

Oblivious Interference Scheduling

Alexander Fanghänel*, Thomas Kesselheim*, Harald Räcke†, Berthold Vöcking*

ABSTRACT

In the *interference scheduling problem*, one is given a set of n communication requests described by pairs of points from a metric space. The points correspond to devices in a wireless network. In the *directed version* of the problem, each pair of points consists of a dedicated sending and a dedicated receiving device. In the *bidirectional version* the devices within a pair shall be able to exchange signals in both directions. In both versions, each pair must be assigned a power level and a color such that the pairs in each color class can communicate simultaneously at the specified power levels. The feasibility of simultaneous communication within a color class is defined in terms of the Signal to Interference Plus Noise Ratio (SINR) that compares the strength of a signal at a receiver to the sum of the strengths of other signals. This is commonly referred to as the “physical model” and is the established way of modelling interference in the engineering community. The objective is to minimize the number of colors as this corresponds to the time needed to schedule all requests.

We study *oblivious power assignments* in which the power value of a pair only depends on the distance between the points of this pair. We prove that oblivious power assignments cannot yield approximation ratios better than $\Omega(n)$ for the directed version of the problem, which is the worst possible performance guarantee as there is a straightforward algorithm that achieves an $O(n)$ -approximation. For the bidirectional version, however, we can show the existence of a universally good oblivious power assignment: For any set of n bidirectional communication requests, the so-called “square root assignment” admits a coloring with at most $\text{polylog}(n)$ times the minimal number of colors. The proof for the exis-

tence of this coloring is non-constructive. We complement it by an approximation algorithm for the coloring problem under the square root assignment. This way, we obtain the first polynomial time algorithm with approximation ratio $\text{polylog}(n)$ for interference scheduling in the physical model.

1. INTRODUCTION

Signals sent by different sources in multipoint radio networks need to be coordinated as simultaneously transmitted signals interfere with each other. If too many signals are sent at the same time in the same region of the network then none of them might get through because the interference is too high. Signals might get scheduled simultaneously, however, if they are exchanged in network regions that are sufficiently far apart so that the interference is relatively low. In order to achieve a good throughput one needs to schedule the signals carefully. The Media Access Control (MAC) layer is responsible for this important task in today’s wireless communication networks. It provides single-hop full-duplex communication channels in multipoint networks to higher layers of the protocol stack. In this paper, we study the task of the MAC layer from an algorithmic point of view. We investigate scheduling algorithms that provide a set of channels between specified pairs of nodes in a wireless network.

Most previous theoretical work (see, e.g., [11, 14, 1]) about scheduling signals or packets in radio networks resort to *graph based vicinity models* of the following flavour. Two nodes in the radio network are connected by an edge in a communication graph if and only if they are in mutual transmission range. Interference is modelled through independence constraints: If a node u transmits a signal to an adjacent node v , then no other node in the vicinity of v , e.g. in the one- or two-hop neighborhood, can transmit. The problem with this modelling approach is that it ignores that neither radio signals nor interference ends abruptly at a boundary.

Some recent theoretical studies [12, 13, 3, 4] use a more realistic model, the so-called *physical model*, which is well-accepted in the engineering community. It is assumed that the strength of a signal diminishes with the distance from its source. More specifically, let $\delta(u, v)$ denote the distance between the nodes u and v . The *loss* between u and v is defined as $\ell(u, v) = \delta(u, v)^\alpha$, where $\alpha \geq 1$ is parameter of the model, the so-called *path-loss exponent*.¹ A signal sent with power p by node u is received by node v at a strength of $p/\ell(u, v)$. Node u can successfully decode this

¹Depending on the environment, it is usually assumed that α has a value between 2 and 5. Our analysis holds for any constant $\alpha \geq 1$.

*Department of Computer Science, RWTH Aachen University, Germany. The authors acknowledge the support of the DFG through the UMIC-excellence cluster at the RWTH.

†Department of Computer Science, University of Warwick, UK. The author acknowledges the support of DIMAP (the Centre for Discrete Mathematics and its Applications).

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Copyright 200X ACM X-XXXXX-XX-X/XX/XX ...\$10.00.

signal if its strength is relatively large in comparison to the strength of other signals received at the same time. This constraint is described in terms of the *Signal to Interference plus Noise Ratio (SINR)* being defined as the ratio between the strength of the signal that shall be received and the sum of the strengths of signals simultaneously sent by other nodes (plus ambient noise). For successfully receiving a signal, it is required that the SINR is at least β with $\beta > 0$ being the second parameter of the model, the so-called *gain*.

In a seminal work [12], Moscibroda and Wattenhofer posted the following problem regarding the physical model: *Assume that we are given a set of directed links between pairs of nodes that indicate communication requests. How much time is required to schedule all these requests?* In this paper, we extend their question towards bidirectional communication requests. In fact, we believe that the bidirectional variant of this problem might be of greater practical relevance as it is the theoretical analog of providing full-duplex communication channels on the MAC layer.

In the *interference scheduling problem* one is given a set of n communication requests each consisting of a pair of points in a metric space. Each pair shall be assigned a power level and a color such that the pairs in each color class can communicate simultaneously at the specified power. The feasibility of simultaneous communication within a color class is described by SINR constraints. We distinguish an unidirectional and a bidirectional version of the problem, depending on whether each pair of nodes consists of a sending and a receiving device or both nodes shall be able to exchange signals in both directions. The exact formulation of the SINR constraints for the undirected and the bidirectional variant can be found in Section 1.1. The objective is to minimize the number of colors, which corresponds to minimizing the time needed to schedule all communication requests. It can be shown via a reduction from 3-Partition that both variants of this problem are strongly NP-hard. We thus seek approximation algorithms.

The interference scheduling problem consists of two correlated subproblems: the *power assignment* and the *coloring*. By far the most literature about MAC layer protocols focuses on scheduling with *uniform power assignment*, in which all pairs send at the same power (see, e.g., [7, 15, 9]). In other studies, a *linear power assignment* is considered, in which the power level for a pair (u, v) is chosen proportional to the loss $\ell(u, v)$. These are examples of *oblivious power assignments* which means the power level assigned to a pair is defined as a function of the loss (or the distance) between the nodes of a pair. The advantage of oblivious power assignments is their simplicity which allows for an immediate implementation in a distributed setting.

1.1 Formal description of the problem

Let the path loss exponent $\alpha \geq 1$ and the gain $\beta > 0$ be fixed. Let V be a set of nodes from a metric space. Let $\delta(u, v)$ denote the distance between two nodes u and v . The loss between u and v is defined as $\ell(u, v) = \delta(u, v)^\alpha$. One is given a set of n requests consisting of pairs $(u_i, v_i) \in V^2$. For every $i \in [n] := \{1, \dots, n\}$, one needs to specify a power level $p_i > 0$ and a color $c_i \in [k] := \{1, \dots, k\}$ such that the number of colors, k , is minimized and the pairs in each color class satisfy the following *SINR constraints* that depend on the specific variant of the problem.

In the directed variant, for every $i \in [n]$, it must hold that

$$\frac{p_i}{\ell(u_i, v_i)} \geq \beta \left(\sum_{\substack{j \in [n] \setminus \{i\} \\ c_j = c_i}} \frac{p_j}{\ell(u_j, v_i)} + \nu \right),$$

where $\nu \geq 0$ expresses ambient noise. In words, for every receiver v_i the strength of the signal received from the corresponding sender u_i needs to be at least as large as β times the sum of the strengths of the signals received from all other senders of the same color plus the noise.

In the bidirectional variant of the problem, for every $i \in [n]$ and $w \in \{u_i, v_i\}$, it must hold that

$$\frac{p_i}{\ell(u_i, v_i)} \geq \beta \left(\sum_{\substack{j \in [n] \setminus \{i\} \\ c_j = c_i}} \frac{p_j}{\min\{\ell(u_j, w), \ell(v_j, w)\}} + \nu \right).$$

In words, for each of the two nodes from a request (u_i, v_i) the strength of the signal received from the communication partner needs to be at least as large as β times the sum of the strengths of the signals sent between other communication partners plus the noise.

Note that this definition implicitly assumes that the communication within a pair follows some unknown protocol ensuring that the signals within a pair do not overlap, as only one end-point of a pair contributes to the interference at another node. Alternatively, one can assume that signals within a pair might overlap. This would increase the interference at other nodes at most by a factor of two. Our results are robust against changes of the interference by constant factors. The relevance of this model in comparison to the directed variant is discussed in Section 6.

The above definition of the problem is derived from modelling interferences via SINR constraints which is the established approach within the engineering community. For our analysis we neglect the ambient noise in the above model, that is, we assume $\nu = 0$ and satisfy the SINR constraints with “>” rather than “ \geq ”. Observe that, under this assumption, any feasible schedule remains feasible when all power levels are multiplied by the same positive factor. Moreover, one can transform a schedule that is feasible under this assumption into a schedule that is feasible for any $\nu > 0$ by multiplying all power levels by a sufficiently large factor. Of course, this might cause problems in practice as it might lead to a high energy requirement. However, this aspect is beyond our analysis.

Finally, let us formally define that a power assignment is called *oblivious* if there is a function $f: \mathbb{R}_{>0} \rightarrow \mathbb{R}_{>0}$ such that, for every $i \in [n]$, $p_i = f(\ell(u_i, v_i))$.

1.2 Our contribution

The question that we study is whether oblivious power assignments are not only easy to implement but also efficient with respect to the number of colors (time steps) that they require in comparison to an optimal schedule. Our answer to this question is different depending on whether one considers the directed or bidirectional version of the problem:

- For any oblivious power assignment p , there exists an instance with n directed communication requests needing $\Omega(n)$ colors when using p , but only one color when using a different power assignment. That is, oblivious power assignments cannot yield approximation ratios better than n for the directed interference scheduling

problem, which corresponds to the worst possible performance guarantee.

- In contrast, there exists a universally good oblivious power assignment for the bidirectional version of the problem: The *square root power assignment* sets the power level for a pair (u, v) equal to $\sqrt{\ell(u, v)}$. We prove that this assignment admits a coloring with at most $\text{polylog}(n)$ times the minimal number of colors, for any set of n bidirectional communication requests.

The negative result for the directed variant is shown by specifying family of request pairs on the line. That is, this result holds already for one-dimensional Euclidean space. In contrast, the positive result about the bidirectional variant holds for request pairs from every metric space.

Let us try to give some intuition about what is the secret behind the square root power assignment. Consider n bidirectional communication requests on the line with $u_i = -2^i$ and $v_i = 2^i$. It is not difficult to see that the uniform assignment allows to execute only $O(1)$ of these nested requests simultaneously as the signals sent between outer pairs are drowned by the signals sent between inner pairs. (The exact number of requests that can be scheduled at the same time depends on the choices for α and β .) Similarly, also the linear assignment allows to schedule only $O(1)$ requests simultaneously since now outer pairs disturb inner pairs. The same is true for any superlinear power assignment, too. The square root assignment, however, allows to schedule a constant fraction of the requests simultaneously as it balances the interference in the right way. Our analysis shows that this kind of balancing effect does not only exist for the line but it is present in any metric space.

The proof for the existence of the coloring in the bidirectional case relies on simulating general metrics by tree metrics and then, as a next step, decomposing tree metrics into star metrics in a hierarchical manner. This existence proof is non-constructive. We make our result constructive by additionally giving an efficient approximation algorithm for the coloring problem under the square root assignment. This way, we obtain the first polynomial time algorithm with approximation ratio $\text{polylog}(n)$ for interference scheduling in the physical model.

1.3 Related Work

The first theoretical studies about interference scheduling in the physical model focus on topologies generated by placing nodes randomly in two-dimensional Euclidean space, see, e.g., [8, 2, 10].

The study of interference scheduling with respect to arbitrary topologies has been initiated by Moscibroda and Wattenhofer [12]. They present the first analysis of the directed interference scheduling problem. However, they cannot handle general request sets but only specific kinds of sets. In particular, they study the question of how many time slots (colors) are needed to schedule a set of communication requests ensuring strong connectivity among n points placed arbitrarily in two-dimensional Euclidean space. On the one hand, they prove that there are configurations requiring $\Omega(n)$ colors when using either uniform or linear power assignments. On the other hand, they show that $O(\log^4 n)$ colors are sufficient to ensure strong connectivity when choosing the right power assignment. This assignment is quite involved and non-oblivious.

The first study for general request sets is presented by Moscibroda *et al.* in [13]. They prove that every set of n directed requests can be scheduled using $O(I_{\text{in}} \cdot \log^2 n)$ colors, where I_{in} is a certain static interference measure depending on the instance. This result enables them to improve the bound for strong connectivity from $O(\log^4 n)$ to $O(\log^3 n)$. However, it does not give any approximation guarantee for general request sets since I_{in} can deviate by a factor that is as large as $\Omega(n)$ from the optimal number of colors.

Chafekar *et al.* [3, 4] study a multi-hop version of the interference scheduling problem on two-dimensional Euclidean instances, that is, they additionally consider the aspect of routing on top of the tasks power assignment and coloring. The considered power assignment is restricted, that is, it is assumed that power levels must be chosen from a specified interval $[p_{\min}, p_{\max}]$. The objective in [3] is to minimize the end-to-end latency, while [4] aims at maximizing throughput. When breaking down the approach in [3] to the directed version of the single-hop interference problem, it yields a schedule using $O(\text{opt}' \cdot \text{polylog}(n, \Delta, \Gamma))$ colors where opt' denotes the minimal number of colors needed for a schedule with slightly smaller power range $[p_{\min}, (1 - \epsilon)p_{\max}]$, Γ denotes the ratio between p_{\max} and p_{\min} , and Δ denotes the aspect ratio, i.e., the ratio between maximum and minimum distance over all pairs of nodes. In a recent work, Fanghänel *et al.* [5] improve on this result and achieve an approximation factor of order $\mathcal{O}(\log n \log \Delta)$. Let us remark that the dependence on the aspect ratio cannot be avoided by both of these approaches as the presented algorithms employ the linear power assignment which, without taking into account other parameters than n , cannot achieve an approximation ratio better than $\Omega(n)$.

2. THE DIRECTED CASE

We already stated in the unidirectional case any oblivious power assignment can have bad performance when compared to an optimal scheme. To prove this we construct a family of instances for a given function f such that using f requires at least $\Omega(n)$ colors or schedule steps while an optimum power assignment needs only $O(1)$ rounds.

THEOREM 1. *Let $f : \mathbb{R}_{>0} \rightarrow \mathbb{R}_{>0}$ be any oblivious power assignment function. For the unidirectional model there exists a family of instances on a line that requires $\Omega(n)$ colors when scheduling with the powers defined by f whereas an optimal schedule has constant length.*

PROOF (SKETCH). We only consider the case that f is asymptotically unbounded, i.e., for every $c > 0$ and every $x_0 > 0$ there exists a value $x > x_0$ with $f(x) > c$.

Consider the following family of instances consisting of n pairs (u_i, v_i) , with distances x_i between two nodes of a pair and χy_i between neighboring pairs. Depending on β , we choose χ as a suitable constant.

Formally, the instance is defined by $u_1, v_1, \dots, u_n, v_n \in \mathbb{R}$ such that

$$u_i = \begin{cases} 0 & \text{if } i = 1 \\ v_{i-1} + \chi y_i & \text{otherwise} \end{cases} \quad \text{and} \quad v_i = u_i + x_i.$$

We now define the distances x_i and y_i between the nodes recursively depending on the function f : $y_i := 2(x_{i-1} + y_{i-1})$. Given x_1, \dots, x_{i-1} and y_i , we choose x_i such that $x_i \geq y_i$ and $f(x_i) \geq (y_i^\alpha f(x_j))/x_j^\alpha$ for all $j < i$. This is always possible

since f is asymptotically unbounded. By this construction it is ensured that a pair k is exposed to high interference by pairs with larger indices. To show this, let $S \subseteq [n]$ be a set of indices of pairs that can be scheduled together in one step; $k = \min S$. For $i \in S \setminus \{k\}$ it holds that

$$\begin{aligned} \delta(u_i, v_k) &= \sum_{j=k+1}^{i-1} x_j + \sum_{j=k+1}^i \chi \cdot y_j \leq 2\chi \sum_{j=k}^i y_j \\ &\leq 2\chi \sum_{j=k}^i \frac{1}{2^{i-j}} y_i \leq 4\chi y_i. \end{aligned}$$

Since all pairs in S can be scheduled in one step the SINR condition is satisfied for pair k :

$$\beta \sum_{i \in S \setminus \{k\}} \frac{p_i}{\ell(u_i, v_k)} \leq \frac{p_k}{\ell(u_k, v_k)} = \frac{f(x_k)}{x_k^\alpha}.$$

Putting these facts together gives

$$\frac{f(x_k)}{\beta x_k^\alpha} \geq \sum_{i \in S \setminus \{k\}} \frac{p_i}{\ell(u_i, v_k)} \geq \sum_{i \in S \setminus \{k\}} \frac{y_i^\alpha \frac{f(x_k)}{x_k^\alpha}}{(4\chi y_i)^\alpha} = \frac{|S| - 1}{(4\chi)^\alpha} \frac{f(x_k)}{x_k^\alpha}.$$

This implies $|S| \leq (4\chi)^\alpha / \beta + 1$, which means there are at least $\frac{\beta}{(4\chi)^\alpha + \beta} n = \Omega(n)$ colors needed when using $p_i = f(\ell(s_i, d_i))$.

On the other hand for these instances there is a power assignment, $p_i = \sqrt{2^i}$, such that there is a coloring using a constant number of colors. This is caused by the fact that for all instances described it holds that $y_i \leq x_i$ and $y_{i+1} \geq 2x_i$. Thus for any link k the interference by the ones with higher index as well as the ones with lower index form a geometric series. This means a constant fraction of all links may have the same color and therefore there is a coloring using a constant number of colors. \square

For bounded, linear and superlinear functions f this proof can be adapted to the bidirectional model. For sublinear functions, however, such an adaptation is not possible. In fact, we will show in the next section that there exists a sublinear function, namely the square-root function, which allows to minimize the number of colors up to a polylogarithmic factor for bidirectional communication.

3. THE SQUARE ROOT POWER ASSIGNMENT HAS A GOOD COLORING

In this section we consider the bidirectional version of the interference scheduling problem. Recall that the square root power assignment sets the power level for a pair (u, v) equal to $\sqrt{\ell(u, v)}$ where $\ell(u, v) = \delta(u, v)^\alpha$ denotes the loss between u and v . In the following, this power assignment is denoted by \bar{p} . We prove the following main theorem.

THEOREM 2. *Let $(u_1, v_1), \dots, (u_n, v_n)$ be a set of request pairs from a metric space for which there is a power assignment p_1, \dots, p_n satisfying the bidirectional SINR constraints with only one color. Then, for \bar{p} , there exists a coloring with $\mathcal{O}(\log^{3.5+\alpha} n)$ colors satisfying the bidirectional SINR constraints.*

Before starting the actual proof we present two main techniques. Section 3.1 shows, that if we are given a valid coloring for gain β , scaling the gain by a constant factor changes the number of colors needed only by a logarithmic factor. In

Section 3.2 we simplify our problem by splitting the communication pairs into single nodes. A so called loss parameter is used to keep track of the loss between the communication partners.

The proof of the main theorem consists of three parts: In Section 3.3 we show how to break the problem for general metrics down to tree metrics. Section 3.4 then reduces the problem from tree metrics to star metrics. In Section 4 we present our analysis for stars, finishing the proof of the theorem.

3.1 Scaling the gain

Consider an instance of the interference scheduling problem in the directed or bidirectional variant with n requests. Suppose both the coloring c and the power assignment p are fixed such that the SINR constraints are satisfied with gain β . The following two propositions show the existence of a coloring c' that for the same power assignment p satisfies the SINR constraints with a more restrictive gain $\beta' > \beta$ and uses only $\mathcal{O}(\beta'/\beta \log n)$ times the number of colors in c . The proofs have been omitted due to space limitations.

PROPOSITION 3. *Let S denote a set of requests with power assignment p satisfying the SINR constraints with gain β . Then there exists a subset S' of S with $|S'| \geq \beta/\beta' |S|$ satisfying the SINR constraints with gain $\beta' > \beta$ for the same power assignment.*

PROPOSITION 4. *Let S denote a set of requests with power assignment p satisfying the SINR constraints with gain β . Then there exists a coloring c' for S with $\mathcal{O}(\beta/\beta' \log |S|)$ colors such that c' together with p satisfy the SINR constraints with gain $\beta' > \beta$.*

3.2 Splitting pairs

For our analysis of the interference scheduling problem we use a slightly modified variant, the *node-loss scheduling problem*. One is given a set of nodes $[n]$ and each node i is associated with a *loss parameter* ℓ_i . For every $i \in [n] := \{1, \dots, n\}$, one needs to specify a power level $p_i > 0$ and a color $c_i \in [k] := \{1, \dots, k\}$ such that the number of colors, k , is minimized and the pairs in each color class satisfy the following *SINR constraints*.

$$\frac{p_i}{\ell_i} \geq \beta \left(\sum_{\substack{j \in [n] \setminus \{i\} \\ c_j = c_i}} \frac{p_j}{\ell(i, j)} + \nu \right)$$

In words, for each node the ratio between the power p_i and the loss ℓ_i needs to be at least β times larger than the sum of the strengths of the signals sent by other nodes plus the noise. Again we neglect the ambient noise in the model, i.e., $\nu = 0$ and fulfil the constraints with “>” rather than “ \geq ”.

A *power assignment* specifies a power level for each node. The *square root power assignment* \bar{p} sets the power level for node $i \in [n]$ equal to $\sqrt{\ell_i}$. For a power assignment $p = p_1, \dots, p_n$ and a set of nodes $U \subseteq [n]$, let

$$I_p(i | U) = \sum_{j \in U \setminus \{i\}} \frac{p_j}{\ell(i, j)}$$

denote the *interference* at node $i \in [n]$ induced by elements of U . We say that U is β -feasible for a power assignment p if $\frac{p_i}{\ell_i} > \beta I_p(i | U)$, for every $i \in U$.

On any given instance, feasible schedule steps for the interference scheduling and the node-loss scheduling problem are related as follows: First, if we have a feasible schedule step S for the node-loss scheduling that schedules a fraction greater than one half of the nodes, we can give a feasible schedule step for a constant fraction of the nodes in the interference scheduling setting by scheduling the pairs with both nodes in S .

Secondly, if we have a set of pairs U that we can schedule in the interference scheduling setting with gain β , the set of all nodes from pairs in U is $\beta/2+\beta$ -feasible for the node-loss scenario, as we show in the following. For a node i let $I'(i)$ denote the interference at this node in the interference scheduling problem, and $I(i)$ denote the interference at this node in the node-loss scheduling problem. If now all nodes from pairs in U transmit, the interference at a single node i is at most twice the interference from the interference scheduling problem plus the interference from the other node of this pair, i.e., p_i/ℓ_i , so

$$I(i) \leq 2I'(i) + \frac{p_i}{\ell_i} \leq \frac{2+\beta}{\beta} \cdot \frac{p_i}{\ell_i},$$

as $I'(i) \leq p_i/\beta\ell_i$. As the results from Section 3.1 can be proven analogously for the node-loss scheduling problem, we can compute a schedule for the node-loss scheduling problem from a schedule for the interference scheduling problem, that is longer by at most a logarithmic factor.

In Section 4 we prove the following result for the node-loss scheduling problem.

LEMMA 5. *Let $\beta' \geq \beta > 0$. Suppose $S([n], \delta, \ell)$ is a star for which there exists a power assignment p such that $[n]$ is β' -feasible under p . Then there is a subset $U \subseteq [n]$ with $|U| \geq (1 - \mathcal{O}((\frac{\beta}{\beta'})^{2/3}))n$ that is β -feasible under the square root assignment \bar{p} .*

The star $S([n], \delta, \ell)$ is defined by a set $[n]$ of nodes placed around a center c , the distances of the nodes δ and their loss parameters ℓ (see Section 4 for details). Using this lemma we now turn to the proof for Theorem 2.

3.3 From general metrics to trees

For this part we utilize the following lemma, which is suitably adapted from a lemma in [6].

LEMMA 6. *Given a finite metric space $([n], \delta)$ there exist $r = \mathcal{O}(\log n)$ edge-weighted trees T_1, \dots, T_r with node-set $[n]$ such that the following holds*

1. *For every pair $(u, v) \in [n]^2$ and for every tree T_i : $\delta(u, v) \leq \delta_{T_i}(u, v)$ where δ_{T_i} denotes the shortest path metric induced by tree T_i .*
2. *For every node $v \in [n]$ there exists a subset $\mathcal{T}_v \subseteq \{T_1, \dots, T_r\}$ with $|\mathcal{T}_v| \geq \frac{9}{10}r$ such that the pairwise distances involving v are only stretched by a logarithmic factor, i.e., $\forall T \in \mathcal{T}_v : \forall u \in [n] : \delta_T(u, v) \leq \mathcal{O}(\log n) \cdot \delta(u, v)$.*

For a tree T_i in the above lemma we call the set of nodes whose distances are at most stretched by the logarithmic factor the *core* of T_i , and denote it with C_i . Suppose that we are given an instance of the node-loss scheduling problem in a metric space $([n], \delta)$. With every tree T_i from the decomposition of Lemma 6 we associate a corresponding

node-loss scheduling instance that only includes nodes in the core of T_i (the loss parameters stay the same).

PROPOSITION 7. *Suppose there exists a β' -feasible set $U \subseteq [n]$ for the node-loss scheduling problem on $([n], \delta)$. Then there exists a tree T_i with a β' -feasible set of size at least $\frac{9}{10} \cdot |U|$ in its core C_i .*

LEMMA 8. *Suppose there is a β' -feasible subset U of core nodes for the node-loss scheduling instance in a tree T_i (for some power assignment p). Then, this set U is β'' -feasible with respect to the original metric for $\beta'' = \Omega(\frac{\beta'}{\log^\alpha n})$.*

3.4 From trees to stars

In this section we extend Lemma 5 to tree metrics.

LEMMA 9. *Suppose we are given an instance $T([n], \delta, \ell)$ of the node-loss scheduling problem on a tree metric for which there exists a power assignment p such that a subset $U \subseteq [n]$ is β' -feasible under p . Then, there exists a subset $U' \subseteq U$ with $|U'| \geq \frac{9}{10}|U|$ that is β -feasible under \bar{p} for $\beta = \Omega(\frac{\beta'}{\log^{2.5} n})$.*

PROOF. In order to show the result we repeatedly make use of Lemma 5, and remove nodes from the set U that cannot be scheduled by the square root power assignment in one round. In the end we show that we did not remove too many nodes from U . For the first round we choose a node c in the tree such that the removal of c partitions the tree into disjoint sub-trees with size at most $n/2$. Such a node can be found in any tree. Now we consider the node-loss scheduling problem on the star metric obtained by selecting c as center and setting the distance δ_v of a node v to the center as the tree-distance $\delta(v, c)$. Note that distances in this star-metric are not smaller than distances in the original tree and that therefore the set U is β' -feasible in this metric.

When applying Lemma 5 with a suitable parameter $\beta'' = \beta/\mathcal{O}(\log^{3/2} n)$ we obtain a subset $U'_1 \subset U$, $|U'_1| \geq (1 - \frac{1}{10 \log n})|U|$ that is β'' -feasible for the square-root power assignment \bar{p} . Here the constant 10 comes from suitably balancing the hidden constant in the \mathcal{O} -notation of Lemma 5 and the hidden constant in the \mathcal{O} -notation of β'' . Of course, this subset may not be feasible for the square root power assignment in the original tree metric $([n], \delta)$, because some nodes of U'_1 are closer in $([n], \delta)$ and hence induce more interference between each other. In order to compensate for this we re-run the algorithm on the forest obtained after splitting the graph at c , i.e., we delete all but one edge incident to c . In each of the trees of this forest we run the above algorithm recursively. For each level i of the recursion, the algorithm returns a set U'_i , $|U'_i| \geq (1 - \frac{1}{10 \log n})|U|$ that is β'' -feasible in the corresponding forest. There are at most $\log n$ recursion levels as the size of a tree reduces by at least a factor of 2 in each iteration. Let $U' := \bigcap_i U'_i$. Then we have $|U'| \geq 9^{\log n}/10$.

Note that a pair $(u, v) \in U' \times U'$ has the correct distance in at least one of the recursions. Therefore, the total interference induced at a node $u \in U'$ (from all the other nodes of U') when using the square-root assignment in the tree metric $([n], \delta)$ is at most the sum of the interferences generated at u in all iterations which is at most $\log n \cdot \frac{1}{\beta'' \sqrt{\ell_u}}$, since u is β'' -feasible in each iteration. This means that the set U' is $\beta = \frac{\beta''}{\log n} = \Omega(\frac{\beta'}{\log^{2.5} n})$ -feasible. \square

3.5 Putting the pieces together

In this section we prove Theorem 2. We are given a set S of request pairs from a metric space $([n], \delta)$ for which there is a power assignment that satisfies the bidirectional SINR constraints with only one color. Let U denote the set of terminal nodes of pairs from S . Following the discussion in Section 3.2 this set is β' -feasible for the node-loss scheduling problem with $\beta' \geq \beta/(2 + \beta)$ (on the same metric $([n], \delta)$).

Now, we apply Proposition 7 to this set U , and obtain a subset $U' \subset U$, $|U'| \geq \frac{9}{10}|U|$ that is β' -feasible and is contained in the core C_i of a tree T_i .

Next, we apply Lemma 9 to this set and obtain a subset U'' , $|U''| \geq \frac{9}{10}|U'|$ that is β'' -feasible for the square-root assignment \bar{p} , where $\beta'' = \Omega(\beta'/\log^{2.5} n)$.

Now, we can use Lemma 8 to conclude that this set is also β''' -feasible for \bar{p} in the original metric, where $\beta''' = \Omega(\beta''/\log^\alpha n)$.

Note that the subset U'' contains at least $\frac{9}{10} \cdot \frac{9}{10}|U| > \frac{8}{10}|U|$ nodes. This means that for at least a $\frac{6}{10}$ -fraction of pairs from the original set S , both end-points are contained in U'' . Let $S' \subset S$ denote a set that contains only these pairs. The pairs in S' fulfill the bidirectional SINR constraints with gain β''' for the power assignment \bar{p} .

Rescaling the gain with Proposition 3, we obtain a subset S'' with $|S''| \geq \beta'''/8\beta$ that fulfills the SINR constraints with gain β . Observe that the size of S'' is $\Omega(1/\beta \log^{2.5+\alpha} n)|S|$. Coloring the requests from S'' with a single color and repeating the process for the remaining request gives that we only need $O(\log^{3.5+\alpha} n)$ colors to color all requests. This completes the proof of the theorem.

4. ANALYSIS FOR STAR METRICS

In this section, we prove Lemma 5. Let $\beta' \geq \beta > 0$. We are given a set $\{(1, \ell_1), \dots, (n, \ell_n)\}$ of node-loss pairs (requests) being β' -feasible under some power assignment p . The nodes $1, \dots, n$ form a star centered around an additional node c . The distance between c and i is denoted by δ_i . Let $d_i = \delta_i^\alpha$, that is, d_i corresponds to the loss between c and i . In the following, this parameter is called *decay* in order to distinguish it from the loss parameter ℓ_i . W.l.o.g., we assume $d_1 \leq d_2 \leq \dots \leq d_n$. Let $a_i = \ell_i/d_i$. We have to show that there exists a subset $U \subseteq [n]$ with $|U| \geq (1 - \mathcal{O}((\frac{\beta}{\beta'})^{2/3}))n$ being β -feasible under the square root power assignment \bar{p} .

We will first prove some helpful properties. These properties will show that the lemma follows relatively easy for the special case in which the loss parameter is relatively large in comparison to the decay, i.e., $a_i > 2^{\alpha+1}/\beta'$, for every $i \in [n]$. We then turn our attention to the case in which the loss parameter is relatively small, i.e., $a_i \leq 2^{\alpha+1}/\beta'$, for every $i \in [n]$. Finally, we will combine the results for these special cases in order to prove the lemma for stars with both small and large loss parameters.

4.1 Helpful properties

Consider two nodes i and i' with $i' < i$. As there exists a power scheme p with a β' -feasible schedule, it holds that

$$\frac{p_{i'}}{\ell_{i'}} > \beta' \frac{p_i}{(\delta_i + \delta_{i'})^\alpha} \quad \text{and} \quad \frac{p_i}{\ell_i} > \beta' \frac{p_{i'}}{(\delta_i + \delta_{i'})^\alpha}.$$

Multiplying these equations we obtain $(\delta_i + \delta_{i'})^{2\alpha} > \beta'^2 \cdot \ell_i \cdot \ell_{i'}$. As $i' < i$ we have $(2\delta_i)^{2\alpha} \geq (\delta_i + \delta_{i'})^{2\alpha}$ and thus

$$d_i^2 = \delta_i^{2\alpha} > \frac{\beta'^2}{4^\alpha} \cdot \ell_i \cdot \ell_{i'}. \quad (1)$$

It follows

$$d_i \geq a_i \cdot \frac{\beta'^2}{4^\alpha} \cdot \ell_{i'}, \quad (2)$$

$$d_i \geq a_i \cdot a_{i'} \cdot \frac{\beta'^2}{4^\alpha} \cdot d_{i'}, \quad (3)$$

$$\ell_i \geq a_i^2 \cdot \frac{\beta'^2}{4^\alpha} \cdot \ell_{i'}. \quad (4)$$

4.2 Stars with large loss parameters

In this section, we assume $a_i > 2^{\alpha+1}/\beta'$, for every $i \in [n]$. We apply Equation 4 with $\ell_{i'} = \ell_{i-1}$ and repeat this for $i - j$ times deriving the following lower bound relating ℓ_i to ℓ_j , for $i > j$,

$$\begin{aligned} \ell_i/\ell_j &\geq a_i^2 \cdot \dots \cdot a_{j+1}^2 \left(\frac{\beta'}{2^\alpha}\right)^{2(i-j)} \\ &> a_i^2 \left(\frac{2^{\alpha+1}}{\beta'}\right)^{2(i-j-1)} \left(\frac{\beta'}{2^\alpha}\right)^{2(i-j)} = a_i^2 \left(\frac{\beta'}{2^{\alpha+1}}\right)^2 2^{2(i-j)}. \end{aligned} \quad (5)$$

Now we solve the equation above for ℓ_j and exchange the indices i and j . This way, for $i < j$,

$$\ell_i < a_j^{-2} \cdot \left(\frac{2^{\alpha+1}}{\beta'}\right)^2 \cdot \ell_j \cdot 2^{2(i-j)}. \quad (6)$$

These inequalities enable us to prove the following result for stars with large loss parameters.

LEMMA 10. *Suppose $a_i > 2^{\alpha+1}/\beta'$, for every $i \in [n]$. If there exists a power scheme p such that $[n]$ is β' -feasible under p then $[n]$ is β -feasible under the square root power assignment \bar{p} with $\beta \leq \beta'/2^{\alpha+2}$.*

PROOF. At node j the received interference is

$$I_{\bar{p}}(j) \leq \sum_{i=1}^{j-1} \frac{\sqrt{\ell_i}}{d_j} + \sum_{i=j+1}^n \frac{\sqrt{\ell_i}}{d_i} = \sum_{i=1}^{j-1} \frac{\sqrt{\ell_i}}{d_j} + \sum_{i=j+1}^n \frac{a_i}{\sqrt{\ell_i}}.$$

Now applying Equation 6 and Equation 5 gives

$$\begin{aligned} I_{\bar{p}}(j) &< \frac{\sqrt{\ell_j}}{a_j d_j} \cdot \frac{2^{\alpha+1}}{\beta'} \cdot \sum_{i=1}^{j-1} 2^{i-j} + \frac{1}{\sqrt{\ell_j}} \cdot \frac{2^{\alpha+1}}{\beta'} \cdot \sum_{i=j+1}^n 2^{j-i} \\ &< \frac{2}{\sqrt{\ell_j}} \cdot \frac{2^{\alpha+1}}{\beta'}. \end{aligned}$$

The SINR constraint at j is satisfied if $I_{\bar{p}}(j) < 1/\beta\sqrt{\ell_i}$. For $\beta \leq \beta'/2^{\alpha+2}$ this condition is satisfied. \square

4.3 Stars with small loss parameters

Now we assume that all loss parameters are relatively small compared to the decay. In this case, given a β' -feasible power assignment p , we can ensure that the square root power assignment is β -feasible for any $\beta < \beta'$ if a small fraction of the nodes that depends on the ratio between β and β' can be dropped.

LEMMA 11. *Suppose $a_i \leq 2^{\alpha+1}/\beta'$, for every $i \in [n]$. If there exists a power scheme p such that $[n]$ is β' -feasible under p then there exists a subset $U \subseteq [n]$ that is β -feasible under \bar{p} with $|U| = (1 - \mathcal{O}((\frac{\beta}{\beta'})^{2/3}))n$.*

PROOF. We partition the nodes into classes depending on their distance/decay to the center c . W.l.o.g., assume $d_u > 1$,

for every $u \in [n]$. Let $D_j = \{u \mid 2^{j-1} < d_u \leq 2^j\}$, $|D_j| = k_j$ and let m denote the largest index for which D_m is not empty.

CLAIM 12. Let $0 < \mu < 1$. For a $(1 - \mu)$ -fraction of the nodes in class D_j , the loss parameter ℓ_u fulfils $\ell_u \leq \frac{2^{\alpha+j+2}}{\mu\beta'k_j}$.

PROOF. In the given power assignment p , a node v from class D_j induces an interference on node $u \in D_j$ of

$$\frac{p_v}{(\delta_u + \delta_v)^\alpha} \geq \frac{p_v}{(2 \cdot 2^{j/\alpha})^\alpha} = \frac{p_v}{2^{\alpha+j}}.$$

The interference at node u is upper-bounded by $p_u/\beta'\ell_u$ because p satisfies the SINR constraint. Thus, it follows

$$\sum_{v \in D_j \setminus \{u\}} \frac{p_v}{2^{\alpha+j}} \leq \sum_{v \in [n] \setminus \{u\}} \frac{p_v}{\ell(u, v)} \leq \frac{p_u}{\beta'\ell_u}.$$

For nodes u that fulfil $p_u \leq \sum_{v \in D_j \setminus \{u\}} p_v$, we thus get

$$\ell_u \leq \frac{2^{\alpha+j}}{\beta'} \cdot \frac{p_u}{\sum_{v \in D_j \setminus \{u\}} p_v} \leq \frac{2^{\alpha+j+1}}{\beta'} \cdot \frac{p_u}{\sum_{v \in D_j} p_v}.$$

For other nodes,

$$\ell_u \leq a_u \cdot d_u \leq \frac{2^{\alpha+j+1}}{\beta'} \leq \frac{2^{\alpha+j+2}}{\beta'} \cdot \frac{p_u}{\sum_{v \in D_j} p_v}$$

since $p_u > \sum_{v \in D_j \setminus \{u\}} p_v$ implies $2p_u > \sum_{v \in D_j} p_v$. Summing the above inequality over all nodes in the class D_j gives $\sum_{u \in D_j} \ell_u \leq 2^{\alpha+j+2}/\beta'$. This means that, on average, a node has a loss parameter of only $2^{\alpha+j+2}/(\beta'k_j)$. Using the Markov inequality, we get that a fraction of at most μ of the nodes have a loss parameter larger than $2^{\alpha+j+2}/(\mu\beta'k_j)$. \square

Claim 12 is based on properties of p . In the rest of the proof of Lemma 11, we will not consider other properties of p than the one given by the claim. For the time being, let us ignore a μ -fraction of the nodes such that all remaining nodes fulfil the bound in the claim. The μ -fraction dropped will be taken into account at the end of the proof.

When using the square root power assignment, the interference induced at a node $u \in D_j$ by a node $v \in D_i$, $i \leq j$ is at most

$$\frac{\sqrt{\ell_v}}{2^{j-1}} \leq \frac{1}{2^{j-1}} \sqrt{\frac{2^{\alpha+i+2}}{\mu\beta'k_i}} = \frac{1}{2^j} \sqrt{\frac{2^{\alpha+i+4}}{\mu\beta'k_i}}.$$

Summing this over all nodes in the class and then over all classes gives the following bound on the interference generated at u by nodes from classes with lower or equal index:

$$I_{\bar{p}}(u \mid D_1 \cup \dots \cup D_j) \leq \sqrt{\frac{2^{\alpha+4}}{\mu\beta'}} \sum_{i=1}^j \frac{\sqrt{k_i 2^i}}{2^j}.$$

The interference generated by nodes from higher classes can be estimated as

$$I_{\bar{p}}(u \mid D_{j+1} \cup \dots \cup D_m) \leq \sqrt{\frac{2^{\alpha+4}}{\mu\beta'}} \sum_{i=j+1}^m \frac{\sqrt{k_i 2^i}}{2^i}.$$

We now select all nodes for which, both, the interference from classes with lower index and the interference from classes with higher index, is no more than $1/2\beta$ times the strength of the received signal.

We first count the number of nodes that are not selected this way because the interference from classes with lower or equal index is too high, that is, the number of nodes $u \in D_j$ satisfying

$$I_{\bar{p}}(u \mid D_1 \cup \dots \cup D_j) \geq \frac{1}{2\beta} \frac{1}{\sqrt{\ell_u}} \geq \frac{1}{2\beta} \sqrt{\frac{\mu\beta'k_j}{2^{\alpha+j+2}}}$$

as the received signal strength at a node u in class D_j is $\frac{\sqrt{\ell_u}}{\ell_u} \geq \sqrt{\frac{\mu\beta'k_j}{2^{\alpha+j+2}}}$. Together with the above bound on the interference we obtain

$$\begin{aligned} k_j &\leq \left(\frac{2^{\alpha+4}\beta}{\mu\beta'} \right)^2 \left(\sum_{i=1}^j \sqrt{\frac{k_i}{2^{j-i}}} \right)^2 \\ &\leq \left(\frac{2^{\alpha+4}\beta}{\mu\beta'} \right)^2 \left(\sum_{i=1}^j \frac{k_i}{\sqrt{2^{j-i}}} \right) \cdot \left(\sum_{i=1}^j \frac{1}{\sqrt{2^{j-i}}} \right) \\ &\leq \left(\frac{2^{\alpha+6}\beta}{\mu\beta'} \right)^2 \sum_{i=1}^j \frac{k_i}{\sqrt{2^{j-i}}}. \end{aligned}$$

Here the third inequality uses Cauchy-Schwarz ($(\sum a_i b_i)^2 \leq \sum a_i^2 \cdot \sum b_i^2$) with $a_i^2 = k_i/\sqrt{2^{j-i}}$ and $b_i^2 = 1/\sqrt{2^{j-i}}$. Now the number of nodes lost because of too much interference from classes with lower or equal index can be estimated by

$$\begin{aligned} \sum_{\substack{j: \text{class } D_j \\ \text{not scheduled}}} k_j &\leq \sum_{j=1}^m \left(\frac{2^{\alpha+6}\beta}{\mu\beta'} \right)^2 \sum_{i=1}^j \frac{k_i}{\sqrt{2^{j-i}}} \\ &= \left(\frac{2^{\alpha+6}\beta}{\mu\beta'} \right)^2 \sum_{i=1}^m \sum_{j=i}^m \frac{k_i}{\sqrt{2^{j-i}}} \leq \left(\frac{2^{\alpha+8}\beta}{\mu\beta'} \right)^2 \sum_{i=1}^m k_i. \end{aligned}$$

Analogously the number of nodes lost because of too much interference from classes with higher index is at most

$$\begin{aligned} \sum_{\substack{j: \text{class } D_j \\ \text{not scheduled}}} k_j &\leq \sum_{j=1}^{m-1} \left(\frac{2^{\alpha+6}\beta}{\mu\beta'} \right)^2 \sum_{i=j+1}^m \frac{k_i}{\sqrt{2^{i-j}}} \\ &\leq \left(\frac{2^{\alpha+6}\beta}{\mu\beta'} \right)^2 \sum_{i=2}^m \sum_{j=1}^{i-1} \frac{k_i}{\sqrt{2^{i-j}}} \leq \left(\frac{2^{\alpha+8}\beta}{\mu\beta'} \right)^2 \sum_{i=2}^m k_i. \end{aligned}$$

So in total we only lose $\mathcal{O}((\beta/\mu\beta')^2 + \mu)n$ nodes. Choosing $\mu = (\beta/\beta')^{2/3}$ gives the bound in Lemma 11.

4.4 Stars with arbitrary loss parameters

In the following, we use the results for the special cases given in Lemma 10 and Lemma 11 to prove Lemma 5 for stars without restrictions on the ratio a_i between ℓ_i and d_i .

W.l.o.g., assume that $\beta' \geq 2c_0\beta$ and choose $\beta'' = 2c_1\beta$, for suitable large positive constant terms c_0 and c_1 as specified at the end of the proof. We will show that there is a way to remove a subset of $\Theta((\beta''/\beta')^{2/3})n = \Theta((\beta/\beta')^{2/3})n$ many nodes such that the interference at any remaining node i is at most $(c_0/\beta' + c_1/\beta'') \cdot 1/\sqrt{\ell_i} \leq 1/\beta\sqrt{\ell_i}$, that is, the set of the remaining nodes is β -feasible.

Suppose we hypothetically reduce the loss ℓ_i , for every $i \in [n]$ with $\ell_i > d_i \cdot 2^{\alpha+1}/\beta'$, to $d_i \cdot 2^{\alpha+1}/\beta'$. Under this hypothesis, all nodes have small loss parameters so that Lemma 11 shows the existence of a subset $U \subseteq [n]$ that is β'' -feasible (wrt the hypothetical loss parameters) under \bar{p} with $|U| = (1 - \mathcal{O}((\beta'/\beta'')^{2/3}))n$. In the following, we will study the interference caused by the square root power assignment

applied to the nodes in U with respect to the original loss parameters.

Define the set $L \subseteq U$ of *large loss nodes* by $L := \{i \in [n] \mid a_i > 2^{\alpha+1}/\beta'\}$. For a node $i \in U$, we use $\text{pred}(i) := \max\{j \in L \mid j < i\}$ and $\text{succ}(i) := \min\{j \in L \mid j > i\}$ to denote the predecessor and successor, respectively, of i in L . The nodes in L partition the remaining nodes into subsets as follows. For $i \in L$ we define the set $S_i := \{j \in U \mid \text{pred}(i) < j < i\}$.

The interference that is induced by *large-loss nodes* (nodes in L) onto other large-loss nodes can be handled by applying Lemma 10. Similarly, the interference that is induced by *low-loss nodes* (nodes not in L) onto other low-loss nodes can be handled by Lemma 11. In the following two lemmas we will derive bounds for the interference that is induced by small-loss nodes onto large-loss nodes and vice versa.

LEMMA 13. *For every node $i \in L$, $\sum_{j \in L \setminus \{i, \text{succ}(i)\}} I_{\bar{p}}(i \mid S_j) < 2^\alpha / \beta'' \sqrt{\ell_i}$.*

In words, for a node $i \in L$, the interference generated at i by the small-loss nodes in the sets S_j (with exception of S_i and $S_{\text{succ}(i)}$), is less than $1/\sqrt{\ell_i}$, the strength of the signal received at node i , times $2^\alpha/\beta''$.

PROOF. To show the lemma, we split the interference at i from classes S_j into two parts, the interference from classes $S_j, j \in L, j < i$, and interference from classes $S_j, j \in L, j > \text{succ}(i)$. We will show that each of these terms is upper-bounded by $2^\alpha/2\beta''\sqrt{\ell_i}$, which proves the lemma.

The interference at node i due to the classes $S_j, j \in L, j < i$ can be bounded as follows. We first prove an upper bound on the interference at node $\text{pred}(i)$ and then we show that this bound translates to the desired upper bound for the interference at node i . Let $L^{<i} := \{j \in L \mid j < i\}$. The interference at node $\text{pred}(i)$ from sets $S_j, j \in L^{<i}$ is at least

$$\begin{aligned} \sum_{j \in L^{<i}} \sum_{k \in S_j} \frac{\sqrt{\ell_k}}{(\delta_k + \delta_{\text{pred}(i)})^\alpha} &\geq \sum_{j \in L^{<i}} \sum_{k \in S_j} \frac{\sqrt{\ell_k}}{(2 \cdot \delta_{\text{pred}(i)})^\alpha} \\ &= \sum_{j \in L^{<i}} \sum_{k \in S_j} \frac{\sqrt{\ell_k}}{2^\alpha d_{\text{pred}(i)}}. \end{aligned}$$

For every $j \in U$, let $\ell'_j = d_j \cdot a'_j$ with $a'_j = \min\{a_j, 2^{\alpha+1}/\beta'\}$, that is, we decrease the large loss parameters so that all loss parameters $\ell'_j, j \in U$, are small. On the one hand, due to the construction of the set U , the interference at node $\text{pred}(i)$ caused by the nodes from U wrt to the loss parameters ℓ'_j is upper-bounded by $1/\beta''\sqrt{\ell'_{\text{pred}(i)}}$ because the nodes in U are β'' -feasible wrt modified loss parameters. On the other hand, the lower bound on the interference at $\text{pred}(i)$ is valid also for the modified loss parameters as it only sums over the strengths of signals received from nodes with small loss parameters. Consequently,

$$\sum_{j \in L^{<i}} \sum_{k \in S_j} \frac{\sqrt{\ell_k}}{2^\alpha d_{\text{pred}(i)}} \leq \frac{1}{\beta''\sqrt{\ell'_{\text{pred}(i)}}} = \frac{1}{\beta''\sqrt{a'_{\text{pred}(i)}d_{\text{pred}(i)}}}.$$

By multiplying with $d_{\text{pred}(i)}2^\alpha/d_i$, we get

$$\begin{aligned} \sum_{j \in L^{<i}} \sum_{k \in S_j} \frac{\sqrt{\ell_k}}{d_i} &\leq \frac{2^\alpha \sqrt{d_{\text{pred}(i)}}}{\beta'' d_i \sqrt{a'_{\text{pred}(i)}}} \\ &< \frac{4^\alpha}{\beta' \beta'' a'_{\text{pred}(i)} \sqrt{a_i d_i}} = \frac{2^\alpha}{2\beta'' \sqrt{\ell_i}}, \end{aligned}$$

where we used $d_{\text{pred}(i)} \leq \frac{4^\alpha}{\beta'^2} \frac{d_i}{a_i a'_{\text{pred}(i)}} < \frac{4^\alpha}{\beta'^2} \frac{d_i}{a_i a'_{\text{pred}(i)}}$ (Equation 3) for the second step and the identities $a'_{\text{pred}(i)} = 2^{\alpha+1}/\beta'$ and $a_i d_i = \ell_i$ for the last one. Now observe that the left hand term of this equation is an upper bound on the interference at node i due to the classes $S_j, j \in L^{<i}$ such that the desired bound on this interference is shown.

Next we show that the interference at node i due to the classes $S_j, j \in L, j < i$ can be bounded by a similar approach studying the interference at $\text{succ}(i)$ instead of $\text{pred}(i)$. Let $L^{>\text{succ}(i)} := \{j \in L \mid j > \text{succ}(i)\}$. The interference at node $\text{succ}(i)$ from sets $S_j, j \in L^{>\text{succ}(i)}$ is

$$\begin{aligned} \sum_{j \in L^{>\text{succ}(i)}} \sum_{k \in S_j} \frac{\sqrt{\ell_k}}{(\delta_k + \delta_{\text{succ}(i)})^\alpha} &\geq \sum_{j \in L^{>\text{succ}(i)}} \sum_{k \in S_j} \frac{\sqrt{\ell_k}}{(2 \cdot \delta_k)^\alpha} \\ &= \sum_{j \in L^{>\text{succ}(i)}} \sum_{k \in S_j} \frac{\sqrt{\ell_k}}{2^\alpha d_k}. \end{aligned}$$

Analogously to the case above, the interference at node $\text{succ}(i)$ due to the nodes from U wrt to the loss parameters ℓ'_j is upper-bounded by $1/\beta''\sqrt{\ell'_{\text{succ}(i)}}$ so that

$$\sum_{j \in L^{>\text{succ}(i)}} \sum_{k \in S_j} \frac{\sqrt{\ell_k}}{2^\alpha d_k} \leq \frac{1}{\beta''\sqrt{\ell'_{\text{succ}(i)}}} = \frac{1}{\beta''\sqrt{a'_{\text{succ}(i)}d_{\text{succ}(i)}}}.$$

By multiplying with 2^α , we get

$$\begin{aligned} \sum_{j \in L^{<i}} \sum_{k \in S_j} \frac{\sqrt{\ell_k}}{d_k} &\leq \frac{2^\alpha}{\beta''\sqrt{a'_{\text{succ}(i)}d_{\text{succ}(i)}}} \\ &< \frac{4^\alpha}{\beta' \beta'' a'_{\text{succ}(i)} \sqrt{a_i d_i}} = \frac{2^\alpha}{2\beta'' \sqrt{\ell_i}}, \end{aligned}$$

where we used $d_{\text{succ}(i)} \geq a_{\text{succ}(i)} a_i \beta'^2 4^{-\alpha} > a'_{\text{succ}(i)} a_i \beta'^2 4^{-\alpha}$ (Equation 3) for the second step and the identities $a'_{\text{succ}(i)} = 2^{\alpha+1}/\beta'$ and $\ell_i = a_i d_i$ for the last one. Finally, observe that the left hand term of this equation is an upper bound on the interference at node i due to the classes $S_j, j \in L^{>\text{succ}(i)}$ such that the desired bound is shown, which completes the proof of Lemma 13. \square

The following lemma gives a bound on the interference induced by the large-loss nodes onto the small-loss nodes. Its proof has been omitted due to space limitations.

LEMMA 14. *For every $i \in U$ and $j \in S_i$,*

$$I_{\bar{p}}(j \mid L \setminus \{\text{pred}(i), i\}) \leq \frac{2^{2\alpha+2}}{\beta' \sqrt{\ell_j}}.$$

It remains to show how to combine Lemmas 10 to 14 to obtain Lemma 5. The interference at a node $j \in S_i$ for $i \in L$ can be bounded as follows. The interference caused by other nodes with small loss parameter is at most $1/(\beta''\sqrt{\ell_j})$ due to Lemma 11. The interference caused by nodes $L \setminus \{\text{pred}(i), i\}$ is at most $2^{2\alpha+2}/(\beta'\sqrt{\ell_j})$ due to Lemma 14. Finally, the interference caused by nodes i and $\text{pred}(i)$ at node j is at most

$$\begin{aligned} \frac{\sqrt{\ell_{\text{pred}(i)}}}{(\delta_{\text{pred}(i)} + \delta_j)^\alpha} + \frac{\sqrt{\ell_i}}{(\delta_i + \delta_j)^\alpha} &\leq \frac{\sqrt{\ell_{\text{pred}(i)}}}{\beta' \sqrt{\ell_{\text{pred}(i)}} \ell_j} + \frac{\sqrt{\ell_i}}{\beta' \sqrt{\ell_i} \ell_j} \\ &= \frac{2}{\beta' \sqrt{\ell_j}}. \end{aligned}$$

Here, the second step follows because $(\delta_i + \delta_j)^{2\alpha} \geq \beta'^2 \ell_i \ell_j$ due to Equation 1. In total the interference at j is at most $((2^{2\alpha+2} + 2) \frac{1}{\beta'} + \frac{1}{\beta''}) \ell_j^{-1/2}$. Hence, the nodes with small loss parameters are β -feasible if the constant terms c_0 and c_1 relating β' and β'' , respectively, to β satisfy the conditions $c_0 \geq 2^{2\alpha+2} + 2$ and $c_1 \geq 1$.

A slightly more involved argument is required to estimate the interference at a node $i \in L$. The interference due to other nodes in L is at most $2^{\alpha+2}/(\beta' \sqrt{\ell_i})$ by Lemma 10. The interference caused by nodes in the sets S_j , $j \notin \{i, \text{succ}(i)\}$ is at most $2^\alpha/(\beta'' \sqrt{\ell_i})$ by Lemma 13. The sets S_i and $S_{\text{pred}(i)}$, however, may cause large interference at node i . We use the following trick to deal with this problem: If $|S_i \cup \{i\} \cup S_{\text{succ}(i)}| > \beta'/\beta''$ then we do not choose the node i for the set U in Lemma 5. We can afford this because only $\mathcal{O}(\beta''/\beta')n = \mathcal{O}(\beta/\beta')n$ nodes satisfy this condition. Now suppose $|S_i \cup \{i\} \cup S_{\text{succ}(i)}| \leq \beta'/\beta''$. Then the interference at i due to these nodes is bounded by

$$\sum_{j \in S_i \cup S_{\text{succ}(i)}} \frac{\sqrt{\ell_j}}{(\delta_i + \delta_j)^\alpha} \leq \sum_{j \in S_i \cup S_{\text{succ}(i)}} \frac{\sqrt{\ell_j}}{\beta' \sqrt{\ell_i \ell_j}} \leq \frac{1}{\beta'' \sqrt{\ell_i}}.$$

Thus, if i is chosen, then the interference at node i is at most $((2^{\alpha+2} + 2) \frac{1}{\beta'} + (2^\alpha + 1) \frac{1}{\beta''}) \ell_i^{-1/2}$. Hence, the nodes with large loss parameters are β -feasible if $c_0 \geq 2^{\alpha+2} + 2$ and $c_1 \geq 2^\alpha + 1$. This completes the proof of Lemma 5.

5. A COLORING ALGORITHM FOR THE SQUARE ROOT POWER ASSIGNMENT

In the *coloring problem for the square root power assignment*, we are given n bidirectional requests and we seek for a coloring satisfying the SINR constraints with a minimal number of colors.

THEOREM 15. *There exists a randomized polynomial time algorithm solving the coloring problem for the square root power assignment with approximation factor $O(\log n)$.*

Let μ denote the maximal number of requests that can be scheduled with the same color. We will devise an algorithm \mathcal{A} that computes a subset $S \subseteq [n]$ of size $\Omega(\mu)$ with the property that the requests in S can be scheduled with the same color. In order to compute a coloring, algorithm \mathcal{A} is called and the requests in the set S are assigned to the first color class. This procedure is repeated recursively on the remaining requests until all requests have been colored. It is easy to see that such a greedy approach yields an $O(\log n)$ approximation for the optimal number of colors.

We now devise an algorithm \mathcal{A} that has the property described above. In the following, when saying that the SINR constraints are satisfied for a set of requests we mean that they are satisfied when all requests in the set are assigned the same color. The algorithm partitions the set of communication pairs into disjoint classes. W.l.o.g., let us assume $\min_{j \in [n]} \delta(u_j, v_j) = 1$ and let k be the smallest integer such that $\max_{j \in [n]} \delta(u_j, v_j) < 4^{k+1}$. For $0 \leq i \leq k$, class C_i contains the pairs $j \in [n]$ with $4^i \leq \delta(u_j, v_j) < 4^{i+1}$. This implies that the loss in this class is in $[4^{\alpha i}, 4^{\alpha(i+1)})$. For the time being, let us assume that all requests in class C_i have loss $4^{\alpha i}$ so that the square root power assignment sets the power level to $2^{\alpha i}$. We discuss the consequences of this simplifying assumption at the end of the proof.

The algorithm proceeds as follows. For $i = 0$ to k , it chooses a set S_i of sufficiently many (as defined later) requests from C_i taking into account interference caused by the previously selected sets S_0, \dots, S_{i-1} . In particular, S_i satisfies the SINR constraints with gain $\beta/2$ on top of S_0, \dots, S_{i-1} , i.e., the interference constraints for every pair in S_i are satisfied with gain $\beta/2$ taking into account the interference caused by the previously inserted pairs in S_0, \dots, S_{i-1} and the other pairs in S_i . Observe that we relaxed the interference constraints by using the gain $\beta/2$ instead of β . Furthermore, choosing S_i might violate the interference constraints of the previously chosen pairs in S_0, \dots, S_{i-1} . We come back to this aspect later.

Let us first take care that the algorithm chooses sufficiently many pairs. Let s_i^* be the maximal size of a subset of requests from C_i such that the SINR constraints at the nodes from S_i are satisfied with original gain β on top of the pairs in S_0, \dots, S_{i-1} .

LEMMA 16. *There is a polynomial time algorithm choosing S_i such that $|S_i| \geq s_i^*/k_0$, for a suitable constant $k_0 \geq 1$.*

PROOF. Let V denote the set of all nodes of the metric. For a node $w \in V$ and a set of requests S , let

$$I(w | S) = \sum_{j \in S} \frac{\sqrt{\ell(u_j, v_j)}}{\min\{\ell(u_j, w), \ell(v_j, w)\}}$$

be the interference at w caused by the pairs in the set S .

Let $S_0 \cup \dots \cup S_{i-1}$ be fixed. For simplicity of notation, we scale all distances such that the requests in class C_i have distance 1. Let $V' \subseteq V$ denote the subset of nodes with $I(w | S_1 \cup \dots \cup S_{i-1}) < 1/\beta$. Let C'_i denote the subset of requests from C_i only using nodes from V' . S_i^* can take only requests from C'_i as the other pairs exceed the interference threshold. Hence, we only need to take into account nodes from V' and requests from C'_i .

We have to choose a subset $S_i \subseteq C'_i$ of cardinality at least $s_i^*/k_0 = |S_i^*|/k_0$, for a suitable constant k_0 . We will choose S_i such that $I(w | S_i) < 1/\beta$, for every node w from any pair of S_i . This implies $I(w | S_1 \cup \dots \cup S_{i-1} \cup S_i) < 2/\beta$ as required in the description of the algorithm. The following claim gives a necessary condition that S_i needs to satisfy.

CLAIM 17. *Let T be any subset of C'_i satisfying the SINR constraints with gain β , then for every node w in V' it holds $I(w | T) < 2^\alpha \beta^{-1}$.*

PROOF. If $w = u_k$ or $w = v_k$, for some $k \in T$, then the condition is met directly by the definition of T . Otherwise, let n_k be the node from $\{u_k, v_k\}$ closest to w . Now let $j \in \arg \min_{k \in T} \delta(n_k, w)$, i. e., n_j is the node from T that is closest to w . By the triangle inequality it holds that $\delta(n_i, n_j) \leq \delta(n_i, w) + \delta(n_j, w) \leq 2\delta(n_i, w)$ so that $\ell(n_i, n_j) \leq 2^\alpha \ell(n_i, w)$. As a consequence,

$$I(w | T) \leq \sum_{i \in T} \frac{1}{\ell(n_i, w)} \leq 2^\alpha \sum_{i \in T} \frac{1}{\ell(n_i, n_j)} < 2^\alpha \beta^{-1}. \quad \square$$

The interference constraints from the claim can be described by an ILP with binary variables $x_j \in \{0, 1\}$, for $j \in C'_i$, and a linear SINR constraint for every node $w \in V'$. The objective is to maximize $|T| = \sum_{j \in C'_i} x_j$. We relax the integrality requirement and obtain an LP with variables $x_j \in [0, 1]$. This LP is solved to optimality. Let x' be the

optimal fractional solution and opt' its value. The claim above yields that opt' is an upper bound on s_i^* .

It is possible to compute a feasible subset S_i from x' of cardinality $\Omega(opt')$ via randomized rounding. The details are omitted due to space limitations. \square

The following lemma shows that we have selected $\Omega(\mu)$ requests. The proof is omitted.

LEMMA 18. $\left| \bigcup_{i=0}^k S_i \right| \geq \frac{\mu}{k_0+2}$.

Notice, when the algorithm computes S_i , it ensures that the interference constraints for S_i on top of S_0, \dots, S_{i-1} are satisfied with gain $\beta/2$. The algorithm does not explicitly take care for the additional interference caused by adding the pairs in S_i at the pairs from S_0, \dots, S_{i-1} . The following lemma, however, shows that this increase is bounded by a constant factor.

LEMMA 19. *There is constant $k_1 \geq 1$ such that $\bigcup_{i=0}^k S_i$ satisfies the SINR constraints with gain at most β/k_1 .*

Lemma 16 and 18 show that the algorithm chooses $\Omega(\mu)$ requests. However, these requests might violate the interference constraints with gain β because of the following reasons: a) We assumed that the loss in class C_i is exactly $4^{-\alpha i}$ rather than from the interval $[4^{\alpha i}, 4^{\alpha(i+1)})$. b) The pairs in each set S_i are chosen with respect to a relaxed gain $\beta/2$ instead of β . c) The SINR constraints for the sets in S_0, \dots, S_{i-1} are not explicitly considered when choosing S_i . (a) and (b) obviously increase the interference at most by a constant factor. Lemma 19 shows that the same is true for (c). Hence, the SINR constraints are violated at most by a constant factor so that they can be thinned out by applying Proposition 3. This way, one obtains a feasible set S of cardinality $\Omega(\mu)$. Thus, Theorem 15 is shown.

6. DISCUSSION AND OPEN PROBLEMS

One should remark that the bidirectional model can be simulated by the directed one using twice the number of steps (colors). Our analysis, hence, reveals that solutions with oblivious power assignments cannot compete with solutions using possibly different power levels and colors within a pair. However, they are capable of achieving nearly the same performance as solutions restricted to symmetric power assignments and colorings. Observe that oblivious power assignments use symmetric power assignments by definition. It is thus an interesting question that is left open by our analysis how they compare with solutions using symmetric power levels but asymmetric colorings.

In our analysis, we neglected some aspects that leave room for future research. For example, the presented coloring algorithm for the square root power assignment is centralized. It is an open question, whether there is a distributed coloring procedure that achieves the same kind of performance guarantee.

Another aspect that we did not take into account in this manuscript is energy efficiency. In comparison to the linear power assignment, the square root power assignment uses increased power levels for pairs of nodes of small distance with the objective to increase the performance. In [5], we study linear and, hence, energy efficient power assignments. We prove upper and lower bounds showing that linear power

assignments lose a factor that is logarithmic in the aspect ratio but linear in n against optimal power assignments. A study of the tradeoff between performance and energy efficiency is left for future work.

7. REFERENCES

- [1] Hari Balakrishnan, Christopher L. Barrett, V. S. Anil Kumar, Madhav V. Marathe, and Shripad Thite. The distance-2 matching problem and its relationship to the MAC-layer capacity of ad hoc wireless networks. *IEEE Journal on Selected Areas in Communications*, 22(6):1069–1079, 2004.
- [2] Nikhil Bansal and Zhen Liu. Capacity, delay and mobility in wireless ad-hoc networks. In *Proceedings of the 22nd Conference of the IEEE Communications Society (INFOCOM)*, volume 2, pages 1553–1563, 2003.
- [3] Deepti Chafekar, V. S. Anil Kumar, Madhav V. Marathe, Srinivasan Parthasarathy, and Aravind Srinivasan. Cross-layer latency minimization in wireless networks with SINR constraints. In *Proceedings of the 8th ACM International Symposium Mobile Ad-Hoc Networking and Computing (MOBIHOC)*, pages 110–119, 2007.
- [4] Deepti Chafekar, V. S. Anil Kumar, Madhav V. Marathe, Srinivasan Parthasarathy, and Aravind Srinivasan. Approximation algorithms for computing capacity of wireless networks with SINR constraints. In *Proceedings of the 27th Conference of the IEEE Communications Society (INFOCOM)*, pages 1166–1174, 2008.
- [5] Alexander Fanghänel, Thomas Kesselheim, and Berthold Vöcking. Improved algorithms for latency minimization in wireless networks with energy-efficient power control. In *Proceedings of the 36th International Colloquium on Automata, Languages and Programming (ICALP)*, 2009. to appear.
- [6] Anupam Gupta, Mohammad Taghi Hajiaghayi, and Harald Räcke. Oblivious network design. In *Proceedings of the 17th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 970–979, 2006.
- [7] Piyush Gupta and P. R. Kumar. Critical power for asymptotic connectivity in wireless networks. In William M. McEneaney, George Yin, and Qing Zhang, editors, *Stochastic Analysis, Control, Optimization, and Applications: A Volume in Honor of W. H. Fleming*, Systems & Control: Foundations & Applications, pages 547–566. Birkhäuser, 1998.
- [8] Piyush Gupta and P. R. Kumar. The capacity of wireless networks. *IEEE Transactions on Information Theory*, 46:388–404, 2000.
- [9] Ramin Hekmat and Piet Van Mieghem. Interference in wireless multi-hop ad-hoc networks and its effect on network capacity. *Wireless Networks*, 10(4):389–399, 2004.
- [10] Ulas C. Kozat and Leandros Tassiulas. Network layer support for service discovery in mobile ad hoc networks. In *Proceedings of the 22nd Conference of the IEEE Communications Society (INFOCOM)*, pages 1965–1975, 2003.
- [11] Sven O. Krumke, Madhav V. Marathe, and S. S. Ravi. Models and approximation algorithms for channel assignment in radio networks. *Wireless Networks*,

7(6):575–584, 2001.

- [12] Thomas Moscibroda and Roger Wattenhofer. The complexity of connectivity in wireless networks. In *Proceedings of the 25th Conference of the IEEE Communications Society (INFOCOM)*, pages 1–13, 2006.
- [13] Thomas Moscibroda, Roger Wattenhofer, and Aaron Zollinger. Topology control meets SINR: The scheduling complexity of arbitrary topologies. In *Proceedings of the 7th ACM International Symposium Mobile Ad-Hoc Networking and Computing (MOBIHOC)*, pages 310–321, 2006.
- [14] Subramanian Ramanathan and Errol L. Lloyd. Scheduling algorithms for multi-hop radio networks. *ACM SIGCOMM Computer Communication Review*, 22(4):211–222, 1992.
- [15] Suresh Singh and C. S. Raghavendra. PAMAS—Power aware multi-access protocol with signalling for ad hoc networks. *ACM SIGCOMM Computer Communication Review*, 28(3):5–26, 1998.