Chapter 2
Scheduling and Power Assignments in the Physical Model

Alexander Fanghänel and Berthold Vöcking

Abstract In the interference scheduling problem, one is given a set of n communication requests each of which corresponds to a sender and a receiver in a multipoint radio network. Each request must be assigned a power level and a color such that signals in each color class can be transmitted simultaneously. 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 modeling interference in the engineering community. The objective is to minimize the schedule length corresponding to the number of colors needed to schedule all requests. We study *oblivious power assignments* in which the power value of a request only depends on the path loss between the sender and the receiver, e.g., in a linear fashion. At first, we present a *measure of interference* giving lower bounds for the schedule length with respect to linear and other power assignments. Based on this measure, we devise distributed scheduling algorithms for the linear power assignment achieving the minimal schedule length up to small factors. In addition, we study a power assignment in which the signal strength is set to the square root of the path loss. We show that this power assignment leads to improved approximation guarantees in two kinds of problem instances defined by directed and bidirectional communication request. Finally, we study the limitations of oblivious power assignments by proving lower bounds for this class of algorithms.

2.1 Introduction

Simultaneously transmitted radio signals interfere with each other. Early theoretical approaches (see, e.g., [11, 13, 17]) about scheduling signals or packets in radio networks resort to *graph-based vicinity models* (also known as *protocol model*) of

A. Fanghänel (✉)
RWTH Aachen University, Aachen, Germany
e-mail: fanghaenel@cs.rwth-aachen.de

S. Nikoletseas, J.D.P. Rolim (eds.), *Theoretical Aspects of Distributed Computing in Sensor Networks,* Monographs in Theoretical Computer Science. An EATCS Series,
DOI 10.1007/978-3-642-14849-1_2,

the following flavor. 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 modeled 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 modeling approach is that it ignores that neither radio signals nor interference ends abruptly at a boundary.

Recent theoretical studies [1–4, 6, 7, 9, 14, 15] 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 $d(u, v)$ denote the distance between the nodes u and v. We assume the *path loss radio propagation* model, where a signal sent by node u with power p is received at node v with $p/d(u, v)^\alpha$, 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 $p/d(u, v)^\alpha$. Node v can successfully decode this 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 $\beta > 1$ being the second parameter of the model, the so-called *gain*.

Let us illustrate the physical model with a simple but intriguing example showing the importance of choosing the right power assignment. Suppose there are two pairs of nodes (u_1, v_1) and (u_2, v_2). Two signals shall be sent simultaneously, one from u_1 to v_1 and the other from u_2 to v_2. Suppose the nodes are placed in a nested fashion on a line, that is, the points are located on the line in the order u_1, u_2, v_2, v_1 such that the distance between u_1 and u_2 is two, the distance between u_2 and v_2 is one, and the distance between v_2 and v_1 is two (cf. Fig. 2.1). For simplicity fix $\alpha = 2$ and $\beta = 1$ and neglect the noise.

- At first, let us assume that both u_1 and u_2 send their signal with the same power 1. Then the strength of u_1's signal at node v_1 is $1/25$ while the strength of u_2's signal at the same node is $1/9$. Hence, v_1 cannot decode the signal sent by node u_1 as it is drowned by u_2's signal. That is, the outer pair is blocked by the inner pair when using uniform powers.
- At second, let us assume that signals are sent in a way that the path loss is compensated, that is, both nodes use a strength that is linear in the path loss. In particular, u_1 sends at power 25 and u_2 sends at power 1. Now consider the strengths of the signals received at v_2: The strength of u_2's signal is only 1 while the strength of u_1's signal is $25/9$. Thus, the inner pair is blocked by the outer pair when using powers that are chosen linear in the path loss.
- Finally, let us make an attempt setting the powers equal to the square root of the path loss, that is, u_1 uses power 5 and u_2 uses power 1. Now easy calculus

[1] Depending on the environment, it is usually assumed that α has a value between 2 and 5.

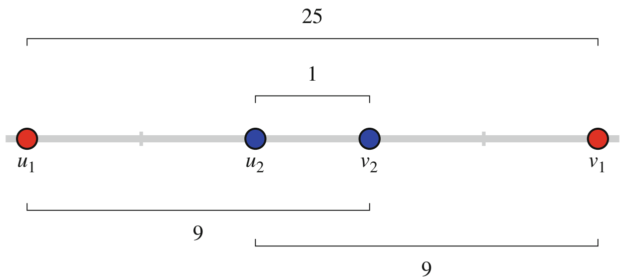

Fig. 2.1 Placement of the nodes and the path loss for $\alpha = 2$. Linear and uniform power assignment both need different schedule steps for each of the requests, the square root power assignment can schedule both requests at once

shows that, at v_1, the strength of u_1's signal is larger than the strength of u_2s and, at v_2, the strength of u_2's signal is larger than the strength of u_1s. Hence, simultaneous communication between the nested pairs is possible when choosing the right power assignment.

In this chapter, we investigate interference scheduling problems like the one in the example above. In general, 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. The objective is to minimize the number of colors, which corresponds to minimizing the time needed to schedule all communication requests. As this problem is NP-hard [1], we are interested in approximation algorithms.

The interference scheduling problem consists of two correlated subproblems: the *power assignment* and the *coloring*. By far, most literature focuses on scheduling with *uniform power assignment*, in which all pairs send at the same power (see, e.g., [8, 12, 18]). In other studies, the *linear power assignment* is considered, in which the power level for a pair (u, v) is chosen proportional to the path loss $d(u, v)^\alpha$. In the example above, we have seen that choosing powers proportional to the square root of the path loss might be an interesting alternative. All these power assignments have the advantage that they are locally computable independent of other requests, which allows for an immediate implementation in a distributed setting. These are examples of *oblivious power assignments* which mean the power level assigned to a pair is defined as a function of the path loss (or the distance) between the nodes of a pair.

2.1.1 Outline

In Sect. 2.2, we formally introduce the physical model with SINR constraints and show a helpful robustness property of this model. In Sect. 2.3, we study scheduling algorithms for the linear power assignment. In particular, we introduce a *measure of interference* giving lower bounds for the schedule length with respect not only to linear but also to other power assignments. Based on this measure, we devise distributed scheduling algorithms for the linear power assignment achieving the minimal schedule length up to small factors. In Sect. 2.4, we study the square root power assignment. We show that this power assignment leads to better approximation guarantees in two kinds of problem instances defined by directed and bidirectional communication request. In Sect. 2.5, we study the limitations of oblivious power assignments by proving lower bounds for this approach. Finally, in Sect. 2.6 we summarize the results from our presentation with pointers to the literature and open problems.

2.2 Notation and Preliminaries

Let the path loss exponent $\alpha \geq 1$ and the gain $\beta > 1$ be fixed. Let V be a set of nodes from a metric space. Let $d(u, v)$ denote the distance between two nodes u and v. One is given a set R of n requests consisting of pairs $(u_i, v_i) \in V^2$, where u_i represents the source and v_i the destination of the signal from the ith request. W.l.o.g., we assume $\min_{i \in R} d(u_i, v_i) = 1$. Let $\Delta = \max_{i \in R} d(u_i, v_i)$ be the *aspect ratio*. We say that a set R of requests is a *nearly equilength* set, if the lengths of the requests in R differ by at most factor 2.

In the *interference scheduling problem* one needs to specify a power level $p_i > 0$ and a color $c_i \in [k] := \{1, \ldots, k\}$ for every $i \in [n] := \{1, \ldots, n\}$ such that the *latency*, i.e., the number of colors k, is minimized and the pairs in each color class satisfy the *SINR constraint*, that is, for every $i \in [n]$, it holds

$$\frac{p_i}{d(u_i, v_i)^\alpha} \geq \beta \left(\sum_{\substack{j \in [n] \setminus \{i\} \\ c_j = c_i}} \frac{p_j}{d(u_j, v_i)^\alpha} + \nu \right) \tag{2.1}$$

The SINR constraint is the central condition for successful communication in the physical model. It characterizes the received strength of the signal emitted from u_i at receiver v_i compared to *ambient noise* ν and the *interference* from signals of all other senders in the same color class. The so-called *scheduling complexity* of R, as introduced by Moscribroda and Wattenhofer [14], is the minimal number of colors (steps) needed to schedule all requests in R.

In this chapter we focus on distance-based power assignments because of their simplicity and locality, which is a striking conceptual advantage in distributed wireless systems. An *oblivious* (or distance-based) power assignment p is given by $p_i = \phi(d(u_i, v_i))$ with a function $\phi : [1, \Delta] \to (0, \infty)$. For uniqueness we assume that ϕ is always scaled such that $\phi(1) = 1$. Examples are the *uniform* $\phi(d(u_i, v_i)) = 1$ or the *linear* $\phi(d(u_i, v_i)) = d(u_i, v_i)^\alpha$ power assignment. Recently, the *square root* assignment $\phi(d(u_i, v_i)) = d(u_i, v_i)^{\alpha/2}$ has attracted some interest [5, 9] as it yields better approximation ratios for request scheduling than the uniform and the linear power assignment.

We define the *relative interference* on a request i from a request set R as

$$\mathrm{RI}_i(R) = c_i \cdot \frac{d(u_i, v_i)^\alpha}{p_i} \cdot \sum_{j \in R} \frac{p_j}{d(u_j, v_i)^\alpha}$$

where

$$c_i = \frac{\beta}{1 - \beta \cdot \nu \cdot p_i / d(u_i, v_i)}$$

denotes a constant that indicates the extent to which the ambient noise approaches the required signal at the receiver of request i. The relative interference describes the received interference at receiver v_i normalized by the received signal strength. The relative interference satisfies the two following properties for a request set R. First, R is SINR feasible iff for every $i \in R$, $\mathrm{RI}_i(R) \leq 1$. Second, the relative interference function is additive, that is, for every partition $R = R_1 \dot\cup R_2$ and every request i it holds $\mathrm{RI}_i(R) = \mathrm{RI}_i(R_1) + \mathrm{RI}_i(R_2)$.

We denote with an r*-signal* set or schedule one where each requests relative interference is at most $1/r$.

2.2.1 Robustness of the Physical Model

The main criticism of graph-based models is that they are too simplistic to model real wireless networks. The physical model requires simplifying assumptions, too, as (2.1) models no obstructions, perfectly isotropic radios and a constant ambient noise level.

In the following proof (from [10]) we show that there are only minor changes in the schedule length, if there are minor changes in the signal requirements. This justifies the analytic study of the physical model despite its simplifying assumptions.

Proposition 1 *Let R be a r-signal schedule under a power assignment p. Then there exists a r'-signal schedule R' for p that is at most $\lceil 2r'/r \rceil^2$ times longer than R, for $r' > r$.*

Proof Let R be a r-signal schedule and T be a single schedule step. We show that we can decompose T in at most $\lceil 2r'/r \rceil$ slots $T_1, T_2, \ldots$ that are r'-signal sets. We now process the requests in T by increasing index. For request i, assign it to the first set T_j, in which the relative interference on i is at most $1/2r'$. Since every request had at most a relative interference of $1/r$, it follows from the additivity of relative interference that there are at most

$$\left\lceil \frac{1/r}{1/2r'} \right\rceil = \left\lceil \frac{2r'}{r} \right\rceil$$

such sets. In each of these sets T_j the relative interference from requests with lower index is at most $1/2r'$. Now, for each of these sets we repeat this process, processing the requests in T_{ji}, now in reverse order. Using the same arguments T_j is split into at most $\lceil 2r'/r \rceil$ sets. In that way we make sure that the requests in each set have a relative interference of at most $1/2r'$ from requests with higher index, which bounds the total relative interference on each request by $1/r'$, while using at most $\lceil 2r'/r \rceil^2$ times more slots than the original schedule. □

2.3 Scheduling with the Linear Power Assignment

In the first part we focus on the linear power assignment, i.e., the power for a request pair (u_i, v_i) is equal to $d(u_i, v_i)^\alpha$ and, hence, linear in the path loss. The linear power assignment has the advantage of being energy efficient as the minimal transmission power required to transmit along a distance $d(u_i, v_i)$ is proportional to $d(u_i, v_i)^\alpha$.

We first present a measure of interference I, which allows us to lower bound the schedule for general metrics using the linear power assignment by $\Omega(I)$. If we allow any power assignment, the schedule length can be bounded by $\Omega(I/\log \Delta \log n)$. For $\alpha > 2$, embedding the instance in the Euclidean space improves this bound to $O(I/\log \Delta)$.

These results are complemented by a simple and efficient algorithm computing a schedule using $O(I \cdot \log n)$ steps. A more sophisticated algorithm computes a schedule using $O(I + \log^2 n)$ steps. This gives a constant factor approximation of the optimal schedule using the linear power assignment for dense instances, i.e., if $I \geq \log^2 n$.

2.3.1 Measure of Interference and Lower Bounds

We first present an instance-based *measure of interference* I, which allows us to lower bound the number of steps needed for scheduling a request set R in terms of I.

Definition 1 (Measure of Interference) Let $R \subseteq V \times V$ be a set of requests. For $w \in V$ define

$$I_w(R) = \sum_{(u,v)\in R} \min\left(1, \frac{d(u,v)^\alpha}{d(u,w)^\alpha}\right)$$

Using this function we define the measure of interference induced by the requests in R:

$$I = I(R) = \max_{w\in V} I_w(R)$$

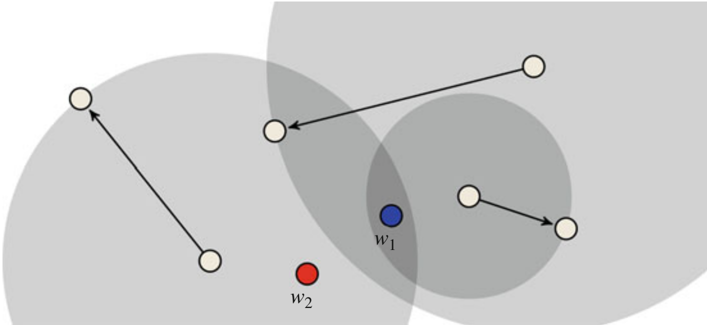

Fig. 2.2 An example for the measure of interference with three requests. *Gray* circles mark the areas where the interference from a sender is at least 1. For the *red* node I_{w_2} is 1 plus the interference from the two rightmost senders (each less than 1). The interference is maximal at the *blue* node w_1, i. e., $I_{w_1} = 3$, so the measure of interference I for this instance is $I = 3$

An example of the measure of interference is illustrated in Fig. 2.2.

Observe that I is subadditive, i.e., for $R = R_1 \cup R_2$ it holds

$$\begin{aligned} I(R) &= \max_{w\in V} I_w(R) \le \max_{w\in V} \{I_w(R_1) + I_w(R_2)\} \\ &\le \max_{w\in V} I_w(R_1) + \max_{w\in V} I_w(R_2) = I(R_1) + I(R_2) \end{aligned}$$

Theorem 1 *Let T be the minimum schedule length for a set of requests R with the linear power assignment. Then we have $T = \Omega(I)$.*

Proof Let there be a schedule of length T when using the linear power assignment. Then there exist sets of requests $R_1, \ldots, R_T$ each of which satisfies the SINR constraint for this power assignment. As I is subadditive we have $I\left(\bigcup_{t=1}^{T} R_t\right) \le \sum_{t=1}^{T} I(R_t)$. Thus it suffices to show that $I(R_t) = O(1)$ for every $t \in \{1, \ldots, T\}$, as this implies $T = \Omega(I)$.

Let $R_t = \{(u_1, v_1), \ldots, (u_{\bar{n}}, v_{\bar{n}})\}$ and let $w \in V$. Furthermore, let v_j be the receiver from R_t that is closest to w, i.e., $j \in \arg\min_{i\in[\bar{n}]} d(v_i, w)$. Possibly $w = v_j$.

We distinguish between two kinds of requests. We define a set U of indices of requests whose senders u_i lie within a distance of at most $\frac{1}{2}d(v_j, w)$ from w, i.e., $U = \{i \in [\bar{n}] \mid d(u_i, w) \le \frac{1}{2}d(v_j, w)\}$. Using the triangle inequality we can conclude for all $i \in U$:

$$d(u_i, v_j) \le d(u_i, w) + d(w, v_j) \le \frac{3}{2}d(v_j, w) \tag{2.2}$$

In addition, we have

$$d(v_j, w) \le d(v_i, w) \le d(v_i, u_i) + d(u_i, w) \le d(v_i, u_i) + \frac{1}{2}d(v_j, w)$$

Here the first equation holds since v_j is the closest receiver to w, the second equation holds by triangle inequality and the third step follows from the definition of U. This implies

$$d(v_j, w) \leq 2d(u_i, v_i) \tag{2.3}$$

Combining (2.2) and (2.3) we get $d(u_i, v_j) \leq 3d(u_i, v_i)$. Thus it holds

$$|U \setminus \{j\}| = \sum_{\substack{i \in U \\ i \neq j}} \frac{d(u_i, v_i)^\alpha}{d(u_i, v_i)^\alpha} \leq \sum_{\substack{i \in U \\ i \neq j}} \frac{d(u_i, v_i)^\alpha}{\frac{1}{3^\alpha} d(u_i, v_j)^\alpha} \leq \frac{3^\alpha}{\beta}$$

Hence,

$$I_w(U) = \sum_{i \in U} \min \left\{ 1, \frac{d(u_i, v_i)^\alpha}{d(u_i, w)^\alpha} \right\} \leq \frac{3^\alpha}{\beta} + 1$$

Next we upper bound $I_w(R_t \setminus U)$. For all $i \in [\bar{n}] \setminus U$ it holds that

$$d(u_i, v_j) \leq d(u_i, w) + d(w, v_j) \leq d(u_i, w) + 2d(u_i, w) = 3d(u_i, w)$$

by applying triangle inequality and the definition of U. As a consequence

$$I_w(R_t \setminus U) \leq \sum_{\substack{i \in [\bar{n}] \setminus U \\ i \neq j}} \frac{d(u_i, v_i)^\alpha}{d(u_i, w)^\alpha} \leq \sum_{\substack{i \in [\bar{n}] \setminus U \\ i \neq j}} \frac{d(u_i, v_i)^\alpha}{\frac{1}{3^\alpha} d(u_i, v_j)^\alpha} \leq \frac{3^\alpha}{\beta}$$

Thus

$$I_w(R_t) \leq I_w(U) + I_w(R_t \setminus U) = \frac{2 \cdot 3^\alpha}{\beta} + 1 = O(1)$$

□

Theorem 2 *Let T denote the optimal schedule length using any power assignment. Then we have $T = \Omega\,(I/\log \Delta \cdot \log n)$.*

Proof We use a similar technique as in the proof of Theorem 1. However, we have to deal with an unknown power assignment. Since there is a schedule of length T in this power assignment, there exist sets of requests $R_1, \ldots, R_T$ each of which satisfies the SINR constraint for this power assignment. We divide such a set R_t into $\log \Delta$ classes $C_{t,j} = \{(u, v) \in R_t \mid 2^{j-1} \leq d(u, v) < 2^j\}$. Again, by using the subadditivity of I, it suffices to show that $I(C_{t,j}) = O(\log n)$ for such a class. Fix $C_{t,j}$ and let $C_{t,j} = \{(u_1, v_1), \ldots, (u_{\bar{n}}, v_{\bar{n}})\}$. Further, for notational simplicity we write $L = 2^{j-1}$.

As an important fact we can bound the number of requests whose senders are located around a node within a distance of at most ℓ.

Lemma 1 *For all* $w \in V$, $\ell \geq L$ *we have for* $K_\ell(w) = \{i \in [\bar{n}] \mid d(u_i, w) \leq \ell\}$*:*

$$|K_\ell(w)| \leq \frac{1}{\beta}\left(\frac{4\ell}{L}\right)^\alpha + 1$$

Proof Let p be the power assignment that allows all requests to be served in a single time slot. Let furthermore (u_k, v_k) be the request with $k \in K_L(w)$ that is transmitted with minimal power p_k. As the SINR condition is satisfied for request (u_k, v_k), we get

$$\frac{1}{\beta}\frac{p_k}{d(u_k, v_k)^\alpha} \geq \sum_{\substack{i \in K_\ell(w) \\ i \neq k}} \frac{p_i}{d(u_i, v_k)^\alpha} \geq \sum_{\substack{i \in K_\ell(w) \\ i \neq k}} \frac{p_i}{(2\ell + 2L)^\alpha} \geq \frac{(|K_\ell(w)| - 1) \cdot p_k}{(2\ell + 2L)^\alpha}$$

So

$$|K_\ell(w)| - 1 \leq \frac{1}{\beta}\left(\frac{2\ell + 2L}{d(u_k, v_k)}\right)^\alpha \leq \frac{1}{\beta}\left(\frac{4\ell}{L}\right)^\alpha$$

□

Now, let $w \in V$. We prove $I_w(C_{t,j}) = O(\log n)$. W. l. o. g., let $u_1, \ldots, u_{\bar{n}}$ be ordered by increasing distance to w. Note that for all $\ell > 0$ we have $K_\ell(w) = \{1, \ldots, x\}$ for some $x \in \mathbb{N}$ by this definition.

For $k \leq \log \bar{n} + 1$ let $S_k = [2^k] \setminus [2^{k-1}]$. Furthermore, let ℓ_k be defined as $\ell_k = \min_{i \in S_k} d(u_i, w)$. For the value of $I_w(C_{t,j})$ follows from these definitions:

$$\begin{aligned} I_w(C_{t,j}) &= \sum_{i=1}^{\bar{n}} \min\left\{1, \frac{d(u_i, v_i)^\alpha}{d(u_i, w)^\alpha}\right\} \\ &\leq \sum_{k=1}^{\log \bar{n}+1} \sum_{i \in S_k} \frac{d(u_i, v_i)^\alpha}{d(u_i, w)^\alpha} + \sum_{i \in K_L(w)} 1 \leq (2L)^\alpha \sum_{k=1}^{\log \bar{n}+1} \frac{|S_k|}{\ell_k^\alpha} + |K_L(w)| \end{aligned}$$

As the distances are increasing, we have $\ell_k \geq d(u_i, w)$ for all $i \leq 2^{k-1}$. In other words $[2^{k-1}] \subseteq K_{\ell_k}(w)$.

Since we add up the interference induced by requests from $K_L(w)$ separately, we may assume $\ell_k \geq L$ for all k and thus apply Lemma 1 on $|K_{\ell_k}(w)|$, which gives

$$2^{k-1} = |[2^{k-1}]| \leq |K_{\ell_k}(w)| \leq \left(\frac{4\ell_k}{L}\right)^\alpha + 1$$

Consequently, we have

$$\ell_k^\alpha \geq (2^{k-1} - 1)\left(\frac{L}{4}\right)^\alpha$$

Using the above results for ℓ_k^α and $|K_L(w)|$ we can bound $I_w(C_{t,j})$ by

$$(2L)^\alpha \sum_{k=1}^{\log \bar{n}+1} \frac{2^{k-1}}{(2^{k-1}-1)\left(\frac{L}{4}\right)^\alpha} + \left(\frac{4^\alpha}{\beta} + 1\right) \leq 8^\alpha \sum_{k=1}^{\log \bar{n}+1} 2 + \frac{4^\alpha}{\beta} + 1 = O(\log n)$$

□

Earlier results restricted the instances often to the Euclidean plane and required α to be strictly greater than 2. Under these assumptions we can use geometric arguments to get an even better bound of $\Omega(I/\log \Delta)$ on the optimal schedule length, as we show in the following.

Theorem 3 *Let the instance be located in the Euclidean plane, let $\alpha > 2$, and let T denote the optimal schedule length using any power assignment. Then we have $T = \Omega\,(I/\log \Delta)$.*

Proof Again, we divide the requests into $\log \Delta \cdot T$ classes $C_{t,i}$. This time, we have to prove $I_w(C_{t,i}) = O(1)$. Let us remark that in the Euclidean plane a ring of inner radius $L \cdot r$ and width L can be covered by $8(r+1)$ circles of radius L. If x is the center of such a circle, we get from Lemma 1 that $|K_L(x)| \leq \frac{4^\alpha}{\beta}$. Thus we have $|K_{L(r+1)}(w) \setminus K_{Lr}(w)| \leq 8(r+1)\frac{4^\alpha}{\beta} \leq 16r\frac{4^\alpha}{\beta} = r\frac{4^{\alpha+2}}{\beta}$ for $r \geq 1$. We can bound $I_w(C_{t,j})$ by

$$I_w(C_{t,j}) \leq \sum_{r=1}^{\infty} |K_{L(r+1)}(w) \setminus K_{Lr}(w)| \cdot \frac{(2L)^\alpha}{(Lr)^\alpha} + |K_L(w)|$$

Using the above result we get

$$I_w(C_{t,j}) \leq 2^\alpha \frac{4^{\alpha+2}}{\beta} \sum_{r=1}^{\infty} r^{1-\alpha} + \frac{4^\alpha}{\beta} \leq \frac{4^\alpha}{\beta}\left(2^\alpha 4^2 \frac{\alpha-1}{\alpha-2} + 1\right) = O(1)$$

□

2.3.2 *Upper Bounds for the Linear Power Assignment*

The measure of interference enables us to design randomized algorithms using the linear power assignment, i.e., the power for the transmission from u to v is $c \cdot d(u,v)^\alpha$ for some fixed $c \geq \beta\nu$. As a key fact, we can simplify the SINR constraint in this setting as follows. If R is a set of requests that can be scheduled in one time slot, we have for all nodes v' with $(u', v') \in R$

$$\sum_{\substack{(u,v)\in R \\ (u,v)\neq(u',v')}} \frac{c \cdot d(u,v)^{\alpha}}{d(u,v')^{\alpha}} \leq \frac{c}{\beta} - \nu$$

Since $\beta > 1$ we can write equivalently

$$I_{v'}(R) = \sum_{(u,v)\in R} \min\left\{1, \frac{d(u,v)^{\alpha}}{d(u,v')^{\alpha}}\right\} \leq \frac{1}{\beta} - \frac{\nu}{c} \tag{2.4}$$

For simplicity of notation we replace $\frac{1}{\beta} - \frac{\nu}{c}$ by $\frac{1}{\beta'}$ in the following proofs.

The idea of our basic algorithm (Algorithm 1) is that each sender decides randomly in each time slot if it tries to transmit until it is successful. The probability of transmission is set to $\frac{1}{2\beta' I}$ and is not changed throughout the process.

Algorithm 1 A simple single-hop algorithm

1: **while** packet has not been successfully transmitted **do**
2: try transmitting with probability $\frac{1}{2\beta' I}$
3: **end while**

Theorem 4 *Algorithm 1 generates a schedule of length at most $O(I \log n)$ whp.*

Proof Let us first consider the probability of success for a fixed request (u_k, v_k) in a single step of the algorithm. Let X_i, $i \in [n]$, be the 0/1 random variable indicating if sender u_i tries to transmit in this step. Assume a sender u_k tries to transmit in this step, i.e., $X_k = 1$. To make this attempt successful, the interference constraint (2.4) has to be satisfied. We can express this event as $Z \leq 1/\beta'$ where Z is defined by

$$Z = \sum_{\substack{i\in[n] \\ i\neq k}} \min\left\{1, \frac{d(u_i,v_i)^{\alpha}}{d(u_i,v_k)^{\alpha}}\right\} X_i$$

We have $\mathbf{E}[Z] \leq 1/2\beta'$ and thus we can use Markov's inequality to bound the probability that this packet cannot be transmitted successfully by

$$\mathbf{Pr} Z \geq \frac{1}{\beta'} \leq \mathbf{Pr} Z \geq 2\mathbf{E}[Z] \leq \frac{1}{2}$$

To make the transmission successful the two events $X_k = 1$ and $Z \leq 1/\beta' \frac{1}{\beta'}$ have to occur. Since they are independent it holds that

$$\mathbf{Pr} X_k = 1, Z \leq \frac{1}{\beta'} = \mathbf{Pr} X_k = 1 \cdot \mathbf{Pr} Z \leq \frac{1}{\beta'} \geq \frac{1}{2\beta' I}\left(1 - \frac{1}{2}\right) = \frac{1}{4\beta' I}$$

The probability for packet k not to be successfully transmitted in $(k_0 + 1)4\beta' I \ln n$ independent repeats of such a step is therefore at most

$$\left(1 - \frac{1}{4\beta' I}\right)^{(k_0+1)4\beta' I \ln n} \leq e^{-(k_0+1)\ln n} = n^{-(k_0+1)}$$

Applying a union bound we get an overall bound on the probability that one of n packets is not successfully transmitted in these independent repeats by n^{-k_0}. This means all senders are successful within $O(I \log n)$ steps whp.

An obvious disadvantage of the basic algorithm is that the probability of transmission stays the same throughout the process. To improve it, one idea could be to increase the probability of transmission after some transmissions have successfully taken place. This is why we need the following weighted Chernoff bound that can deal with dependent random variables.

Lemma 2 *Let $X_1, \ldots, X_n$ be 0/1 random variables for which there is a $p \in [0, 1]$ such that for all $k \in [n]$ and all $a_1, \ldots, a_{k-1} \in \{0, 1\}$*

$$\mathbf{Pr} X_k = 1 \mid X_1 = a_1, \ldots X_{k-1} = a_{k-1} \leq p \tag{2.5}$$

Let furthermore $w_1, \ldots, w_n$ be reals in $(0, 1]$ and $\mu \geq p \sum w_i$. Then the weighted Chernoff bound

$$\mathbf{Pr} \sum_{i=1}^{n} w_i X_i \geq (1 + \delta)\mu \leq \left(\frac{e^\delta}{(1 + \delta)^{(1+\delta)}}\right)^\mu$$

holds.

Proof (*Sketch*). To show this bound, a standard proof for the weighted Chernoff bound [16] can be adapted. By using the definition of expectation and repeatedly applying (2.5), one can show that

$$\mathbf{E}\left[e^{tX}\right] \leq \prod_{i=1}^{n} \left(p e^{t w_i} + 1 - p\right)$$

although random variables are no more independent. In the original proof no other step makes use of the independence. □

We can now use this bound to analyze Algorithm 2. This algorithm assigns random delays to all packets. The maximum delay is decreased depending on I^{curr}, which denotes the measure of interference that is induced by the requests that have not been scheduled at this point.

The algorithm works as follows: During one iteration of the outer *while* loop by repeatedly assigning random delays to the packets the measure of interference is

Algorithm 2 An $O(I + \log^2 n)$ whp algorithm

```
while I^curr ≥ log n do
    J := I^curr
    while I^curr ≥ J/2 do
        if packet i has not been successfully transmitted then
            assign a delay 1 ≤ δ_i ≤ 16eβ'J i. u. r.
            try transmission after waiting the delay
        end if
    end while
end while
execute algorithm Algorithm 1
```

reduced to a half of its initial value. This is repeated until we have $I^{\text{curr}} < \log n$ and the basic algorithm is applied.

Our first observation is that reducing I^{curr} by factor 2 takes $O(I^{\text{curr}})$ scheduling steps whp.

Lemma 3 *During one iteration of the outer* while *loop of Algorithm 2, the inner* while *loop is executed at most* $k_0 + 2$ *times with probability at least* $1 - n^{-k_0}$ *for all constants* k_0.

Proof Let us first consider a single iteration of this loop. We assume all senders are taking part as if none has been successful during this iteration of the outer *while* loop yet. We only benefit from any previous success.

Observe, if the senders of a set S are transmitting and there is a collision for packet i we have

$$\sum_{\substack{j \in S \\ j < i}} \min\left\{1, \frac{d(u_j, v_j)^\alpha}{d(u_j, v_i)^\alpha}\right\} > \frac{1}{2\beta'} \quad \text{or} \quad \sum_{\substack{j \in S \\ j > i}} \min\left\{1, \frac{d(u_j, v_j)^\alpha}{d(u_j, v_i)^\alpha}\right\} > \frac{1}{2\beta'}$$

In the first case let $Y_i^< = 1$, in the second one $Y_i^> = 1$. We now show that the random variables $Y_1^<, \ldots, Y_n^<$ fulfill (2.5) for $p = \frac{1}{8e}$. Let us fix $k \in [n]$ and $a_1, \ldots, a_{k-1} \in \{0, 1\}$. We have to show $\mathbf{Pr} Y_k^< = 1 \mid Y_1^< = a_1, \ldots, Y_{k-1}^< = a_{k-1} \leq p$.

Since the delays δ_i are drawn independently they can be considered as if they were drawn one after the other in the order $\delta_1, \delta_2, \ldots$. Then the value of $Y_i^<$ would already be determined after drawing δ_i by definition. In other words, the values of $\delta_1, \ldots, \delta_{k-1}$ already determine the values of $Y_1^<, \ldots, Y_{k-1}^<$. It follows that there is a subset $M \subseteq [16e\beta' J]^{k-1}$ of delay values such that $Y_1^< = a_1, \ldots, Y_{k-1}^< = a_{k-1}$ iff $(\delta_1, \ldots, \delta_{k-1}) \in M$.

Now let X_i be a 0/1 random variable for $i \in [k-1]$ such that $X_i = 1$ iff $\delta_i = \delta_k$. We can observe that we have for all $(b_1, \ldots, b_{k-1}) \in [16e\beta' J]^{k-1}$:

$$\mathbf{E}\left[X_i \mid \delta_1 = b_1, \ldots, \delta_{k-1} = b_{k-1}\right] = \frac{1}{16e\beta' J}$$

Define furthermore

$$Z_k^< = \sum_{i=1}^{k-1} \min\left\{1, \frac{d(u_i, v_i)^\alpha}{d(u_i, v_k)^\alpha}\right\} X_i$$

with $\mathbf{E}\left[Z_k^< \mid \delta_1 = b_1, \ldots, \delta_{k-1} = b_{k-1}\right] \leq \frac{1}{16e\beta'}$. Now it holds that

$$\begin{aligned}
&\mathbf{Pr}\left[\, Y_k^< = 1 \;\mid\; \delta_1 = b_1, \ldots, \delta_{j-1} = b_{k-1} \,\right] \\
&\qquad = \mathbf{Pr} Z_k^< > \frac{1}{2\beta'} \;\middle|\; \delta_1 = b_1, \ldots, \delta_{k-1} = b_{k-1} \\
&\qquad \leq 2\beta' \mathbf{E}\left[Z_k^< \mid \delta_1 = b_1, \ldots, \delta_{k-1} = b_{k-1}\right] \\
&\qquad = \frac{1}{8e} = p
\end{aligned}$$

We now apply the law of alternatives:

$$\begin{aligned}
&\mathbf{Pr} Y_k^< = 1 \mid Y_1^< = a_1, \ldots, Y_{k-1}^< = a_{k-1} \\
= &\sum_{(b_1,\ldots,b_{k-1})\in M} \mathbf{Pr} \delta_1 = b_1, \ldots, \delta_{k-1} = b_{k-1} \mid Y_1^< = a_1, \ldots, Y_{k-1}^< = a_{k-1} \\
&\quad \cdot \mathbf{Pr} Y_k^< = 1 \mid \delta_1 = b_1, \ldots, \delta_{k-1} = b_{k-1} \\
\leq\; & p
\end{aligned}$$

Thus, for $w \in V$, we may apply Lemma 2 on $I_w^<$ defined as follows:

$$I_w^< = \sum_{i=1}^{n} \min\left\{1, \frac{d(u_i, v_i)^\alpha}{d(u_i, w)^\alpha}\right\} Y_i^<$$

This random variable indicates the remaining measure of interference that is caused by these collisions. Setting $\delta = 2e - 1$ and $\mu = \frac{J}{8e}$ Lemma 2 states

$$\mathbf{Pr} I_w^< \geq \frac{J}{4} \leq 2^{-\frac{J}{4}} \leq n^{-1}$$

Now consider the situation after $k_0 + 2$ iterations of the inner *while* loop. Since these are independent repeats we have

$$\mathbf{Pr} I_w^< \geq \frac{J}{4} \leq n^{-(k_0+2)}$$

With a symmetric argument this also applies to $I_j^>$. For a sender that has not been successful we have $Z_j^< + Z_j^> \geq 1$. This means we have the bound $I_w^{\text{curr}} \leq I_w^< + I_w^>$. For the remaining measure of interference $I^{\text{curr}} = \max_{w\in V} I_w^{\text{curr}}$ we can conclude

$$\begin{aligned}\mathbf{Pr}I^{\text{curr}} \geq \frac{J}{2} &\leq \sum_{w\in V} \mathbf{Pr}I_w^{\text{curr}} \geq \frac{J}{2} \\ &\leq \sum_{w\in V} \mathbf{Pr}I_w^{<} \geq \frac{J}{4} \text{ or } I_w^{<} \geq \frac{J}{4} \\ &\leq n\left(n^{-(k_0+2)} + n^{-(k_0+2)}\right) \\ &\leq n^{-k_0}\end{aligned}$$

□

Using the previous lemma, we can bound the numbers of steps that are generated in the *while* loops.

Theorem 5 *Algorithm 2 generates a schedule of length at most $O(I + \log^2 n)$ steps whp.*

Proof Let T_k denote the number of scheduling steps generated in the kth execution of the outer *while* loop. As shown in the previous lemma, it holds that

$$\mathbf{Pr} v_k \geq (k_0 + 3)\, 16e\beta' \frac{1}{2^{k-1}} I \leq \frac{1}{n^{k_0+1}}$$

Let furthermore U denote the number of scheduling steps generated in the execution of Algorithm 1. As shown in Lemma 4, it holds that

$$\mathbf{Pr} U \geq (k_0 + 2)\, 4\beta' \ln n \log n \leq \frac{1}{n^{k_0+1}}$$

Thus the total number of steps generated in the *while* loops $\sum_k v_k + U$ can be estimated by

$$\begin{aligned}&\mathbf{Pr}\sum_k v_k + U \geq (k_0 + 3)\, 32e\beta' I + (k_0 + 2)\, 4\beta' \ln n \log n \\ &\leq \mathbf{Pr}\bigvee_k v_k \geq (k_0 + 3)\, 16e\beta' \frac{1}{2^{k-1}} I \vee U \geq (k_0 + 2)\, 4\beta' \ln n \log n \\ &\leq \sum_k \mathbf{Pr} v_k \geq (k_0 + 3)\, 16e\beta' \frac{1}{2^{k-1}} I + \mathbf{Pr} U \geq (k_0 + 2)\, 4\beta' \ln n \log n \\ &\leq \sum_k \frac{1}{n^{k_0+1}} + \frac{1}{n^{k_0+1}} \\ &\leq (\log n + 1)\, \frac{1}{n^{k_0+1}} \\ &\leq \frac{1}{n^{k_0}}\end{aligned}$$

This means the total number of steps upper bounded by

$$(k_0 + 3)\, 32e\beta' I + (k_0 + 2)\, 4\beta' \ln n \log n = O(I + \log^2 n)$$

with probability at least $1 - \frac{1}{n^{k_0}}$. □

In sufficiently dense instances, i.e., $I \geq \log^2 n$, this algorithm yields a constant-factor approximation for the optimal schedule compared to the linear power assignment with high probability. Compared to the optimal power assignment the approximation factor then is $O(\log \Delta \cdot \log n)$ whp for general metrics, respectively. $O(\log \Delta)$ for the two-dimensional Euclidean plane.

Algorithm 1 can be implemented in a distributed way losing a factor $\log n$ in the following way. In contrast to the centralized problem, the nodes do not know the correct value of I, thus, they do not know their transmission probability. Now in the distributed setting the algorithm processes in each *while* iteration $\log n$ steps, where in each of these steps the transmission probability is halfed, that is, starting by $1/2\beta'$ down to $1/2\beta' n$.

Algorithm 2 can be modified analogously, leading to a schedule of length $O(\log n \cdot (I + \log^2 n))$ whp.

2.4 Scheduling with the Square Root Power Assignment

The scheduling algorithms for the linear power assignment presented in Sect. 2.3 achieve an approximation factor of order $\log \Delta \text{polylog}\, n$ in comparison to an optimal solution with respect to general power assignments. In this section, we show that the square root power assignment admits schedules beating this bound achieving an approximation factor of order $\log \log \Delta \text{polylog}\, n$. Furthermore, we present a bidirectional variant of the interference scheduling problem in which the square root power assignment yields an approximation of order $\text{polylog}\, n$ and is, hence, independent of the aspect ratio.

2.4.1 Scheduling Directed Requests

In this section we show how to achieve an $O(\log \log \Delta \log^3 n)$ approximation on the interference scheduling problem using square root power. To prove this result we first show two properties that make use of the following definitions. We call a set R of requests *well separated*, if the length of any pair of requests differs by a factor of either at most 2 or at least $16n^{2/\alpha}$. We say that two requests (u_i, v_i) and (u_j, v_j) are τ*-close* under the square root power assignment if $\max\{\text{RI}_i(j), \text{RI}_j(i)\} \geq \tau$.

Lemma 4 *Let R be a well-separated SINR-feasible set of requests. Let (u_0, v_0) be a request that is shorter than the requests in R by at least a factor of $16n^{2/\alpha}$. If all the requests in R are $1/2n$-close to (u_0, v_0) under the square root power assignment, then $|R| = O(\log \log \Delta)$.*

Proof Let R' be a maximum 3^α-signal subset of R, let n' denote the number of requests in R' and w.l.o.g. let the requests in R' be labeled in increasing order of length. From Proposition 1 we know $\left|R'\right| = n' \geq |R|/9^\alpha$. As all the requests in R' are $1/2n$-close to (u_0, v_0), R' consists of two types of requests:

- Requests j for which the ratio between j's interference and the received signal from u_0 at receiver v_0 is at least $1/2n$ (or $\sqrt{d(u_0, v_0) \cdot d(u_j, v_j)}^\alpha \geq d(u_j, v_0)^\alpha \cdot \frac{1}{2n}$) and
- Requests j for which the ratio between u_0's interference and the received signal from j's sender at v_j is at least $1/2n$ (or $\sqrt{d(u_0, v_0) \cdot d(u_j, v_j)}^\alpha \geq d(u_0, v_j)^\alpha \cdot \frac{1}{2n}$).

We only consider the former type, the argument is almost identical for the latter type and will be left to the reader.

Let $j, j' \in R'$, w.l.o.g. assume $j \geq j'$. As they are $1/2n$-close to (u_0, v_0), it holds $\sqrt{d(u_0, v_0) \cdot d(u_j, v_j)}^\alpha \geq d(u_j, v_0)^\alpha \cdot \frac{1}{2n}$ (and analogously for j'). So we get

$$d(u_j, v_0) \leq \sqrt{d(u_0, v_0) \cdot d(u_j, v_j)}(2n)^{1/\alpha}$$

and

$$d(u_{j'}, v_0) \leq \sqrt{d(u_0, v_0) \cdot d(u_{j'}, v_{j'})}(2n)^{1/\alpha}$$

With triangle inequality we can conclude

$$\begin{aligned} d(u_{j'}, v_j) &\leq d(u_{j'}, v_i) + d(v_i, u_j) + d(u_j, v_j) \\ &\leq d(u_j, v_j) + 2^{1+1/\alpha} n^{1/\alpha} \sqrt{d(u_0, v_0) \cdot d(u_j, v_j)} \end{aligned}$$

Applying $\alpha \geq 1$ and $d(u_j, v_j) \geq 16n^{2/\alpha} d(u_0, v_0)$ to this inequality, we get

$$d(u_{j'}, v_j) \leq d(u_j, v_j) + 2^{1+1/\alpha} n^{1/\alpha} \sqrt{d(u_0, v_0) \cdot d(u_j, v_j)} \leq 2d(u_j, v_j)$$

For technical simplicity, we use the more relaxed $d(u_{j'}, v_j) < 3d(u_j, v_j)$ in the following. Using the same arguments as above we get

$$d(u_j, v_{j'}) \leq d(u_{j'}, v_{j'}) + 2^{1+1/\alpha} n^{1/\alpha} \sqrt{d(u_0, v_0) \cdot d(u_j, v_j)}$$

Multiplying this inequality with $d(u_{j'}, v_j) < 3d(u_j, v_j)$ it follows

$$(u_{j'}, v_j) \cdot d(u_j, v_{j'}) < 3d(u_j, v_j) d(u_{j'}, v_{j'}) + 12n^{1/\alpha} d(u_j, v_j) \sqrt{d(u_0, v_0) \cdot d(u_j, v_j)}$$

Since R' is a 3^α-signal set, we have $d(u_{j'}, v_j) \cdot d(u_j, v_{j'}) \geq 9d(u_j, v_j) \cdot d(u_{j'}, v_{j'})$. Again, applying the well separation, the last two inequalities yield (with canceling a $6d(u_j, v_j)$ factor)

$$d(u_{j'}, v_{j'}) < 2n^{1/\alpha}\sqrt{d(u_0, v_0) \cdot d(u_j, v_j)} \tag{2.6}$$

This equation implies $d(u_j, v_j) > 2d(u_{j'}, v_{j'})$. By well separation of R it follows $d(u_j, v_j) \geq 16n^{2/\alpha}d(u_{j'}, v_{j'})$. Now it follows from (2.6)

$$d(u_{i+1}, v_{i+1}) \geq \frac{d(u_i, v_i)^2}{4d(u_0, v_0)n^{2/\alpha}} \geq \frac{2d(u_i, v_i)^2}{d_1}$$

for any $i \in \{2, \ldots, n'\}$. Let $\lambda_i = d(u_i, v_i)/d(u_1, v_1)$. Then $\lambda_{i+1} \geq 2\lambda_i^2$ and by induction $\lambda_{n'} \geq 2^{2^{n'-1}-1}$. Hence, $n' = |R'| \leq \lg\lg \lambda_{n'} + 2 = \lg\lg \Delta + 2$, which proves the lemma. □

Lemma 5 *Let R be a well-separated set of requests. If any subset of R containing only nearly equilength requests can be scheduled with the linear power assignment using at most c colors, then all requests in R can be scheduled with* $O(c \log\log \Delta)$ *colors using the square root power assignment.*

Proof In the following we show that a single step from a schedule of R can be scheduled in $O(\log\log \Delta)$ steps. Let $R = R_1 \dot\cup R_2 \dot\cup \ldots \dot\cup R_t$ denote the decomposition of R in length groups, such that the length of the requests in each group differs by at most factor 2 and in different groups by at least factor $16n^{2/\alpha}$. First we transform the schedules for each length group in an r-signal schedule, with $r = 2^{\alpha/2}$. This changes the number of schedule steps by at most factor $(r + 1)^2$ (by Proposition 1). Let $T = \bigcup_i T_i$ be a single schedule step from the schedule of R and let T_i denote the requests in T from length group R_i. W.l.o.g., let the requests in T be ordered by decreasing length.

Lemma 4 states that for each request i there are at most $O(\log\log \Delta)$ longer requests in T that are $1/2n$-close to i. Let $p = O(\log\log \Delta)$ denote this bound. Now process the requests $i \in T$ by decreasing length: Assign i to a step T'_j with $j \in [p + 1]$ that does not contain a $1/2n$-close request for i.

It remains to show that this assignment yields a feasible schedule. Consider a request $i \in T'_j$ that originally came from set R_k. The relative interference on i from nearly equilength requests in $T'_j \cap R_k$ under the linear power assignment is at most $1/r$, since each length group is an r-signal set. We first analyze the influence from changing the power assignment from linear to square root in a length class. It holds for two requests a and b for the linear power assignment

$$\mathrm{RI}_a\left((u_b, v_b)\right) = c_a \cdot \frac{d(u_a, v_a)^\alpha}{p_a} \cdot \frac{p_b}{d(u_b, v_a)^\alpha} = c_a \cdot \frac{d(u_b, v_b)^\alpha}{d(u_b, v_a)^\alpha}$$

and for the square root power assignment

$$\mathrm{RI}_a\left((u_b, v_b)\right) = c_a \cdot \sqrt{d(u_a, v_a)^{\alpha}} \cdot \frac{\sqrt{d(u_b, v_b)}^{\alpha}}{d(u_b, v_a)^{\alpha}}$$

Since the requests in the same length class differ by at most factor 2 combining these two bounds yields that changing the power in a feasible schedule from the linear power assignment to the square root power assignment changes the relative interference by a factor of at most $2^{\alpha/2}$ in such nearly equilength request sets. Thus, the relative interference on i from requests in the same length class is at most $1/2$. On the other hand, the relative interference on i from requests not in the same length class is at most $1/2n$ each, by construction, which is at most $1/2$ in total. The relative interference on each link is not greater than one, which gives us an SINR-feasible schedule. □

Theorem 6 *Suppose there exists a ρ-approximate algorithm for the interference scheduling problem on nearly equilength request sets using uniform power assignment. Then there exists an $O(\rho \cdot \log\log\Delta \cdot \log n)$-approximate algorithm for the interference scheduling problem using the square root power assignment.*

Proof Let R be the set of requests. We partition R into $k = \left\lceil \frac{2}{\alpha} \log 16n \right\rceil$ well-separated sets as follows. Let $R_1, R_2, \ldots$ denote length groups with $R_i = \{j \in R \mid d(u_i, v_i) \in [2^{i-1}, 2^i)\}$. Then, partition R into classes $B_i = \cup_j R_{i+j\cdot k}$, for $i = 1, 2, \ldots, k$. Now the theorem follows from applying Lemma 5 on each of the classes B_i separately. □

Recall that Algorithm 2 had an approximation ratio of $O(\log\Delta \log^2 n)$ in general metrics. For nearly equilength request sets this ratio reduces to $O(\log^2 n)$, which gives the following result.

Corollary 1 *The interference scheduling problem in general metrics has an approximation factor of $O(\log\log\Delta \cdot \log^3 n)$ for the square root power assignment.*

For instances embedded in the Euclidean plane the approximation factor of Algorithm 2 is $O(\log\Delta \log n)$ which reduces to $O(\log n)$ for nearly equilength request sets.

Corollary 2 *For $\alpha > 2$, the interference scheduling problem in the two-dimensional Euclidean space has an approximation factor of $O(\log\log\Delta \cdot \log^2 n)$ for the square root power assignment.*

2.4.2 Scheduling Bidirectional Requests

Most communication protocols used in practice rely on bidirectional point-to-point communication. This is reflected by the following variant of the physical model in which requests are undirected, that is, each of the two nodes of a request acts as both sender and receiver. The SINR constraint is adapted as follows. For every request pair $(u_i, v_i) \in R$ and $w \in \{u_i, v_i\}$, it must hold

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

In every request set that fulfills this condition the two nodes of a request can exchange messages in both directions, as long as only one of them acts as sender at any given time.

In this setting, bounded, linear, and superlinear power functions still can have schedule lengths of $\Omega(n)$, while the optimal schedule has constant length. This can be shown by a straightforward adaption of the proof for Theorem 8. For sublinear assignments this adaption is not possible. In fact, we show in the following that the square root power assignment guarantees an approximation factor of $O(\log^3 n)$.

First, we need the following technical lemma.

Lemma 6 *Let (u_i, v_i) and (u_j, v_j) be two requests. If they can be scheduled simultaneously, then*

$$\min\{d(w_i, w_j)\}^2 \geq \beta^{2/\alpha} \cdot d(u_i, v_i) \cdot d(u_j, v_j)$$

Proof Let $w_1 \in \{u_i, v_i\}$ and $w_2 \in \{u_j, v_j\}$, such that $\min\{d(w_i, w_j)\} = d(w_1, w_2)$. The SINR constraint gives

$$\frac{p_i}{d(u_i, v_i)^\alpha} \leq \beta \frac{p_j}{d(w_1, w_2)^\alpha}$$

and

$$\frac{p_j}{d(u_j, v_j)^\alpha} \leq \beta \frac{p_i}{d(w_1, w_2)^\alpha}$$

From multiplying both equations follows

$$d(w_1, w_2)^2 \geq \beta^{2/\alpha} \cdot d(u_i, v_i) \cdot d(u_j, v_j)$$

□

Lemma 7 *Let R be a set of requests that can be scheduled with an arbitrary power assignment and let i be a request. Then there is at most a constant number of requests $j \in R$ with $d(u_j, v_j) \geq n^{2/\alpha} \cdot d(u_i, v_i)$ that cause a relative interference of at least $1/2n$ on i under the square root power assignment.*

Proof In the following we show that for fixed β there is at most one request $j \in R$ with $d(u_j, v_j) \geq n^{2/\alpha} \cdot d(u_i, v_i)$ that causes a relative interference of at least $1/2n$ on i under the square root power assignment. By Proposition 1 this yields the claimed result.

Assume that there are two requests $j, j' \in R$ with $d(u_j, v_j)$ and $d(u_{j'}, v_{j'})$ at least $n^{2/\alpha} \cdot d(u_i, v_i)$ that cause a relative interference of more than $1/2n$ on i under the square root power assignment. W.l.o.g, let $d(u_j, v_j) \geq d(u_{j'}, v_{j'})$. For $k \in \{j, j'\}$ and $w \in \{u_i, v_i\}$ let $d_m = \min\{d(u_k, w), d(v_k, w)\}$. The relative interference under the square root power assignment implies

$$\left(\frac{\sqrt{d(u_k, v_k)d(u_i, v_i)}}{d_m}\right)^{\alpha} \geq \frac{1}{2n}$$

This implies

$$d_m \leq (2n)^{1/\alpha}\sqrt{d(u_k, v_k) \cdot d(u_i, v_i)} \leq (2n)^{1/\alpha}\sqrt{d(u_j, v_j) \cdot d(u_i, v_i)}$$

To avoid notational clutter, let $d(u_j, v_{j'})$ be the minimal distance between j and j'. Applying triangle inequality we get

$$\begin{aligned} d(u_j, v_{j'}) &\leq 2d_m \leq 2(2n)^{1/\alpha}\sqrt{d(u_i, v_i) \cdot d(u_j, v_j)} \\ &\leq 2(2n)^{1/\alpha}\sqrt{\frac{d(u_{j'}, v_{j'})}{n^{2/\alpha}} \cdot d(u_j, v_j)} \leq 2^{1+1/\alpha}\sqrt{d(u_j, v_j) \cdot d(u_{j'}, v_{j'})} \end{aligned}$$

Thus

$$d(u_j, v_{j'})^2 \leq \left(2^{\alpha+1}\right)^{2/\alpha} d(u_j, v_j) \cdot d(u_{j'}, v_{j'})$$

From Lemma 6 follows for $\beta < 2^{\alpha+1}$ there is at most one request $j \in R$ with $d(u_j, v_j) \geq n^{2/\alpha} \cdot d(u_i, v_i)$ that causes a relative interference of at least $1/2n$ on i under the square root power assignment. □

We now can use an almost identical approach like shown in Lemma 5 and Theorem 6 for the unidirectional case.

Lemma 8 *Let R be a request set where the length of every pair of links differs by at most factor 2 or at least $n^{2/\alpha}$. If any subset of R containing only nearly equilength requests can be scheduled with the linear power assignment using at most c colors, then all requests in R can be scheduled with $O(c)$ colors.*

Theorem 7 *Suppose there exists a ρ-approximate algorithm for the bidirectional interference scheduling problem on equilength requests. Then there exists an algorithm for the bidirectional interference scheduling problem with approximation factor $O(\rho \log n)$ for the square root power assignment.*

We omit the proofs for these lemmas, as the arguments are analogous to the unidirectional case. For scheduling the equilength request sets we again can use Algorithm 2.

Corollary 3 *The bidirectional interference scheduling problem in general metrics has an approximation factor of* $O(\log^3 n)$ *for the square root power assignment.*

Corollary 4 *For* $\alpha > 2$*, the bidirectional interference scheduling problem in the two-dimensional Euclidean space has an approximation factor of* $O(\log^2 n)$ *for the square root power assignment.*

2.5 The Gap of Oblivious Power Schemes

Our upper bounds on the approximation ratios of oblivious scheduling algorithms for directed communication requests depend on the aspect ratio. In this section, we show that the dependence on the aspect ratio is unavoidable. To prove this we construct a family of instances for a given oblivious power assignment 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 8 *Let* $f : \mathbb{R}_{>0} \to \mathbb{R}_{>0}$ *be any oblivious power assignment function. 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 We distinguish three cases. In the first case, we assume that f is asymptotically unbounded, that is, for every $c > 0$ and every $x_0 > 0$ there exists a value $x > x_0$ with $f(x) > c$.

We consider the following family of instances. They consist of n pairs (u_i, v_i) on a line, 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 that is large enough to get along with different values of β.

Formally, this kind of instance can be defined by $u_1, v_1, \ldots, u_n, v_n \in \mathbb{R}$ such that

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

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 $x_1, \ldots, x_{i-1}$ and y_i, we choose x_i such that $x_i \geq y_i$ and

$$f(x_i) \geq y_i^{\alpha} \frac{f(x_j)}{x_j^{\alpha}} \qquad \text{for all } j < i$$

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

$$d(u_i, v_k) = \sum_{j=k+1}^{i-1} x_j + \sum_{j=k+1}^{i} \chi \cdot y_j \le 2\chi \sum_{j=k}^{i} y_j \le 2\chi \sum_{j=k}^{i} \frac{1}{2^{i-j}} y_i \le 4\chi y_i$$

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

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

Putting these facts together

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

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

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 $y_i \le x_i$ and $y_{i+1} \ge 2x_i$. Thus for any link k the interference by the ones with higher index as well as the ones with lower index forms 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.

In the second case, we assume that f is asymptotically bounded from above by some value $c > 0$ but does not converge to 0. In this case, there exists a value $b \in (0, c]$ such that for every $x_0 > 0$ there exists a value $x > x_0$ with $f(x) \in [b, 2b]$. Let $\chi > 1$ be a suitable constant. We choose n numbers $x_1, \ldots, x_n$ satisfying the properties (a) $f(x_i) \in [b, 2b]$, for $1 \le i \le n$, and (b) $x_i \ge \chi x_{i-1}$, for $2 \le i \le n$. We set $u_i = -x_i/2$ and $v_i = x_i/2$. This instance corresponds to nested pairs on the line, whereas the power assignment is similar to the uniform power assignment, which already indicates the desired result.

To be more precise, let $S \subseteq [n]$ be a set of indices of requests that can be scheduled together and let $k = \max S$. For $i \in S$ it holds $d(u_i, v_k) \le x_k/2$. The SINR condition for k states

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

This yields

$$\frac{1}{\beta} \cdot \frac{f(x_k)}{x_k^\alpha} \geq \sum_{i \in S \setminus \{k\}} \frac{p_i}{d(u_i, v_k)^\alpha} \geq \sum_{i \in S \setminus \{k\}} \frac{b}{(x_k/2)^\alpha} = (|S| - 1) \cdot \frac{2^\alpha b}{x_k^\alpha}$$

Since $2b \geq f(x_k)$, we have $|S| \leq 1/\beta \cdot 2^{1-\alpha} + 1$. It follows again that at least $\Omega(n)$ colors are needed to schedule these instances using $p_i = f(d(u_i, v_i))$.

In contrast, if χ is chosen sufficiently large than the square root power assignment can schedule all these requests simultaneously.

Finally, in the third case, $\lim f(x) = 0$, we again construct a sequence of nested pairs analogously to second case but replacing condition (a) by the condition $f(x_i) \leq f(x_{i-1})$. Analogously to the second case, the power assignment defined by f allows only for scheduling a constant number of pairs simultaneously while the square root assignment can schedule all pairs simultaneously. □

The last result shows that the dependence on Δ is necessary for nontrivial results. The following theorem shows that there is a gap of at least $\Omega\left(\sqrt{\log\log \Delta}\right)$ between oblivious and optimal power assignments.

Theorem 9 *An instance of the interference scheduling problem exists such that every schedule using an oblivious power function needs at least* $\Omega\left(\sqrt{\log\log \Delta}\right)$ *more steps than the optimal schedule.*

Proof In this proof we construct an instance that can be scheduled in a constant number of rounds by a non-oblivious power assignment, but every oblivious power assignment needs at least $\Omega\left(\sqrt{\log\log \Delta}\right)$ steps. The instance consist of two nearly identical requests sets, only the role of sender and receiver in each request is exchanged. More formally, let $x_1 = 1$, $y_i = x_i^2$, and $x_{i+1} = 2y_i$ for every $i \in [n]$. Let the request set R_1 consist of the requests (u_i, v_i) described by

$$u_i = \begin{cases} 0 & \text{if } i = 1 \\ -\sum_{j=2}^{i} x_j & \text{otherwise} \end{cases} \qquad \text{and} \qquad v_i = \sum_{j=1}^{i} y_i$$

and let R_2 consist of requests $\left(u_i', v_i'\right)$ with

$$u_i' = M + \sum_{j=1}^{i} y_i \qquad \text{and} \qquad v_i' = \begin{cases} M & \text{if } i = 1 \\ M - \sum_{j=2}^{i} x_j & \text{otherwise} \end{cases}$$

where M denotes a constant large enough that interferences between requests from R_1 and R_2 become negligible. Since for all $i \in [n]$ holds $d(u_i, v_i) = d\left(u_i', v_i'\right)$, every oblivious power assignment uses the same power p_i for request (u_i, v_i) and $\left(u_i', v_i'\right)$.

Let T denote the schedule under an arbitrary, fixed oblivious power assignment. In this schedule there must be a step where at least n/T requests from R_1 are scheduled. Let $M \subseteq [n]$ denote their indices. Let $i, j \in M$ with $i < j$. The SINR constraint states

$$\beta \frac{p_i}{d(u_i, v_j)^\alpha} \leq \frac{p_j}{d(u_j, v_j)^\alpha}$$

Using $d(u_i, v_j) \leq x_j$ and $d(u_j, v_j) \geq y_j = x_j^2$ we get

$$\beta \frac{p_i}{x_j^\alpha} \leq \frac{p_j}{x_j^{2\alpha}}$$

which implies $p_i \leq p_j/\beta x_j^\alpha$. With $d\left(u'_j, v'_i\right) \leq 2x_j$, the interference from $\left(u'_j, v'_j\right)$ on $\left(u'_i, v'_i\right)$ is

$$\beta \frac{p_j}{d\left(u'_j, v'_i\right)^\alpha} \geq \beta \frac{p_j}{(2x_j)^\alpha} \geq \frac{\beta^2 p_i}{2^\alpha} > \frac{p_i}{d\left(u'_i, v'_i\right)}$$

Thus, for every $i \neq j, i, j \in M$, the requests $\left(u'_i, v'_i\right)$ and $\left(u'_j, v'_j\right)$ cannot be scheduled in the same step. In fact, for every $i \in M$, $\left(u'_i, v'_i\right)$ must be assigned to a different schedule step. This yields $T \geq |M|$ and it follows $T \geq \sqrt{n} = \sqrt{\Omega(\log \log \Delta)}$. □

2.6 Summary and Open Problems

We have studied the interference scheduling problem with a focus on oblivious power assignments, i.e., the power for a signal is defined as a function of the path loss. Examples of such power assignments are the uniform, the linear, and the square root power assignment. The major advantage of these power assignments is their simplicity. In particular, they can be computed for every request without taking into account other requests. In our study we investigated the approximation factors with respect to the schedule length that can be achieved with oblivious power assignments.

The linear power assignment is of special interest as it is energy efficient in the sense that signals are sent at a power level that is only a constant factor larger than the power level needed to drown out ambient noise. In Sect. 2.3, we presented lower and upper bounds for the linear power assignment from [5]. The key to both the lower and upper bounds is the measure of interference I. On the one hand, we have shown that $\Omega(I)$ is a lower bound on the schedule length when using linear power assignments. On the other hand, we have presented distributed scheduling algorithms for the linear power assignment computing schedules of length $O(I \log n)$ and $O(I + \log^2 n)$, respectively. For dense instances this gives a constant factor approximation of the optimal schedule for linear power assignment.

Similar results have been achieved recently for the uniform power assignment. In [6] it is presented an algorithm that achieves a constant factor approximation guarantee with respect to the number of requests that can be scheduled simultaneously. A straight forward extension of this approach yields an approximation factor of $O(\log n)$ with respect to the schedule length for the uniform power assignment.

How do these results compare to the schedule length for general power assignments? – In Sect. 2.3, we show a lower bound of $\Omega(I/\log \Delta \log n)$ for schedules with general power assignments, where Δ denotes the aspect ratio of the metric. When restricting to the two-dimensional Euclidean space the bound improves to $\Omega(I/\log \Delta)$. Thus, the best known scheduling algorithms for the linear and the uniform power assignments achieve approximation ratios of order $\log \Delta \text{polylog} n$ in comparison to the optimal power assignment.

In Sect. 2.4, we present an analysis showing that the square root power assignment can achieve significantly better approximation ratio in terms of the aspect ratio than the linear and the uniform power assignment: For directed communication requests the approximation ratio of the square root power assignment is of order $O(\log \Delta \text{polylog}\, n)$ and for bidirectional requests even of order only $O(\text{polylog}\, n)$. Both of these ratios compare the schedule length of the square root power assignment with the schedule length for general power assignments. The result for directed communication requests is from [9] and the result for bidirectional requests was first shown in [5] and then improved in [9].

Finally, in Sect. 2.5 we study lower bounds for oblivious power assignments. We show that the dependence on the aspect ratio cannot be avoided for directed communication requests and present a lower bound of order $\Omega\left(\sqrt{\log \log \Delta}\right)$ on the approximation ratio holding for every oblivious power assignment. In particular, one cannot achieve approximation factors better than $\Omega(n)$ for directed communication requests with unbounded aspect ratio when restricting to oblivious power assignments.

We want to conclude with the major open problems about interference scheduling in the physical model: Devise a polynomial time constant factor approximation algorithm or approximation scheme for the interference scheduling problem with general power assignments or show that such an approximation is not possible. Present improved distributed algorithms beating the currently best known approximation ratios for oblivious power assignments.

References

1. M. Andrews and M. Dinitz. Maximizing capacity in arbitrary wireless networks in the SINR model: Comple xity and game theory. In *Proceedings of the 28th Conference of the IEEE Communications Society (INFOCOM)*, Rio de Janeiro, Brazil, 2009.

2. C. Avin, Z. Lotker, and Y. A. Pignolet. On the power of uniform power: Capacity of wireless networks with bounded resources. In *Proceedings of the 17th Annual European Symposium on Algorithms (ESA)*, Copenhagen, Denmark, 2009.
3. D. Chafekar, V. S. Anil Kumar, M. V. Marathe, S. Parthasarathy, and A. Srinivasan. Cross-layer latency minimization in wireless networks with SINR con. In *Proceedings of the 8th ACM International Symposium Mobile Ad-Hoc Networking and Computing (MOBIHOC)*, pages 110–119, Montreal, Quebec, Canada, 2007.
4. D. Chafekar, V. S. Anil Kumar, M. V. Marathe, S. Parthasarathy, and A. 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, Phoenix, AZ, USA, 2008.
5. A. Fanghänel, T. Kesselheim, H. Räcke, and B. Vöcking. Oblivious interference scheduling. In *Proceedings of the 28th Annual ACM Symposium on Principles of Distributed Computing (PODC)*, Calgary, Alberta, Canada, 2009.
6. O. Goussevskaia, M. M. Halldórsson, R. Wattenhofer, and E. Welzl. Capacity of arbitrary wireless networks. In *Proceedings of the 28th Conference of the IEEE Communications Society (INFOCOM)*, Rio de Janeiro, Brazil, 2009.
7. O. Goussevskaia, Y. A. Oswald, and R. Wattenhofer. Complexity in geometric SINR. In *Proceedings of the 8th ACM International Symposium Mobile Ad-Hoc Networking and Computing (MOBIHOC)*, pages 100–109, Montreal, Quebec, Canada, 2007.
8. P. Gupta and P. R. Kumar. Critical power for asymptotic connectivity in wireless networks. In W. M. McEneaney, G. Yin, and Q. 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.
9. M. M. Halldórsson. Wireless scheduling with power control. In *Proceedings of the 17th Annual European Symposium on Algorithms (ESA)*, Copenhagen, Denmark, 2009.
10. M. M. Halldórsson and R. Wattenhofer. Computing Wireless Capacity. In press.
11. V. S. Anil Kumar, M. V. Marathe, S. Thite, H. Balakrishnan, C. L. Barrett. 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.
12. R. Hekmat and P. Van Mieghem. Interference in wireless multi-hop ad-hoc networks and its effect on network capacity. *Wireless Networks*, 10(4):389–399, 2004.
13. S. O. Krumke, M. V. Marathe, and S. S. Ravi. Models and approximation algorithms for channel assignment in radio networks. *Wireless Networks*, 7(6):575–584, 2001.
14. T. Moscibroda and R. Wattenhofer. The complexity of connectivity in wireless networks. In *Proceedings of the 25th Conference of the IEEE Communications Society (INFOCOM)*, pages 1–13, Barcelona, Catalunya, Spain, 2006.
15. T. Moscibroda, R. Wattenhofer, and A. 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, Florence, Italy, 2006.
16. P. Raghavan. Probabilistic construction of deterministic algorithms: approximating packing integer programs. *Journal of Computer and System Sciences*, 37(2):130–143, 1988.
17. S. Ramanathan and E. L. Lloyd. Scheduling algorithms for multi-hop radio networks. *ACM SIGCOMM Computer Communication Review*, 22(4):211–222, 1992.
18. S. 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.