

Online Capacity Maximization in Wireless Networks ^{*}

Alexander Fanghänel [†] Sascha Geulen [‡] Martin Hoefer [§] Berthold Vöcking [¶]

Abstract

In this paper we study a dynamic version of capacity maximization in the physical model of wireless communication. In our model, requests for connections between pairs of points in Euclidean space of constant dimension d arrive iteratively over time. When a new request arrives, an online algorithm needs to decide whether or not to accept the request and to assign one out of k channels and a transmission power to the request. Accepted requests must satisfy constraints on the signal-to-interference-plus-noise (SINR) ratio. The objective is to maximize the number of accepted requests.

Using competitive analysis we study algorithms using distance-based power assignments, for which the power of a request relies only on the distance between the points. Such assignments are inherently local and particularly useful in distributed settings. We first focus on the case of a single channel. For request sets with spatial lengths in $[1, \Delta]$ and duration in $[1, \Gamma]$ we derive a lower bound of $\Omega(\Gamma \cdot \Delta^{d/2})$ on the competitive ratio of any deterministic online algorithm using a distance-based power assignment. Our main result is a near-optimal deterministic algorithm that is $O(\Gamma \cdot \Delta^{(d/2)+\varepsilon})$ -competitive, for any constant $\varepsilon > 0$.

Our algorithm for a single channel can be generalized to k channels. It can be adjusted to yield a competitive ratio of $O(k \cdot \Gamma^{1/k'} \cdot \Delta^{(d/2k'')+\varepsilon})$ for any factorization (k', k'') such that $k' \cdot k'' = k$. This illustrates the effectiveness of multiple channels when dealing with unknown request sequences. In particular, for $\Theta(\log \Gamma \cdot \log \Delta)$ channels this yields an $O(\log \Gamma \cdot \log \Delta)$ -competitive algorithm. Additionally, we show how this approach can be turned into a randomized algorithm, which is $O(\log \Gamma \cdot \log \Delta)$ -competitive even for a single channel.

1 Introduction

Determining the capacity of wireless networks is a major challenge in networking. Most studies in this area rely on the *physical model* taking into account that the strength of a signal fades with the distance from the sender. A node can successfully receive a signal if the *signal to interference plus noise ratio (SINR)* is above some threshold, that is, if the signal's strength is sufficiently large in comparison to the sum of other signals received simultaneously plus ambient noise.

Only very recently, we have seen significant progress in understanding the algorithmic aspects of the scheduling problems arising in the physical model [3–5, 10, 11, 13, 17]. Previous work focuses on offline optimization problems of the following kind. Suppose one is given a set of n requests for connections between pairs of points in Euclidean space of constant dimension d . One has to

^{*}This work has been supported by DFG through UMIC Research Centre at RWTH Aachen University and grant Ho 3831/3-1.

[†]Dept. of Computer Science, RWTH Aachen University, Germany. fanghaenel@cs.rwth-aachen.de.

[‡]Dept. of Computer Science, RWTH Aachen University, Germany. sgeulen@cs.rwth-aachen.de.

[§]Dept. of Computer Science, RWTH Aachen University, Germany. mhoefer@cs.rwth-aachen.de.

[¶]Dept. of Computer Science, RWTH Aachen University, Germany. voecking@cs.rwth-aachen.de.

specify a subset of requests and a power assignment to each pair such that the requests can be scheduled simultaneously, that is, the chosen requests with the power assignments satisfy the SINR constraint. The objective is to maximize the number of chosen requests. This variant is sometimes referred to the *throughput version of the capacity maximization problem*. A related problem is to minimize the number of batches such that the requests in each batch satisfy the SINR constraint. In this case batches can be mapped to orthogonal channels or time slots.

Most of the previous work focuses on power assignments that are *distance-based*, i. e., the power assignment is a function of the distance between the two nodes of a request. Prominent examples are the linear power assignment in which the power is chosen proportional to the loss in power between the nodes of a request (and, hence, depends polynomially on the distance) and the uniform power assignment in which all requests get assigned the same power. Such assignments are inherently local and, hence, particularly useful in distributed settings. The linear power assignment has the additional advantage of being energy-minimal.

The best known offline results for the uniform power assignment are achieved in [10] and [14]. They present algorithms that achieve $O(1)$ approximation guarantees with respect to the number of requests that can be scheduled simultaneously when restricting to the uniform power assignment. This approach can easily be extended to obtain an $O(\log n)$ approximation ratio on the number of batches for the uniform power assignment [10]. Similarly, in [9] an algorithm is presented that achieves an $O(\log n)$ approximation on the number of batches when restricting to linear power assignments. Very recently, Kesselheim [15] presented a constant factor approximation algorithm for maximizing the number of requests using arbitrary power assignments.

Let us remark that offline approximation ratios restricted to uniform or linear power assignments can often be translated into approximation ratios with respect to general power assignments by spending an additional factor of order $\log \Delta$, where Δ is the ratio between the largest and the smallest distance among all request pairs. To see this, consider an optimum solution with general powers and cluster the accepted requests into $O(\log \Delta)$ classes such that in class i requests have lengths in $[2^{i-1}, 2^i)$. For the restricted instance of all requests with lengths in $[2^{i-1}, 2^i)$ the optimum solution using uniform power assignment yields at least a constant fraction of the size of class i if requests are located in the Euclidean plane. This can be derived from, e.g., results in [1] and shows the result for uniform powers. As uniform and linear power assignments differ only by a constant factor if request lengths are within a factor of 2, the same result holds for linear assignments using, e.g., results of [3]. More generally, using our insights below one can easily extend this argument to hold for arbitrary fading metrics (for a definition see below).

Consequently, the algorithms from [10] and [14] using uniform power assignments achieve an $O(\log \Delta)$ approximation ratio in comparison to general power assignments. The same approximation ratio has been achieved independently in [1]. It follows from the analysis presented in [8,16] that this approximation ratio is best possible for algorithms using uniform or linear power assignments.

A drawback of the previous work is that it neglects the dynamic nature of request scheduling in wireless networks. The focus of our paper lies exactly on this aspect. We study request scheduling in wireless networks as an online problem, that is, requests arrive one by one. When a new request arrives, an online algorithm needs to decide whether or not to accept the request and to assign a power rate. In the multi-channel version, accepted requests must also be assigned to one out of k available channels. Decisions about acceptance as well as power and channel assignments cannot be revoked later.

1.1 Online Request Scheduling

In our online model we receive an unknown number of n *communication requests* sequentially over time. Each request $1 \leq i \leq n$ consists of a point pair. For a *directed* request there is a *sender* s_i and a *receiver* r_i that strive to establish an uninterrupted connection. For *undirected* requests, both points are receiver and sender at the same time. In this paper we consider sets of directed and undirected requests, as well as mixed sets of requests. We assume that points come from a metric space with a distance function $d(x, y)$. We use short notation for $d_{ii} = d(s_i, r_i)$, the *distance* between sender s_i and receiver r_i . We also refer to d_{ii} as the *length* of request i . More generally, for two different directed requests we use $d_{ij} = d(s_i, r_j)$. We denote by $\Delta = (\max_i d_{ii})/(\min_i d_{ii})$ the so called *aspect ratio*. Further, each request pair i comes with a parameter t_i , which denotes the duration of the request. We denote by $\Gamma = (\max_i t_i)/(\min_i t_i)$, where w.l.o.g. we let $\min_i t_i = 1$ and $\max_i t_i = \Gamma$. We assume requests lie in \mathbb{R}^d of constant dimension d , and the distance function is an l_p -norm or the l_{max} -norm. Finally, we show how to extend most of our results to doubling metrics [2, 6].

Requests arrive sequentially over time and are assumed to be characterized by the physical model [12]. The goal is to accept the maximum number of requests that can successfully communicate simultaneously. For each request an online algorithm must make a decision whether to accept the request or not. For an accepted request i it needs to set a *power level* p_i and a *channel* $k_i \in \{1, \dots, k\}$ for the sender s_i to emit a signal. For undirected requests we assume that both points emit signals with the same power and on the same channel. The algorithm iteratively expands the sets S_1, \dots, S_k of accepted requests on the corresponding channels. Decisions on acceptance, power levels, and channels of a request cannot be revoked later on. If a request is accepted, the algorithm must ensure that it remains successful throughout the time. The criterion of “successful” for an accepted directed request i is the following *SINR constraint*:

$$\frac{p_i}{d_{ii}^\alpha} \geq \beta \left(\sum_{j \in S_{k_i}, j \neq i} \frac{p_j}{d_{ji}^\alpha} + N_{k_i} \right). \quad (1)$$

This constraint is the central condition for successful communication in the physical model. It characterizes the strength at r_i of the signal emitted by s_i compared to *ambient noise* N_{k_i} and the *interference* from signals of all other senders *on the same channel* k_i . In this expression α is the *path loss exponent* that characterizes the decay of a signal over a distance. In this paper we consider a Euclidean *fading metric* [13], i. e., we require that $\alpha > d$, where we treat both α and d as constants. The constant β is called the *gain*.

For a successful undirected request the SINR constraint has to be satisfied at both points of the pair. Similarly, when considering another receiver i , both points of j are senders using power p_j and create interference. For notational simplicity, however, we will treat them as two directed requests in the right-hand side of (1). An online algorithm has to ensure that (1) is satisfied for all $i \in S = S_1 \cup \dots \cup S_k$ throughout.

For simplicity we assume that noise is absent, $N_1 = \dots = N_k = 0$. Our algorithms will satisfy the SINR constraint with strict inequality. This allows to scale powers up sufficiently to satisfy the constraints also when there is noise. Clearly, such a scaling might be wasteful or infeasible in practice, but this aspect is beyond our analysis. When there is no noise, we can scale all distances such that $\min_i d_{ii} = 1$ and $\max_i d_{ii} = \Delta$.

In this paper we are particularly interested in distance-based power assignments because of their simplicity and locality, which is a striking conceptual advantage in distributed wireless systems. A

distance-based power assignment p is given by $p_i = \phi(d_{ii})$ with a function $\phi : [1, \Delta] \rightarrow (0, \infty)$. For uniqueness we assume ϕ is always scaled such that $\phi(1) = 1$. Examples are *uniform* $\phi(d_{ii}) = 1$ or *linear* $\phi(d_{ii}) = d_{ii}^\alpha$ assignments. Recently, a *square-root* assignment $\phi(d_{ii}) = d_{ii}^{\alpha/2}$ has attracted some interest [8,13] as it yields better approximation ratios for the offline version of request scheduling than uniform and linear power assignments. We generalize these three classes to *polynomial* assignments of the form $\phi(d_{ii}) = d_{ii}^{r\alpha}$ with parameter $r \in \mathbb{R}$.

For the analysis of our online algorithms we make use of the following definitions. Let $A(\omega)$ denote the number of request pairs an online algorithm A accepts, and let $\text{OPT}(\omega)$ denote the number of requests in an optimal offline solution on an input sequence ω . An online algorithm is *c-competitive* (or “yields competitive ratio c ”) if there exists a constant a , such that for every input ω we have $A(\omega) \geq (\text{OPT}(\omega)/c) + a$. We call algorithm A *strictly c-competitive* if it is c -competitive with $a = 0$. Note that all algorithms presented in this paper are strictly competitive. For the lower bounds we do not need to rely on strictness.

1.2 Our Results

Our first contribution are lower bounds for deterministic online algorithms choosing requests for a single channel. We show that any deterministic online algorithm using a polynomial power assignment with parameter r cannot yield a competitive ratio better than $\Omega(\Gamma \cdot \Delta^{d \cdot \max\{r, 1-r\}})$. For uniform and linear power assignments, this results in a lower bound of $\Omega(\Gamma \cdot \Delta^d)$; for the square root power assignment, it yields a lower bound of $\Omega(\Gamma \cdot \Delta^{d/2})$. In fact, we can show that the $\Omega(\Gamma \cdot \Delta^{d/2})$ lower bound on the competitive ratio is not restricted to polynomial power assignments: In the case of directed requests, this bound holds for any distance-based power assignment and, in the case of undirected requests, the same bound holds even for general power assignments.

Our lower bounds reveal an exponential gap between the approximation guarantees achievable by deterministic online and offline algorithms. The main difficulty of the online scenario turns out to be that requests cannot be ordered by length. This has been a crucial ingredient to all existing deterministic offline algorithms with polylogarithmic approximation guarantee [1, 10, 13].

Our second contribution is a deterministic online algorithm for a single channel that almost matches the lower bounds. All following results hold for directed and undirected requests. Algorithm SAFE-DISTANCE works for polynomial power assignments with $r \in [0, 1]$. For uniform and linear power assignments, it achieves a competitive ratio of $O(\Gamma \cdot \Delta^d)$. For the square-root power assignment, we extend the basic idea and obtain algorithm MULTI-CLASS SAFE-DISTANCE. For any constant $\epsilon > 0$ it achieves a competitive ratio of $O(\Gamma \cdot \Delta^{d/2+\epsilon})$.

Let us explicitly point out that these competitive ratios compare the performance of online algorithms with polynomial power assignments to optimal offline algorithms with general power assignments. Combining the upper bound for the square root power assignment with the lower bounds above shows that this power assignment achieves nearly the best possible competitive ratio among all (distance-based) power assignments (in case of directed requests) and is superior to any other polynomial power assignment.

Our third contribution is an illustration of the power of multiple channels for deterministic online algorithms. We generalize algorithm MULTI-CLASS SAFE-DISTANCE and its analysis from 1 to k channels and achieve an exponential reduction in the competitive ratio. We prove that algorithm MULTI-CLASS SAFE-DISTANCE using $k = k' \cdot k''$ channels is only $O(k \cdot \Gamma^{1/k'} \cdot \Delta^{(d/2k'')+\epsilon})$ -competitive. In particular, with just a logarithmic number of channels we obtain a deterministic algorithm with logarithmic competitive ratio. This algorithm is only constant-competitive against

an optimum solution that uses only one channel. By randomly choosing a channel, we thus obtain a randomized algorithm for a single channel that is $O(\log \Gamma \cdot \log \Delta)$ -competitive with respect to the expected number of accepted requests.

Finally, we show the robustness of our results by extending all upper bounds from Euclidean to doubling metrics. This allows to introduce features such as obstacles in our model, which locally disturb Euclidean distances but do not affect the global structure of the metric.

Outline For technical reasons, we present our results in a different order than listed above. In Section 2 we first analyze algorithm SAFE-DISTANCE before stating the general lower bound in Theorem 2.5. In Section 3 we give the near-optimal algorithm MULTI-CLASS SAFE-DISTANCE (Section 3.1), the generalization to k channels (Section 3.2) and the randomized algorithm (Section 3.3). In Section 4 we reach the full level of generality by describing the adjustments to requests with duration (Section 4.1) and to doubling metrics (Section 4.2). Finally, we conclude the paper in Section 5 with some open problems.

2 A Simple Algorithm and a Lower Bound

In the following we first analyze the spatial aspect of the problem and assume that requests last forever, i. e., for all requests i , $t_i = \infty$. We begin by analyzing a simple online algorithm for the case of a single channel and any polynomial power assignment. Subsequently, we show a general lower bound. Our analysis of the online algorithm introduces a number of critical observations that we use in later sections.

The main idea of the algorithm is to accept a new request only if it keeps a *safe distance* σ from every other previously accepted request. In particular, we accept incoming request i only if $\min\{d_{ij}, d_{ji}\} \geq \sigma$ for every other previously accepted request $j \in S$. We call this algorithm SAFE-DISTANCE. For the choice of σ there is a conflict between correctness and competitive ratio. A larger σ blocks out a larger portion of the space, in which an optimal algorithm knowing the request sequence might be able to accept requests. If σ is too small, then at some point the interference at an accepted request can get too large and the SINR constraint becomes violated.

We strive to choose σ as small as possible to ensure correctness of SAFE-DISTANCE. To bound the interference at accepted requests we construct a worst-case scenario. We consider a receiver r_i from a single accepted request and bound the maximum number of senders that can be at a certain distance from r_i . In the following we show that for $r \in [0, 1]$ the choice of

$$\sigma = \max \left\{ 2\Delta, \Delta \cdot 18d \cdot \sqrt[3]{2\beta/(\alpha - d)} \right\}$$

is sufficient to yield the following result.

Theorem 2.1. *SAFE-DISTANCE is $O(\Delta^d)$ -competitive for any polynomial power assignment with $r \in [0, 1]$ and a single channel.*

Proof. We first show that SAFE-DISTANCE is correct, i. e., for an accepted request i the SINR constraint of i never becomes violated. In particular, we will underestimate the distances of accepted senders of other requests to overestimate the interference at receiver r_i . However, even under such pessimistic conditions the SINR constraint at r_i will remain valid.

Consider a receiver r_i of an accepted request i . To estimate the interference at r_i we have to count how many senders may be placed at which distance. Using $\sigma \geq 2\Delta$ and the choice rule of

the algorithm it is straightforward to verify that senders of any two different accepted requests are at least a distance of $\sigma - \Delta \geq \sigma/2$ apart. We segment all of \mathbb{R}^d into d -dimensional hypercubes with length $\sigma/3d$, which we call *sectors*. The greatest distance within a sector is $\sigma d/3d = \sigma/3 < \sigma/2$. Each sector can contain senders from at most one request, so there are at most 2 senders in every sector. Without loss of generality, we assume that sectors are created such that r_i lies in a corner point of 2^d sectors. We divide the set of sectors into *layers*. The first layer consists of the 2^d sectors incident to r_i . The second layer are all sectors that are not in the first layer but share at least a point with a sector from the first layer, and so on. In this construction there are exactly $(2\ell)^d$ sectors from layers 1 through ℓ , and their union is a hypercube of side length $2\ell\sigma/3d$ with r_i in the center. Therefore, there are exactly $2^d(\ell^d - (\ell - 1)^d)$ sectors in layer ℓ .

Due to the algorithm there can be no sender at a distance smaller than σ from r_i . The sector of smallest layer that is at a distance at least σ from r_i can be reached along the volume diagonal of the layer hypercubes. There can be no sender in all sectors from layers 1 through ℓ' , where ℓ' is bounded by $\sigma \leq \ell'(\sigma/3)$, which yields $\ell' \geq 3$. For bounding the interference assume that in all sectors of layer $\ell \geq 3$ there are 2 senders. Note that all senders in sectors from a layer ℓ have a distance at least $(\ell - 1)\sigma/3d$ to r_i . To bound the interference that is created at r_i , we use the following technical lemma.

Lemma 2.2. *For $\alpha > d \geq 1$ it holds that*

$$2^d \cdot \sum_{\ell=3}^{\infty} \frac{\ell^d - (\ell - 1)^d}{(\ell - 1)^\alpha} < \frac{6^d}{\alpha - d} .$$

Proof. We observe

$$\begin{aligned} & 2^d \cdot \sum_{\ell=3}^{\infty} \frac{\ell^d - (\ell - 1)^d}{(\ell - 1)^\alpha} \\ & \leq 2^d \cdot \sum_{\ell=3}^{\infty} \frac{2^d \ell^{d-1}}{(\ell - 1)^\alpha} \\ & = 2^{2d} \cdot \sum_{\ell=3}^{\infty} \frac{\ell^{d-1}}{(\ell - 1)^\alpha} . \end{aligned}$$

We now bound $\ell^{d-1}/(\ell - 1)^\alpha$, where we assume that $\epsilon = \alpha - d > 0$. This yields

$$\begin{aligned} & 2^{2d} \cdot \sum_{\ell=3}^{\infty} \frac{\ell^{d-1}}{(\ell - 1)^\alpha} \\ & = 2^{2d} \cdot \sum_{\ell=3}^{\infty} \frac{\ell^{d-1}}{(\ell - 1)^{d-1}} \cdot \frac{1}{(\ell - 1)^{1+\epsilon}} \\ & = 2^{2d} \cdot \sum_{\ell=3}^{\infty} \left(1 + \frac{1}{\ell - 1}\right)^{d-1} \cdot \frac{1}{(\ell - 1)^{1+\epsilon}} \\ & < 6^d \cdot \sum_{\ell=3}^{\infty} \frac{1}{(\ell - 1)^{1+\epsilon}} \\ & = 6^d \cdot \sum_{\ell=2}^{\infty} \ell^{-1-\epsilon} . \end{aligned}$$

The assumption $\epsilon > 0$ yields a constant value for the expression, which is $6^d \cdot (\zeta(1 + \epsilon) - 1)$. We estimate this value by $\sum_{\ell=2}^{\infty} \ell^{-1-\epsilon} < \int_{\ell=1}^{\infty} \ell^{-1-\epsilon} = 1/\epsilon$, which proves the lemma. \square

This yields

$$\begin{aligned} I &= \sum_{j \in S, j \neq i} \frac{d_{ii}^{r\alpha}}{d_{ji}^\alpha} \\ &\leq 2\Delta^{r\alpha} \sum_{\ell=3}^{\infty} 2^d (\ell^d - (\ell-1)^d) \cdot \frac{1}{((\ell-1)\sigma/3d)^\alpha} \\ &< 2\Delta^{r\alpha} \left(\frac{3d}{\sigma}\right)^\alpha \cdot \frac{6^d}{\alpha-d} . \end{aligned}$$

Note that the SINR constraint is satisfied if $p_i/d_{ii}^\alpha \geq \Delta^{r\alpha}/\Delta^\alpha \geq \beta I$, or

$$2\beta\Delta^{r\alpha} \cdot \left(\frac{3d}{\sigma}\right)^\alpha \cdot \frac{6^d}{\alpha-d} \leq \Delta^{(r-1)\alpha} .$$

This yields a lower bound for the distance of

$$\sigma \geq \Delta \cdot 3d \cdot \sqrt[\alpha]{\frac{2\beta 6^d}{\alpha-d}} , \quad (2)$$

which can be verified to hold for our choice of σ .

To bound the competitive ratio we need the following *Density Lemma*, which is an extension of Lemma 3 in [1] to both senders and receivers, and to metric spaces of arbitrary dimension d . The proof requires some adjustments from [1].

Lemma 2.3 (Density Lemma). *Consider a sector A with side-length $x \geq 1$ and any feasible solution with arbitrary power assignment. There can be only $(d+1)^\alpha(x+1)^d/\beta$ requests with a receiver in A and only $(d+1)^\alpha(x+1)^d/\beta$ requests with a sender in A .*

Proof. We first assume $x = 1$ and consider the number of receivers and senders in A separately.

Receivers: We first prove the lemma for the receivers. Let us assume that the transmission powers in the solution are such that there is a constant \bar{p} such that the signal strength received by a receiver $p_i/d_{ii}^\alpha = \bar{p}$ for any request with $r_i \in A$. Consider another request with $r_j \in A$. The interference of j at r_i is $p_j/d_{ji}^\alpha \geq p_j/(d(r_i, r_j) + d_{jj})^\alpha$. Due to the size of the sector we have that $d(r_i, r_j) \leq d$. Also $d_{jj} \geq 1$, which implies

$$\frac{p_j}{(d(r_i, r_j) + d_{jj})^\alpha} \geq \frac{1}{(d+1)^\alpha} \cdot \frac{p_j}{d_{jj}^\alpha} \geq \frac{\bar{p}}{(d+1)^\alpha} .$$

Thus, if more than $(d+1)^\alpha/\beta$ such connections are present, the SINR constraint *for all of them* is violated.

Now consider a solution with arbitrary powers. Here we artificially reduce powers such that all connections experience a minimal signal strength \bar{p} and then increase powers to their original value. The increase lowers SINR ratios for the requests that continue to have a signal strength of \bar{p} . Hence, if more than $(d+1)^\alpha/\beta$ receivers are present in A , *at least one* SINR constraint is violated.

Senders: For bounding the number of senders in A we use a similar approach. This time, however, we first assume that all senders have the same power. For two requests i and j this yields $p_j/d_{ji}^\alpha \geq p_j/(d(s_i, s_j)+d_{ii})^\alpha$. We have that $d(s_i, s_j) \leq d$. Also $d_{jj} \geq 1$, so $p_j/(d(s_j, s_i)+d_{ii})^\alpha \geq \frac{1}{(d+1)^\alpha} \cdot \frac{p_j}{d_{ii}^\alpha}$ as before. Thus, for the SINR constraint it is necessary that

$$\frac{p_i}{d_{ii}^\alpha} \geq \frac{\beta}{(d+1)^\alpha} \cdot \sum_{j \neq i} \frac{p_j}{d_{ii}^\alpha}.$$

Using $p_i = p_j$ for all requests i and j , there can be at most $(d+1)^\alpha/\beta$ senders in A , otherwise the SINR constraint *for all* requests is violated. A similar observation as before generalizes the argument to arbitrary powers.

This proves the lemma for $x = 1$. If $x > 1$, we overestimate A to a sector of side-length $\lfloor x + 1 \rfloor$, divide this into sectors of length 1, apply the above arguments, and the bound follows. \square

The density lemma allows a simple way to bound the number of connections the optimum solution can accept in the blocked area. First consider a sender s_i of a request accepted by SAFE-DISTANCE. The sender blocks a hypersphere of radius σ for receivers of other requests. We overestimate its size by a sector of side-length 2σ centered at s_i . By the density lemma, the optimum solution can accept at most $(d+1)^\alpha(2\sigma+1)^d/\beta$ requests, which is $O(\Delta^d)$ for fixed α, β , and d . For the receiver r_i there is a similar estimation. This time we bound the number of senders in a hypersphere around r_i , which is $O(\Delta^d)$ for fixed α, β , and d . Finally, note that σ is chosen to maximize conceptual simplicity and does not optimize the involved constants in the competitive ratio. \square

We can use similar arguments to show a result for any other polynomial power assignment. As safe distance we pick $\sigma^+ = \Delta^r \cdot \sigma$ if $r > 1$, and $\sigma^- = \Delta^{1-r} \cdot \sigma$ if $r < 0$.

Corollary 2.4. SAFE-DISTANCE is $O(\Delta^{d \cdot \max\{r, 1-r\}})$ -competitive for a polynomial power assignment with $r \notin (0, 1)$ and a single channel.

Proof. In the case $r > 1$ we note for correctness of the algorithm that the interference at an accepted receiver r_i is again bounded by

$$I = \sum_{j \in S, j \neq i} \frac{d_{jj}^{r\alpha}}{d_{ji}^\alpha} \leq \Delta^{r\alpha} \sum_{j \in S, j \neq i} \frac{1}{d_{ji}^\alpha} < 2\Delta^{r\alpha} \cdot \left(\frac{3d}{\sigma^+}\right)^\alpha \cdot \frac{6^d}{\alpha - d}.$$

The SINR constraint now requires that $p_i/d_{ii}^\alpha = d_{ii}^{(r-1)\alpha} \geq 1 \geq \beta I$. This yields a lower bound of

$$\sigma^+ \geq \Delta^r \cdot 3d \cdot \sqrt[\alpha]{\frac{2\beta 6^d}{\alpha - d}}. \quad (3)$$

Bounding the competitive ratio can be done as before and proves the result for the case $r > 1$.

If $r < 0$, then the interference is maximized with requests of length 1 in each sector. The interference is thus bounded by

$$I = \sum_{j \in S, j \neq i} \frac{d_{jj}^{r\alpha}}{d_{ji}^\alpha} \leq \sum_{j \in S, j \neq i} \frac{1}{d_{ji}^\alpha} < 2 \cdot \left(\frac{3d}{\sigma^-}\right)^\alpha \cdot \frac{6^d}{\alpha - d}.$$

The SINR constraint now requires that $p_i/d_{ii}^\alpha = d_{ii}^{(r-1)\alpha} \geq \Delta^{(r-1)\alpha} \geq \beta I$. This yields a lower bound

$$\sigma^- \geq \Delta^{1-r} \cdot 3d \cdot \sqrt[\alpha]{\frac{2\beta 6^d}{\alpha - d}} . \quad (4)$$

The corollary follows. \square

As it turns out, the competitive ratio of SAFE-DISTANCE is asymptotically best possible for polynomial power assignments with $r \notin (0, 1)$. This includes both the uniform and linear power assignment. Next, we bound the competitive ratio for any deterministic online algorithm using polynomial power assignments. This can be generalized to a lower bound for any distance-based power assignment.

Theorem 2.5. *Every deterministic online algorithm using polynomial power assignments has a competitive ratio of $\Omega(\Delta^{d \cdot \max\{r, 1-r\}})$. Every deterministic online algorithm is $\Omega(\Delta^{d/2})$ -competitive (1) using arbitrary power assignments in the case of undirected requests and (2) using distance-based power assignments in the case of only directed requests.*

Proof. The main observation in the proof is that every deterministic online algorithm has to accept the first request that arrives, otherwise it risks having an unbounded competitive ratio. While this is true only for strictly competitive algorithms, we can repeat the following instance sufficiently often and keep a sufficiently large distance between the instances. In this way we can neglect the constant a from the competitive ratio.

We first consider the case that all requests are directed and polynomial power assignment. Let the first request have length Δ . From the SINR constraint we bound the minimum distance every other successful request has to keep to sender s_1 or receiver r_1 . This yields a blocked area in which the online algorithm is not able to accept any request. We then count the maximum number of requests that can be placed into this area, and which the optimum solution can accept simultaneously. The next Proposition yields a bound on the minimum distance between two requests with a polynomial power assignment.

Proposition 2.6. *Consider two directed successful requests i and j with polynomial power assignment. The distance between s_i and r_j must be at least $d_{ij} \geq \sqrt[\alpha]{\beta} \cdot d_{ii}^r \cdot d_{jj}^{1-r}$.*

Proof. Consider the SINR constraint for request j when only requests i and j are accepted. It reads

$$d_{jj}^{\alpha(r-1)} \geq \beta(d_{ii}^r/d_{ij}^\alpha) ,$$

and rearranging yields the result. \square

Now suppose the online algorithm has accepted the first request of length Δ . The adversary subsequently presents requests of length 1. If the sender of one such request is closer than $\sqrt[\alpha]{\beta} \cdot \Delta^{1-r}$ to r_1 , the online algorithm cannot accept the request. The same holds if the receiver is closer than $\sqrt[\alpha]{\beta} \cdot \Delta^r$ to s_1 . Thus, there are two hyperspherical areas blocked around sender and receiver of request 1. Let us consider the case $r \leq 0.5$ and the hypersphere around the receiver. All subsequent arguments follow similarly for $r > 0.5$ and the sender.

The adversary can place requests, all of equal length $d_{ii} = 1$, into the hypersphere of radius $\sqrt[\alpha]{\beta} \cdot \Delta^{1-r}$ around r_1 . Similar to the proof of Theorem 2.1 we divide the space into sectors of length $2\sigma_1$, where

$$\sigma_1 = \max \left\{ 2, 18d \cdot \sqrt[\alpha]{2\beta/(\alpha - d)} \right\} .$$

We again assume that r_1 is located on the boundary of d sectors. How many sectors are completely enclosed by the blocked hypersphere around r_1 ? The side-length of the maximum hypercube that is contained is $2\Delta^{1-r} \sqrt[d]{\beta}/d$. There are at least $\frac{2\Delta^{1-r} \sqrt[d]{\beta}}{d\sigma_1} - 1$ sectors along each dimension within the hypercube, a number in $\Theta(\Delta^{1-r})$. This obviously yields a total number of $\Omega(\Delta^{(1-r)d})$ sectors, in which the online algorithm must not accept any request. However, we observe that σ_1 is chosen using the formula for σ with ratio 1. It is possible to locate one request of length 1 in each sector such that receivers and senders of two different requests are at least a distance of σ_1 apart. By Theorem 2.1 it is possible to accept all these $\Omega(\Delta^{(1-r)d})$ small requests simultaneously, which proves the theorem for case $r \leq 0.5$. For $r > 0.5$ we can place requests in the hypersphere around s_1 to derive a similar result.

To extend the previous arguments to arbitrary distance-based power assignments, we observe that the previous lower bound uses only requests of length 1 and Δ . Let ϕ be the function of the distance-based power assignment, then $\phi(\Delta)$ is the power of the first request. The lower bound for this power assignment behaves exactly as for a polynomial assignment with $r = (\log \phi(\Delta))/(\alpha \log \Delta)$.

Note that when a power assignment is not distance-based, it might assign different powers to small requests based on whether they are near the sender or the receiver of the first request. This, however, does not help if the requests are undirected. In this case we create the same instance using only undirected requests. Then we get a blocked area of at least $\Omega(\Delta^{d/2})$ for any polynomial power assignment around both points of the first request. Using the normalization of powers as before we observe that there is a blocked area of size $\Omega(\Delta^{d/2})$ for any small request, *no matter which power we assign to it*. This proves the theorem. \square

3 Improved Competitive Ratios

3.1 A Near-Optimal Algorithm for the Square-Root Assignment

Algorithm 1 MULTI-CLASS SAFE-DISTANCE

```

1: Initialize accepted requests  $S = \emptyset$ .
2: while a new request  $i$  arrives do
3:   Set  $p_i = \sqrt{d_{ii}^{\alpha}}$  and temporarily accept  $S' \leftarrow S \cup i$ 
4:   for all  $j \in S$  do
5:     Let  $\mathcal{C}_x$  and  $\mathcal{C}_y$  be the length classes of requests  $i$  and  $j$ , respectively
6:     if  $\min\{d_{ij}, d_{ji}\} \leq \min\{\sigma(\mathcal{C}_x), \sigma(\mathcal{C}_y)\}$  then
7:       decline request:  $S' \leftarrow S$ .
8:     end if
9:   end for
10:  Update:  $S \leftarrow S'$ .
11: end while

```

In this section we extend algorithm SAFE-DISTANCE to achieve a competitive ratio, which is close to the best-possible ratio for any distance-based power assignment. The algorithm uses the square-root power assignment, and the main idea of the algorithm is to block areas based on the distances of the involved requests. In particular, we classify requests into m length classes, where class \mathcal{C}_x contains requests i with $d_{ii} \in [\Delta^{a_x}, \Delta^{a_{x-1}}]$ with $a_x = 1/2^x$, for $x = 1, \dots, m-1$ and

$[1, \Delta^{a_{m-1}}]$ for class \mathcal{C}_m . With each class we associate a safe distance $\sigma(\mathcal{C}_x)$ chosen as

$$\sigma(\mathcal{C}_x) = \max \left\{ 2\Delta^{a_{x-1}}, \Delta^{0.5+a_x} \cdot 18d \cdot \sqrt[\alpha]{2\beta m \cdot \left(2 + \frac{1}{\alpha-d}\right)} \right\} .$$

This yields the following result.

Theorem 3.1. *For any constant $\varepsilon > 0$, MULTI-CLASS SAFE-DISTANCE is $O(\Delta^{d/2+\varepsilon})$ -competitive for a single channel.*

Proof. We first show that the algorithm is correct. We again treat a single accepted request and bound the interference from other accepted requests. This time, however, we have to consider the class the request is contained in. Suppose a request i is from class \mathcal{C}_x . To show that it is successful we have to estimate the distances d_{ji} for other requests. We will bound the interference from requests of each class separately and apply the construction outlined in Theorem 2.1. For requests of class \mathcal{C}_y we assume a worst-case placement and divide the space into sectors of side-length $\sigma(\mathcal{C}_y)/3d$. This again shows that no sector can contain more than two senders. The consideration of layers allows to bound the joint interference from all senders. For a class $y \geq x$, the minimum distance from r_i to each sender is at least $\sigma(\mathcal{C}_y)$. Thus, there is no sender in layers 1 and 2, and we can apply previous arguments to bound the interference. For classes with $y < x$ the minimum distance between r_i and any sender from this class is only $\sigma(\mathcal{C}_x) < \sigma(\mathcal{C}_y)$. Senders can be closer to r_i creating more interference. In particular, there can be senders in sectors of layers 1 and 2. For these senders we explicitly bound the distance using $\sigma(\mathcal{C}_x)$.

$$\begin{aligned} I &\leq \sum_{y=1}^m \sum_{j \in \mathcal{C}_y, j \neq i} \frac{d_{jj}^{\alpha/2}}{d_{ji}^\alpha} \leq \sum_{y \geq x} \sum_{j \in \mathcal{C}_y, j \neq i} \frac{\Delta^{\alpha/2^y}}{d_{ji}^\alpha} + \sum_{y < x} \sum_{j \in \mathcal{C}_y} \frac{\Delta^{\alpha/2^y}}{d_{ji}^\alpha} \\ &< \sum_{y \geq x} 2\Delta^{\alpha/2^y} \cdot \left(\frac{3d}{\sigma(\mathcal{C}_y)}\right)^\alpha \cdot \frac{6^d}{\alpha-d} + \underbrace{\sum_{y < x} \Delta^{\alpha/2^y} \sum_{j \in \mathcal{C}_y} \frac{1}{d_{ji}^\alpha}}_{I^{<x}} . \end{aligned}$$

Using Lemma 2.2, the definition of $\sigma(\mathcal{C}_x)$, and $y \geq 1$ we see that

$$\begin{aligned} I^{<x} &\leq 2 \sum_{y < x} \Delta^{\alpha/2^y} \cdot \left(\frac{2^d}{\sigma(\mathcal{C}_x)^\alpha} + \left(\frac{3d}{\sigma(\mathcal{C}_y)}\right)^\alpha \cdot \left(4^d + 2^d \sum_{\ell=3}^{\infty} \frac{\ell^d - (\ell-1)^d}{(\ell-1)^\alpha}\right) \right) \\ &< 2 \sum_{y < x} \Delta^{\alpha/2^y} \cdot \left(\frac{2^d}{\sigma(\mathcal{C}_x)^\alpha} + \left(\frac{3d}{\sigma(\mathcal{C}_y)}\right)^\alpha \cdot \left(4^d + \frac{6^d}{\alpha-d}\right) \right) \\ &\leq \sum_{y < x} 2\Delta^{\alpha/2^y} \cdot \left(\frac{3d}{\sigma(\mathcal{C}_x)}\right)^\alpha \cdot 6^d \cdot \left(2 + \frac{1}{\alpha-d}\right) \\ &< \sum_{y < x} \frac{\Delta^{\alpha/2^y}}{\beta m \cdot \Delta^{\alpha/2 + \alpha/2^x}} \\ &\leq \frac{x-1}{\beta m \cdot \Delta^{\alpha/2^x}} . \end{aligned}$$

For the total interference we use $x \geq 1$ and bound as follows

$$\begin{aligned}
I &< \sum_{y \geq x} 2\Delta^{\alpha/2^y} \cdot \left(\frac{3d}{\sigma(\mathcal{C}_y)}\right)^\alpha \cdot \frac{6^d}{\alpha - d} + \frac{x-1}{\beta m \cdot \Delta^{\alpha/2^x}} \\
&\leq \frac{m-x+1}{\beta m \cdot \Delta^{\alpha/2}} + \frac{x-1}{\beta m \cdot \Delta^{\alpha/2^x}} \\
&\leq \frac{1}{\beta \cdot \Delta^{\alpha/2^x}}.
\end{aligned}$$

As request i is in class \mathcal{C}_x , the minimum signal strength is $p_i/d_{ii}^\alpha \geq 1/\Delta^{\alpha/2^x} > \beta I$, which proves correctness of the algorithm.

For bounding the competitive ratio we consider the number of requests from the optimum solution that are blocked per accepted request. We consider blocked requests from each class separately. Obviously, the largest blocked areas are generated by a request from class 1. It blocks a hypersphere of radius $\sigma(\mathcal{C}_x)$ for requests from class \mathcal{C}_x , which we overestimate by the corresponding sector of side-length $2\sigma(\mathcal{C}_x)$. We must take into account that requests from class \mathcal{C}_x are bounded from below in distance. The proof of the density lemma can be adjusted to show that there can be only $(d+1)^\alpha/\beta$ receivers and senders in a sector of side-length h when each request has distance at least $d_{ii} \geq h$. There are only $(d+1)^\alpha(x/h+1)^d/\beta$ requests of minimum length h in a sector of side-length x . In the blocked area of \mathcal{C}_x we can schedule at most $(d+1)^\alpha(2\sigma(\mathcal{C}_x)/\Delta^{1/2^x}+1)^d/\beta$ requests. Assuming that d , α , and β are constants, this number is in $O(m\Delta^{d/2})$ for each $x = 1, \dots, m-1$. For class \mathcal{C}_m it is in $O(m\Delta^{d/2+d/2^m})$. Hence, the total number of requests blocked per accepted request is $O(m^2\Delta^{d/2+d/2^m})$. In order to obtain a bound for a constant ε , we apply MULTI-CLASS SAFE-DISTANCE using $m = \log d/\varepsilon$ length classes. This proves the theorem. \square

3.2 Multiple Channels

In this section we show how to generalize the algorithms above to k channels and decrease their competitive ratio. We propose a *k-channel adjustment*, in which we separate the problem by using certain channels only for specific request lengths. All requests with length in $[\Delta^{(i-1)/k}, \Delta^{i/k}]$ are assigned to channel i , for $i = 1, \dots, k$, where we assign requests of length $\Delta^{i/k}$ arbitrarily to channel i or $i+1$. For each channel i we apply an algorithm outlined above, which makes decisions about acceptance and power of requests assigned to channel i . Using this separation, we effectively reduce the aspect ratio to $\Delta^{1/k}$ on each channel. If the optimum solution has to adhere to the same length separation on the channels, this would yield a denominator k in the exponents of Δ of the competitive ratios. Obviously, the optimum solution is not tied to our separation, but the possible improvement due to this degree of freedom can easily be bounded by a factor k . This yields the following corollary.

Corollary 3.2. MULTI-CLASS SAFE-DISTANCE with *k-channel adjustment* is $O(k\Delta^{(d/2k)+\varepsilon})$ -competitive using the square-root power assignment. SAFE-DISTANCE with *k-channel adjustment* is $O(k\Delta^{d/k})$ -competitive for any polynomial assignment with $r \in [0, 1]$, and $O(k\Delta^{\max\{r, 1-r\} \cdot d/k})$ -competitive for $r \notin [0, 1]$.

3.3 A Randomized Algorithm

In the previous section for $k = \Theta(\log \Delta)$, the length differences on each channel reduce to a constant factor, e.g., for suitable k the requests on channel j are of length $[2^{j-1}, 2^j]$. This implies that

we approximate the requests on each channel by a constant factor. Thus, we obtain an $O(\log \Delta)$ -competitive algorithm against an optimum that can use $k = \Theta(\log \Delta)$ channels. Similarly, if the optimum was restricted to use only one channel, we would obtain a constant factor approximation algorithm. This is the main insight for designing our randomized algorithm RANDOM SAFE-DISTANCE. We virtually set up $\Theta(\log \Delta)$ channels, pick one channel uniformly at random, and then run our algorithm restricted to this channel. This yields a $O(\log \Delta)$ -competitive randomized algorithm, even for the case of a single channel. Using an additional k -channel adjustment in this case shows a similar result for k channels. We have the following corollary.

Corollary 3.3. *RANDOM SAFE-DISTANCE with k -channel adjustment is $O(\log \Delta)$ -competitive for any polynomial power assignment and any number k of channels.*

Note that for polynomial assignments with $r \notin (0, 1)$ and one channel the logarithmic ratio is asymptotically optimal. This follows with a simple example from [8]. There are $n = \Theta(\log \Delta)$ nested request pairs on the line with exponentially increasing distance. The optimum power assignment can successfully schedule $\Omega(\log \Delta)$ requests. Using any polynomial assignment with $r \notin (0, 1)$ there can be only $O(1)$ successful requests. Thus, using such a power assignment even an optimal offline algorithm knowing all requests is $\Omega(\log \Delta)$ -competitive. A similar observation holds with results of [8] in the case of directed request sets and any distance-based power assignment. In this case, however, the lower bound is only $\Omega(\log \log \Delta)$. Closing this gap remains as an open problem.

4 Extensions

4.1 Requests with Duration

In the previous sections we assumed that requests last forever, analyzing only the spatial aspect of the problem. We now show how our results extend when each request i has a duration t_i . After time t_i an accepted request stops sending and leaves (thus, no longer causing interference).

We show the modification for the algorithm SAFE-DISTANCE for $r \in [0, 1]$. We adapt the algorithm in the following way. It accepts a given request i if and only if the safe distance σ holds to all previously accepted requests that are active at some point in time in i 's duration.

Our first observation is the following. If we consider a fixed point in time, an optimal solution OPT can have at most $O(\Delta^d)$ more requests than our algorithm, as this corresponds to the spatial problem. Now let i be a request accepted by SAFE-DISTANCE with smallest duration possible, that is, $t_i = 1$. Each request contained in an optimal solution that interferes with i is active at least either when i starts or when it stops sending. So it is sufficient to count the accepted requests in OPT at both of these points in time to upper bound the number of requests blocked by i , which is $2 \cdot O(\Delta^d)$. Furthermore, a request i with $t_i \leq \Gamma$ can be split into at most Γ requests of duration 1, thus blocking at most $(\Gamma + 1) \cdot O(\Delta^d)$ requests. The argumentation is similar for other polynomial power assignments and results in an additional factor of Γ in all previously shown bounds (c.f. Section 1.2).

In the case of multiple channels, for $k = k' \cdot k''$, clustering of requests w.r.t. similar length and duration values can be used to improve the ratio for our algorithm MULTI-CLASS SAFE-DISTANCE to $O\left(k \cdot \Gamma^{1/k'} \Delta^{(d/2k'')+\varepsilon}\right)$. Choosing $k = \log \Gamma \cdot \log \Delta$, RANDOM SAFE-DISTANCE becomes $O(\log \Gamma \cdot \log \Delta)$ -competitive.

4.2 Doubling Metrics

All of our algorithms can be adjusted to work in more generalized metric spaces. In particular, we consider doubling metrics [7]. Let (\mathcal{V}, d) be a metric space and $B(x, r) = \{y \in \mathcal{V} \mid d(x, y) \leq r\}$ a ball of radius r around a point x . Consider an ϵ -covering of such a ball, i. e., a set of balls of radius ϵr such that their union contains $B(x, r)$. The *doubling dimension* of a metric space is the minimum number d such that for any ball $B(x, 2r)$ with $x \in \mathcal{V}$ and $r > 0$ there is a covering with 2^d balls of radius r . A metric with constant d is called a *doubling metric*. We again assume that α and d are constants, and that we have a fading metric with $\alpha > d$. A slight adjustment of the constants involved in the definition of the safe distance then yields similar bounds on the performance of SAFE-DISTANCE, MULTI-CLASS SAFE-DISTANCE, and RANDOM SAFE-DISTANCE in this more general scenario.

Theorem 4.1. *All bounds on the competitive ratios of SAFE-DISTANCE, MULTI-CLASS SAFE-DISTANCE, and RANDOM SAFE-DISTANCE continue to hold for doubling fading metrics. In particular, for $k = k' \cdot k''$, algorithm MULTI-CLASS SAFE-DISTANCE with k -channel adjustment is $O\left(k \cdot \Gamma^{1/k'} \cdot \Delta^{(d/2k'')+\epsilon}\right)$ -competitive for the square-root power assignment. RANDOM SAFE-DISTANCE with k -channel adjustment is $O(\log \Gamma \cdot \log \Delta)$ -competitive for any polynomial power assignment and any number k of channels.*

Proof. We discuss the respective results mentioned in the theorem:

Algorithm SAFE-DISTANCE: Let us first consider an adjusted algorithm SAFE-DISTANCE that uses the uniform power assignment and keeps a distance of at least

$$\tau = \max \left\{ 2\Delta, \Delta \cdot 20 \cdot \sqrt[\alpha]{\frac{2\beta}{2^\alpha - 2^d}} \right\} .$$

Then no two senders can be closer than $\tau/2$. Thus, in a ball of radius $\tau/5$ there can be at most two senders. We first require correctness of the algorithm and derive a lower bound on τ . We structure the space into balls of radius $2^\ell \cdot \tau/5$, for $\ell = 1, 2, \dots$. A ball of size ℓ can be covered by at most 2^d many balls of layer $\ell - 1$. Applying this argument recursively, the ball can be covered by $2^{\ell d}$ of radius $\tau/10$. Note that there can be at most $2^{\ell d+1}$ many different senders in such a ball, because the number of balls of radius r required for covering is at most the number of points with mutual distance $2r$ that can be placed in an area. We now overestimate the number of senders and at a distance by using concentric balls around a receiver r_i . We consider an annulus $B(r_i, 2^\ell \cdot \tau/5) - B(r_i, 2^{\ell-1} \cdot \tau/5)$, and assume that $2^{\ell d+1}$ senders are located in this area, which all have a distance of $2^{\ell-1} \cdot \tau/5$ to r_i . As there is a minimum distance of τ of any sender to r_i , we start to count at $\ell = 2$. This yields an upper bound for the interference of

$$\begin{aligned} I &< \sum_{\ell=2}^{\infty} \frac{2^{(\ell+1)d+1}}{(2^\ell \cdot \tau/5)^\alpha} = 2^{d+1} \cdot \left(\frac{5}{\tau}\right)^\alpha \cdot \sum_{\ell=2}^{\infty} (2^{d-\alpha})^\ell \\ &< 2^{d+1} \cdot \left(\frac{5}{\tau}\right)^\alpha \cdot \left(\frac{2^\alpha}{2^\alpha - 2^d}\right) . \end{aligned} \tag{5}$$

For the last inequality we have used that $\alpha > d$. This yields $2^{d-\alpha} < 1$, and the sum amounts to less than $1/(1 - 2^{d-\alpha})$. This allows to derive a lower bound of τ on our safe distance, which is satisfied by our choice, and proves correctness.

For bounding the competitive ratio we adjust the Density Lemma in a straightforward way and note that in a ball of radius 1 there can be only $3^\alpha/\beta$ many senders and receivers. To cover a ball of radius τ , we need at most $2^{\lceil \log_2 \tau \rceil d}$ many balls of radius 1. Thus, for α , β and d being constants, there are at most $O(\Delta^d)$ many requests that are blocked in the optimum by any accepted request of the online algorithm.

Note that the previous proof can be generalized easily to any polynomial power assignment, resulting in similar bounds as shown in Corollary 2.4.

Algorithm MULTI-CLASS SAFE-DISTANCE: For algorithm MULTI-CLASS SAFE-DISTANCE we use the same distribution of request lengths into classes \mathcal{C}_x for $x = 1, \dots, m$ as before. The safe distances $\tau(\mathcal{C}_x)$ used by the algorithm can be estimated similarly. In particular, we use

$$\tau(\mathcal{C}_x) = \max \left\{ 2\Delta^{a_{x-1}}, \Delta^{0.5+a_x} \cdot 20 \cdot \sqrt[\alpha]{2\beta m \cdot \left(2 + \frac{1}{2^\alpha - 2^d}\right)} \right\}.$$

The construction to show correctness is the same extension that we used to extend SAFE-DISTANCE to MULTI-CLASS SAFE-DISTANCE as before. Here, however, we use the bounds of equation (5), which yields

$$I < \sum_{y \geq x} 2\Delta^{\alpha/2^y} \cdot \left(\frac{20}{\tau(\mathcal{C}_y)}\right)^\alpha \cdot \frac{1}{2^\alpha - 2^d} + \underbrace{\sum_{y < x} \Delta^{\alpha/2^y} \sum_{j \in \mathcal{C}_y} \frac{1}{d_{ji}^\alpha}}_{I^{<x}}.$$

Using a minimum distance of $\tau(\mathcal{C}_x)$ for the requests from the smallest balls, we derive similarly as before

$$\begin{aligned} I^{<x} &\leq 2 \sum_{y < x} \Delta^{\alpha/2^y} \cdot \left(\frac{2^d}{\tau(\mathcal{C}_x)^\alpha} + \left(\frac{5}{\tau(\mathcal{C}_y)}\right)^\alpha \cdot \left(4^d + \sum_{\ell=2}^{\infty} \frac{2^{(\ell+1)d}}{2^{\ell\alpha}}\right)\right) \\ &< \sum_{y < x} 2\Delta^{\alpha/2^y} \cdot \left(\frac{2^d}{\tau(\mathcal{C}_x)^\alpha} + \left(\frac{5}{\tau(\mathcal{C}_y)}\right)^\alpha \cdot \left(4^d + \frac{2^{\alpha+d}}{2^\alpha - 2^d}\right)\right) \\ &\leq \sum_{y < x} 2\Delta^{\alpha/2^y} \cdot \left(\frac{20}{\tau(\mathcal{C}_x)}\right)^\alpha \cdot \left(2 + \frac{1}{2^\alpha - 2^d}\right). \end{aligned}$$

Thus, by using the definition of $\tau(\mathcal{C}_x)$ and noting $y \geq 1$ we see that $I^{<x} < \frac{x-1}{\beta m \cdot \Delta^{\alpha/2^x}}$. For the total interference we use $x \geq 1$ and bound as follows

$$I < \sum_{y \geq x} 2\Delta^{\alpha/2^y} \cdot \left(\frac{20}{\tau(\mathcal{C}_y)}\right)^\alpha \cdot \frac{1}{2^\alpha - 2^d} + \frac{x-1}{\beta m \cdot \Delta^{\alpha/2^x}} \leq \frac{1}{\beta \cdot \Delta^{\alpha/2^x}},$$

which proves correctness of the algorithm. Estimation of the competitive ratio can be done similarly as before. We use the adjustment of the Density Lemma outlined above for SAFE-DISTANCE to bound the maximum number of connections from OPT blocked by MULTI-CLASS SAFE-DISTANCE. This results in a competitive ratio of $O(\Delta^{(d/2)+\epsilon})$.

Channels and RANDOM SAFE-DISTANCE: The generalization to multiple channels and the randomized algorithm are independent of the metric and apply directly without adjustment.

□

5 Conclusion

In this paper we have studied a throughput optimization problem in wireless networks using the physical interference model. The lack of information about arrival of requests in practice naturally gives rise to the online scenario studied in this paper, where requests arrive iteratively over time and must be scheduled or dropped. We here use competitive analysis to derive near-optimal deterministic online algorithms. We also show how our algorithms can be extended to requests with duration and more generalized interference scenarios. In addition, we observe that randomized algorithms can be significantly more powerful than deterministic ones.

A natural open problem that arises from our work is to obtain optimal randomized algorithms for the problem. In particular, it would be interesting to see if the bound of $\log \Delta$ is tight or can be beaten using more intelligent algorithms with distance-based or arbitrary power assignments.

In addition, many of our lower bounds hold only with respect to arbitrary power assignments but not with respect to approximating the problem within a given power assignment. For example, can we obtain a smaller competitive ratio than $O(\log \Delta)$ using, e.g., uniform power assignments if we compare to the optimum using uniform power?

Finally, one could also consider adaptive adversaries that, e.g., adjust to the choice of request class made initially in our randomized algorithm. While such adversaries are not necessarily of practical relevance, showing good performance under these conditions can highlight the robustness of an algorithm.

References

- [1] Matthew Andrews and Michael Dinitz. Maximizing capacity in arbitrary wireless networks in the SINR model: Complexity and game theory. In *Proc. 28th IEEE Conf. Computer Communications (INFOCOM)*, 2009.
- [2] Patrice Assouad. Binary contraction of graphs. *Disc. Math.*, 47:315–319, 1983.
- [3] Chen Avin, Zvi Lotker, and Yvonne Anne Pigolet. On the power of uniform power: Capacity of wireless networks with bounded resources. In *Proc. 17th European Symposium on Algorithms (ESA)*, pages 373–384, 2009.
- [4] Hari Balakrishnan, Christopher Barrett, Anil Kumar, Madhav Marathe, and Shripad Thite. The distance-2 matching problem and its relationship to the MAC-layer capacity of ad hoc wireless networks. *IEEE J. Sel. Area Comm.*, 22(6):1069–1079, 2004.
- [5] Deepti Chafekar, Anil Kumar, Madhav Marathe, Srinivasan Parthasarathy, and Aravind Srinivasan. Approximation algorithms for computing capacity of wireless networks with SINR constraints. In *Proc. 27th IEEE Conf. Computer Communications (INFOCOM)*, pages 1166–1174, 2008.
- [6] Kenneth Clarkson. Nearest neighbor queries in metric spaces. *Disc. Comput. Geom.*, 22(1):63–93, 1999.
- [7] Kenneth Clarkson. Nearest-neighbor search and metric space dimensions. In G. Shakhnarovich, T. Darrall, and P. Indyk, editors, *Nearest-Neighbor Methods for Learning and Vision: Theory and Practice*, pages 15–59. MIT Press, 2006.

- [8] Alexander Fanghänel, Thomas Kesselheim, Harald Räcke, and Berthold Vöcking. Oblivious interference scheduling. In *Proc. 28th Symp. Principles of Distributed Computing (PODC)*, pages 220–229, 2009.
- [9] Alexander Fanghänel, Thomas Kesselheim, and Berthold Vöcking. Improved algorithms for latency minimization in wireless networks. In *Proc. 36th Intl. Coll. Automata, Languages and Programming (ICALP)*, volume 2, pages 208–219, 2009.
- [10] Olga Goussevskaia, Magnús Halldórsson, Roger Wattenhofer, and Emo Welzl. Capacity of arbitrary wireless networks. In *Proc. 28th IEEE Conf. Computer Communications (INFOCOM)*, 2009.
- [11] Olga Goussevskaia, Yvonne Anne Oswald, and Roger Wattenhofer. Complexity in geometric SINR. In *Proc. 8th Symp. Mobile Ad Hoc Networking and Computing (MobiHoc)*, pages 100–109, 2007.
- [12] Piyush Gupta and P. R. Kumar. The capacity of wireless networks. *IEEE Trans. Inf. Theory*, 46(2):388–404, 2000.
- [13] Magnús Halldórsson. Wireless scheduling with power control. In *Proc. 17th European Symposium on Algorithms (ESA)*, pages 361–372, 2009.
- [14] Magnús Halldórsson and Roger Wattenhofer. Computing wireless capacity. Unpublished manuscript, 2010.
- [15] Thomas Kesselheim. A constant-factor approximation for wireless capacity maximization with power control in the SINR model. In *Proc. 22nd Symp. Discrete Algorithms (SODA)*, pages 1549–1559, 2011.
- [16] Thomas Moscibroda, Stefan Schmid, and Roger Wattenhofer. On the topologies formed by selfish peers. In *Proc. 25th Symp. Principles of Distributed Computing (PODC)*, pages 133–142, 2006.
- [17] Thomas Moscibroda, Roger Wattenhofer, and Aaron Zollinger. Topology control meets SINR: The scheduling complexity of arbitrary topologies. In *Proc. 7th Symp. Mobile Ad Hoc Networking and Computing (MobiHoc)*, pages 310–321, 2006.