On the Power of Uniform Power: Capacity of Wireless Networks with Bounded Resources

Chen Avin[1], Zvi Lotker[1], Yvonne-Anne Pignolet[2]

[1]Ben Gurion University of the Negev, Israel, [2]ETH Zurich, Switzerland
{avin,zvilo}@cse.bgu.ac.il, pignolet@tik.ee.ethz.ch

Abstract. The *throughput capacity* of arbitrary wireless networks under the physical *Signal to Interference Plus Noise Ratio* (SINR) model has received much attention in recent years. In this paper, we investigate the question of how much the worst-case performance of uniform and non-uniform power assignments differ under constraints such as a bound on the area where nodes are distributed or restrictions on the maximum power available. We determine the maximum factor by which a non-uniform power assignment can outperform the uniform case in the SINR model. More precisely, we prove that in one-dimensional settings the capacity of a non-uniform assignment exceeds a uniform assignment by at most a factor of $O(\log L_{\max})$ when the length of the network is $L_{\max}$. In two-dimensional settings, the uniform assignment is at most a factor of $O(\log P_{\max})$ worse than the non-uniform assignment if the maximum power is $P_{\max}$. We provide algorithms that reach this capacity in both cases. Due to lower bound examples in previous work, these results are tight in the sense that there are networks where the lack of power control causes a performance loss in the order of these factors. As a consequence, engineers and researchers may prefer the uniform model due to its simplicity if this degree of performance deterioration is acceptable.

1 Introduction

The great success of wireless networks is mainly due to the fact that any device can exchange information with any other device in its reception range. However, this advantage is also the most problematic characteristic of wireless networks. Since the communication medium is shared by all participants, it is necessary to address the problem of interference. Simultaneous communication attempts cause interference and might even prevent the correct reception of a signal. To tap the full potential of a network, algorithms that coordinate the transmission of messages are necessary. To reach the *throughput capacity* of a network, we have to solve the problem of assigning time slots, frequencies and (depending on hardware capabilities) transmitting power levels to a set of n pairs of wireless transmitters (senders) and receivers distributed in a given area.

When attempting to solve this and related problems, we must first choose the appropriate interference model. A standard interference model that captures some of the key characteristics of wireless communication and is sufficiently concise for rigorous reasoning is the physical SINR model [7]. It describes interference as continuous property, decreasing polynomially with the distance from the sender. In this model, a message is received successfully if the ratio between the strength of the sender signal at the receiving location and the sum of interferences created by all other simultaneous senders plus ambient noise is larger than some hardware-defined threshold. The fading speed depends on the value of the so-called path-loss exponent α.

The analysis of problems in the SINR model is intricate, due to the non-binary and accumulative features of interference. Only recently have some theoretical guarantees for SINR-based algorithms been provided. One of problems under scrutiny is the *scheduling problem*. Given a set of n pairs of senders and receivers along with the power level of the transmitters, the goal is to devise a scheduling scheme that minimizes the total number of rounds that will satisfy all the communication requests of every pair. In addition to the timing, the signal strengths of the transmitting nodes greatly influence the performance of wireless networks, since the number of simultaneous transmissions can be increased if the nodes are able to emit signals of different power levels. Thus, power control constitutes an additional aspect of interest. Orthogonally to the scheduling problem, it is necessary to address the *power assignment problem*, i.e., determining a power assignment for each sender of a given set of communication pairs in such a way that the total number of communication requests in one round is maximized. The two problems are often combined, and many algorithms addressing the problem of joint power control and scheduling of a set of links have been devised (see related work section). Since *power control*, i.e., the possibility to assign a different power level to each sender, may play a major role in the complexity of the problems or the performance of the algorithms, we distinguish between two settings: *non-uniform power* (i.e., power control), where each transmitter can transmit with a different power, and *uniform power*, where there is only one power level. It has been demonstrated [12, 13] that uniform power has significant performance disadvantages compared to the non-uniform case. However, examples of situations in which power control algorithms outperform uniform energy assignment schemes usually position the nodes in an area of exponential size in the number of nodes and require transmission power levels that differ by a factor exponential in the number of nodes.

On the other hand, a uniform power assignment has several important advantages due to its simplicity. Most importantly, the production cost of wireless devices that always transmit at the same power is lower. Therefore, the uniform power assignment has been widely adopted in practical systems. The lack of power control implies that a device only has to decide whether or not it should send a message at the certain point of time, and not at which power level. As a consequence, there are fewer possibilities to consider which makes reaching a decision much simpler. Moreover, recently a study of *SINR diagrams*[1] [1]

[1] The SINR diagram of a set of transmitters divides the plane into $n+1$ *regions* or reception zones, one region for each transmitter that indicates the set of locations

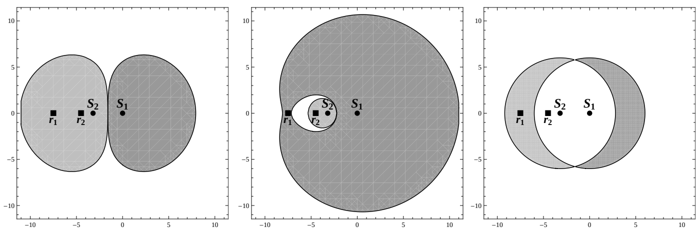

a) *b)* *c)*

Fig. 1. Impact of the choice of the interference model and power assignment. Given two communication pairs, $l_1 = (s_1, r_1)$ and $l_2 = (s_2, r_2)$, the shaded areas indicate where the signal of a sender can be received (the area in the lighter gray belongs to sender s_2). White areas imply that the received signal power is too weak for reception. *a)* Uniform power: only node r_2 receives a message from its sender, the interference is to high at r_1. *b)* Non-uniform power: both transmissions are successful. *c)* Unit Disk Graph model: neither r_1 nor r_2 receive a message from their corresponding senders.

showed that the reception zones of all senders are *convex* for a uniform scheme but not necessarily for non-uniform power assignments. This finding suggests that designing algorithms may be much simpler for uniform networks than for non-uniform networks.

In this paper, we compare the uniform and non-uniform cases and study the trade-off involved between the two, i.e., simplicity vs. performance. As mentioned above, in the absence of restrictions, the performance of the non-uniform model clearly exceeds the uniform model. However, by taking a closer look, we notice that this conclusion is based on examples that involve unbounded resources. Of course, resources are restricted in reality, e.g., the maximum available power for a transmitter or the space where nodes are distributed may be limited. This observation has motivated us to ask the following question:

> *In a resource-constrained setting (i.e., bounded area, bounded power), what is the worst-case performance difference between the uniform and non-uniform case?*

In a nutshell, we show that with bounded resources the two cases are not significantly different; therefore, engineers and researchers may actually prefer the uniform model due to its simplicity.

1.1 Problem Statement and Overview of Our Results

In order to quantify the gap induced by the ability to adjust the transmission power when the available resources are bounded, we consider the following game between two players, a *non-uniform (power control) player* and a *uniform player*.

in which it can be heard successfully, and one more region that indicates the set of locations in which no sender can be heard. This concept is perhaps analogous to the role played by Voronoi diagrams in computational geometry.

The non-uniform player begins by setting up a *configuration* by selecting n communication pairs where each pair consists of a sender and a receiver and their locations. For these pairs the following two conditions must be met:

1. The distance between a sender and its intended receiver is at least one.
2. There exists a power assignment such that the receivers are able to decode the messages of their senders when all senders transmit simultaneously with the same frequency, i.e., the non-uniform configuration is *feasible*.

After the first player has reached its decision, the uniform player can, choose a subset of these pairs that can transmit simultaneously with uniform power, i.e., a feasible uniform configuration. The non-uniform player tries to position the sender/receivers in such a way that the uniform player can pick only a small subset of the available pairs, without causing too much interference. On the other hand, the uniform player tries to select as many pairs as possible.[2]

Let the size of a configuration be the number of pairs in the configuration. Our first result concerns the one-dimensional case. If the non-uniform player can place its configuration on a interval of length at most L_{max}, then we can state the following:

Theorem 1. *For any feasible non-uniform configuration of size n, there is a feasible uniform configuration of a size of at least $\Omega(n/\log L_{max})$ if $\alpha = 2$.*

This result is tight in the sense that from previously known examples [11] there are feasible non-uniform configurations of size n for which the size of any feasible uniform configuration is at most $O(n/\log L_{max})$. Our proof is constructive and we present an algorithm that achieves our bound. This algorithm, combined with previous algorithms, can be used as an approximation scheme for the *uniform scheduling problem* and the *uniform power assignment problem*. In both cases, one can take the output a non-uniform scheduling and/or power assignment algorithm produces, which is basically a non-uniform configuration, and obtain a uniform configuration by "paying" an additional approximation price of $\log L_{\max}$.

In the two-dimensional case, we consider power assignments rendering the non-uniform player's configuration feasible with transmission power levels in the range of $[1, P_{max}]$. In this case we prove:

Theorem 2. *For any feasible non-uniform configuration of size n, there is a feasible uniform configuration of a size of at least $\Omega(n/\log P_{max})$ if $0 < \alpha$.*

This result is tight since there is a non-uniform configuration of size n for which the uniform player can at most select a configuration with a size of at most $O(n/\log P_{max})$. As for Theorem 1, we devise an algorithm that achieves this bound. Note that even if the ratio of the lowest and the highest power level

[2] Observe that the players are assumed to have unlimited computational power, since the problem of selecting the largest subset of nodes transmitting with fixed power levels has been shown to be NP-hard [6].

is constant, it is not immediately obvious that in this case the capacity for the uniform power assignment is in the same order as in setting with power control; namely because there are infinitely many possible power assignments and nested pairs of links are feasible.

It follows from Theorem 1 and Theorem 2 that if we bound either $P_{\max}$ or $L_{\max}$ to n in the one-dimensional case, then a uniform power assignment is at most $\log n$ worse than the best non-uniform power assignment. In the two-dimensional case, this is true only if we bound $P_{\max}$.

The number of links able to transmit simultaneously crucially depends on the path-loss exponent α. The faster the signal strength falls, the smaller an amount of interference is caused. In [15], measurements of indoor and outdoor path-loss exponents at various frequencies are reported, ranging from 1.6 to 6. Most existing work relies on the assumption that $\alpha > 2$, exploiting the fact that in this case the interference of far away nodes can be bounded easily. For $\alpha \leq 2$ the situation changes dramatically and different arguments are necessary. In this paper, the results for two dimension hold for all $\alpha > 0$, the results for one dimension for $\alpha = 2$.

2 Related Work

The study of the capacity of wireless metworks has been initiated by the seminal work of Gupta and Kumar [7]. The authors bounded the throughput capacity in the best-case (i.e., optimal configurations) for the physical models for $\alpha > 2$. More recently, a worst-case view point was adopted [10] by proving lower bounds. We also use this approach in the current paper. The fact that interference is continuous and accumulative as well as the geometric constraints render the scheduling task difficult in the physical model, even if the transmission power of the nodes is fixed. See [5, 6, 9] for the analysis of such scheduling algorithms. The complexity of connectivity of a uniform power network is examined in [2].

Depending on the hardware, nodes are able to adjust their transmission power. This capability can increase the number of links that are able to transmit successfully at the same time. To exploit this fact, efficient power control algorithms are necessary. For a given set of links, the highest achievable signal to noise ratio can be computed in polynomial time [16], yet the complexity of the problem of joint scheduling and power control in the physical model taking into account the geometry of the problem is unknown. Nevertheless many algorithms and heuristics have been suggested, see [11] for a classification and more detailed discussion of these approaches. Very recent work, [3, 4, 8] gives upper and lower bounds for power-controlled oblivious scheduling.

Non-uniform power assignment can clearly outperform a uniform assignment [13, 12] and increase the capacity of the network, therefore the majority of the work on capacity and scheduling addressed non-uniform power. As we discussed earlier, the study of the uniform case is still worthwhile, due to its simplicity. To the best of our knowledge, the gap between these two models has not been investigated under restricted resources.

The proof of the scheduling algorithm for fixed power levels in [5] can be adapted to conclude that the number of links that are able to communicate concurrently with uniform power is bounded logarithmically in the ratio of the highest and the lowest power, yet the analysis of their algorithm depends on the fact that $\alpha > 2$. In addition, the authors adopt a different viewpoint: Given a set of links, they try to find an approximation of the shortest schedule without power control. In contrast, we are concerned with the lower bound of the size of the largest subset of links that are able to communicate simultaneously with uniform power under the assumption that the original set was feasible with a non-uniform assignment.

Very few papers have been devoted to the one-dimensional case, as the capacity is more restricted than in two dimensions, especially for randomly distributed nodes. Nevertheless, Moscibroda et al. [12, 11] showed that the capacity of one-dimensional networks can be linear in the number of the links, at the expense of exponentially long links and and exponentially high power. Hence, we are among the first to study the capacity for uniform and non-uniform power assignments in one dimension.

Another paper published at ESA addresses power control and scheduling in the SINR model. It proposes an oblivious $O(\log n \log\log \Lambda)$-approximation algorithm for the scheduling problem, where Λ is the ratio between the longest and the shortest link length. Moreover it considers the approximation ratio uniform power algorithms can achieve. Using a different tool set from ours, [8] shows that if the Assouad dimension A of the underlying metric is strictly less than α, uniform power assignments are at most a $O(\log(\Lambda))$-factor worse than unconstrained power control. In other words, this more general result works well for $\alpha > A$, while our results for the two-dimensional case hold for any $\alpha > 0$. We believe that the techniques of [8] and the approaches of this paper are complementary and their combination might help to understand the remaining open problems in the SINR interference model.

3 Model and Preliminaries

Let (M, d) be a metric space and $V \subseteq M$ a finite set of $|V|$ *nodes*. A node v_j successfully receives a message from node v_i depending on the set of concurrently transmitting nodes and the applied interference model. In this paper, we adopt the physical SINR model [7], where the successful reception of a transmission depends on the strength of the received signal, the interference caused by nodes transmitting simultaneously, and the ambient noise level. The received power $P_{r_i}(s_i)$ of a signal transmitted by a sender s_i at an intended receiver r_i is $P_{r_i}(s_i) = P(s_i) \cdot g(s_i, r_i)$, where $P(s_i)$ is the transmission power of s_i and $g(s_i, r_i)$ is the propagation attenuation (link gain) modeled as $g(s_i, r_i) = d(s_i, r_i)^{-\alpha}$. The *path-loss exponent* $\alpha \geq 1$ is a constant typically between 1.6 and 6. The exact value of α depends on external conditions of the medium (humidity, obstacles, etc.) and on the exact sender-receiver distance. Measurements for indoor and outdoor path-loss exponents can be found in [15].

Given a sender and a receiver pair $l_i = (s_i, r_i)$, we use the notation $I_{r_i}(s_j) = P_{r_i}(s_j)$ for any other sender s_j concurrent to s_i in order to emphasize that the signal power transmitted by s_j is perceived at r_i as interference. The *total interference* $I_{r_i}(L)$ experienced by a receiver r_i is the sum of the interference power values created by the set L of nodes transmitting simultaneously (except the intending sender s_i), i.e., , $I_{r_i}(L) := \sum_{l_j \in L \setminus \{l_i\}} I_{r_i}(s_j)$. Finally, let N denote the ambient noise power level. Then, r_i receives s_i's transmission if and only if

$$\text{SINR}(l_i) = \frac{P_{r_i}(s_i)}{N+I_{r_i}(L)} = \frac{P(s_i)g(s_i,r_i)}{N+\sum_{j\neq i} P(s_j)g(s_j,r_i)} = \frac{\frac{P(s_i)}{d(s_i,r_i)^\alpha}}{N+\sum_{j\neq i}\frac{P(s_j)}{d(s_j,r_i)^\alpha}} \geq \beta,$$

where $\beta \geq 1$ is the minimum SINR required for a successful message reception. In the sequel we assume $\beta = 1$ we set $N = 0$ and ignore the influence of noise in the calculation of the SINR, for the sake of simplicity. However, this has no significant effect on the results: by scaling the power of all senders, the influence of ambience noise can be made arbitrarily small. Observe that for real scenarios with upper bounds on the maximum transmission power this is not possible, however, for our asymptotic calculations we can neglect this term.

For a uniform power assignment, we say a set of links $L = \{l_1, \ldots, l_n\}$ is a *uniformly feasible configuration of size* n if $P(s_i) = 1$ and $\text{SINR}(l_i) \geq \beta$ for all links $l_i \in L$. If the power level of a device is adjustable, we denote a set L to be a *PC feasible configuration of size* n if there exists a power assignment such that $\text{SINR}(l_i) \geq \beta$ for all links $l_i \in L$.

Zander [16] showed that the *maximum achievable SINR* (denoted SINR*) for wireless networks can be computed efficiently. Solving the eigenvalue problem for the matrix $Z = \left[\frac{g(s_i,r_j)}{g(s_i,r_i)}\right]$ yields an eigenvalue λ^* for which all elements of the corresponding eigenvector have the same sign. Then, the maximum achievable SINR, is given by $\text{SINR}^* = 1/(\lambda^* - 1)$. Furthermore, the corresponding eigenvector $\mathbf{P}^*$ is a power vector reaching this maximum for all links, i.e., they all have the same SINR level. Since we defined $\beta = 1$, this means that the largest eigenvalue of Z has to be less than 2, otherwise the successful concurrent transmission of all links is impossible.[3]

Theorem 3 (Zander [16]). *A set of senders can transmit simultaneously if the largest eigenvalue of the normalized link gain matrix is less than 2.*

We use the following equations from linear algebra repeatedly.

Theorem 4 (Eigenvalue relationships). *Given an n-by-n matrix with real or complex entries where $\lambda_1, \ldots, \lambda_n$ are the (complex and distinct) eigenvalues of A, then it holds for $k \in \mathbb{N}$ that for the trace of A^k, $\text{tr}(A^k) = \sum \lambda_i^k$. In contrast, the determinant of A is the product of its eigenvalues; i.e.,* $\det(A) = \prod \lambda_i$.

4 One Dimension: Length Constraint

In this section, we aim at determining the advantage of power control in one-dimensional settings. We prove that at most a factor of $\log L_{max}$ more links can

[3] [16] ignores the influence of noise. See [14] for an approach that handles noise as well.

be scheduled with power control than without, if the senders and receivers are located on a line of length L_{max}. Moreover, we present an algorithm that given a configuration feasible with power control, selects a subset of these links that can be scheduled with uniform power and contains at least a $1/\log L_{max}$-fraction of the links in the original configuration. As a first step, we show that, even when power is adjustable, two links transmitting concurrently must not cross, otherwise at least one of the receiver cannot decode the message.

Lemma 1 (Crossings). *Two senders s_1 and s_2 cannot transmit successfully at the same time if their respective receiver is closer to the other sender, i.e., if $d(s_1, r_1) > d(s_1, r_2)$ and $d(s_2, r_2) > d(s_2, r_1)$. [Proof in full version, applies Theorem 3]*

Nested pairs however are possible. But, as soon as there are more than two nested pairs, they cannot be too close to each other, since the interference is too high otherwise. More precisely, we can show that no matter how we position three nested sender/receiver pairs in an interval of length three and regardless of the power levels we assign to them, they cannot transmit simultaneously.

Lemma 2 (Nestings). *Three nested communication pairs need at least an interval of two times the shortest link distance, otherwise they cannot transmit simultaneously if $\alpha = 2$.*

Proof. (Sketch) We compute the normalized gain matrix Z for three nested links sending in the same direction. We show that for $\alpha = 2$ the following holds

- $\mathrm{tr}(Z)) = 3$
- $\mathrm{tr}(Z^2) > 5$
- $\det(Z) > 0$

Applying Theorem 4 we can derive three conditions the eigenvalues of Z have to satify and we can show that they conflict with the requirement that the largest eigenvalue has be less than two (otherwise no feasible solution, Theorem 3). Thus there is no configuration with three links transmitting in the same direction simultaneously if the longest link is at most twice as long as the shortes. For the other scenarios, where at least one link transmits in the other direction, $\mathrm{tr}(Z)) = 3$ and $\mathrm{tr}(Z^2))$ exceeds 9 and the arguments carry over.

We know from Lemma 2 that three nested links require an interval of more than three times the shortest link distance. Let the shortest distance be one. If we want to add three additional nested links that include the first three links we need at least an interval of length 3^2, if we neglect the interference from the three inner most links. We can repeat this procedure at most $O(\log L_{max})$ times, before we cover the entire interval length L_{max}. Thus Corollary 1 follows.

Corollary 1 (Nestings). *At most $O(\log L_{max})$ links can be nested on a interval of length L_{max}, otherwise there exists no power assignment allowing them to transmit simultaneously if $\alpha = 2$.*

Apart from crossing and nested links, we need to consider parallel links as well. We can show the following lemma.

Lemma 3. *Let $\zeta(x) = \sum_i^{\infty} i^{-x}$ the Riemann zeta function. Out of m parallel links (no crossings, no nestings) where the length of the longest link is at most twice the length of the shortest link, there are at least $k = \lceil 2^{1/\alpha-1}\zeta(\alpha)^{(1/\alpha)} \rceil$ senders that can transmit successfully at the same time with a uniform power assignment if $\alpha > 1$.*

Since $\zeta(x)$ konverges for all real $x > 1$, we have now all the ingredients necessary to conclude how much the capacity in a uniform power setting suffers from the lack of power control. We rephrase Theorem 1:

Theorem 1. *Consider an interval of length L_{max}. Given a PC-feasible configuration of size n, there exists a uniformly feasible configuration of size at least $\Omega(n/\log L_{max})$ if $\alpha = 2$.*

Proof. Given a set of links, we divide the links into $\log L_{max}$ length classes such that the class k contains the links of length in the interval $[2^{k-1}, 2^k]$. We pick the class containing the largest number of links. At least half of them transmit in the same direction. If there are any nested links, we know thanks to Lemma 2 that there are at most two nested links in the same length class. As a next step we can apply Lemma 3 by picking the first and every $\lceil 2^{1/\alpha-1}\zeta(\alpha)^{(1/\alpha)} \rceil^{th}$ of these links and let them transmit concurrently with a uniform power assignment. Due to this procedure a power control algorithm can schedule at most $O(\log n)$ more links simultaneously and the claim follows. □

5 Power Restriction

Even if the transmission power of a device is adjustable, there is typically a bounded range for the power or, even more restricted, a set of available power levels. In these situations, a uniform assignment can achieve a substantial fraction of the capacity a power control scheme can reach. In the following we examine the difference between the two strategies if the first player can place the nodes in arbitrary positions and use transmission power levels in the range $[1, P_{max}]$. We rephrase Theorem 2:

Theorem 2. *Given a PC-feasible configuration of size n in the two-dimensional Euclidean space, there exists a uniformly feasible configuration of size at least $\Omega(n/\log P_{max})$*

We construct an algorithm that given a PC-feasible set of n links placed in the two-dimensional Euclidean plane returns a uniformly feasible set that contains at least $\Omega(n/\log P_{max})$ links. Before we start with the description of the algorithm we elaborate on the problems that our algorithm faces. Since the senders in the resulting set of links have to be able to transmit with the same power and since there is typically background noise we are forced to increase the power levels (of a selected subset of transmitters) to the same power level. Consequently, the interference raises and we need to push the interference back to the level it was. We can do this by further reducing the number of senders. The main challenge is to keep the number of senders as large as possible. So far we have only considered the senders. In order to ensure that the signal to interference ratio is high enough at the receivers as well, we use the fact that from far away the position of the sender and the receiver almost coincide.

Proof (Sketch). Our algorithm (cf. Algorithm 2 for a description in pseudo code) starts by computing an optimal power assignment for the given set of links. Then, we divide the links into classes of similar power levels, i.e. we build subsets of links where the highest transmission power is at most twice the lowest power. Among these sets, we choose the one of greatest cardinality, T_1. The intuition for this step is that we now have a fairly homogeneous set of links and the difference between a uniform power assignment and arbitrary power levels is negligible. We remove some links of this set T_1 to guarantee that all remaining senders can transmit concurrently. To this end, we start with the longest link l_1 and we first clean an area around its receiver, i.e., we delete links with senders too close to r_1 from the set T_1. We assign l_1 to the candidate set T_2 and repeat this procedure with the remaining links until no links are left. These steps ensure that the interference at s_j and at r_j is roughly the same. Next, we remove more links to guarantee that the remaining set is uniformly feasible. We pick one link l_j of the candidate set T_2 and we partition the plane into six sectors of $60°$ around s_j. In each sector we remove the links whose senders are closest to s_j. Thereafter we add l_i to the set T_3 and recursively repeat the partitioning and removal with the remaining links. Before finally declaring the set of the surviving links to be simultaneously schedulable, we repeat the procedure with T_3, this time considering the links in the opposite sequence. We show now that the configuration S produced by this algorithm is uniformly feasible[4] and at most a factor of $O(1/\log P_{max})$ smaller than the original set L. Line 4 determines the largest set of similar power levels. A set of links where the feasible power levels differ by at most a factor of two is very similar to a uniformly feasible set. In particular, the number of senders in close proximity to a receiver is limited. This can be shown by adapting Lemma 4.2 from [5] to this case.

Lemma 4 (Extension of Lemma 4.2 from [5]). *For any link l_i of a PC-feasible set L the number of senders within distance $c \cdot d(s_i, r_i)$ from the receiver r_i is at most $2c^\alpha/\beta$ if $P_{max}/P_{min} < 2$. [Proof in full version]*

[4] This algorithm can be generalized to higher dimensions at the expense of a higher constant in its approximation guarantee. The only adjustments affect the lines 11-12 and 16-17, where cones instead of sectors are considered. We omit the explicit treatment of higher dimensional cases to increase the clarity of the arguments and more than three dimension are unlikely to be of practical importance.

Thus we can conclude that this algorithm runs in polynomial time and the original set size is reduced by a factor of O($1/\log P_{max}$) in line 3, by $O(\frac{\beta}{2\mu^\alpha})$ in lines 4-7 (due to Lemma 4), by $O(\frac{1}{6\nu})$ in lines 8-12 and by $O(\frac{1}{6\nu})$ in lines 13-17. Consequently, the ratio between the input and the output set is $\frac{L}{S} \le O(\frac{n\beta}{72\mu^\alpha\nu^2\log P_{max}}) \in O\left(\frac{n}{\log P_{max}}\right)$.

Algorithm 2 2D $\log(p_{max})$-approximation

Require: PC-feasible set $L = \{l_1, \ldots, l_n\}$
Ensure: uniformly feasible set $S \subset L$
1: set $\mu := 1 + 2\alpha$ and $\nu := \left\lceil (2 \cdot \frac{\mu}{\mu-1})^\alpha \right\rceil - 1$
2: determine power assignment for L (Zander)
3: partition L into subsets $L_i := \{l_j | 2^i \le P(s_j) < 2^{i+1}\}$, $i = 1, \ldots, \log P_{max} - 1$
4: set $T_1 := \arg\max_{0 \le i < \log(p_{max})} |L_i|, T_2 := \emptyset, T_3 := \emptyset$
5: **repeat**
6: move longest link l_j from T_1 to T_2
7: remove l_i from T_1 if $d(s_i, r_j) < \mu d(s_j, r_j)$
8: **until** $T_1 = \emptyset$
9: **repeat**
10: move longest link l_j from T_2 to T_3
11: partion plane into 6 sectors of 60° around s_j
12: in each sector remove the ν links with the closest senders to s_j from T_2
13: **until** $T_2 = \emptyset$
14: **repeat**
15: move shortest link l_j from T_3 to S
16: partion plane into 6 sectors of 60° around s_j
17: in each sector remove the ν links with the closest senders to s_j from T_3
18: **until** $T_3 = \emptyset$
19: **return** S;

It remains to demonstrate that the resulting set is indeed uniformly feasible. By setting the transmission power to two for every sender, the strength of the interference at the receivers is at most doubled. As a consequence we have to reduce the number of simultaneous transmissions such that the interference is halved in order to obtain a uniformly feasible set. Clearing a disk around each receiver and removing close senders diminishes the interference by half. We prove this in two steps. First we show how much the interference experienced at the senders is decreased and then we derive the resulting amount of interference at their respective receivers.

Lemma 5. *$I_{s_i}(S) < \frac{1}{\nu+1} I_{s_i}(L)$ and $I_{r_i}(S) < (\frac{\mu}{\mu-1})^\alpha I_{s_i}(S)$ for all $l_i \in S$. [Proof in full version]*

As a consequence the algorithm has reduced the interference at the receivers by at least $\Omega(\frac{1}{\nu+1} \cdot \left(\frac{\mu}{\mu-1}\right)^\alpha) = \frac{1}{\lceil 2\left(\frac{\mu}{\mu-1}\right)^\alpha \rceil} \cdot \left(\frac{\mu}{\mu-1}\right)^\alpha \ge \frac{1}{2} \cdot \left(\frac{\mu-1}{\mu}\right)^\alpha \cdot \left(\frac{\mu}{\mu-1}\right)^\alpha = \frac{1}{2}$. Therefore all transmitters that survive can transmit at power 2, while their receivers are guaranteed to be able to decode the message successfully. This concludes the proof that there always exists a uniform power $O(\log P_{max})$-approximation of a power control problem. □

6 Conclusion

In this paper we show that for limited resources, e.g., an upper bound on the maximum transmission power or the maximum distance between a sender and a receiver, a uniform power assignment provides a log-approximation for the achievable capacity by a non-uniform power assignment. These results can be understood in two ways. We can design and solve algorithmic problems in the uniform power model instead of the non-uniform power model and lose a log factor in the solution. The result presented in this paper suggest the following methodology for solving algorithmic problems in the non uniform power models. First solve the same problem in the uniform power model (this task is usually simpler and less general). Use this solution as a guide line for the general case involving power control and try to eliminate the logarithmic factor.

References

1. C. Avin, Y. Emek, E. Kantor, Z. Lotker, D. Peleg, and L. Roditty. SINR Diagrams: Towards Algorithmically Usable SINR Models of Wireless Networks. In *Proc. 28th Symposium on Principles of Distributed Computing (PODC)*, 2009.
2. C. Avin, Z. Lotker, F. Pasquale, and Y. A. Pignolet. A Note on Uniform Power Connectivity in the SINR Model. In *Proc.5th Intl Workshop on Algorithmic Aspects of Wireless Sensor Networks (ALGOSENSOR)*, 2009.
3. A. Fanghänel, T. Kesselheim, H. Räcke, and B. Vöcking. Oblivious Interference Scheduling. In *Proc. 28th Symposium on Principles of Distributed Computing (PODC)*, 2009.
4. A. Fanghänel, T. Kesselheim, and B. Vöcking. Improved Algorithms for Latency Minimization in Wireless Networks. In *Proc. 36th Intl. Coll. on Automata, Languages and Programming (ICALP)*, 2009.
5. O. Goussevskaia, M. Halldorsson, R. Wattenhofer, and E. Welzl. Capacity of Arbitrary Wireless Networks. In *28th Annual IEEE Conference on Computer Communications (INFOCOM)*, 2009.
6. O. Goussevskaia, Y. A. Oswald, and R. Wattenhofer. Complexity in Geometric SINR. In *ACM International Symposium on Mobile Ad Hoc Networking and Computing (MOBIHOC), Montreal, Canada*, September 2007.
7. P. Gupta and P. R. Kumar. The Capacity of Wireless Networks. *IEEE Trans. Inf. Theory*, 46(2):388–404, 2000.
8. M. Halldórsson. Wireless Scheduling with Power Control. In *Proc. 17th annual European Symposium on Algorithms (ESA)*, pages 368–380, 2009.
9. M. Halldorsson and R. Wattenhofer. Wireless Communication is in APX. In *Proc. 36th International Colloquium on Automata, Languages and Programming (ICALP)*, July 2009.
10. T. Moscibroda. The worst-case capacity of wireless sensor networks. In *Proc. 6th Conference on Information Processing in Sensor Networks (IPSN)*, 2007.
11. T. Moscibroda, Y. A. Oswald, and R. Wattenhofer. How Optimal are Wireless Scheduling Protocols? In *Proc. 26th IEEE Conference on Computer Communications (INFOCOM)*, 2007.
12. T. Moscibroda and R. Wattenhofer. The Complexity of Connectivity in Wireless Networks. In *Proc. 25th Conference of the IEEE Computer and Communications Societies (INFOCOM)*, 2006.
13. T. Moscibroda, R. Wattenhofer, and Y. Weber. Protocol Design Beyond Graph-based Models. In *Proc. 5th ACM SIGCOMM Workshop on Hot Topics in Networks (HOTNETS)*, 2006.
14. S. Pillai, T. Suel, and S. Cha. The Perron-Frobenius theorem. *IEEE Signal Processing Magazine*, 22(2):62–75, 2005.
15. T. Rappaport. *Wireless communications*. Prentice Hall, 2002.
16. J. Zander. Performance of optimum transmitter power control in cellular radio systems. *IEEE Trans. Veh. Technol.*, 41, 1992.