Probabilistic Analysis of Power Assignments*

Maurits de Graaf[1,2] and Bodo Manthey[1]

[1] University of Twente, Department of Applied Mathematics, Enschede, Netherlands
m.degraaf/b.manthey@utwente.nl
[2] Thales Nederland B. V., Huizen, Netherlands

Abstract. A fundamental problem for wireless ad hoc networks is the assignment of suitable transmission powers to the wireless devices such that the resulting communication graph is connected. The goal is to minimize the total transmit power in order to maximize the life-time of the network. Our aim is a probabilistic analysis of this power assignment problem. We prove complete convergence for arbitrary combinations of the dimension d and the distance-power gradient p. Furthermore, we prove that the expected approximation ratio of the simple spanning tree heuristic is strictly less than its worst-case ratio of 2.

Our main technical novelties are two-fold: First, we find a way to deal with the unbounded degree that the communication network induced by the optimal power assignment can have. Minimum spanning trees and traveling salesman tours, for which strong concentration results are known in Euclidean space, have bounded degree, which is heavily exploited in their analysis. Second, we apply a recent generalization of Azuma-Hoeffding's inequality to prove complete convergence for the case $p \geq d$ for both power assignments and minimum spanning trees (MSTs). As far as we are aware, complete convergence for $p > d$ has not been proved yet for any Euclidean functional.

1 Introduction

Wireless ad hoc networks have received significant attention due to their many applications in, for instance, environmental monitoring or emergency disaster relief, where wiring is difficult. Unlike wired networks, wireless ad hoc networks lack a backbone infrastructure. Communication takes place either through single-hop transmission or by relaying through intermediate nodes. We consider the case that each node can adjust its transmit power for the purpose of power conservation. In the assignment of transmit powers, two conflicting effects have to be taken into account: if the transmit powers are too low, the resulting network may be disconnected. If the transmit powers are too high, the nodes run out of energy quickly. The goal of the power assignment problem is to assign transmit powers to the transceivers such that the resulting network is connected and the sum of transmit powers is minimized [12].

* A full version with all proofs is available at http://arxiv.org/abs/1403.5882.

1.1 Problem Statement and Previous Results

We consider a set of vertices $X \subseteq [0,1]^d$, which represent the sensors, $|X| = n$, and assume that $\|u - v\|^p$, for some $p \in \mathbb{R}$ (called the *distance-power gradient* or *path loss exponent*), is the power required to successfully transmit a signal from u to v. This is called the power-attenuation model, where the strength of the signal decreases with $1/r^p$ for distance r, and is a simple yet very common model for power assignments in wireless networks [14]. In practice, we typically have $1 \leq p \leq 6$ [13].

A power assignment $\mathsf{pa} : X \to [0, \infty)$ is an assignment of transmit powers to the nodes in X. Given pa, we have an edge between two nodes u and v if both $\mathsf{pa}(x), \mathsf{pa}(y) \geq \|x - y\|^p$. If the resulting graph is connected, we call it a *PA graph*. Our goal is to find a PA graph and a corresponding power assignment pa that minimizes $\sum_{v \in X} \mathsf{pa}(v)$. Note that any PA graph $G = (X, E)$ induces a power assignment by $\mathsf{pa}(v) = \max_{u \in X : \{u,v\} \in E} \|u - v\|^p$.

PA graphs can in many aspects be regarded as a tree as we are only interested in connectedness, but it can contain more edges in general. However, we can simply ignore edges and restrict ourselves to a spanning tree of the PA graph.

The minimal connected power assignment problem is NP-hard for $d \geq 2$ and APX-hard for $d \geq 3$ [4]. For $d = 1$, i.e., when the sensors are located on a line, the problem can be solved by dynamic programming [11]. A simple approximation algorithm for minimum power assignments is the minimum spanning tree heuristic (MST heuristic), which achieves a tight worst-case approximation ratio of 2 [11]. This has been improved by Althaus et al. [1], who devised an approximation algorithm that achieves an approximation ratio of 5/3. A first average-case analysis of the MST heuristic was presented by de Graaf et al. [6]: First, they analyzed the expected approximation ratio of the MST heuristic for the (non-geometric, non-metric) case of independent edge lengths. Second, they proved convergence of the total power consumption of the assignment computed by the MST heuristic for the special case of $p = d$, but not of the optimal power assignment. They left as open problems, first, an average-case analysis of the MST heuristic for random geometric instances and, second, the convergence of the value of the optimal power assignment.

1.2 Our Contribution

In this paper, we conduct an average-case analysis of the optimal power assignment problem for Euclidean instances. The points are drawn independently and uniformly from the d-dimensional unit hypercube $[0,1]^d$. We believe that probabilistic analysis is better-suited for performance evaluation in wireless ad hoc networks than worst-case analysis, as the positions of the sensors – in particular if deployed in areas that are difficult to access – are subjected to randomness.

Roughly speaking, our contributions are as follows:

1. We show that the power assignment functional has sufficiently nice properties in order to apply Yukich's general framework for Euclidean functionals [16] to obtain concentration results (Section 3).

2. Combining these insights with a recent generalization of the Azuma-Hoeffding bound [15], we obtain concentration of measure and complete convergence for all combinations of d and $p \geq 1$, even for the case $p \geq d$ (Section 4). In addition, we obtain complete convergence for $p \geq d$ for minimum-weight spanning trees. As far as we are aware, complete convergence for $p \geq d$ has not been proved yet for such functionals. The only exception we are aware of are minimum spanning trees for the case $p = d$ [16, Sect. 6.4].
3. We provide a probabilistic analysis of the MST heuristic for the geometric case. We show that its expected approximation ratio is strictly smaller than its worst-case approximation ratio of 2 [11] for any d and p (Section 5).

Our main technical contributions are two-fold: First, we introduce a transmit power redistribution argument to deal with the unbounded degree that graphs induced by the optimal transmit power assignment can have. The unboundedness of the degree makes the analysis of the power assignment functional PA challenging. The reason is that removing a vertex can cause the graph to fall into a large number of components and it might be costly to connect these components without the removed vertex. In contrast, the degree of any minimum spanning tree, for which strong concentration results are known in Euclidean space for $p \leq d$, is bounded for every fixed d, and this is heavily exploited in the analysis. (The concentration result by de Graaf et al. [6] for the power assignment obtained from the MST heuristic also exploits that MSTs have bounded degree.)

Second, we apply a recent generalization of Azuma-Hoeffding's inequality by Warnke [15] to prove complete convergence for the case $p \geq d$ for both power assignments and minimum spanning trees. We introduce the notion of *typically smooth* Euclidean functionals, prove convergence of such functionals, and show that minimum spanning trees and power assignments are typically smooth. In this sense, our proof of complete convergence provides an alternative and generic way to prove complete convergence, whereas Yukich's proof for minimum spanning trees is tailored to the case $p = d$. In order to prove complete convergence with our approach, one only needs to prove convergence in mean, which is often much simpler than complete convergence, and typically smoothness. Thus, we provide a simple method to prove complete convergence of Euclidean functionals along the lines of Yukich's result that, in the presence of concentration of measure, convergence in mean implies complete convergence [16, Cor. 6.4].

2 Definitions and Notation

Throughout the paper, d (the dimension) and p (the distance-power gradient) are fixed constants. For three points x, y, v, we by $\overline{xv}$ the line through x and v, and we denote by $\angle(x, v, y)$ the angle between $\overline{xv}$ and $\overline{yv}$.

A *Euclidean functional* is a function F^p for $p > 0$ that maps finite sets of points from the unit hypercube $[0, 1]^d$ to some non-negative real number and is translation invariant and homogeneous of order p [16, page 18]. From now on, we omit the superscript p of Euclidean functionals, as p is always fixed and clear from the context.

PA_B is the canonical boundary functional of PA (we refer to Yukich [16] for boundary functionals of other optimization problems): given a hyperrectangle $R \subseteq \mathbb{R}^d$ with $X \subseteq R$, this means that a solution is an assignment $\mathsf{pa}(x)$ of power to the nodes $x \in X$ such that

- x and y are connected if $\mathsf{pa}(x), \mathsf{pa}(y) \geq \|x - y\|^p$,
- x is connected to the boundary of R if the distance of x to the boundary of R is at most $\mathsf{pa}(x)^{1/p}$, and
- the resulting graph, called a *boundary PA graph*, is either connected or consists of connected components that are all connected to the boundary.

Then $\mathsf{PA}_B(X, R)$ is the minimum value for $\sum_{x \in X} \mathsf{pa}(x)$ that can be achieved by a boundary PA graph. Note that in the boundary functional, no power is assigned to the boundary. It is straight-forward to see that PA and PA_B are Euclidean functionals for all $p > 0$ according to Yukich [16, page 18].

For a hyperrectangle $R \subseteq \mathbb{R}^d$, let $\operatorname{diam} R = \max_{x,y \in R} \|x - y\|$ denote the diameter of R. For a Euclidean functional F, let $\mathsf{F}(n) = \mathsf{F}(\{U_1, \ldots, U_n\})$, where $U_1, \ldots, U_n$ are drawn uniformly and independently from $[0,1]^d$. Let $\gamma_{\mathsf{F}}^{d,p} = \lim_{n \to \infty} \frac{\mathbb{E}\left(\mathsf{F}(n)\right)}{n^{\frac{d-p}{d}}}$. (In principle, $\gamma_{\mathsf{F}}^{d,p}$ need not exist, but it does exist for all functionals considered in this paper.)

A sequence $(R_n)_{n \in \mathbb{N}}$ of random variables *converges in mean* to a constant γ if $\lim_{n \to \infty} \mathbb{E}(|R_n - \gamma|) = 0$. The sequence $(R_n)_{n \in \mathbb{N}}$ *converges completely to a constant* γ if we have $\sum_{n=1}^{\infty} \mathbb{P}\left(|R_n - \gamma| > \varepsilon\right) < \infty$ for all $\varepsilon > 0$ [16, page 33].

Besides PA, we consider two other Euclidean functions: $\mathsf{MST}(X)$ denotes the length of the minimum spanning tree with lengths raised to the power p. $\mathsf{PT}(X)$ denotes the total power consumption of the assignment obtained from the MST heuristic, again with lengths raised to the power p. The MST heuristic proceeds as follows: First, we compute a minimum spanning tree of X. Then let $\mathsf{pa}(x) = \max\{\|x - y\|^p \mid \{x, y\} \text{ is an edge of the MST}\}$. By construction and a simple analysis, we have $\mathsf{MST}(X) \leq \mathsf{PA}(X) \leq \mathsf{PT}(X) \leq 2 \cdot \mathsf{MST}(X)$ [11].

For $n \in \mathbb{N}$, let $[n] = \{1, \ldots, n\}$.

3 Properties of the Power Assignment Functional

After showing that optimal PA graphs can have unbounded degree and providing a lemma that helps solving this problem, we show that the power assignment functional fits into Yukich's framework for Euclidean functionals [16].

3.1 Degrees and Cones

As opposed to minimum spanning trees, whose maximum degree is bounded from above by a constant that depends only on the dimension d, a technical challenge is that the maximum degree in an optimal PA graph cannot be bounded by a constant in the dimension. This holds even for the simplest case of $d = 1$ and $p > 1$. We conjecture that the same holds also for $p = 1$, but proving this seems to be more difficult and not to add much.

Lemma 3.1. *For all $p > 1$, all integers $d \geq 1$, and for infinitely many n, there exist instances of n points in $[0,1]^d$ such that the unique optimal PA graph is a tree with a maximum degree of $n - 1$.*

The unboundedness of the degree of PA graphs make the analysis of the functional PA challenging. The technical reason is that removing a vertex can cause the PA graph to fall into a non-constant number of components. The following lemma is the crucial ingredient to get over this "degree hurdle".

Lemma 3.2. *Let $x, y \in X$, let $v \in [0,1]^d$, and assume that x and y have power $\mathsf{pa}(x) \geq \|x - v\|^p$ and $\mathsf{pa}(y) \geq \|y - v\|^p$, respectively. Assume further that $\|x - v\| \leq \|y - v\|$ and that $\angle(x, v, y) \leq \alpha$ with $\alpha \leq \pi/3$. Then the following holds:*

(a) $\mathsf{pa}(y) \geq \|x - y\|^p$, i.e., y has sufficient power to reach x.
(b) If x and y are not connected (i.e., $\mathsf{pa}(x) < \|x - y\|^p$), then $\|y - v\| > \frac{\sin(2\alpha)}{\sin(\alpha)} \cdot \|x - v\|$.

For instance, $\alpha = \pi/6$ results in a factor of $\sqrt{3} = \sin(\pi/3)/\sin(\pi/6)$. In the following, we invoke this lemma always with $\alpha = \pi/6$, but this choice is arbitrary as long as $\alpha < \pi/3$, which causes $\sin(2\alpha)/\sin(\alpha)$ to be strictly larger than 1.

3.2 Deterministic Properties

In this section, we state properties of the power assignment functional. Subadditivity (Lemma 3.3), superadditivity (Lemma 3.4), and growth bound (Lemma 3.5) are straightforward.

Lemma 3.3 (subadditivity). PA *is subadditive [16, (2.2)] for all $p > 0$ and all $d \geq 1$, i.e., for any point sets X and Y and any hyperrectangle $R \subseteq \mathbb{R}^d$ with $X, Y \subseteq R$, we have $\mathsf{PA}(X \cup Y) \leq \mathsf{PA}(X) + \mathsf{PA}(Y) + O\big((\operatorname{diam} R)^p\big)$.*

Lemma 3.4 (superadditivity). PA_B *is superadditive for all $p \geq 1$ and $d \geq 1$ [16, (3.3)], i.e., for any X, hyperrectangle $R \subseteq \mathbb{R}^d$ with $X \subseteq R$ and partition of R into hyperrectangles R_1 and R_2, we have $\mathsf{PA}_B^p(X, R) \geq \mathsf{PA}_B^p(X \cap R_1, R_1) + \mathsf{PA}_B^p(X \cap R_2, R_2)$.*

Lemma 3.5 (growth bound). *For any $X \subseteq [0,1]^d$ and $0 < p$ and $d \geq 1$, we have $\mathsf{PA}_B(X) \leq \mathsf{PA}(X) \leq O\left(\max\left\{n^{\frac{d-p}{d}}, 1\right\}\right)$.*

The following lemma shows that PA is smooth, which roughly means that adding or removing a few points does not have a huge impact on the function value. Its proof requires Lemma 3.2 to deal with the fact that optimal PA graphs can have unbounded degree.

Lemma 3.6. *The power assignment functional* PA *is smooth for all* $0 < p \leq d$ *[16, (3.8)], i.e.,* $\left|\mathsf{PA}^p(X \cup Y) - \mathsf{PA}^p(X)\right| = O\left(|Y|^{\frac{d-p}{d}}\right)$ *for all point sets* $X, Y \subseteq [0,1]^d$.

Proof. One direction is straightforward: $\mathsf{PA}(X \cup Y) - \mathsf{PA}(X)$ is bounded by $\Psi = O\big(|Y|^{\frac{d-p}{d}}\big)$, because the optimal PA graph for Y has a value of at most Ψ by Lemma 3.5. Then we can take the PA graph for Y and connect it to the tree for X with a single edge, which costs at most $O(1) \leq \Psi$ because $p \leq d$.

For the other direction, consider the optimal PA graph T for $X \cup Y$. The problem is that the degrees $\deg_T(v)$ of vertices $v \in Y$ can be unbounded (Lemma 3.1). (If the maximum degree were bounded, then we could argue in the same way as for the MST functional.) The idea is to exploit the fact that removing $v \in Y$ also frees some power. Roughly speaking, we proceed as follows: Let $v \in Y$ be a vertex of possibly large degree. We add the power of v to some vertices close to v. The graph obtained from removing v and distributing its energy has only a constant number of components.

To prove this, Lemma 3.2 is crucial. We consider cones rooted at v with the following properties:

- The cones have a small angle α, meaning that for every cone C and every $x, y \in C$, we have $\angle(x, v, y) \leq \alpha$. We choose $\alpha = \pi/6$.
- Every point in $[0,1]^d$ is covered by some cone.
- There is a finite number of cones. (This can be achieved because d is a constant.)

Let $C_1, \ldots, C_m$ be these cones. By abusing notation, let C_i also denote all points $x \in C_i \cap (X \cup Y \setminus \{v\})$ that are adjacent to v in T. For C_i, let x_i be the point in C_i that is closest to v and adjacent to v (breaking ties arbitrarily), and let y_i be the point in C_i that is farthest from v and adjacent to v (again breaking ties arbitrarily). (For completeness, we remark that then C_i can be ignored if $C_i \cap X = \emptyset$.) Let $\ell_i = \|y_i - v\|$ be the maximum distance of any point in C_i to v, and let $\ell = \max_i \ell_i$.

We increase the power of x_i by ℓ^p/m. Since the power of v is at least ℓ^p and we have m cones, we can account for this with v's power because we remove v. Because $\alpha = \pi/6$ and x_i is closest to v, any point in C_i is closer to x_i than to v. According to Lemma 3.2(a), every point in C_i has sufficient power to reach x_i. Thus, if x_i can reach a point $z \in C_i$, then there is an established connection between them.

From this and increasing x_i's power to at least ℓ^p/m, there is an edge between x_i and every point $z \in C_i$ that has a distance of at most $\ell/\sqrt[p]{m}$ from v. We recall that m and p are constants.

Now let $z_1, \ldots, z_k \in C_i$ be the vertices in C_i that are not connected to x_i because x_i has too little power. We assume that they are sorted by increasing distance from v. Thus, $z_k = y_i$. We can assume that no two z_j and $z_{j'}$ are in the same component after removal of v. Otherwise, we can simply ignore one of the edges $\{v, z_j\}$ and $\{v, z_{j'}\}$ without changing the components.

Since z_j and z_{j+1} were connected to v and they are not connected to each other, we can apply Lemma 3.2(b), which implies that $\|z_{j+1}-v\| \geq \sqrt{3}\cdot\|z_j-v\|$. Furthermore, $\|z_1 - v\| \geq \ell/\sqrt[p]{m}$ by assumption. Iterating this argument yields $\ell = \|z_k-v\| \geq \sqrt{3}^{k-1}\|z_1-v\| \geq \sqrt{3}^{k-1}\cdot\ell/\sqrt[p]{m}$. This implies $k \leq \log_{\sqrt{3}}(\sqrt[p]{m})+1$. Thus, removing v and redistributing its energy as described causes the PA graph to fall into at most a constant number of components. Removing $|Y|$ points causes the PA graph to fall into at most $O(|Y|)$ components. These components can be connected with costs $O(|Y|^{\frac{d-p}{d}})$ by choosing one point per component and applying Lemma 3.5. □

Lemma 3.7. PA_B *is smooth for all* $1 \leq p \leq d$ *[16, (3.8)].*

Crucial for convergence of PA is that PA, which is subadditive, and PA_B, which is superadditive, are close to each other. Then both are approximately both subadditive and superadditive. The following lemma states that indeed PA and PA_B do not differ too much for $1 \leq p < d$.

Lemma 3.8. PA *is point-wise close to* PA_B *for* $1 \leq p < d$ *[16, (3.10)], i.e.,* $\left|\mathsf{PA}^p(X) - \mathsf{PA}_B^p(X,[0,1]^d)\right| = o\left(n^{\frac{d-p}{d}}\right)$ *for every set* $X \subseteq [0,1]^d$ *of* n *points.*

3.3 Probabilistic Properties

For $p > d$, smoothness is not guaranteed to hold, and for $p \geq d$, point-wise closeness is not guaranteed to hold. But similar properties typically hold for random point sets, namely smoothness in mean (Definition 3.10) and closeness in mean (Definition 3.12). In the following, let $X = \{U_1,\ldots,U_n\}$. Recall that $U_1,\ldots,U_n$ are drawn uniformly and independently from $[0,1]^d$. We need the following bound on the longest edge of an optimal PA graph.

Lemma 3.9 (longest edge). *For every constant* $\beta > 0$*, there exists a constant* $c_{\text{edge}} = c_{\text{edge}}(\beta)$ *such that, with a probability of at least* $1-n^{-\beta}$*, every edge of an optimal PA graph and an optimal boundary PA graph* PA_B *is of length at most* $r_{\text{edge}} = c_{\text{edge}}\cdot(\log n/n)^{1/d}$.

Yukich gave two different notions of smoothness in mean [16, (4.13) and (4.20) & (4.21)]. We use the stronger notion, which implies the other.

Definition 3.10 (smooth in mean [16, (4.20), (4.21)]). *A Euclidean functional* F *is called* smooth in mean *if, for every constant* $\beta > 0$*, there exists a constant* $c = c(\beta)$ *such that the following holds with a probability of at least* $1 - n^{-\beta}$*:*

$$\left|\mathsf{F}(n) - \mathsf{F}(n \pm k)\right| \leq ck\cdot\left(\tfrac{\log n}{n}\right)^{p/d} \quad \text{and} \quad \left|\mathsf{F}_B(n) - \mathsf{F}_B(n \pm k)\right| = ck\cdot\left(\tfrac{\log n}{n}\right)^{p/d}.$$

for all $0 \leq k \leq n/2$.

Lemma 3.11. PA_B *and* PA *are smooth in mean for all* $p > 0$ *and all* d.

Definition 3.12 (close in mean [16, (4.11)]). *A Euclidean functional* F *is close in mean to its boundary functional* F_B *if* $\mathbb{E}\left(|\mathsf{F}(n) - \mathsf{F}_B(n)|\right) = o(n^{\frac{d-p}{d}})$.

Lemma 3.13. PA *is close in mean to* PA_B *for all* d *and* $p \geq 1$.

4 Convergence

4.1 Standard Convergence

Our findings of Sections 3.2 yield complete convergence of PA for $p < d$ (Theorem 4.1). Together with the probabilistic properties of Section 3.3, we obtain convergence in mean in a straightforward way for all combinations of d and p (Theorem 4.2). In Sections 4.2 and 4.3, we prove complete convergence for $p \geq d$.

Theorem 4.1. *For all d and p with $1 \leq p < d$, there exists a constant $\gamma_{\mathsf{PA}}^{d,p}$ such that $\frac{\mathsf{PA}^p(n)}{n^{\frac{d-p}{d}}}$ converges completely to $\gamma_{\mathsf{PA}}^{d,p}$.*

Theorem 4.2. *For all $p \geq 1$ and $d \geq 1$, there exists a constant $\gamma_{\mathsf{PA}}^{d,p}$ such that $\lim_{n\to\infty} \frac{\mathbb{E}(\mathsf{PA}^p(n))}{n^{\frac{d-p}{d}}} = \lim_{n\to\infty} \frac{\mathbb{E}(\mathsf{PA}_B^p(n))}{n^{\frac{d-p}{d}}} = \gamma_{\mathsf{PA}}^{d,p}$.*

4.2 Concentration with Warnke's Inequality

McDiarmid's or Azuma-Hoeffding's inequality are powerful tools to prove concentration of measure for a function that depends on many independent random variables, all of which have only a bounded influence on the function value. If we consider smoothness in mean (see Lemma 3.11), then we have the situation that the influence of a single variable is typically very small (namely $O((\log n/n)^{p/d})$), but can be quite large in the worst case (namely $O(1)$). Unfortunately, this situation is not covered by McDiarmid's or Azuma-Hoeffding's inequality. Fortunately, Warnke [15] proved a generalization specifically for the case that the influence of single variables is typically bounded and fulfills a weaker bound in the worst case.

The following theorem is a simplified version (personal communication with Lutz Warnke) of Warnke's concentration inequality [15, Theorem 2], tailored to our needs.

Theorem 4.3 (Warnke). *Let $U_1, \ldots, U_n$ be a family of independent random variables with $U_i \in [0,1]^d$ for each i. Suppose that there are numbers $c_{\text{good}} \leq c_{\text{bad}}$ and an event Γ such that the function $\mathsf{F} : ([0,1]^d)^n \to \mathbb{R}$ satisfies*

$$\max_{i\in[n]} \max_{x\in[0,1]^d} |\mathsf{F}(U_1,\ldots,U_n) - \mathsf{F}(U_1,\ldots,U_{i-1},x,U_{i+1},\ldots,U_k)| \leq \begin{cases} c_{\text{good}} & \text{if } \Gamma \text{ holds and} \\ c_{\text{bad}} & \text{otherwise.} \end{cases} \tag{1}$$

Then, for any $t \geq 0$ and $\gamma \in (0,1]$ and $\eta = \gamma(c_{\text{bad}} - c_{\text{good}})$, we have

$$\mathbb{P}\big(|\,\mathsf{F}(n) - \mathbb{E}(\mathsf{F}(n))| \geq t\big) \leq 2\exp\big(-\tfrac{t^2}{2n(c_{\text{good}}+\eta)^2}\big) + \tfrac{n}{\gamma} \cdot \mathbb{P}(\neg\Gamma). \tag{2}$$

Next, we introduce *typical smoothness*, which means that, with high probability, a single point does not have a significant influence on the value of F, and we apply Theorem 4.3 for typically smooth functionals F. The bound of $c \cdot (\log n/n)^{p/d}$ in Definition 4.4 below for the typical influence of a single point is somewhat arbitrary, but works for PA and MST. This bound is also essentially the smallest possible, as there can be regions of diameter $c' \cdot (\log n/n)^{1/d}$ for some small constant $c' > 0$ that contain no or only a single point. It might be possible to obtain convergence results for other functionals for weaker notions of typical smoothness.

Definition 4.4 (typically smooth). *A Euclidean functional* F *is* typically smooth *if, for every* $\beta > 0$*, there exists a constant* $c = c(\beta)$ *such that*

$$\max_{x\in[0,1]^d, i\in[n]} \big|\mathsf{F}(U_1,\ldots,U_n) - \mathsf{F}(U_1,\ldots,U_{i-1},x,U_{i+1},\ldots,U_n)\big| \leq c \cdot \left(\tfrac{\log n}{n}\right)^{p/d}$$

with a probability of at least $1 - n^{-\beta}$.

Theorem 4.5 (concentration of typically smooth functionals). *Let* $p, d \geq 1$*. Assume that* F *is typically smooth. Then*

$$\mathbb{P}\big(|\,\mathsf{F}(n) - \mathbb{E}(\mathsf{F}(n))| \geq t\big) \leq O(n^{-\beta}) + \exp\big(-\tfrac{t^2 n^{\frac{2p}{d}-1}}{C(\log n)^{2p/d}}\big)$$

for an arbitrarily large constant $\beta > 0$ *and another constant* $C > 0$ *that depends on* β.

Choosing $t = n^{\frac{d-p}{d}} / \log n$ yields a nontrivial concentration result that suffices to prove complete convergence of typically smooth Euclidean functionals.

Corollary 4.6. *Let* $p, d \geq 1$*. Assume that* F *is typically smooth. Then*

$$\mathbb{P}\big(|\,\mathsf{F}(n) - \mathbb{E}(\mathsf{F}(n))| > n^{\frac{d-p}{d}} / \log n\big) \leq O\big(n^{-\beta} + \exp\big(-\tfrac{n}{C(\log n)^{2+\frac{2p}{d}}}\big)\big) \tag{3}$$

for any constant β *and* C *depending on* β *as in Theorem 4.5.*

4.3 Complete Convergence for $p \geq d$

In this section, we show that typical smoothness (Definition 4.4) suffices for complete convergence. This implies complete convergence of MST and PA by Lemma 4.8 below.

Theorem 4.7. *Let* $p, d \geq 1$*. Assume that* F *is typically smooth and* $\mathsf{F}(n)/n^{\frac{d-p}{d}}$ *converges in mean to* $\gamma_{\mathsf{F}}^{d,p}$*. Then* $\mathsf{F}(n)/n^{\frac{d-p}{d}}$ *converges completely to* $\gamma_{\mathsf{F}}^{d,p}$.

Although similar in flavor, smoothness in mean does not immediately imply typical smoothness or vice versa: the latter makes only a statement about *single* points at *worst-case* positions. The former only makes a statement about adding and removing *several* points at *random* positions. However, the proofs of smoothness in mean for MST and PA do not exploit this, and we can adapt them to yield typical smoothness.

Lemma 4.8. PA *and* MST *are typically smooth.*

Corollary 4.9. *For all d and p with $p \geq 1$, $\mathsf{MST}(n)/n^{\frac{d-p}{d}}$ and $\mathsf{PA}(n)/n^{\frac{d-p}{d}}$ converge completely to constants $\gamma^{d,p}_{\mathsf{MST}}$ and $\gamma^{d,p}_{\mathsf{PA}}$, respectively.*

5 Average-Case Ratio of the MST Heuristic

In this section, we show that the average-case approximation ratio of the MST heuristic for power assignments is strictly better than its worst-case ratio of 2. First, we prove that the average-case bound is strictly (albeit marginally) better than 2 for any combination of d and p. Second, we show a simple improved bound for the 1-dimensional case.

5.1 The General Case

The idea behind showing that the MST heuristic performs better on average than in the worst case is as follows: the weight of the PA graph obtained from the MST heuristic can not only be upper-bounded by twice the weight of an MST, but it is in fact easy to prove that it can be upper-bounded by twice the weight of the heavier half of the edges of the MST [6]. Thus, we only have to show that the lighter half of the edges of the MST contributes $\Omega(n^{\frac{d-p}{d}})$ to the value of the MST in expectation.

For simplicity, we assume that the number $n = 2m + 1$ of points is odd. The case of even n is similar but slightly more technical. We draw points $X = \{U_1, \ldots, U_n\}$ as described above. Let $\mathsf{PT}(X)$ denote the power required in the power assignment obtained from the MST. Furthermore, let H denote the m heaviest edges of the MST, and let L denote the m lightest edges of the MST. We omit the parameter X since it is clear from the context. Then we have

$$\mathsf{H} + \mathsf{L} = \mathsf{MST} \leq \mathsf{PA} \leq \mathsf{PT} \leq 2\,\mathsf{H} = 2\,\mathsf{MST} - 2\,\mathsf{L} \leq 2\,\mathsf{MST} \tag{4}$$

since the weight of the PA graph obtained from an MST can not only be upper bounded by twice the weight of a minimum-weight spanning tree, but it is easy to show that the PA graph obtained from the MST is in fact by twice the weight of the heavier half of the edges of a minimum-weight spanning tree [6]. We can show that $\mathbb{E}(\mathsf{L}) = \Omega(n^{\frac{d-p}{d}})$. This yields the following result.

Theorem 5.1. *For any $d \geq 1$ and any $p \geq 1$, we have*

$$\gamma^{d,p}_{\mathsf{MST}} \leq \gamma^{d,p}_{\mathsf{PA}} \leq 2(\gamma^{d,p}_{\mathsf{MST}} - C) < 2\gamma^{d,p}_{\mathsf{MST}}$$

for some constant $C > 0$ that depends only on d and p.

By exploiting that PA converges completely, we can obtain a bound on the expected approximation ratio from the above result.

Corollary 5.2. *For any $d \geq 1$ and $p \geq 1$ and sufficiently large n, the expected approximation ratio of the MST heuristic for power assignments is bounded from above by a constant strictly smaller than 2.*

5.2 An Improved Bound for the One-Dimensional Case

The case $d = 1$ is much simpler than the general case, because the MST is just a Hamiltonian path starting at the left-most and ending at the right-most point. Furthermore, we also know precisely what the MST heuristic does: assume that a point x_i lies between x_{i-1} and x_{i+1}. The MST heuristic assigns power $\mathsf{PA}(x_i) = \max\{|x_i - x_{i-1}|, |x_i - x_{i+1}|\}^p$ to x_i. The example that proves that the MST heuristic is no better than a worst-case 2-approximation shows that it is bad if x_i is very close to either side and good if x_i is approximately in the middle between x_{i-1} and x_{i+1}. By analyzing $\gamma_{\mathsf{MST}}^{1,p}$ and $\gamma_{\mathsf{PA}}^{1,p}$ carefully, we obtain the following theorem.

Theorem 5.3. *For all* $p \geq 1$*, we have* $\gamma_{\mathsf{MST}}^{1,p} \leq \gamma_{\mathsf{PA}}^{1,p} \leq (2 - 2^{-p}) \cdot \gamma_{\mathsf{MST}}^{1,p}$.

The high probability bounds for the bound of $2 - 2^{-p}$ of the approximation ratio of the power assignment obtained from the spanning tree together with the observation that in case of any "failure" event we can use the worst-case approximation ratio of 2 yields the following corollary.

Corollary 5.4. *The expected approximation ratio of the MST heuristic is at most* $2 - 2^{-p} + o(1)$.

6 Conclusions and Open Problems

We have proved complete convergence of Euclidean functionals that are *typically smooth* (Definition 4.4) for the case that the distance-power gradient p is larger than the dimension d. The case $p > d$ appears naturally in the case of transmission questions for wireless networks. As examples, we have obtained complete convergence for the MST and the PA functional. To prove this, we have used a recent concentration of measure result by Warnke [15]. His concentration inequality might be of independent interest to the algorithms community. As a technical challenge, we have had to deal with the fact that the degree of an optimal power assignment graph can be unbounded.

To conclude this paper, let us mention some problems for further research:

1. Is it possible to prove complete convergence of other functionals for $p \geq d$? The most prominent one would be the traveling salesman problem (TSP).
2. Is it possible to prove improved bounds on the approximation ratio of the MST heuristic?
3. Can our findings about power assignments be generalized to other problems in wireless communication, such as the k-station network coverage problem of Funke et al. [5], where transmit powers are assigned to at most k stations such that X can be reached from at least one sender, or power assignments in the SINR model [7, 9]? Interestingly, in the SINR model the MST turns out to be a good solution to schedule all links within a short time [8, 10]. More general, can this framework also be exploited to analyze other approximation algorithms for geometric optimization problems? As far as we are aware, besides partitioning heuristics [2, 16], the only other algorithm analyzed within this framework is Christofides' algorithm for the TSP [3].

References

1. Althaus, E., Calinescu, G., Mandoiu, I.I., Prasad, S.K., Tchervenski, N., Zelikovsky, A.: Power efficient range assignment for symmetric connectivity in static ad hoc wireless networks. Wireless Networks 12(3), 287–299 (2006)
2. Bläser, M., Manthey, B., Rao, B.V.R.: Smoothed analysis of partitioning algorithms for Euclidean functionals. Algorithmica 66(2), 397–418 (2013)
3. Bläser, M., Panagiotou, K., Rao, B.V.R.: A probabilistic analysis of christofides' algorithm. In: Fomin, F.V., Kaski, P. (eds.) Proc. of the 13th Scandinavian Workshop on Algorithm Theory (SWAT). Lecture Notes in Computer Science, vol. 7357, pp. 225–236. Springer (2012)
4. Clementi, A.E.F., Penna, P., Silvestri, R.: On the power assignment problem in radio networks. Mobile Networks and Applications 9(2), 125–140 (2004)
5. Funke, S., Laue, S., Lotker, Z., Naujoks, R.: Power assignment problems in wireless communication: Covering points by disks, reaching few receivers quickly, and energy-efficient travelling salesman tours. Ad Hoc Networks 9(6), 1028–1035 (2011)
6. de Graaf, M., Boucherie, R.J., Hurink, J.L., van Ommeren, J.K.: An average case analysis of the minimum spanning tree heuristic for the range assignment problem. Memorandum 11259 (revised version), Department of Applied Mathematics, University of Twente (2013)
7. Halldórsson, M.M., Holzer, S., Mitra, P., Wattenhofer, R.: The power of nonuniform wireless power. In: Proc. of the 24th Ann. ACM-SIAM Symp. on Discrete Algorithms (SODA). pp. 1595–1606. SIAM (2013)
8. Halldórsson, M.M., Mitra, P.: Wireless connectivity and capacity. In: Rabani, Y. (ed.) Proc. of the 23rd Ann. ACM-SIAM Symp. on Discrete Algorithms (SODA). pp. 516–526. SIAM (2012)
9. Kesselheim, T.: A constant-factor approximation for wireless capacity maximization with power control in the SINR model. In: Proc. of the 22nd Ann. ACM-SIAM Symp. on Discrete Algorithms (SODA). pp. 1549–1559. SIAM (2011)
10. Khan, M., Pandurangan, G., Pei, G., Vullikanti, A.K.S.: Brief announcement: A fast distributed approximation algorithm for minimum spanning trees in the SINR model. In: Aguilera, M.K. (ed.) Proc. of the 26th Int. Symp. on Distributed Computing (DISC). Lecture Notes in Computer Science, vol. 7611, pp. 409–410. Springer (2012)
11. Kirousis, L.M., Kranakis, E., Krizanc, D., Pelc, A.: Power consumption in packet radio networks. Theoretical Computer Science 243(1-2), 289–305 (2000)
12. Lloyd, E.L., Liu, R., Marathe, M.V., Ramanathan, R., Ravi, S.S.: Algorithmic aspects of topology control problems for ad hoc networks. Mobile Networks and Applications 10(1-2), 19–34 (2005)
13. Pahlavan, K., Levesque, A.H.: Wireless Information Networks. Wiley (1995)
14. Rappaport, T.S.: Wireless Communication. Prentice Hall (2002)
15. Warnke, L.: On the method of typical bounded differences. Computing Research Repository 1212.5796 [math.CO], arXiv (2012)
16. Yukich, J.E.: Probability Theory of Classical Euclidean Optimization Problems, Lecture Notes in Mathematics, vol. 1675. Springer (1998)