ARTICLE IN PRESS

Discrete Applied Mathematics xxx (xxxx) xxx

Contents lists available at ScienceDirect

Discrete Applied Mathematics

journal homepage: www.elsevier.com/locate/dam

Probabilistic properties of highly connected random geometric graphs

Bodo Manthey *, Victor M.J.J. Reijnders

University of Twente, Department of Applied Mathematics, Enschede, The Netherlands

ARTICLE INFO

Article history:
Received 18 April 2018
Received in revised form 14 November 2018
Accepted 10 April 2019
Available online xxxx

Keywords:
Euclidean optimization problems
Power assignments
Probabilistic analysis

ABSTRACT

We study the probabilistic properties of reliable networks of minimum costs in d-dimensional Euclidean space, with reliability in terms of k-edge-connectivity in graphs. We show that this problem fits into Yukich's framework for Euclidean functionals for arbitrary k, dimension d and distant-power gradient p with $p < d$. With this framework, several theorems on convergence and concentration of the value of optimal solutions follow.

These results are then extended to optimal k-edge-connected power assignment graphs, where we assign transmit power to nodes, and two nodes are connected if they both have sufficient transmit power. This variant models wireless networks.

Finally, we devise a partitioning heuristic to find approximate solutions quickly, and we analyze its performance in the framework of smoothed analysis.

1. Introduction

The design of fault-tolerant networks is an important issue in today's research due to their numerous applications. The goal is to find cheap and reliable networks with some specific characteristics. Reliability is generally expressed in terms of the connectivity of the network. For example, we might want to have multiple paths between each pair of nodes to account for possible failures in a link. Applications of this type of problems can be found in the design of reliable communication and transportation networks [2,11,12].

Wireless ad hoc networks have also received significant attention in recent studies [5,10,18]. Instead of direct connections between nodes, communication can also take place by relaying through intermediate nodes. Here we assign a transmission power to each node. As the transmission range is directly related to the power used by a node, the goal is to find a fault tolerant network with minimum total power usage. Possible applications are environmental monitoring, emergency disaster relief where wiring is difficult, communication between mobile computers for conferencing and home networking, wireless sensor networks, multi-hop extensions of cellular telecommunication systems, and networks of vehicles [3,10]. Metricom Inc.'s Ricochet network and the Army Near-Term Digital Radio network are examples of fully operational multi-hop wireless networks [18].

In this paper, we study the probabilistic properties of the value of optimal solutions of the problems mentioned above within Yukich's framework for Euclidean optimization problems. This yields results both about convergence and concentration of the value of optimal solutions as a function of the number n of nodes, dimension d of the underlying Euclidean space, distance-power gradient p, and connectivity requirement k.

* Corresponding author.
E-mail addresses: b.manthey@utwente.nl (B. Manthey), v.m.j.j.reijnders@utwente.nl (V.M.J.J. Reijnders).

https://doi.org/10.1016/j.dam.2019.04.007
0166-218X/

Please cite this article as: B. Manthey and V.M.J.J. Reijnders, Probabilistic properties of highly connected random geometric graphs, Discrete Applied Mathematics (2019), https://doi.org/10.1016/j.dam.2019.04.007.

Finding a cheapest k-edge-connected network is NP-hard [9], and so is finding a minimal power wireless network [6]. As we still want to have reasonably good solutions in acceptable computation time, we need to find a good approximation algorithm. Partitioning algorithms are a simple, easy-to-implement type of heuristic that show good performance for Euclidean optimization problems [4]. We devise a partitioning heuristic for the design of optimal k-edge-connected networks. Furthermore, we analyze it in the framework of smoothed analysis [16,22] in order to explain its performance.

The rest of this paper is organized as follows. In Section 2 we give the relevant definitions. We summarize related work in Section 3. The properties of k-edge-connected graphs are in Section 4. The partitioning heuristic and its smoothed analysis are presented in Section 5. We extend these results to k-edge-connected power graphs in Section 6. We conclude with some open problems (Section 7).

2. Definitions

Let $V \subseteq \mathbb{R}^d$ be a finite set of nodes, where $d \in \mathbb{N}$ is an arbitrary constant. In the rest of the paper, $n = |V|$ is the number of nodes. For two nodes $u, v \in V$, let (u, v) denote the edge connecting u and v, and let $|(u, v)|$ denote the Euclidean distance between u and v.

A graph $G = (V, E)$ is called k-edge-connected if G is connected after removal of any set of at most $k-1$ of its edges. An alternative characterization of k-edge-connectedness is that there exist k edge-disjoint paths between every pair of nodes. In this paper, we also call any complete graph k-edge-connected, even if it contains fewer than k nodes. (Otherwise, no k-edge-connected graphs below a certain size exist, which would cause technical issues.)

We study k-edge-connected graphs of minimal costs. The cost of an edge (u, v) is $|(u, v)|^p$, where $p > 0$ denotes the distant-power gradient. For a given graph $G = (V, E)$, the costs of this graph are the sum of its edge costs, i.e., $\sum_{(u,v)\in E} |(u, v)|^p$. Then $k\text{-EC}^p(V)$ is the minimum cost of any k-edge-connected graph on V with costs computed with distant-power gradient p.

In the remainder of this paper, $k \in \mathbb{N}$, $d \in \mathbb{N}$, and $p > 0$ are assumed to be arbitrary, but fixed constant.

Besides the model above, where we pay per edge, and which could be viewed as modeling wired networks, we also consider a model where we assign transmit power $\text{PA}(v)$ to node v for each v. Two nodes u and v are connected by an edge if both have sufficient transmit power to reach each other, i.e., if $\text{PA}(u), \text{PA}(v) \geq |(u, v)|^p$. The costs of such a power assignment is the sum of all transmit powers, i.e., $\sum_{v\in V} \text{PA}(v)$. The graph resulting from a power assignment is called the corresponding power assignment graph.

$k\text{-ECPA}^p(V)$ denotes the minimum costs of any power assignment whose corresponding power assignment graph is k-edge-connected.

Both $k\text{-ECPA}^p$ and $k\text{-EC}^p$ are Euclidean functionals. This means that they map a finite point set to a non-negative real number, are translation invariant, and scaling all points by a factor of $\alpha > 0$ changes the costs by a factor of α^p [24].

Following Yukich [24], for a Euclidean functional L^p, we write $L^p(V, R)$ to denote the functional on $V \cap R$, where R is some hyperrectangle. Usually, $R = [0, 1]^d$, and we omit R if it is clear from the context. In the following, we denote by $\operatorname{diam} R$ the diameter of R.

In order to fit $k\text{-EC}^p$ and $k\text{-ECPA}^p$ into Yukich's framework for Euclidean optimization problems [24], we have to define corresponding canonical boundary functionals, an idea first articulated in Redmond's thesis [19]. Roughly speaking, in these functionals, the entire boundary of the rectangle is considered as one additional node that can be used. To distinguish between a functional and its boundary functional, we refer to the former as the original functional.

Given a hyperrectangle $R \subseteq \mathbb{R}^d$ and a point set $V \subseteq R$, a boundary graph is a graph with nodes V plus the boundary ∂R of R as an additional node. We view the boundary as a single node. For a node to be connected to the boundary, it is sufficient that this node is connected to an arbitrary point of ∂R. A boundary graph is called k-edge-connected if the graph restricted to V is k-edge-connected, or if the graph on $V \cup \{\partial R\}$ is k-edge-connected. Here, any edge connecting $v \in V$ to ∂R counts as up to k independent edges. In Fig. 1, you can find an example of a 3-edge-connected boundary graph. We denote by $k\text{-EC}^p_B$ the boundary functional corresponding to $k\text{-EC}^p$. This means that $k\text{-EC}^p_B(V, R)$ is the minimum total length of a k-edge-connected graph in terms of the sum of the edge lengths on $V \cup \partial R$ raised to the power p. A node v is connected to the boundary ∂R by adding edge (v, v_∂) where $v_\partial = \arg\min_{w\in\partial R} |(v, w)|$.

Similarly, $k\text{-ECPA}^p_B$ is the boundary functional of $k\text{-ECPA}^p$. Here, v is connected to the boundary if $\text{PA}(v) \geq |(v, v_\partial)|^p$.

3. Related work

For a survey about properties of k-edge-connected graphs, we refer to Kammer and Täubig [14].

Finding k-edge-connected networks of minimum costs is NP-hard for $k \geq 2$, both in the classical and the power assignment model [6,9]. Therefore, a considerable amount of research has been focused on approximation algorithms.

Khuller and Vishkin [15] proved that the problem of finding a minimum-cost k-edge connected graph can be approximated within a factor of 2. Czumaj and Lingas [7] gave a polynomial-time approximation scheme for this problem for the Euclidean case, where points are contained in $\mathbb{R}^d$ for some fixed d.

Althaus et al. [1] devised several approximation algorithms for the problem of finding a connected power assignment graph. Fuchs [8] proved NP-hardness of the power assignment problem for simple connectivity for $d = 2$ and APX-hardness for $d \geq 3$. Santi et al. [21] studied the connectivity of power assignment graphs in the Euclidean case for

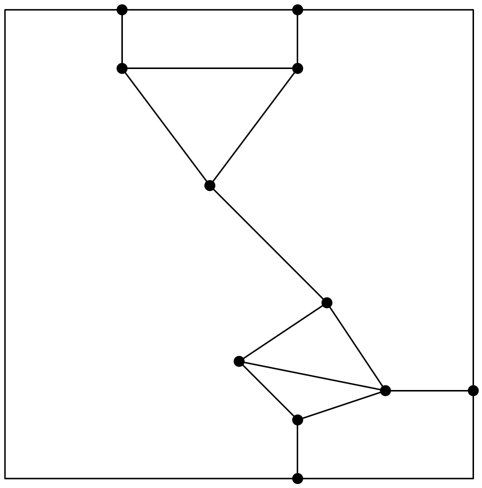

Fig. 1. A 3-edge-connected boundary graph on 7 nodes.

$d \in \{1, 2, 3\}$ under the restriction that every node is assigned the same power r. They derived bounds for r to achieve connectivity with high probability. De Graaf and Manthey [10] analyzed connectivity in power assignment graphs. They proved properties similar to our results of Section 4, but only for simple connectivity.

Călinescu and Wan [5] analyzed several approximation algorithms for finding cheap k-edge-connected power assignment graphs and obtained an approximation ratio of $2k$ for this problem.

4. Properties of k-EC^p

In this section, we show that k-EC^p fits into Yukich's framework for Euclidean functionals [24].

In the following, we make heavy use of the following lemma, which states that the union of two k-edge-connected graphs with non-empty intersection is also k-edge-connected.

Lemma 4.1 (*Matula [17]*)**.** *Let $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$ be k-edge-connected graphs with $V_1 \cap V_2 \neq \emptyset$. Then the graph $H = (V_1 \cup V_2, E_1 \cup E_2)$ is k-edge-connected.*

4.1. Yukich's framework

For a functional to fit into Yukich's framework, we have to prove that it is subadditive, has a superadditive boundary functional, is pointwise close to its boundary functional, and is smooth. With these properties, we obtain several convergence and concentration results. These properties are worked out in the remainder of this section.

First, we prove that k-EC^p is geometrically subadditive. Roughly speaking, this shows that the function value of a set is not larger – up to an additive error term – than the sum of the function values of a partition of this set.

Lemma 4.2. *For $p \geq 1$, k-EC^p is geometrically subadditive, i.e. for all finite sets V, all rectangles R and all partitions of R into rectangles R_1 and R_2 we have*

$$k\text{-EC}^p(V, R) \leq k\text{-EC}^p(V, R_1) + k\text{-EC}^p(V, R_2) + C_1(\operatorname{diam} R)^p,$$

where $C_1 = C_1(d, p)$ is a constant.

Proof. Consider the graphs (V_1, E_1) and (V_2, E_2) with $V_i = V \cap R_i$ that realize the optimal solutions of k-EC$^p(V, R_1)$ and k-EC$^p(V, R_2)$, respectively. Without loss of generality, we assume that $|V_1| \geq |V_2|$. We distinguish three cases.

1. $|V_1|, |V_2| \geq k + 1$. We join (V_1, E_1) and (V_2, E_2) by k vertex-disjoint edges $e_1, \ldots, e_k$. This results in a k-edge-connected graph on V. We have $|e_i| \leq \operatorname{diam} R$. Thus, the cost of this k-edge-connected graph is bounded from above by

$$\begin{aligned} k\text{-EC}^p(V, R) &\leq k\text{-EC}^p(V, R_1) + k\text{-EC}^p(V, R_2) + \sum_{i=1}^{k} |e_i|^p \\ &\leq k\text{-EC}^p(V, R_1) + k\text{-EC}^p(V, R_2) + kd^{p/2}(\operatorname{diam} R)^p. \end{aligned}$$

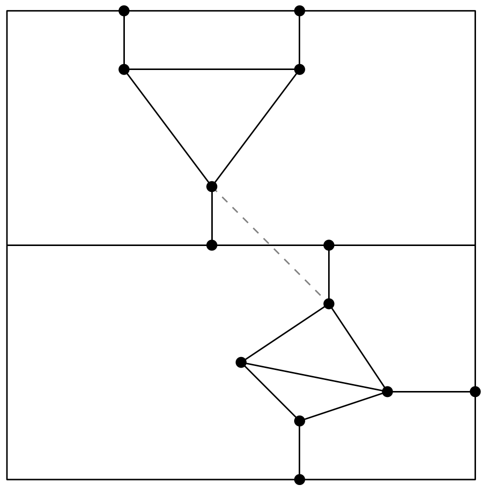

Fig. 2. Two 3-edge-connected boundary graphs after partitioning the graph from Fig. 1 into two rectangles.

2. $|V_1| \geq k+1$ and $|V_2| \leq k$. We know that (V_2, E_2) is complete, as this is the only way a graph with $|V_2| \leq k$ can be k-edge-connected. For each node $v_i \in V_2$, we add k edges $e_{i_1}, \dots e_{i_k}$ with endpoints in V_1. Each pair of nodes of V_1 has at least k edge-disjoint paths in E_1. By adding e_{i_j}, each pair in V has k edge-disjoint paths as well. At most k^2 edges are added this way. As each edge has a length of at most $\operatorname{diam} R$ and k is constant, we can bound the costs of the k-edge-connected graph obtained in a similar way as in the first case.
3. $|V_1|, |V_2| \leq k$. Both (V_1, E_1) and (V_2, E_2) are complete, and we know that a complete graph is always k-edge-connected. If we make the combined graph complete, then it is k-edge-connected by definition. We have to add at most k^2 edges in this way, and $k\text{-EC}^p(V, R)$ can be bounded as in the previous cases. □

Subadditivity directly yields an upper bound for the growth of k-EC^p as a function of $|V|$.

Lemma 4.3 (*Growth Bound*)**.** *Let $0 < p < d$. Then there exists a constant $C = C(d, p)$ such that for all cubes $R \subset \mathbb{R}^d$ and all $V \subset R$, we have*

$$k\text{-EC}^p(V, R) \leq C(\operatorname{diam} R)^p |V|^{(d-p)/d}.$$

Proof. This lemma follows directly by combining Lemma 4.2 and a result by Yukich [24, Lemma 3.3]. □

Unfortunately, likewise to the Euclidean functionals considered by Yukich [24], k-EC^p is not superadditive. If it were superadditive, then this together with subadditivity makes the functional nearly additive in the sense that $k\text{-EC}^p(V, R) \approx k\text{-EC}^p(V, R_1) + k\text{-EC}^p(V, R_2)$. We could then approximate the optimal solution value of the whole set by the sum of optimal solutions on its partitions.

The standard way to superadditivity is via the canonical boundary functional of k-EC^p_B, which is superadditive according to the following lemma.

Lemma 4.4. *For $p \geq 1$, k-EC^p_B is superadditive, i.e. for all finite sets V, all rectangles $R \subseteq \mathbb{R}^d$ and all partitions of R into rectangles R_1 and R_2 we have*

$$k\text{-EC}^p_B(V, R) \geq k\text{-EC}^p_B(V, R_1) + k\text{-EC}^p_B(V, R_2).$$

Proof. Let V and R together with a partition of R into rectangles R_1 and R_2 be given. Let $V_i = V \cap R_i$ for $i \in \{1, 2\}$. Let E be the edge set of a boundary k-edge-connected graph on V or $V \cup \{\partial R\}$ of minimum costs. For each edge $(u_1, u_2) \in E$ with $u_i \in V_i$, we add two edges: One edge connecting u_1 to the closest point ∂u_1 on the boundary of R_1, and one edge connecting u_2 to the closest point ∂u_2 on the boundary of R_2 (see Fig. 2).

The two induced subgraphs are k-edge-connected, as both still have k edge-disjoint paths between any pair of their nodes, albeit via ∂R_i for $i \in \{1, 2\}$.

The sum of the costs of these two graphs does not exceed the costs of k-EC$^p(V, R)$ by the triangle inequality and because $p \geq 1$. □

Next, we show that k-EC^p and k-EC^p_B are pointwise close. Together with subadditivity of k-EC^p and superadditivity of k-EC^p_B, this yields that both are approximately subadditive and superadditive.

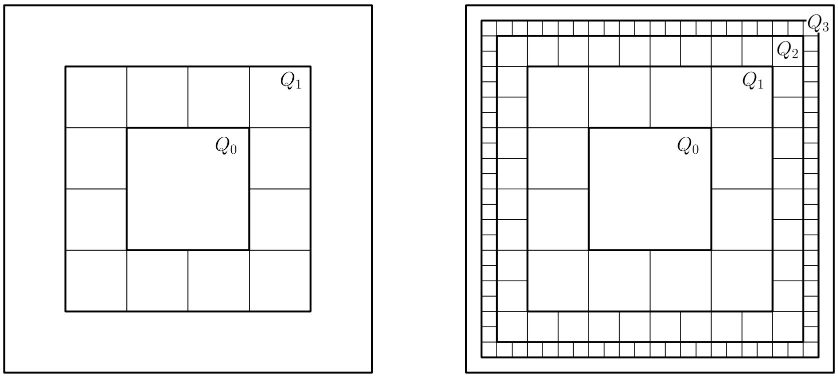

(a) Subdivision up to Q_1. (b) Subdivision up to Q_3.

Fig. 3. Dyadic subdivision of $[0,1]^2$.

Lemma 4.5. *For $1 \le p < d$ and all rectangles $R \subseteq \mathbb{R}^d$, k-EC^p is pointwise close to k-EC^p_B, i.e. for all finite sets $V \subset R$, we have*

$$\left|k\text{-EC}^p(V,R) - k\text{-EC}^p_B(V,R)\right| = o(\mathrm{diam}(R)^p \cdot |V|^{(d-p)/d}).$$

Proof. By scaling, we can restrict ourselves to $R = [0,1]^d$. Let $V \subseteq R$ be a finite set of points. Clearly, k-EC$^p(V,R) \ge k$-EC$^p_B(V,R)$, as each solution for a k-edge-connected graph is also a solution for a k-edge-connected boundary graph. Thus, we only have to prove k-EC$^p(V,R) \le k$-EC$^p_B(V,R) + o(|V|^{(d-p)/p})$. If $|V| \le k$, then this holds because $|V|$ is constant.

Thus, we assume that $|V| \ge k+1$ from now on. We first prove the following claim.

Claim 4.6. *Let $V \subseteq [0,1]^d$, $|V| = n$, and $1 \le p < d$. Consider a graph $G = (V,E)$ that realizes the optimal solution of k-EC$^p_B(V,R)$ for $R = [0,1]^d$. Then the sum of the pth powers of the lengths of the edges connecting nodes in V with the boundary of R is bounded by $O(n^{(d-p-1)/(d-1)})$.*

Proof. The proof is almost identical to a proof by Yukich [24, Lemma 3.8] and depends on a dyadic subdivision of $[0,1]^d$. Let Q_0 by the cube of side length $1/3$ and centered within $R = [0,1]^d$. Let Q_1 be the cube of side length $2/3$, also centered within R. We partition $Q_1 - Q_0$ into subcubes of side length $1/6$. The number of such subcubes is bounded by $C6^{d-1}$ for some constant $C = C(d)$.

We continue with this subdivision recursively. This means that at the jth stage we define cube Q_j of side length $1 - 2(3 \cdot 2^j)^{-1}$ and partition $Q_j - Q_{j-1}$ into subcubes of side length $(3 \cdot 2^j)^{-1}$. In Fig. 3, this subdivision is shown for $d = 2$. The number of such subcubes is bounded from above by $C3^{d-1}2^{j(d-1)}$. We carry out this recursion until the ℓ-th stage, where ℓ is the unique integer satisfying $2^{(\ell-1)(d-1)} \le n \le 2^{\ell(d-1)}$.

This procedure produces nested cubes $Q_1 \subseteq Q_2 \subseteq \ldots \subseteq Q_\ell$. It produces a dyadic covering of the cube until the moat $R - Q_\ell$ has a width of $O(n^{-1/(d-1)})$. We use these properties to prove Claim 4.6 as follows.

This dyadic subdivision partitions the largest cube Q_ℓ into at most

$$\sum_{j=0}^{k} C3^{d-1}2^{j(d-1)} \le Cn$$

subcubes, each with a side length equal to the distance between the subcube and the boundary of R. Furthermore, by partitioning each subcube into $(k2^y)^d$ congruent subcubes, where y is the least integer satisfying $2^y \ge d^{1/2}$, we obtain a partition $\mathcal{P}$ of Q_ℓ consisting of at most Cn subcubes with the property that k times the diameter of each subcube is less than the distance to the boundary.

Now let $G = (V,E)$ be a graph that realizes an optimal solution of k-EC$^p_B(V,R)$. We observe that in G, each subcube Q in $\mathcal{P}$ contains at most k points in V that are rooted to the boundary. If there were more than k points in $V \cap Q$ rooted to the boundary, we can do the following. We know that these points do not have edges directly between them as they already have k edge-disjoint paths between them, and edges between them would then not all be in G. So we can take one of them and connect it to all the other points rooted to the boundary (of which there are at least k), while removing the connection to the boundary. As the diameter of each subcube is less than $1/k$ times the distance to the boundary, this relinking gives us a cheaper solution, contradicting the optimality of G.

The sum of the pth powers of the lengths of the edges connecting nodes in $V \cap (Q_j - Q_{j-1})$ with the boundary is thus bounded by the product of the number of subcubes in $Q_j - Q_{j-1}$ and the pth power of the common diameter of the subcubes, namely

$$C3^{d-1}2^{j(d-1)} \cdot (3 \cdot 2^j)^{-p}.$$

Summing over all $1 \leq j \leq \ell$ yields a bound for the sum of the pth power of the lengths of the edges connecting points to the boundary in $V \cap Q_\ell$:

$$\sum_{j=1}^{\ell} C 3^{d-1} 2^{j(d-1)} \cdot (3 \cdot 2^j)^{-p} \leq \begin{cases} C \max\{n^{(d-p-1)/(d-1)}, \log n\} & \text{if } 1 \leq p \leq d-1 \\ C & \text{if } d-1 < p < d. \end{cases} \tag{1}$$

The $\log n$ term is needed to cover the case $p = d-1$. The sum of the pth powers of the lengths of the edges connecting nodes in $V \cap (R - Q_\ell)$ with the boundary is at most the product of $n = |V|$ and the pth power of the width of the moat $R - Q_\ell$, i.e. at most

$$n \cdot Cn^{-p/(d-1)} = Cn^{(d-p-1)/(d-1)}. \tag{2}$$

Combining (1) and (2) proves the claim. □

We can now continue with the proof of Lemma 4.5. Consider the set $U \subset V$ of all nodes connected to the boundary. Let $\mathcal{B} \subseteq \partial R$ be the set of points on the boundary to which these nodes are connected. Then $|U| \geq |\mathcal{B}|$. As we want to remove all edges to the boundary to get to a solution for our original functional, we have to add new edges to maintain k-edge-connectivity. Recall that edges to the boundary can count as up to k edges. As in the proof of Claim 4.6, we know that there are no edges between the points of U. To get a good bound on the increase of costs that incurs by adding these edges, we first prove the following claim.

Claim 4.7. *Fix $1 \leq p < d$ and let $G = (V, E)$ be a k-edge-connected graph realizing the optimal solution for k-$\mathrm{EC}^p(V, R)$, where $R = [0, 1]^d$. Then there exists a constant $c = c(k, d)$ such that the degree of every node $v \in V$ is bounded by c.*

Proof. Let us assume to the contrary that there exists a node $v \in V$ for which the degree is not bounded by c. Consider cones originating from v with aperture $\pi/6$:

$$C_r = \left\{ x \;\middle|\; \frac{(x - v) \cdot r}{\|x - v\|_2} \geq \cos(\pi/6) \right\}.$$

Here, r denotes a unit vector in the direction of the center of the cone. The aperture of $\pi/6$ implies that for every two $x, y \in C_r$, we have $\angle(x, y, v) \leq \pi/3$. Since the aperture of the cones that we consider is constant, a constant number of cones (where the exact number depends only on the dimension d) cover $\mathbb{R}^d$ (see also de Graaf and Manthey [10]).

We consider a cone C with an unbounded number of points connected to v and look at the point y that is farthest from v and connected to v. Such a cone has to exist as the number of cones is bounded. We distinguish two cases.

In the first case, the degree of y is greater than c as well, and y is connected to all nodes in C that are also connected to v. This means that both are connected to more than k nodes in C. In this case, we can remove the edge between v and y as we would still have more than k edge-disjoint paths between v and y (and other points will not be affected). Removing an edge would only lower the cost for the graph, so this would contradict the optimality of the solution.

In the other case, we can find a node z to which y is not connected, but v is connected to z. We consider the triangle $\Delta(v, z, y)$. As we know that $\angle(z, v, y) < \pi/3$, either $\angle(z, y, v) > \pi/3$ or $\angle(v, z,)y > \pi/3$ (or both). Using that $|(v, y)| \geq |(v, z)|$, we can see that $|(y, z)| < |(v, y)|$. If we then replace the edge from v to y by the edge from y to z, the number of edge-disjoint paths from y to z or v cannot decrease. But as $|(y, z)| < |(v, y)|$, we also have $|(y, z)|^p < |(v, y)|^p$. Thus, we have lowered the cost and still have a k-edge-connected graph. This again would contradict the optimality of G. □

With this lemma we can continue the proof of Lemma 4.5. We can also get a bound on the length of each edge we have to add. Using the triangle inequality for powers of metrics, we obtain

$$\sum_{(u,v)\in M} |(u, v)|^p \leq 2^{p-1} \left(\sum_{(u,\cdot)\in M} |(u, u_B)|^p + \sum_{(u,v)\in M} |(u_B, v)|^p \right) \tag{3}$$

$$\begin{aligned}
&\leq 4^{p-1} \left(\sum_{(u,\cdot)\in M} |(u, u_B)|^p + \sum_{(u,v)\in M} |(u_B, v_B)|^p + \sum_{(\cdot,v)\in M} |(v_B, v)|^p \right) \\
&\leq 4^p \sum_{(u,\cdot)\in M} |(u, u_B)|^p + 4^{p-1} k\text{-}\mathrm{EC}^p(\mathcal{B}, R) \\
&\leq 4^p C_1 \sum_{u\in U} |(u, u_B)|^p + 4^{p-1} C_2 |V|^{(d-p-1)/(d-1)} && \text{(by Lemma 4.3)} \\
&\leq 4^p C_1 C_3 |V|^{(d-p-1)/(d-1)} + o(|V|^{(d-p)/d}) && \text{(by Claim 4.6)} \\
&\leq o(|V|^{(d-p)/d}).
\end{aligned}$$

Now we have changed a graph that realizes an optimal solution for $k\text{-EC}^p_B(V, R)$, where $R = [0, 1]^d$, into a k-edge-connected graph without connection to the boundary. The cost of this new graph provides an upper bound for $k\text{-EC}^p(V, R)$. Observing that we have increased the costs by at most $o(|V|^{(d-p)/d})$ concludes the proof. □

We have shown geometric subadditivity, superadditivity, and pointwise closeness. The last property that we need is smoothness. Roughly speaking, smoothness means that adding or removing a few nodes does not change the function value by much.

Lemma 4.8. *For $1 \le p < d$, $k\text{-EC}^p$ is smooth, i.e. for all finite sets $U, V \subseteq [0, 1]^d$ we have*

$$\left|k\text{-EC}^p(U \cup V) - k\text{-EC}^p(U)\right| = O(|V|^{(d-p)/d}).$$

Proof. Subadditivity (Lemma 4.2) and the growth bound (Lemma 4.3) together yield

$$k\text{-EC}^p(U \cup V) \le k\text{-EC}^p(U) + k\text{-EC}^p(V) + O(1) \le k\text{-EC}^p(U) + O\big(|V|^{(d-p)/d}\big).$$

It remains to be shown that

$$k\text{-EC}^p(U) - k\text{-EC}^p(U \cup V) \le O(|V|^{(d-p)/d}).$$

We start with a graph $G = (U \cup V, E)$ that realizes $k\text{-EC}^p(U \cup V)$. After removal of V, we show that we can modify the remaining graph to obtain a k-edge-connected graph on U without increasing the cost by more than $O(|V|^{(d-p)/d})$. Let $N_V \subseteq U$ be the set of direct neighbors of nodes in V. By Claim 4.7 we know that $|N_V| \le c|V|$ for some constant c, that only depends on k and d. Let $m = |N_V|$.

We now compute a k-edge-connected graph T of minimum cost on N_V. By Lemma 4.3, the cost of this graph is bounded by $O(m^{(d-p)/d})$. Let F be the set of edges of the subgraph of G induced by U. Clearly, the graph with edge set $T \cup F$ has weight at most $k\text{-EC}^p(U \cup V) + O(|V|^{(d-p)/d})$. It remains to be proved that it is k-edge-connected. Fix any two nodes $u, v \in V$. We distinguish three cases:

1. Both u and v are in N_V. As T is k-edge-connected graph, there are k edge-disjoint paths connecting u to v.
2. Only one of u and v is in N_V. Without loss of generality, let $v \in N_V$ and $u \in U \setminus N_V$. Consider k edge-disjoint paths $P_1, \dots, P_k$ from u to v in $(U \cup V, E)$. Let q_i be the first node of N_V that P_i reaches. The nodes $q_1, \dots, q_k$ are not necessarily distinct, and we can have $q_i = v$. However, since T is k-edge-connected, there exist k edge-disjoint paths within T connecting q_i to v for each $i \in \{1, \dots, k\}$.
3. Both u and v are not in N_V. Take any $x \in N_V$, then Item 2. yields that there are k edge-disjoint paths from u to x and from x to v. By Lemma 4.1, we know there are also k edge-disjoint paths from u to v. □

Smoothness for the boundary functional now follows with an almost identical proof.

Lemma 4.9. *For $1 \le p < d$, $k\text{-EC}^p_B$ is smooth.*

Proof. The proof is similar to the proof of Lemma 4.8. For removing V, we ignore the possible connections to the boundary and create a boundary k-edge-connected graph on N_V for $k\text{-EC}^p_B$. □

4.2. Limit theorems

For the theorems on the convergence of optimal solutions, we need the notion of complete convergence. Let $(X_n)_{n\in\mathbb{N}}$ be a sequence of random variables. Then X_n converges completely (c.c.) to a constant C if and only if for all $\epsilon > 0$ we have

$$\sum_{n=1}^{\infty} P\big(|X_n - C| > \epsilon\big) < \infty.$$

This notion of convergence was first introduced by Hsu and Robbins [13].

As we have geometric subadditivity, superadditivity, pointwise closeness, and smoothness, several limit theorems directly follow, which are given in this section. The results show that the functional on random points is highly concentrated around its expected value.

Theorem 4.10. *Let V be a set of n points drawn independently and uniformly from $[0, 1]^d$ and $R = [0, 1]^d$. Fix $1 \le p < d$ and $k \in \mathbb{N}$. Then there exists a positive constant $\alpha = \alpha(k\text{-EC}^p, d, k)$ such that*

$$\lim_{n\to\infty} \frac{k\text{-EC}^p(V, R)}{n^{(d-p)/p}} = \alpha \quad \text{c.c., and}$$

$$\lim_{n\to\infty} \frac{k\text{-EC}^p_B(V, R)}{n^{(d-p)/p}} = \alpha \quad \text{c.c.}$$

Proof. As k-EC^p is a smooth subadditive Euclidean functional and pointwise close to its smooth superadditive boundary functional, we can use a result by Yukich [24, Theorem 4.1] to directly get this theorem. □

The following concentration of measure result is due to Rhee [20]. It states that the solution value is not far from its expected value.

Theorem 4.11. *Fix $d \geq 2$, $1 \leq p < d$ and $k \in \mathbb{N}$. Let V be a set of n points drawn independently and uniformly from $[0, 1]^d$. Then there exist constants $c_1 = c_1(k\text{-EC}^p, d)$ and $c_2 = c_2(k\text{-EC}^p, d)$ such that for all $t > 0$ we have*

$$\mathbb{P}\left(\left|k\text{-EC}^p\left(\{X_1, \dots, X_n\}\right) - \mathbb{E}\left[k\text{-EC}^p\left(\{X_1, \dots, X_n\}\right)\right]\right| > t\right) \leq c_1 \exp\left(\frac{-c_2 t^{2d/(d-p)}}{n}\right).$$

Proof. As k-EC^p is a smooth subadditive Euclidean functional which is pointwise close to its superadditive boundary functional, we can use Rhee's theorem [20] to obtain this result. □

5. Partitioning heuristics

Partitioning heuristics are a generic approach to design heuristics for Euclidean optimization problems: The d-dimensional Euclidean space is divided into a number of cells such that each cell contains only a small number of points. This allows us to compute quickly optimal solutions for the set of points in each cell. Finally, the solutions of the individual cells are combined to obtain a solution of the whole set of points. We describe this more formally in Algorithm 5.1.

Algorithm 5.1 (*Partitioning Scheme*)**. Input:** set $V \subseteq [0, 1]^d$ of n points and number $s > k$.

1. Partition $[0, 1]^d$ into $\ell = \sqrt[d]{n/s}$ stripes of dimension $d - 1$ such that each stripe contains exactly $n/\ell = (n^{d-1}s)^{1/d}$ points.
2. Keep partitioning each $i+1$-dimensional stripe into ℓ stripes of dimension i such that each stripe contains exactly $n/\ell^i = (n^{d-i}s^i)^{1/d}$ points. Stop at $i = 1$ so that each 2-dimensional stripe is partitioned into ℓ cells with $n/\ell^d = s$ points. In this way we end up with $\ell^d = n/s$ cells.
3. Compute a graph achieving the optimal solution of k-EC^p for each cell.
4. Join the graphs to obtain a k-edge-connected graph on V as follows: Choose k points of each cell. Connect these k points to the k points of an adjacent cell such that the graph becomes k-edge-connected.

Overall, we obtain the following additive upper bound on the approximation performance.

Theorem 5.2. *Let $s > k$, and let $1 \leq p < d$. The partitioning heuristic (Algorithm 5.1) for k-EC^p can be implemented to run in time $2^{O(s^2)} + O(n)$. Furthermore, let PSE$^p(V)$ denote the cost of the k-edge-connected graph computed by Algorithm 5.1. Then PSE$^p(V) \leq k$-EC$^p(V) + O\left((n/s)^{\frac{d-p}{d}}\right)$.*

Proof. First, the graph that we get as an output from Algorithm 5.1 is k-edge-connected because of Lemma 4.1. Thus, a feasible solution is computed.

Second, a simple brute-force search shows that an optimal k-edge-connected graph of s points can be computed in time $2^{O(s^2)}$. The joining can easily be done in linear time.

Third, let PSE$^p(V)$ be the cost of the k-edge-connected graph on V computed by Algorithm 5.1. By using subadditivity, PSE$^p(V)$ is bounded from above by k-EC$^p(V) + C_1$ plus the costs of joining the optimal solutions in the cells (Step 4.). Here, $C_1 = C_1(d, p) > 0$ is the constant of Lemma 4.2.

The joining (Step 4.) can be implemented to yield additional costs of at most $O((n/s)^{(d-p)/d})$ using Lemma 4.3. We do this by lacing cells together in a snakelike succession. We create a matching of k nodes between two succeeding cells. (Both the choice of the k nodes of each cell and the exact matching is arbitrary.) This yields a k-edge-connected graph on the cells that we lace together and completes the proof. □

The running-time of the partitioning algorithm is polynomial for $s = \Omega(n/\sqrt{\log n})$.

The above theorem does not yield a good approximation ratio in the worst case. If the value of an optimum solution is small compared to the additive error term, then the approximation guarantee is poor. However, typically, this is not the case, and partitioning heuristics work quite well on typical instances.

In order to explain this, partitioning heuristics for Euclidean optimization problems have been analyzed in the framework of smoothed analysis for Euclidean optimization problems introduced by Bläser et al. [4]. They use the so-called one-step model to construct smoothed instances: Let $\phi \geq 1$ be a perturbation parameter. For each of the n points, an adversary specifies a probability density function $[0, 1]^d \to [0, \phi]$. Then the points are drawn independently according to their respective probability density function. The smoothed performance is then the maximum expected performance that the adversary can achieve by choosing the density functions. The parameter ϕ limits the power of the adversary: If $\phi = 1$, then the adversary can only choose the uniform distribution on the unit hypercube $[0, 1]^d$. For larger ϕ, the adversary can concentrate more and more probability mass and, thus, is able to specify (worst-case) instances more accurately.

For ease of presentations, we restrict ourselves to the case $d = 2$ and $p = 1$ in the remainder of this section, because Bläser et al. [4] also stated their results only for $d = 2$. It is quite straightforward to generalize the results to larger values of d, but it seems to be non-trivial to generalize them to $p > 1$. We obtain the following result.

Theorem 5.3. *For $p = 1$, $d = 2$, and $s = \Theta(\sqrt{\log n})$, Algorithm 5.1 has polynomial running-time and achieves a smoothed approximation ratio of $1 + O\left(\sqrt{\frac{\phi}{\sqrt{\log n}}}\right)$.*

Proof. The polynomial running-time follows from the discussion above. The approximation ratio follows from a result by Bläser et al. [4, Theorem 3.8] together with Theorem 5.2. □

6. Extension to k-ECPA

The results for k-EC^p are easily copied to k-ECPA^p by making some small adjustments to the proofs. We start with subadditivity.

Lemma 6.1. *For $p \geq 1$, k-ECPA^p is a geometric subadditive functional.*

Proof. We follow the proof of Lemma 4.2. We need to increase $\mathrm{PA}(v)$ accordingly for the edges that need to be added in the power assignment graph. We define $d_{uv} = \max\{0, |(u, v)| - \mathrm{PA}(u)\}$. Then we can add an edge (u, v) by increasing the transmit powers of u and v by a total of $(\mathrm{PA}(u) + d_{uv})^p + (\mathrm{PA}(v) + d_{vu})^p - \mathrm{PA}(u)^p - \mathrm{PA}(v)^p \leq 2(\operatorname{diam} R)^p$. The rest of the proof follows. □

Proving superadditivity of k-ECPA^p_B is even easier than proving it of k-EC^p_B as the power of nodes that lose a connection is already large enough to reach the boundary.

Lemma 6.2. *For $p \geq 1$, k-ECPA^p_B is superadditive.*

Proof. We apply the same construction as in the proof of Lemma 4.4. Notice that since $|a| \leq |c|$ and $|b| \leq |c|$, the edges to the boundary are already included without changing the power assignment. This means that $k\text{-ECPA}^p_B(V, R) \geq k\text{-ECPA}^p_B(V, R_1) + k\text{-ECPA}^p_B(V, R_2)$ and therefore k-ECPA^p_B is superadditive for $p \geq 1$. □

Proving pointwise closeness of k-ECPA^p and k-ECPA^p_B is a bit more difficult as not all claims we used in the proof of Lemma 4.5 hold for these functionals. We first prove corresponding auxiliary results.

Claim 6.3. *Let $V \subset [0, 1]^d$, $|V| = n$, and $1 \leq p < d$. Consider a graph realizing the optimal solution of $k\text{-ECPA}^p_B(V, R)$, where $R = [0, 1]^d$. The sum of the pth powers of the lengths of the edges connecting nodes in V with the boundary of R is bounded by $O(n^{(d-p-1)/(d-1)})$.*

Proof. The proof is similar to that of Claim 4.6, except that for power assignment, in case more than k points are rooted to the boundary, all of them are also connected to each other (as the diameter of the subcube is less than the distance to the boundary). Without changing the solution, we can remove one of the roots to the boundary. This also gives us that we only have to take into account at most k points rooted to the boundary in each subcube. The rest of the proof follows. □

As we know the degree of nodes in an optimal power assignment graph can be unbounded [10], we cannot show a claim similar to Claim 4.7 for k-ECPA^p as well. We do however have the following claim.

Claim 6.4. *Fix $1 \leq p < d$, let $V \subset [0, 1]^d$ be V is a finite subset and let R be a d-dimensional rectangle. Then we have that $k\text{-ECPA}^p(V, R) \leq 2k\text{-EC}^p(V, R)$.*

Proof. Consider a graph $G = (V, E)$ achieving the optimal solution for $k\text{-EC}^p(V, R)$. Then for each node $v \in V$ we take the longest edge from v and we set power $\mathrm{PA}(v)$ to the length of this edge. This means that all edges that were in E, are also in the graph resulting from power assignment PA. So the power assignment graph resulting from power assignment PA is also k-edge-connected, and has cost no more than twice $k\text{-EC}^p(V, R)$. An optimal solution for $k\text{-ECPA}^p(V, R)$ can only have costs lower than or equal to that of PA, so $k\text{-ECPA}^p(V, R) \leq 2k\text{-EC}^p(V, R)$ follows. □

Lemma 6.5. *For $1 \leq p < d$, k-ECPA^p is pointwise close to k-ECPA^p_B.*

Proof. The proof for power assignments is the same as the proof of Lemma 4.5, except that we use Claim 6.4 and Eq. (3) to get the following result for a k-edge-connected power assignment graph on U

$$k\text{-ECPA}^p(U, R) \leq 2k\text{-EC}^p(U, R) = 2 \sum_{(u,v)\in M} \mathrm{d}(u, v)^p \leq o(|V|^{(d-p)/d}), \tag{4}$$

where R is $[0,1]^d$ and M is the set of edges used create a k-edge-connected graph on U. This gives us for k-ECPA^p that we have $k\text{-ECPA}^p(V,R) \le k\text{-ECPA}^p_B(V,R) + o(|V|^{(d-p)/p})$, and therefore k-ECPA^p_B and k-ECPA^p are pointwise close. □

If we try to extend the proof of Lemma 4.8 to power assignments, we run into trouble with the possible unbounded degree of power assignment. So instead of trying to bound the number of nodes connected to one node, we bound the number of k-edge-connected components connected to one node. To do this, we use another lemma [10, Lemma 3.2]. With this lemma, we can prove smoothness for k-ECPA^p.

Lemma 6.6. *For $1 \le p < d$, k-ECPA^p is smooth.*

Proof. To prove smoothness, we follow de Graaf and Manthey [10, Proof of Lemma 3.6]. Consider a graph $G = (U \cup V, E)$ that realizes an optimal solution for k-ECPA$^p(U \cup V, R)$. The problem for smoothness is that the degree $\deg_G(v)$ of nodes $v \in V$ can be unbounded. The idea is to exploit the fact that removing $v \in V$ also frees some power. Roughly speaking, we proceed as follows: Let $v \in V$ be a node of possibly large degree. We add the power of v to some nodes close to v. The graph obtained from removing v and distributing its energy has only a constant number of k-edge-connected components (either separate nodes, or components with $k+1$ or more nodes). To prove this, we consider cones again, as done by de Graaf and Manthey [10, Lemma 3.2] and in our Claim 4.7. We consider cones rooted at v with the following properties:

- The cones have a small angle α with v; for all cones C and for all $x, y \in C$, we have $\angle(x, v, y) < \alpha = \pi/6$.
- Every point in $[0,1]^d$ is covered by some cone.
- The number of cones m is finite (as d is a constant).

Let $C_1, \ldots, C_m$ be these cones (m is a constant, depending only on the dimension d). By abusing notation, let C_i also denote all points $C_i \cap (U \cup V \setminus \{x\})$ that are adjacent to v in G. Let $x_{i_1}, \ldots, x_{i_k}$ be the k points in C_i closest to v and let y_i be farthest from v. (We ignore C_i if $C_i \cap U = \emptyset$.) Let $\ell_i = |(y_i, v)|$ be the maximum distance of a node in C_i to v and let $\ell = \max_i \ell_i$. We note that $\ell \le \sqrt[p]{\mathrm{PA}(v)}$.

We increase the power of the k closest points in each C_i by $\ell^p/(mk)$. Since the power of v is at least ℓ^p and we have m cones, we can account for this increase by the removal of v. As $\alpha = \pi/6$, and as $x_{i_1}, \ldots, x_{i_k}$ are the closest points to v, any point in C_i is closer to $x_{i_1}, \ldots, x_{i_k}$ than to v. According to de Graaf and Manthey [10, Lemma 3.2(a)], every point in $C_i \setminus \{x_{i_1}, \ldots, x_{i_k}\}$ has sufficient power to reach all nodes in $x_{i_1}, \ldots, x_{i_k}$. There is now an edge between $x_{i_1}, \ldots, x_{i_k}$ and every point $z \in C_i$ that has a distance of at most $\ell/\sqrt[p]{mk}$. Now let z_1 be the first node not connected to all $x_{i_1}, \ldots, x_{i_k}$ as it has too little power. The second part of de Graaf and Manthey's lemma [10, Lemma 3.2(b)] implies that if x and y are not connected, then $|(y,v)| > 2\cos(\alpha)|(x,v)| = \sqrt{3}|(x,v)|$. We section the cone in pieces such that if x and y are in one piece, we have $|(y,v)| \le \sqrt{3}|(x,v)|$. In this way, adjacent pieces are scaled by a factor $\sqrt{3}$. We cover all nodes in the cone in this way, starting with the piece containing all $x_{i_1}, \ldots, x_{i_k}$ up to z_1. Let us denote the number of pieces we need to cover the cone by h. We can state that $\ell \ge \ell_i = |(y_i, v)| \ge \sqrt{3}^{h-1}|(z_1, v)| \ge \ell/\sqrt[p]{mk}$. This implies that $h \le \log_{\sqrt{3}}(\sqrt[p]{mk}) + 1$.

For each piece, we have two options.

1. The number of nodes in there is smaller than or equal to k. In this case we just count all of them as separate k-edge-connected components.
2. The number of nodes in there is larger than or equal to $k+1$. As $|(y,v)| \le \sqrt{3}|(x,v)|$ for all x and y in one piece, we know all nodes in this piece are connected to each other by Lemma 3.2 in [10]. A complete graph on $k+1$ nodes is k-edge-connected, so we can count this whole piece as one k-edge-connected component.

This means that the number of k-edge-connected components per cone after removing k and redistributing the power is bounded by $h \cdot k = k \log_{\sqrt{3}}(\sqrt[p]{mk}) + k$, which is a constant number. As the number of cones is finite as well, the total number of k-edge-connected components is bounded by a constant as well. If we remove $|V|$ points, the graph falls apart into at most $O(|V|)$ k-edge-connected components. Creating a k-edge-connected graph on these components increases the costs no more than $O(|V|^{(d-p)/d})$, meaning that we can follow the proof from Lemma 4.8. This proves that k-ECPA^p is also smooth for $k \ge 2$. □

Now the limit theorems also hold for k-ECPA^p (Theorems 4.10 and 4.11), as well as the results for the partitioning algorithm (Theorems 5.2 and 5.3).

7. Conclusions and open problems

In this paper, we have looked at fault tolerant networks in terms of k-edge-connected graphs. We studied both a standard (wired) model and a model for wireless networks. We analyzed the corresponding Euclidean functionals k-EC^p and k-ECPA^p on random inputs. We fitted k-EC^p into Yukich's framework for Euclidean functionals and obtained probabilistic results for k-EC^p. With Yukich's framework, we derived several concentration results. We derived a partitioning heuristic for k-EC^1, for which we proved an additive approximation guarantee. We analyzed its approximation ratio in the framework of smoothed analysis. Finally, we transferred these results to k-ECPA^p.

We conclude this paper with a few open problems for future research.

We have only looked at k-edge-connected graphs. But we can also consider connectivity in terms of k-vertex-connected graphs. While subadditivity, superadditivity of the corresponding boundary functional, and pointwise closeness are relatively straightforward to prove, we feel that the main technical difficulty in proving smoothness is the lack of a counterpart of Lemma 4.1.

Another extension would be the case $p \geq d$. As is the case with other Euclidean functionals, our functionals also lack smoothness and closeness. Instead, we require closeness in mean and smoothness in mean to deal with $p \geq d$. With these properties, complete convergence could also be proved for $p \geq d$ using Warnke's concentration of measure result [23], as has been shown by de Graaf and Manthey [10].

References

[1] Ernst Althaus, Gruia Călinescu, Ion I. Mandoiu, Sushil Prasad, Nickolay Tchervenski, Alexander Zelikovsky, Power efficient range assignment for symmetric connectivity in static ad hoc wireless networks, Wirel. Netw. 12 (3) (2006) 287–299.
[2] Fatiha Bendali, I. Diarrassouba, Ali Ridha Mahjoub, Mohamed Didi Biha, Jean Mailfert, A branch-and-cut algorithm for the k-edge connected subgraph problem, Networks 55 (1) (2010) 13–32.
[3] Christian Bettstetter, On the minimum node degree and connectivity of a wireless multihop network, in: Proc. of the 3rd ACM Int. Symp. on Mobile Ad Hoc Networking & Computing (MobiHoc), ACM, 2002, pp. 80–91.
[4] Markus Bläser, Bodo Manthey, B.V. Raghavendra Rao, Smoothed analysis of partitioning algorithms for euclidean functionals, Algorithmica 66 (2) (2013) 397–418.
[5] Gruia Călinescu, Peng-Jun Wan, Range assignment for biconnectivity and k-edge connectivity in wireless ad hoc networks, Mob. Netw. Appl. 11 (2) (2006) 121–128.
[6] Andrea E.F. Clementi, Paolo Penna, Riccardo Silvestri, On the power assignment problem in radio networks, Mob. Netw. Appl. 9 (2) (2004) 125–140.
[7] Artur Czumaj, Andrzej Lingas, On approximability of the minimum-cost k-connected spanning subgraph problem, in: Robert Endre Tarjan, Tandy J. Warnow (Eds.), Proc. of the 10th Ann. ACM-SIAM Symp. on Discrete Algorithms (SODA), ACM/SIAM, 1999, pp. 281–290.
[8] Bernhard Fuchs, On the hardness of range assignment problems, Networks 52 (4) (2008) 183–195.
[9] Michael R. Garey, David S. Johnson, Computers and Intractability: A Guide To the Theory of NP-Completeness, 1979.
[10] Maurits de Graaf, Bodo Manthey, Probabilistic analysis of power assignments, Random Struct. Algorithms 51 (3) (2017) 483–505.
[11] Martin Grötschel, Clyde L. Monma, Integer polyhedra arising from certain network design problems with connectivity constraints, SIAM J. Discrete Math. 3 (4) (1990) 502–523.
[12] Martin Grötschel, Clyde L. Monma, Mechthild Stoer, Polyhedral approaches to network survivability, in: Fred Roberts, Frank Hwang, Clyde L. Monma (Eds.), Proc. Workshop on Reliability of Computer and Communication Networks, in: Discrete Mathematics and Theoretical Computer Science, vol. 5, American Mathematical Society, 1991, pp. 121–141.
[13] Pao-Lu Hsu, Herbert Robbins, Complete convergence and the law of large numbers, Proc. Natl. Acad. Sci. 33 (2) (1947) 25–31.
[14] Frank Kammer, Hanjo Täubig, Connectivity, in: Ulrik Brandes, Thomas Erlebach (Eds.), Network Analysis: Methodological Foundations, in: Lecture Notes in Computer Science, vol. 3418, Springer, 2005, pp. 143–177.
[15] Samir Khuller, Uzi Vishkin, Biconnectivity approximations and graph carvings, J. ACM 41 (2) (1994) 214–235.
[16] Bodo Manthey, Heiko Röglin, Smoothed analysis: Analysis of algorithms beyond worst case, Inf. Technol. 53 (6) (2011) 280–286.
[17] David W. Matula, The cohesive strength of graphs, in: G. Chartrand, S.F. Kapoor (Eds.), The Many Facets of Graph Theory, Springer, 1969, pp. 215–221.
[18] Ram Ramanathan, Regina Rosales-Hain, Topology control of multihop wireless networks using transmit power adjustment, in: 19th Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM 2000), Vol. 2, IEEE, 2000, pp. 404–413.
[19] Charles Redmond, Boundary Rooted Graphs and Euclidean Matching Algorithms, (Ph.D. thesis), Lehigh University, Bethlehem, PA, USA, 1993.
[20] Wansoo T. Rhee, A matching problem and subadditive euclidean functionals, Ann. Appl. Probab. 3 (3) (1993) 794–801.
[21] Paolo Santi, Douglas M. Blough, Feodor Vainstein, A probabilistic analysis for the range assignment problem in ad hoc networks, in: Proceedings of the 2nd ACM International Symposium on Mobile Ad Hoc Networking & Computing, ACM, 2001, pp. 212–220.
[22] Daniel A. Spielman, Shang-Hua Teng, Smoothed analysis: an attempt to explain the behavior of algorithms in practice, Commun. ACM 52 (10) (2009) 76–84.
[23] Lutz Warnke, On the method of typical bounded differences, Combin. Probab. Comput. 25 (2) (2016) 269–299.
[24] Joseph E. Yukich, Probability Theory of Classical Euclidean Optimization Problems, in: Lecture Notes in Mathematics, vol. 1675, Springer, Berlin, 1998.

Please cite this article as: B. Manthey and V.M.J.J. Reijnders, Probabilistic properties of highly connected random geometric graphs, Discrete Applied Mathematics (2019), https://doi.org/10.1016/j.dam.2019.04.007.