

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
b.manthey@utwente.nl / v.m.j.j.reijnders@utwente.nl

In this paper, we study the probabilistic properties of reliable networks of minimum costs in d -dimensional Euclidean space. We study 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 vertices, and two vertices 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 for these type of problems can be found in the design of reliable communication and transportation networks [2, 10, 11].

Wireless ad hoc networks have also received significant attention in recent studies [6, 8, 17]. 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, 8]. Metricom Inc's Ricochet network and the Army Near-Term Digital Radio network are examples of fully operational multi-hop wireless networks [17].

In this paper, we study the probabilistic properties of the value of the optimal solution 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 vertices, dimension d of the underlying Euclidean space, distance-power gradient p , and connectivity requirement k .

Finding a cheapest k -edge-connected network is NP-hard [9], and so is finding a minimal power wireless network [5]. 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 on 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 [15,21] 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.

2 Definitions

Let $V \subseteq \mathbb{R}^d$ be a finite set of vertices, where $d \in \mathbb{N}$ is an arbitrary constant. In the rest of the paper, $n = |V|$ is the number of vertices. 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 vertices. In this paper, we also call any complete graph k -edge-connected, even if it contains fewer than k vertices. (Otherwise, no k -edge-connected graphs below a certain size exist, which would cause technical issues.)

In this paper, we study k -edge-connected graphs of minimal costs, where costs are defined as the sum of all edge lengths. 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 is 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 a 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 the vertices v . Two vertices u and v are connected by an edge if both have sufficient transmit power, 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 of 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 [22].

Following Yukich [22], 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 order to fit $k\text{-EC}^p$ and $k\text{-ECPA}^p$ into Yukich’s framework for Euclidean optimization problems [22], we have to define corresponding canonical boundary functionals, an idea first articulated in Redmond’s thesis [18]. Roughly speaking, in these functionals, the entire boundary of the rectangle is considered as one additional vertex 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 and a point set $V \subseteq R$, a boundary graph is a graph with nodes V plus the boundary ∂R of R as additional node. 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 independent edges. We denote by $k\text{-EC}_B^p$ the boundary functional corresponding to $k\text{-EC}^p$. This means that $k\text{-EC}_B^p(V, R)$ is the minimum total length of a k -edge-connected graph in terms of summed edge lengths on $V \cup \partial R$ in d -dimensional rectangle R with p th power-weighted edges. A vertex v is connected to the boundary ∂R by adding edge (v, v_∂) where $v_\partial = \arg \min_{w \in \partial R} |(v, w)|$.

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

3 Related Work

For a survey about properties of k -edge-connected graphs, we refer to Kammer and Täubig [13].

Finding k -edge-connected networks of minimum costs is NP-hard for $k \geq 2$, both in the classical and the power assignment model [5, 9]. Therefore, a considerable amount of research has been focused on approximation algorithms.

Khuller and Vishkin [14] proved that the problem of finding a minimum-cost k -edge connected graph can be approximated within a factor of 2 approximation algorithm. 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. Santi et al. [20] studied the connectivity of power assignment graphs in the Euclidean case for $d \in \{1, 2, 3\}$ under the restriction that every vertex is assigned the same power r . They derived bounds for r to achieve connectivity with high probability. De Graaf and Manthey [8] analyzed connectivity in power assignment graphs. They proved properties similar to our results of Section 4 for simple connectivity.

Călinescu and Wan [6] 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 $k\text{-EC}^p$

In this section we show that $k\text{-EC}^p$ fits into Yukich’s framework for Euclidean functionals [22].

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 [16]). *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

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

Lemma 4.2. *For $p \geq 1$, $k\text{-EC}^p$ 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*

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

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 $k\text{-EC}^p(V, R_1)$ and $k\text{-EC}^p(V, R_2)$, 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, \dots, e_k . This results in a k -edge-connected graph on V . We have $|e_i| \leq \text{diam } R$. Thus, the costs of this k -edge-connected graph is bounded from above by

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

2. $|V_1| \geq k + 1$ and $|V_2| \leq k$. We know that (V_2, E_2) is complete. For each vertex $v_i \in V_2$ we add k edges e_{i_1}, \dots, e_{i_k} with endpoints in V_1 . Each pair of vertices from V_1 has at least k edge-disjoint paths in E_1 , and by adding e_{i_j} each pair in V has k edge-disjoint paths as well. At most k^2 edges have been added this way, and we can bound the costs of the k -edge-connected graph thus 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. So if we make the combined graph complete as well, it is k -edge-connected. We need to add at most k^2 edges 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 $k\text{-EC}^p$ 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*

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

Proof. This lemma follows directly by combining Lemma 4.2 and a result by Yukich [22, Lemma 3.3]. □

Unfortunately, as the Euclidean functionals considered by Yukich [22], $k\text{-EC}^p$ is not superadditive. If it were superadditive, then this together with subadditivity makes the functional nearly additive in the sense that

$$k\text{-EC}^p(F, R) \approx k\text{-EC}^p(F, R_1) + k\text{-EC}^p(F, R_2).$$

We could then approximate the optimal solution value of the whole set by the sum of optimal solutions on its partitions.

The way to superadditivity is via the canonical boundary functional of $k\text{-EC}_B^p$, which is superadditive according to the following lemma.

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

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

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 $V \cup \{\text{partial}R\}$ 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 .

The two induced subgraphs are k -edge-connected.

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

Next, we show that $k\text{-EC}^p$ and $k\text{-EC}_B^p$ are pointwise close. This yields that both are approximately subadditive and superadditive.

Lemma 4.5. *For $1 \leq p < d$, $k\text{-EC}^p$ is pointwise close to $k\text{-EC}_B^p$, i.e. for all finite sets $V \subset [0, 1]^d$ we have*

$$|k\text{-EC}^p(V) - k\text{-EC}_B^p(V)| = o(|V|^{(d-p)/d}).$$

Proof. Let $V \subseteq [0, 1]^d$ be a finite set of points. Clearly, $k\text{-EC}^p(V) \geq k\text{-EC}_B^p(V)$. Thus, we only have to prove $k\text{-EC}^p(V) \leq k\text{-EC}_B^p(V) + o(|V|^{(d-p)/p})$. If $|V| \leq k$, then this holds as $|V|$ is constant.

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

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

Proof. The proof follows ideas from Yukich [22] and depends on a dyadic subdivision of $[0, 1]^d$. Let Q_0 be the cube of edge length $1/3$ and centered within $[0, 1]^d$. Let Q_1 be the cube of edge length $2/3$, also centered within $[0, 1]^d$. We partition $Q_1 - Q_0$ into subcubes of edge 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 j -th stage we define cube Q_j of edge length $1 - 2(3 \cdot 2^j)^{-1}$ and partition $Q_j - Q_{j-1}$ into subcubes of edge length $(3 \cdot 2^j)^{-1}$. 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 $2^{(\ell-1)(d-1)} \leq n \leq 2^{\ell(d-1)}$.

This procedure produces nested cubes $Q_1 \subseteq Q_2 \subseteq \dots \subseteq Q_\ell$. It produces a dyadic covering of the cube until the moat $[0, 1]^d - Q_\ell$ 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

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

subcubes, each with an edge length equal to the distance between the subcube and the boundary of $[0, 1]^d$. Furthermore, by partitioning each subcube into $(k2^y)^d$ congruent subcubes, where y is the least integer satisfying $2^y \geq d^{1/2}$, we obtain a partition \mathcal{P} of Q_ℓ consisting of at most Cn subcubes with the property that the k times the diameter of each subcube is less than the distance to the boundary.

Suppose $V \subseteq [0, 1]^d$ is a finite point set and $G = (V, E)$ is the graph realizing the optimal solution of $k\text{-EC}_B^p(V)$. We observe that in G , each subcube Q in \mathcal{P} contains at most k points in V which are rooted to the boundary. If there were more than k point in $V \cap Q$ rooted to the boundary, we can do the following. We know that these points will 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 (which 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 p -th powers of the lengths of the edges connecting vertices 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 p -th 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 $1 \leq j \leq \ell$ gives a bound for the sum of the p -th power of the lengths of the edges connecting points to the boundary in $V \cap Q_\ell$:

$$\sum_{j=1}^{\ell} C3^{d-1}2^{j(d-1)} \cdot (3 \cdot 2^j)^{-p} \leq \begin{cases} C \max\{n^{(d-p-1)/(d-1)}, \log n\} & \text{if } 1 \leq p \leq d-1 \\ C & \text{if } d-1 < p < d. \end{cases} \quad (1)$$

The $\log n$ term is needed to cover the case $p = d-1$. The sum of the p -th powers of the lengths of the edges connecting vertices in $V \cap ([0, 1]^d - Q_\ell)$ with the boundary is at most the product of $n = |V|$ and the p -th power of the width of the moat $[0, 1]^d - Q_\ell$, i.e. at most

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

Combining (1) and (2) proves the claim. \square

We can now continue with the proof of Lemma 4.5. Consider $U \subset V$ the set of all vertices connected to the boundary. Let $\mathcal{B} \subseteq \partial[0, 1]^d$ be the set of points on the boundary to which vertices 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 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 lemma.

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

Proof. Let us assume to the contrary that there exists a vertex $v \in V$ for which the degree is not bounded by c . We then divide $[0, 1]^d$ into cones originating from v and with the property that for every two points x and y in a cone we have $\angle(x, y, v) < \pi/3$. In this way, the number of cones we create is finite since d is constant, and every point is covered by a cone. We consider a cone C with an unbounded number of points connected to v and look at the point y that is furthest from v and connected to v . Such a cone has to exist as the number of cones is bounded. Let us consider two cases.

In the first case, the degree of y is greater than c as well, and y is connected to all vertices in C that are also connected to v . This means both are connected to more than k vertices in C . In that 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 vertex to which y is not connected, but v is. Let us call this vertex z and 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 . \square

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 need to add. By using the triangle inequality for $p = 1$ we get $|(u, v)| \leq |(u, u_B)| + |(u_B, v_B)| + |(v_B, v)|$. Using the triangle inequality for powers of metrics, we obtain

$$\begin{aligned}
\sum_{(u,v) \in M} |(u, v)|^p &\leq 2^{p-1} \left(\sum_{(u,\cdot) \in M} |(u, u_B)|^p + \sum_{(u,v) \in M} |(u_B, v)|^p \right) & (3) \\
&\leq 4^{p-1} \left(\sum_{(u,\cdot) \in M} |(u, u_B)|^p + \sum_{(u,v) \in M} |(u_B, v_B)|^p + \sum_{(\cdot, v) \in M} |(v_B, v)|^p \right) \\
&\leq 4^p \sum_{(u,\cdot) \in M} |(u, u_B)|^p + 4^{p-1} k\text{-EC}^p(\mathcal{B}, [0, 1]^d) \\
&\leq 4^p C_1 \sum_{u \in U} |(u, u_B)|^p + 4^{p-1} C_2 |V|^{(d-p-1)/(d-1)} & \text{(by Lemma 4.3)} \\
&\leq 4^p C_1 C_3 |V|^{(d-p-1)/(d-1)} + o(|V|^{(d-p)/d}) & \text{(by Claim 4.6)} \\
&\leq o(|V|^{(d-p)/d}).
\end{aligned}$$

Now we have changed a graph that achieves the optimal solution for $k\text{-EC}_B^p(V, [0, 1]^p)$ into a k -edge-connected graph. The cost of this new graph provides an upper bound for $k\text{-EC}^p(V)$. Observing that we have increased the costs by at most $o(|V|^{(d-p)/d})$ concludes the proof. \square

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 vertices does not change the function value by much.

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

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

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

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

It remains to be shown that

$$k\text{-EC}^p(U, [0, 1]^d) - k\text{-EC}^p(U \cup V, [0, 1]^d) \leq O(|V|^{(d-p)/d}).$$

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 vertices in V . By Lemma 4.7 we know that $|N_V| \leq c|V|$ 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 $k\text{-EC}^p(U \cup V, [0, 1]^d) + O(|V|^{(d-p)/d})$. It remains to be proved that it is k -edge-connected. Fix any two vertices $u, v \in V$. 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 vertex 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.

Theorem 4.9. *For $1 \leq p < d$, $k\text{-EC}_B^p$ is smooth.*

Proof. The proof is similar to the proof of Theorem 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 $k\text{-EC}_B^p$. □

4.2 Limit Theorems

For the theorems on the convergence of optimal solutions, we need the notion of complete convergence. Let X_n , for $n \in \mathbb{N}$, 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

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

This notion of convergence was first introduced by Hsu and Robbins [12].

As we have geometric subadditivity, superadditivity, pointwise closeness, and smoothness, several limit theorems directly follow. These 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 $[0, 1]^d$. Fix $1 \leq p < d$ and $k \in \mathbb{N}$. Then there exists a positive constant $\alpha = \alpha(k\text{-EC}^p, d, k)$ such that*

$$\begin{aligned} \lim_{n \rightarrow \infty} \frac{k\text{-EC}^p(V)}{n^{(d-p)/p}} &= \alpha \quad \text{c.c., and} \\ \lim_{n \rightarrow \infty} \frac{k\text{-EC}_B^p(V)}{n^{(d-p)/p}} &= \alpha \quad \text{c.c.} \end{aligned}$$

Proof. As $k\text{-EC}^p$ is a smooth subadditive Euclidean functional and pointwise close to its smooth superadditive boundary functional, we can use a result by Yukich [22, Theorem 4.1] to directly get this theorem. \square

The following famous limit theorem is due to Rhee [19]. This theorem states that the solution value is not far from its expected value when we look at randomly placed vertices.

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 exists 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*

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

Proof. As $k\text{-EC}^p$ is a smooth subadditive Euclidean functional which is pointwise close to its superadditive boundary functional, we can use the theorem in Rhee [19] to instantly get this result. \square

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 the following 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 $\text{PSE}^p(V)$ denote the cost of the k -edge-connected graph computed by Algorithm 5.1. Then $\text{PSE}^p(V) \leq k\text{-EC}^p(V) + O((n/s)^{\frac{d-p}{d}})$.*

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 $\text{PSE}^p(V)$ be the cost of the k -edge-connected graph on V computed by Algorithm 5.1. By using subadditivity, $\text{PSE}^p(V)$ is bounded from above by $k\text{-EC}^p(V) + C_1$ 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 Theorem 4.3). We do this by lacing cells together in a snakelike succession. We create a minimal length matching of k vertices between two succeeding cells. It is easy to check that this yields a k -edge-connected graph on the cells we lace together. This yields the last claim of the theorem. \square

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

The above theorem does not yield an 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 semi-random instances: Let $\phi \geq 1$ be a perturbation parameter. For each of the n points, an adversary specifies a probability density function $[0, 1]^d \rightarrow [0, \phi]$. 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 smoothed analysis result.

Theorem 5.3. *For $p = 1$, $d = 2$, and $s = \Theta(\sqrt{\log n})$, Algorithm 5.1 has polynomial running-time and achieve 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. \square

6 Extension to k -ECPA

The results for k -ECP are easily copied to k -ECPA^p by making some small adjustments to the proofs. We start off 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 $\text{PA}(v)$ accordingly for the edges that need to be added in the power assignment graph. We define $d_{uv} = \max\{0, |(u, v)| - \text{PA}(u)\}$. This denotes the increase in power needed for u to reach v . Then we can add an edge (u, v) by a change of power equal to $(\text{PA}(u) + d_{uv})^p + (\text{PA}(v) + d_{vu})^p - \text{PA}(u)^p - \text{PA}(v)^p \leq 2(\text{diam } R)^p$. The rest of the proof follows. \square

Proving superadditivity of k -ECPA_B^p is even easier than proving it of k -ECP_B^p as the power of vertices that lose a connection is already large enough to reach to boundary.

Lemma 6.2. *For $p \geq 1$, k -ECPA_B^p 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 -ECPA_B^p(V, R) \geq k -ECPA_B^p(V, R_1) + k -ECPA_B^p(V, R_2) and therefore k -ECPA_B^p is superadditive for $p \geq 1$. \square

Proving pointwise closeness of k -ECPA^p and k -ECPA_B^p is a bit more difficult as not all theorems 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 realising the optimal solution of k -ECPA_B^p($V, [0, 1]^d$). The sum of the p -th powers of the lengths of the edges connecting vertices in V with the boundary of $[0, 1]^d$ 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. \square

As we know the degree of vertices in an optimal power assignment graph can be unbounded [8], we cannot show Lemma 4.7 for k -ECPA^p as well. We do however have the following lemma:

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

Proof. Consider a graph $G = (V, E)$ achieving the optimal solution for k -EC^p(V, R). Then for each vertex $v \in V$ we take the longest edge from v and we set power 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 costs no more than twice k -EC^p(V, R). An optimal solution for k -ECPA^p(V, R) can only have costs lower or equal to that of PA, so k -ECPA^p(V, R) $\leq 2k$ -EC^p(V, R) follows. \square

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

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

$$k\text{-ECPA}^p(U, [0, 1]^d) \leq 2k\text{-EC}^p(U, [0, 1]^d) = 2 \sum_{(u,v) \in M} d(u,v)^p \leq o(|V|^{(d-p)/d}), \quad (4)$$

where 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 -ECPA^p($V, [0, 1]^d$) $\leq k$ -ECPA_B^p($V, [0, 1]^d$) + $o(|V|^{(d-p)/p})$, and therefore k -ECPA_B^p and k -ECPA^p are pointwise close. \square

If we try to extend the proof of Theorem 4.8 to power assignments, we get into trouble with the possible unbounded degree of power assignment. So instead of trying to bound the number of vertices connected to one vertex, we bound the number of k -edge-connected components connected to one vertex. To do this, we use another lemma [8, Lemma 3.2]. With this lemma, we can prove smoothness for k -ECPA^p.

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

Proof. To prove smoothness, we follow a proof by de Graaf and Manthey [8, Lemma 3.6]. Consider a graph $G = (U \cup V, E)$ achieving the optimal solution for k -ECPA^p($U \cup V, R$). The problem for smoothness is that the degree $\deg_G(v)$ of vertices $v \in V$ 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 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 k -edge-connected components (either separate vertices, or components with $k + 1$ or more vertices). To prove this, Lemma 3.2 in [8] is crucial. 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, \dots, C_m be these cones. 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}, \dots, 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 vertex in C_i to v and let $\ell = \max_i \ell_i$. We note that $\ell \leq \sqrt[p]{\text{PA}(v)}$.

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}, \dots, x_{i_k} are the closest points to v , any point in C_i is closer to x_{i_1}, \dots, x_{i_k} than to v . According to Lemma 3.2(a) in [8], every point in $C_i \setminus \{x_{i_1}, \dots, x_{i_k}\}$ has sufficient power to reach all vertices in x_{i_1}, \dots, x_{i_k} . There is now an edge between x_{i_1}, \dots, 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 vertex not connected to all x_{i_1}, \dots, x_{i_k} as it has too little power. We can use Lemma 3.2(b) in [8] which 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 $|(y, v)| \leq \sqrt{3}|(x, v)|$. In this way, adjacent pieces are scaled by a factor $\sqrt{3}$. We cover all vertices in the cone in this way, starting with the piece containing all x_{i_1}, \dots, 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 $\ell \geq \ell_i = |(y_i, v)| \geq \sqrt{3}^{h-1}|(z_1, v)| \geq \ell/\sqrt[p]{mk}$. This implies that $h \leq \log_{\sqrt{3}}(\sqrt[p]{mk}) + 1$.

For each piece, we have two options.

1. The number of vertices in there is smaller or equal to k . In this case we just count all of them as separate k -edge-connected components.
2. The number of vertices in there is larger or equal to $k + 1$. As $|(y, v)| \leq \sqrt{3}|(x, v)|$ for all x and y in one piece, we know all vertices in this piece are connected to each other by Lemma 3.2 in [8]. A complete graph on $k + 1$ vertices 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 fall 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 we can follow the proof from Theorem 4.8. This then proves that k -ECPA ^{p} is also smooth for $k \geq 2$. \square

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 -EC¹, for which we proved an additive approximation guarantee. We analyzed its approximation ratio in the framework of smoothed analysis. Finally, we transferred results to k -ECPA ^{p} .

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

In this paper we have looked only 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 is 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 $p \geq d$. This requires closeness in mean and smoothness in mean. With these properties, complete convergence could also be proved for $p \geq d$ using Warnke's result, as has been shown by de Graaf and Manthey [8].

References

- [1] Ernst Althaus, Gruia Călinescu, Ion I. Mandoiu, Sushil Prasad, Nickolay Tchernovski, and Alexander Zelikovsky. Power efficient range assignment for symmetric connectivity in static ad hoc wireless networks. *Wireless Networks*, 12(3):287–299, 2006.
- [2] Fatiha Bendali, I. Diarrassouba, Ali Ridha Mahjoub, Mohamed Didi Biha, and Jean Mailfert. A branch-and-cut algorithm for the k -edge connected subgraph problem. *Networks*, 55(1):13–32, 2010.
- [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)*, pages 80–91. ACM, 2002.
- [4] Markus Bläser, Bodo Manthey, and B. V. Raghavendra Rao. Smoothed analysis of partitioning algorithms for euclidean functionals. *Algorithmica*, 66(2):397–418, 2013.
- [5] Andrea E. F. Clementi, Paolo Penna, and Riccardo Silvestri. On the power assignment problem in radio networks. *Mobile Networks and Applications*, 9(2):125–140, 2004.
- [6] Gruia Călinescu and Peng-Jun Wan. Range assignment for high connectivity in wireless ad hoc networks. In *Proc. 2nd Int. Conf. on Ad-Hoc, Mobile, and Wireless Networks*, volume 2865 of *Lecture Notes in Computer Science*, pages 235–246. Springer, 2003.
- [7] Artur Czumaj and Andrzej Lingas. On approximability of the minimum-cost k -connected spanning subgraph problem. In Robert Endre Tarjan and Tandy J. Warnow, editors, *Proc. of the 10th Ann. ACM-SIAM Symp. on Discrete Algorithms (SODA)*, pages 281–290. ACM/SIAM, 1999.
- [8] Maurits de Graaf and Bodo Manthey. Probabilistic analysis of power assignments. *Random Structures & Algorithms*, to appear.
- [9] Michael R. Garey and David S. Johnson. Computers and intractability: A guide to the theory of NP-completeness, 1979.
- [10] Martin Grötschel and Clyde L. Monma. Integer polyhedra arising from certain network design problems with connectivity constraints. *SIAM Journal on Discrete Mathematics*, 3(4):502–523, 1990.

- [11] Martin Grötschel, Clyde L. Monma, and Mechthild Stoer. Polyhedral approaches to network survivability. In Fred Roberts, Frank Hwang, and Clyde L. Monma, editors, *Proc. Workshop on Reliability of Computer and Communication Networks*, volume 5 of *Series in Discrete Mathematics and Theoretical Computer Science*, pages 121–141. American Mathematical Society, 1991.
- [12] Pao-Lu Hsu and Herbert Robbins. Complete convergence and the law of large numbers. *Proceedings of the National Academy of Sciences*, 33(2):25–31, 1947.
- [13] Frank Kammer and Hanjo Täubig. Connectivity. In Ulrik Brandes and Thomas Erlebach, editors, *Network Analysis: Methodological Foundations*, volume 3418 of *Lecture Notes in Computer Science*, pages 143–177. Springer, 2005.
- [14] Samir Khuller and Uzi Vishkin. Biconnectivity approximations and graph carvings. *Journal of the ACM*, 41(2):214–235, 1994.
- [15] Bodo Manthey and Heiko Röglin. Smoothed analysis: Analysis of algorithms beyond worst case. *it – Information Technology*, 53(6):280–286, 2011.
- [16] David W. Matula. The cohesive strength of graphs. In G. Chartrand and S. F. Kapoor, editors, *The Many Facets of Graph Theory*, pages 215–221. Springer, 1969.
- [17] Ram Ramanathan and 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)*, volume 2, pages 404–413. IEEE, 2000.
- [18] Charles Redmond. *Boundary Rooted Graphs and Euclidean Matching Algorithms*. PhD thesis, Lehigh University, Bethlehem, PA, USA, 1993.
- [19] Wansoo T. Rhee. A matching problem and subadditive euclidean functionals. *Annals of Applied Probability*, 3(3):794–801, 1993.
- [20] Paolo Santi, Douglas M. Blough, and 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*, pages 212–220. ACM, 2001.
- [21] Daniel A. Spielman and Shang-Hua Teng. Smoothed analysis: an attempt to explain the behavior of algorithms in practice. *Communications of the ACM*, 52(10):76–84, 2009.
- [22] Joseph E. Yukich. *Probability theory of classical Euclidean optimization problems*, volume 1675 of *Lecture Notes in Mathematics*. Springer, Berlin, 1998.