Approximating Survivable Networks with Minimum Number of Steiner Points*

Lior Kamma†
The Open University of Israel
lior.kamma@gmail.com

Zeev Nutov
The Open University of Israel
nutov@openu.ac.il

Key-words: Sensor networks, Unit-disc graphs, Node-connectivity, Approximation algorithms.

Abstract

Given a graph $H = (U, E)$ and connectivity requirements $r = \{r(u, v) : u, v \in R \subseteq U\}$, we say that H satisfies r if it contains $r(u, v)$ pairwise internally-disjoint uv-paths for all $u, v \in R$. We consider the Survivable Network with Minimum Number of Steiner Points (SN-MSP) problem: given a finite set V of points in a normed space $(M, \|\cdot\|)$ and connectivity requirements, find a minimum size set $S \subset M \setminus V$ of additional points, such that the unit disc graph induced by $U = V \cup S$ satisfies the requirements. In the (node-connectivity) Survivable Network Design Problem (SNDP) we are given a graph $G = (V, E)$ with edge costs and connectivity requirements, and seek a min-cost subgraph H of G that satisfies the requirements. Let $k = \max_{u,v \in V} r(u, v)$ denote the maximum connectivity requirement. We will show a natural transformation of an SN-MSP instance (V, r) into an SNDP instance $(G = (V, E), c, r)$, such that an α-approximation algorithm for the SNDP instance implies an $\alpha \cdot O(k^2)$-approximation algorithm for the SN-MSP instance. In particular, for the case of uniform requirement $r(u, v) = k$ for all $u, v \in V$, we obtain for SN-MSP ratio $O(k^2 \ln k)$, which solves an open problem from [3].

1 Introduction

1.1 Problem definition and motivation

Network design problems require finding a minimum cost (sub-)network that satisfies prescribed properties, often connectivity requirements. Classic examples with $0, 1$ connectivity requirements are: Shortest Path, Minimum Spanning Tree, Minimum Steiner Tree/Forest, and others. Examples of problems with high connectivity requirements are: Min-Cost k-Flow, k-Edge/Node-Connected Spanning Subgraph, Steiner Network, and others. Such problems were studied extensively, see [1,

*This research was supported by The Open University of Israel's Research Fund (grant no. 100685).
†Part of this work was done as a part of author's M.Sc. Thesis at The Open University of Israel.

15, 13, 4, 17, 5, 6, 9, 19, 21, 18] for only a small sample of papers in the area.

We consider node-connectivity only. Let $H = (U, E)$ be an undirected graph, possibly with parallel edges. For $u, v \in U$ let $\kappa_H(u, v)$ denote the maximum number of pairwise internally disjoint uv-paths in H. Given non-negative integer connectivity requirements $r = \{r(u, v) : u, v \in R \subseteq U\}$ on a set R of terminals, we say that H *satisfies* r if $\kappa_H(u, v) \geq r(u, v)$ for all $u, v \in R$. In the Survivable Network Design Problem (SNDP) we are given a graph $G = (V, E)$ with edge-costs $\{c_e : e \in E\}$ and *node-connectivity* requirements $r = \{r(u, v) : u, v \in R \subseteq V\}$. The goal is to find a minimum cost subgraph H of G that satisfies r. SNDP problems were extensively studied, especially the k-Connected Subgraph problem when $r(u, v) = k$ for all $u, v \in V$, see [4, 17, 19]. For some recent work on other SNDP problems see [6, 18].

Problems of designing fault tolerant (highly connected) wireless sensor networks were studied extensively, c.f. [3, 2, 7, 20, 22, 16]. In a common scenario, we are given a set U of sensors/transmitters in some metric space, usually $\mathbb{R}^2$ or $\mathbb{R}^3$. Every sensor $u \in U$ can transmit to a known distance $d(u)$, and any other sensor v can receive messages from u if, and only if, v belongs to the disc with radius $d(u)$ and center u. Thus two sensors u, v can communicate with each other if, and only if, the distance $d(u, v)$ between them is at most $\min\{d(u), d(v)\}$. The resulting communication network (graph) $G[U]$ has node-set U and edge-set $\{uv : u, v \in U, d(u, v) \leq \min\{d(u), d(v)\}\}$.

In a network, the reliability of communication between two terminals u, v is usually measured by the node-connectivity between them; in sensor networks it models sensor failures. One way to satisfy the connectivity requirements between sensor terminals is by increasing the transmission range of some sensors. However, the energy needed to transmit through a distance d might be proportional to d^4 [8]. Since energy is a limited source, an alternative way to increase the connectivity is by adding new sensors. Thus the problem of adding a minimum number of new sensors to satisfy the requirements arises. We consider the simplest case, when all sensors have the same transmission radius, which can be assumed to be 1. This motivates the following definition.

Definition 1.1 *Given a finite set of points $U \subset M$ in a metric space (M, d), the* unit disc graph $G[U]$ induced by U *has node set U and edge set $\{uv : u, v \in U,\ 0 < d(u, v) \leq 1\}$.*

We consider SNDP problems on unit-disc graphs, where the goal is to add a minimum number of Steiner points (transmitters) to satisfy the connectivity requirements between the terminals. Namely, given a metric space, (M, d), a finite set $V \subseteq M$ of terminals, and connectivity requirements $\{r(u, v) : u, v \in V\}$, we wish to adjust the network to satisfy the requirements between the terminals by adding a minimum number of transmitters (Steiner points). Formally, our problem is as follows.

Survivable Network with Minimum Number of Steiner Points (SN-MSP)
Instance: A finite set V of points in a metric space (M, d) and pairwise connectivity requirements $r = \{r(u, v) : u, v \in R \subseteq V\}$.
Objective: Find a minimum size set of points $S \subset M \setminus V$ such that $G[V \cup S]$ satisfies r.

Note that the size of a solution to an SN-MSP instance may not be polynomial in the input size, as the solution size may depend on the maximum distance $\max_{u,v\in V} d(u,v)$ between the terminals. Hence we will say that an algorithm for SN-MSP is polynomial if its running time is polynomial in the size of the input and the optimal solution size. For simplicity of exposition, let us even assume that $\max_{u,v\in V} d(u,v)$ is polynomial in the input size.

An important special case is the case of *uniform requirements*, when $r(u,v) = k$ for all $u, v \in V$. We call this particular case k-Connectivity with Minimum Number of Steiner Points (k-C-MSP). In SNDP problems, also the following types of requirements are often considered in the literature, c.f. [6, 18].

- *Rooted requirements:* there is $s \in V$ such that $r(u,v) > 0$ implies $u = s$ or $v = s$; in *rooted-uniform* requirements $r(s,v) = k$ for all $v \in V - \{s\}$.
- *Subset uniform requirements:* there is $R \subseteq V$ such that $r(u,v) = k$ for all $u, v \in R$, and $r(u,v) = 0$ otherwise; (k-C-MSP is the case of *uniform requirements* when $R = V$).

1.2 Our results

Given an instance of SNDP or of SN-MSP, let $k = \max\limits_{u,v\in R} r(u,v)$ denote the maximum connectivity requirement. As in practical networks k is rather small, we focus on obtaining approximation ratios that depend on k only. For $k = 1$, SN-MSP with uniform requirements is the Steiner Tree with Minimum Number of Steiner Points problem (ST-MSP). In the Euclidean plane, this problem admits a 2.5-approximation algorithm [7]. On graphs with unit edge lengths ST-MSP includes the Set-Cover problem [15], and thus has an $\Omega(\ln |V|)$-approximation threshold. Hence for SN-MSP one cannot expect in arbitrary metric spaces a ratio that depends on k only. We will consider instances of SN-MSP defined on a *normed space* $(M, \|\cdot\|)$, when the metric d is induced by the norm $\|\cdot\|$.

One can easily reduce SN-MSP to an SNDP variant with unit weights on the nodes rather than with costs on the edges; this reduction invokes a constant loss factor in the approximation ratio. In this reduction however, uniform requirements in SN-MSP instance become subset uniform requirements in the SNDP instance. The currently best known ratios for SNDP with node weights are: $O(k^2 \log |V|)$ for rooted requirements, $O(k^3 \log |V|)$ for subset uniform requirements, and $O(k^4 \log^2 |V|)$ for general requirements [18]. The factor $O(\log |V|)$ in these ratios is unavoidable even for $k = 1$, as even for $k = 1$ the problem includes the Set-Cover problem [15].

Obtaining an approximation ratio that depends on k only for k-C-MSP in $\mathbb{R}^2$ was posed as an open problem in [3]. We will prove a much more general result. Our ratios are expressed in terms of k and the following parameter that depends on the normed space.

Definition 1.2 *Given a metric space (M,d) let $\Delta = \Delta(M,d)$ be the minimum integer h such that for any $V \subseteq M$ contained in a unit ball, $G[V]$ has a dominating set of size at most h.*

On graphs with unit edge lengths we have $\Delta = 1$, but note that this is *not* a normed space. It is known that $\Delta = 5$ in $\mathbb{R}^2$ and $\Delta = 11$ in $\mathbb{R}^3$. In [23] it is proved that for $M = \mathbb{R}^\ell$ with the norm $\|(x_1, x_2, \ldots, x_\ell)\|_p = \left(\sum_{i=1}^{\ell} |x_i|^p\right)^{1/p}$, Δ is at most the Hadwiger number of the unit ball. (The Hadwiger number of an open convex set X is the maximal number of disjoint translations of X which share a boundary point with X). Thus, for the Euclidean $\mathbb{R}^\ell$, $\Delta \leq 2^{0.401\ell(1+o(1))}$, by [14].

Let $\rho(k) = 2\lceil k/2\rceil(\Delta k/2 + \lceil k/2\rceil + 1)$ (so $\rho(k) = k^2(\Delta+1)/2 + k$ if k is even and $\rho(2) = 2(\Delta+2)$). Our main result is the following.

Theorem 1.1 *An α-approximation algorithm for* SNDP *(on multigraphs) implies an $\alpha \cdot \rho(k)$-approximation algorithm for* SN-MSP*, and this is so also for subset uniform, uniform, rooted, rooted subset uniform, and rooted uniform requirements.*

In SNDP problems, the input graph is usually assumed to be simple, while in Theorem 1.1 it may have parallel edges. One novelty in our approach is considering SNDP on multigraphs, and proving that the best known ratios for SNDP with different requirement types remain the same on multigraphs. Specifically, we will prove the following statement in Section 2.

Lemma 1.2 *There exists an approximation ratio preserving reduction from* SNDP *on multigraphs to* SNDP *on simple graphs, for approximation ratios that do not depend on $|V|$. The reduction is requirement type preserving for uniform, rooted, and subset uniform requirements. In the case of rooted uniform requirements, the problem on multigraphs admits a 2-approximation algorithm.*

The best known values of α are as follows. For k-Connected Subgraph on simple graphs, an $O(\log k)$-approximation algorithm for $k = O(\sqrt{n})$ [4] was obtained long time ago. This ratio was recently extended to almost all values of n, k in [19]; specifically, the ratio in [19] is $O(\ln \frac{|V|}{|V|-k} \cdot \ln k)$ (which is $O(\ln k)$ unless $k = |V| - o(|V|)$). For other SNDP problems, the currently best known approximation ratios are: 2 for rooted uniform requirements [11], $O(k \ln k)$ for rooted requirements [18], $O(k^2 \ln k)$ for subset uniform requirements [18], and $O(k^3 \ln |R|)$ for general requirements [5]. By substituting the currently best known values of α in Theorem 1.1, we obtain the following.

Corollary 1.3 *k-C-MSP admits an approximation ratio $O\left(\ln \frac{|V|}{|V|-k} \cdot \ln k\right) \cdot \rho(k) = O(k^2 \ln k)$. Other* SN-MSP *problems admit the following approximation ratios: $2\rho(k) = O(k^2)$ for rooted uniform requirements, $O(k \ln k) \cdot \rho(k) = O(k^3 \ln k)$ for rooted requirements, and $O(k^2 \ln k) \cdot \rho(k) = O(k^4 \ln k)$ for* subset uniform requirements*,*

Corollary 1.3 solves an open problem of Bredin, Demaine, Hajiaghayi, and Rus [3], by giving the first non-trivial approximation algorithm for k-C-MSP with $k \geq 2$. In [3] the problem of adding a minimum size set S of Steiner points such that the entire graph $G[V \cup S]$ is k-connected was considered (note that in k-C-MSP we require k-connectivity only between terminals). For this problem in $\mathbb{R}^2$, [3] gave a reduction that invokes a loss of $O(k^4)$. They also conjectured that for k-C-MSP an adaptation of their reduction can be used to reduce the instance to an SNDP instance with subset uniform requirements, thus leading to an approximation ratio that depends on k only, and with a loss of $O(k^3)$, provided existence of such an approximation for SNDP with subset uniform requirements. Here we prove a stronger result. Note that even if the conjecture of [3] were proved, it leads to ratio $O(k^3 \cdot k^2 \log k)$, which is much worse than the ratio $O(k^2 \cdot \log k)$ proved in this paper. The reason is not only the worse reduction factor, but also since [3] reduces instances with uniform requirements into instances with subset uniform requirements, while our reduction preserves the requirements type; consequently, our results for k-C-MSP rely on algorithms for k-Connected Subgraph only, and not on recently discovered algorithms for SNDP with subset uniform requirements [18]. Furthermore, our algorithm works for arbitrary normed spaces, and for various connectivity requirement types.

We also note that Theorem 1.1 together with the $O(k^3 \ln |R|)$-approximation algorithm for SNDP of [5] implies the ratio $O(k^3 \ln |R|) \cdot \rho(k) = O(k^5 \ln |R|)$ for SN-MSP with arbitrary requirements. It is an open question whether in this case, a ratio that depends on k only can be achieved.

2 Preliminaries

In this section we give some generic statements that are used later, most of them on connectivity of graphs, and also prove Lemma 1.2.

We start with two (essentially known) statements on k-connected graphs. Recall that a graph $H = (V, E)$ is k-connected if $\kappa_H(u, v) \geq k$ for all $u, v \in V$. A theorem of Whitney (c.f. [10, Theorem 7.5]) states that for $|V| \geq k+1$, H is k-connected if, and only if, $H \setminus Q$ is connected for any $Q \subseteq V$ with $|Q| \leq k-1$. The following two statements can be easily deduced from Whitney's Theorem. For completeness of exposition, we provide a direct proof for the first one, using the following undirected node-connectivity version of Menger's Theorem, c.f. [10, Theorem 7.5]: *If u, v are non-adjacent nodes in a graph H, then $\kappa_H(u, v) = \min\{|Q| : Q \subseteq V \setminus \{u, v\}, \kappa_{H \setminus Q}(u, v) = 0\}$.*

Lemma 2.1 *Let $H = (V, E)$ be a graph on at least $k+1$ nodes. If $\kappa_H(u, v) \geq k$ holds for every pair of non-adjacent nodes $u, v \in V$, then H is k-connected.*

Proof: Suppose to the contrary that H is not k-connected. Then there is $ab \in E$ such that $\kappa_H(a, b) \leq k-1$. Consider the graph H' obtained by removing all ab-edges from H. Note that $\kappa_{H'}(a, b) \leq \kappa_H(a, b) - 1 \leq k-2$. By Menger's Theorem, we can disconnect a, b by removing a set $Q \subseteq V \setminus \{a, b\}$ such that $|Q| = \kappa_{H'}(a, b)$. Thus there exists a partition $\{A, Q, B\}$ of V such that $a \in A$, $b \in B$, $|Q| = \kappa_{H'}(a, b) \leq k-2$, and there is no edge between A and B in H'. Since $|Q| \leq k-2$ and $|V| \geq k+1$, $|A| \geq 2$ or $|B| \geq 2$, say $|A| \geq 2$. Consider the partition $\{A', Q', B\}$ of V, where $Q' = Q \cup \{a\}$ and $A' = A \setminus \{a\}$. Note that $A' \neq \emptyset$, that $|Q'| = |Q| + 1 \leq k-1$, and that there is no edge between A' and B in H. Now let $z \in A'$. Then $zb \notin E$, and hence by Menger's Theorem we have $\kappa_H(z, b) \leq |Q'| \leq k-1$. This contradicts the assumption of the lemma. $\square$

Lemma 2.2 *Let H be a k-connected graph on at least k nodes. Then the graph obtained from H by adding a new node and joining it to some k nodes of H is also k-connected.* $\square$

Now we prove Lemma 1.2, which is restated here for the convenience of the reader.

Lemma 1.2 *There exists an approximation ratio preserving reduction from* SNDP *on multigraphs to* SNDP *on simple graphs, for approximation ratios that do not depend on $|V|$. The reduction is requirement type preserving for uniform, rooted, and subset uniform requirements. In the case of rooted uniform requirements, the problem on multigraphs admits a 2-approximation algorithm.*

Proof: Given an SNDP instance (with parallel edges), insert a new node into every edge, and divide (arbitrarily) the cost of the edge between the corresponding two new edges. Clearly, the obtained graph is simple. It is easy to see that an α-approximation for the modified instance implies an α-approximation for the original instance and that this transformation is requirement type preserving for subset uniform, rooted, and rooted subset uniform requirements. It remains therefore to consider uniform and rooted uniform requirements.

We now consider the case of uniform requirements, when feasible solutions are k-connected spanning subgraphs of G. Let $H = (V, E)$ be a minimally k-connected multi-graph (so $H - e$ is not k-connected for every $e \in E$).

If $|V| \geq k + 1$ then H is simple, by Lemma 2.1; thus we can keep for every maximal set of pairwise parallel edges of G only the cheapest one. Now suppose that $|V| \leq k$. We claim that then H has exactly $k+2-|V|$ edges between every pair of it nodes; thus an optimal solution is found by taking the $k + 2 - |V|$ cheapest edges in G between every pair of nodes. Note that if $|V| \leq k$ and if H has exactly $k + 2 - |V|$ edges between every pair of its nodes, then H is k-connected. Hence it is sufficient to prove that there are at least $k + 2 - |V|$ edges between every two nodes of H. To see this, consider a set of k internally disjoint uv-paths in H. At most $|V| - 2$ of these paths may not be edges between u, v, thus at least $k - (|V| - 2)$ of these paths are edges between u, v.

Finally, for rooted uniform requirements, we note that the existing 2-approximation algorithm in [11] does not have the restriction that G is simple, and hence works also for multi-graphs. $\square$

Clearly, a k-connected graph on q nodes has at least $\lceil kq/2 \rceil$ edges. For any $q \geq k+1$ this bound is achievable by so called Harary graphs [12]. We use the construction of Harary for k even.

Lemma 2.3 (Harary [12]) *Let $V = \{1, 2, \ldots, q\}$ be a set of nodes and let $k \leq q - 1$ be even. Then the graph on V with edge set $E(V, k) = \{ij : 1 \leq i < j \leq q, \min\{j - i, q + i - j\} \leq k/2\}$ is k-connected and has $kq/2$ edges.*

For a subset C of nodes of a graph G let $\Gamma_G(C)$ denote the set of neighbors of C in G. The following lemma plays a key role in the proof of Theorem 1.1, but it also of independent interest and may have other applications. To understand the implications of this lemma, consider the following scenario. We are given a set V of terminals in a graph G and an integer k. We want to remove from

G the nonterminal nodes and to add a "small" number of new edges, such that in the resulting graph J we will have $\kappa_J(u,v) \geq \min\{\kappa_G(u,v),k\}$ for all $u,v \in V$. We may do it by repeatedly removing a connected component C of the graph $G \setminus V$ and adding new edges on $\Gamma_G(C)$; note that $\Gamma_G(C) \subseteq V$, since C is a connected component of $G \setminus V$. Formally, our statement is the following.

Lemma 2.4 *Let V be a subset of nodes of a graph G, let k be an integer, and let C be a connected component of $G \setminus V$. Let J_C be a set of new edges on $\Gamma_G(C)$ such that the following holds.*

(i) *If $|\Gamma_G(C)| \leq k$ then J_C has $\min\{\ell_{uv}, k-|I_{uv}|\}$ uv-edges for any $u,v \in \Gamma_G(C)$, where I_{uv} is the set of uv-edges in G and ℓ_{uv} is the maximum number of internally disjoint uv-paths in the subgraph of G induced by $\{u,v\} \cup C$.*

(ii) *If $|\Gamma_G(C)| \geq k+1$ then the graph induced by $\Gamma_G(C)$ in $G \cup J_C$ is k-connected.*

Let $J = (G \setminus C) \cup J_C$. Then $\kappa_J(u,v) \geq \min\{\kappa_G(u,v),k\}$ for all $u,v \in V$.

Proof: The case $|\Gamma_G(C)| \leq k$ easily follows from the following construction. Let $u,v \in G \setminus C$. Given a set Π of at most k internally disjoint uv-paths in G, for every $P \in \Pi$ do the following. For every maximal $u'v'$-subpath of P that visits C and has all its internal nodes in C, replace this subpath by a $u'v'$-edge e not used by any other path in Π. Such e is chosen to be an edge of G if $\{u',v'\} \neq \{u,v\}$ and $I_{u'v'} \neq \emptyset$ or if $\{u',v'\} = \{u,v\}$ and $\min\{\ell_{uv}, k-|I_{uv}|\} = 0$. Otherwise, e is a new edge added to G. This gives a set of $|\Pi|$ internally disjoint uv-paths that do not visit C. Since the paths in Π are internally disjoint, the set of edges added to G may have parallel edges only between u and v, and by the construction, the number of uv-edges added, if any, can be at most $\min\{\ell_{uv}, |\Pi|-|I_{uv}|\} \leq \min\{\ell_{uv}, k-|I_{uv}|\}$.

Now suppose that $|\Gamma_G(C)| \geq k+1$, so $\Gamma_G(C)$ induces in $G+J_C$ a k-connected graph. Let $u,v \in G-C$. Let I_{uv} be a set of uv-edges in J. Let Q be a minimum size subset of nodes of J such that $J \setminus (Q \cup I_{uv})$ has no uv-path. By Menger's Theorem $\kappa_J(u,v) = |Q|+|I_{uv}|$. Thus if $|Q|+|I_{uv}| \geq k$ then $\kappa_J(u,v) \geq k \geq \min\{\kappa_G(u,v),k\}$. We claim that if $|Q|+|I_{uv}| \leq k-1$ then $G \setminus (Q \cup I_{uv})$ has no uv-path, hence by Menger's Theorem $\kappa_J(u,v) = |Q|+|I_{uv}| \geq \kappa_G(u,v)$. Suppose to the contrary that $G \setminus (Q \cup I_{uv})$ has a uv-path P. Going along P from u to v, let u' be the first and v' the last node in $\Gamma_G(C)$; such u',v' exist since P must contain at least one node from C, as P is not a uv-path in $J \setminus (Q \cup I_{uv})$. As J has k internally disjoint $u'v'$-paths and $|Q|+|I_{uv}| \leq k-1$, the graph $J \setminus (Q \cup I_{uv})$ has at least one $u'v'$-path P'. Replacing the $u'v'$-subpath of P by P' gives a uv-path in $J \setminus (Q \cup I_{uv})$, contradicting the definition of Q. $\square$

We also need the following lemma on dominating sets in unit disc graphs; this is the only place where we use Definition 1.2. Given a graph $G=(U,E)$, we say that $D \subseteq U$ is a *k-dominating set* in G if $|\Gamma_G(u) \cap D| \geq k$ for every $u \in U \setminus D$; note that for $k=1$ we get the usual definition of a dominating set in a graph.

Lemma 2.5 *Let (M, d) be a metric space, let $\Delta \geq 2$ be as in Definition 1.2, let $U \subseteq M$, and let C be a dominating set in $G[U]$. Then for any $k \leq |U \setminus C|$ there is a k-dominating set D in $G[U \setminus C]$ of size $|D| \leq \Delta k |C|$.*

Proof: For every $u \in C$ let D_u be defined as follows. Let $\Gamma_u = \Gamma_{G[U]}(u) \setminus C$. Since Γ_u is contained in the unit ball centered at u, there is a dominating set D_u^1 of size at most Δ in $G[\Gamma_u]$, by the definition of Δ. By the same argument there is a dominating set D_u^2 of size at most Δ in $G[\Gamma_u \setminus D_u^1]$. The set D_u is obtained by repeating the process k times and accumulating the dominating sets. Clearly, $|D_u| \leq k\Delta$ and D_u is a k-dominating set in $G[\Gamma_u]$. Let $D = \bigcup_{u \in C} D_u$. Then $|D| \leq k\Delta|C|$, and since $\bigcup_{u \in C} \Gamma_u = U \setminus C$, D is a k-dominating set in $G[U \setminus C]$. $\square$

3 Proof of the main result

To prove Theorem 1.1, we will prove the following statement.

Lemma 3.1 *There exists a polynomial time algorithm that, given an instance V, r of* SN-MSP*, constructs an instance $G = (V, E), c, r$ of* SNDP *such that the following holds. Any solution of cost C to* SNDP *can be converted in polynomial time to a solution of size $\leq C$ to* SN-MSP*, and for every solution S to* SN-MSP *there exists a solution J of cost $\leq |S| \cdot \rho(k)$ to* SNDP*. Furthermore, the construction preserves the requirement type (subset uniform, uniform, rooted, rooted subset uniform, and rooted uniform).*

Theorem 1.1 easily follows from Lemma 3.1. To see this, consider the following approximation algorithm for SN-MSP.

1. Construct the SNDP instance $(G = (V, E), c, r)$ as in Lemma 3.1.
2. Compute a subgraph $J \subseteq G$ satisfying r using an α-approximation algorithm.
3. Construct from J a feasible solution S to SN-MSP.

Lemma 3.1 ensures that the algorithm runs in polynomial time and computes a feasible solution S to the SN-MSP instance. We prove the approximation ratio. Let J^* be a minimum cost subgraph of G satisfying r, and let S^* be a minimum size set of points such that $G[V \cup S^*]$ satisfies r. Then

$$|S| \leq c(J) \leq \alpha \cdot c(J^*) \leq \alpha \cdot |S^*| \cdot \rho(k) \ .$$

The second inequality is since J is computed using an α-approximation algorithm, and the last inequality is by Lemma 3.1.

In the rest of this section we prove Lemma 3.1. We start by describing the construction of the SNDP instance $G = (V, E), c, r$.

Definition 3.1 *Given a finite set of points $V \subset M$ in a metric space (M, d) and an integer $k \geq 1$, the graph K_V is obtained by connecting every $u, v \in V$ by k parallel edges, one of cost $\lceil d(u,v) \rceil - 1$ and the others of cost $\lceil d(u,v) \rceil$.*

Clearly, given an SN-MSP instance, (V, r), the graph K_V with the corresponding edge costs $c(e)$ can be constructed in polynomial time. The triple (K_V, c, r) will serve as the SNDP instance guaranteed in Lemma 3.1. The construction preserves the requirement types listed in Lemma 3.1. Let J be a subgraph of K_V. Let $u, v \in V$ be connected in J by $j + 1 \leq k$ edges. Place $\lceil d(u,v) \rceil - 1$ new points uniformly on the line segment between u and v, dividing the segment into $\lceil d(u,v) \rceil$ subsegments, each of length $\frac{d(u,v)}{\lceil d(u,v) \rceil} \leq 1$. Since M is a normed space, and thus is also a linear space, this can be done; in fact, for $1 \leq i \leq \lceil d(u,v) \rceil - 1$, the ith point is of the form

$$\left(1 - \frac{i}{\lceil d(u,v) \rceil}\right) u + \frac{i}{\lceil d(u,v) \rceil} v$$

On each subsegment, place uniformly j new points in a similar fashion. Let $S(u,v)$ be the set of added points. Denote by $S(J)$ the union of $S(u,v)$ over all adjacent pairs $u, v \in V$.

Claim 3.2 *$|S(J)| \leq c(J)$ holds for any subgraph J of K_V. Furthermore, if $H = G[V \cup S(J)]$ is the unit disc graph induced by $V \cup S(J)$ then $\kappa_H(u,v) \geq \kappa_J(u,v)$ for all $u, v \in V$.*

Proof: To prove that $|S(J)| \leq c(J)$ it is enough to show that for all $u, v \in V$, $|S(u,v)|$ is at most the sum of the costs of all uv-edges in J. Note that K_V contains one uv-edge of cost $\lceil d(u,v) \rceil - 1$, and $k - 1$ uv-edges of cost $\lceil d(u,v) \rceil$ each. Thus, if J contains $j + 1 \leq k$ uv-edges, the sum of their costs is at least $\lceil d(u,v) \rceil - 1 + j \lceil d(u,v) \rceil = |S(u,v)|$.

We prove that $\kappa_H(u,v) \geq \kappa_J(u,v)$ for all $u, v \in V$. Let $u, v \in V$ and let $H' = G[\{u,v\} \cup S(u,v)]$. Assume that J contains $j + 1 \leq k$ uv-edges. To show that $\kappa_H(u,v) \geq \kappa_J(u,v)$ it is sufficient to show that $\kappa_{H'}(u,v) \geq j + 1$. Let U_0 be the set of points placed uniformly on the uv-segment. The distance between every pair of consecutive points is $\frac{d(u,v)}{\lceil d(u,v) \rceil} \leq 1$. The distance between u and the first point and v and the last point is also at most 1. Thus there is a uv-path in H' with node-set $U_0 \cup \{u, v\}$. Next, we order the sets of j points placed uniformly on each subsegment $1, 2, \ldots, j$ by their distance from u. For $1 \leq i \leq j$ let U_i be the set obtained by taking the ith point on each subsegment. The distance between every pair of consecutive points in U_i is $\frac{d(u,v)}{\lceil d(u,v) \rceil} \leq 1$. The distance between u and the ith point on the first subsegment is $\frac{i \cdot d(u,v)}{(j+1) \cdot \lceil d(u,v) \rceil} \leq 1$. The distance between v and the ith point on the last subsegment is $\frac{d(u,v)}{\lceil d(u,v) \rceil} - \frac{i \cdot d(u,v)}{(j+1) \cdot \lceil d(u,v) \rceil} = \frac{(j+1-i) \cdot d(u,v)}{(j+1) \cdot \lceil d(u,v) \rceil} \leq 1$. Thus for every i, there is a uv-path in H' with node-set $U_i \cup \{u, v\}$. Since $U_0, U_1, \ldots, U_j$ are pairwise disjoint, the result follows. $\square$

Clearly, $S(J)$ can be computed from J in polynomial time. This proves all parts of Lemma 3.1, except the one stating that for every solution S to SN-MSP there exists a solution J of cost $c(J) \leq |S| \cdot \rho(k)$ to SNDP; this will be proved in the rest of this section.

Let S be a feasible solution to an SN-MSP instance, so $G = G[V \cup S]$ satisfies r. The key step in constructing a solution J to SNDP of cost $c(J) \leq |S| \cdot \rho(k)$ is replacing every connected component C of $G \setminus V$ by an edge set J_C as in Lemma 2.4. Obviously, $\Gamma_G(C) \subseteq V$, and thus $J_C \subseteq K_V$. The following lemma shows that there exists such J_C of low cost.

Lemma 3.3 *For every connected component C of $G \setminus V$ there exists an edge set J_C on $\Gamma_G(C)$ as in Lemma 2.4 of edges of K_V such that $c(J_C) \leq \rho(k) \cdot |C|$.*

We prove Lemma 3.3 after the following corollary, which implies the last part of Lemma 3.1.

Corollary 3.4 *Let $\mathcal{C}$ be the set of connected components of $G \setminus V$. For $C \in \mathcal{C}$ let J_C be an edge set as in Lemma 3.3. Then $J = (G \setminus S) \bigcup (\bigcup_{C \in \mathcal{C}} J_C)$ is a subgraph of K_V of cost $c(J) \leq \rho(k) \cdot |S|$ that satisfies r.*

Proof: It is easy to see that for any $u, v \in V$ the number of uv-edges in J is at most k. Hence J is a subgraph of K_V. As $\mathcal{C}$ is a partition of S, we have by Lemma 3.3

$$c(J) \leq \sum_{C \in \mathcal{C}} c(J_C) \leq \sum_{C \in \mathcal{C}} \rho(k) \cdot |C| = \rho(k) \cdot |S| \ .$$

We prove that J satisfies r. Let $\mathcal{C} = \{C_1, C_2, \ldots, C_m\}$. Let $G_0 = G$ and for $j = 1, \ldots, m$ let $G_j = \left(G \setminus (\bigcup_{i=1}^{j} C_i) \right) \bigcup (\bigcup_{i=1}^{j} J_{C_i})$. Using Lemma 2.4, a simple induction shows that for all $1 \leq j \leq m$, G_j satisfies r. In particular, this is so for $J = G_m$. $\square$

Now we prove Lemma 3.3. Let $C \in \mathcal{C}$. We start with the case $|\Gamma_G(C)| \leq k$. In the notation of Lemma 2.4, ℓ_{uv} is the maximum number of internally disjoint uv-paths in the subgraph of G induced by $\{u, v\} \cup C$. Then J_C has at most ℓ_{uv} edges for every $u, v \in \Gamma_G(C)$. Let $u, v \in \Gamma_G(C)$. Since there are ℓ_{uv} internally disjoint uv-paths in the subgraph of G induced by $\{u, v\} \cup C$, there is one such path containing no more than $\lfloor |C|/\ell_{uv} \rfloor$ points in C. Since the distance between two consecutive nodes in the path is at most 1, and due to the triangle inequality, $d(u, v) \leq \lfloor |C|/\ell_{uv} \rfloor + 1$. Thus $c(u, v) \leq \lceil d(u, v) \rceil \leq \lfloor |C|/\ell_{uv} \rfloor + 1$. Consequently, the total cost of uv-edges in J_C is bounded by $\ell_{uv} \cdot \left(\frac{|C|}{\ell_{uv}} + 1 \right) = |C| + \ell_{uv} \leq 2|C|$. Thus as $|\Gamma_G(C)| \leq k$ we have

$$c(J_C) \leq \binom{k}{2} \cdot 2|C| \leq \frac{1}{2} k(k-1) \cdot 2|C| = k(k-1)|C| \leq \rho(k) \cdot |C| \ .$$

This finishes the proof of Lemma 3.3 for the case $|\Gamma_G(C)| \leq k$.

If $|\Gamma_G(C)| = k + 1$, J_C is the complete graph on $\Gamma_G(C)$. By the triangle inequality, $d(u, v) \leq |C| + 1$ for every $u, v \in \Gamma_G(C)$. Hence we get

$$c(J_C) \leq \binom{k+1}{2} \cdot 2|C| \leq \frac{1}{2} k(k+1) \cdot (|C| + 1) \leq k(k+1)|C| \leq \rho(k) \cdot |C| \ .$$

This finishes the proof of Lemma 3.3 for the case $|\Gamma_G(C)| = k + 1$.

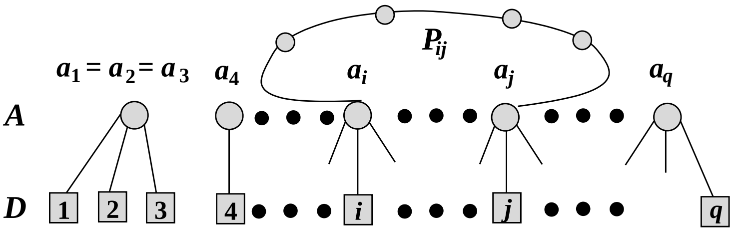

Figure 1: Illustration to the proof of Lemma 3.5.

In the rest of the proof of Lemma 3.3 assume that $|\Gamma_G(C)| \geq k+2 > 2\lceil k/2 \rceil$.

Lemma 3.5 *Let $U \subseteq M$, let C be a connected dominating set in $G[U]$, and let $D = U \setminus C$. Then for any $k \leq |D|-1$ even there is $E_D \subseteq K_D$ such that the graph (D, E_D) is k-connected and $c(E_D) \leq k|D|/2 + k(k+2)|C|/2$.*

Proof: Since C is a connected dominating set in $G[U]$, $G[C]$ is connected and $C \cap \Gamma_{G[U]}(v) \neq \emptyset$ for every $v \in D$. For every $v \in D$ choose some $a_v \in C \cap \Gamma_{G[U]}(v)$. Let $A = \{a_v : v \in D\}$. Let T be a spanning tree in $G[C]$. Order the nodes in A by running DFS on T. Let $1, \ldots, q$ be an order of D such that $i < j$ if a_i precedes a_j in the DFS order of A, see Figure 1. Let $E_D = E(D,k) = \{ij : 1 \leq i < j \leq q, \min\{j-i, q+i-j\} \leq k/2\}$ be as in Lemma 2.3, so the graph (D, E_D) is k-connected and has $|E_D| = k|D|/2$ edges.

We prove that $c(E_D) \leq k|D|/2 + k(k+2)|C|/2$. For $i, j \in D$ such that $i < j$, let $P_{i,j} = P(a_i, a_j)$ denote the unique $a_i a_j$-path in T and let $|P_{i,j}|$ be the number of edges in $P_{i,j}$ ($|P_{i,j}| = 0$ if $a_i = a_j$). By the triangle inequality we have

$$d(i,j) \leq d(i,a_i) + d(a_i,a_j) + d(a_j,j) \leq 1 + |P_{ij}| + 1 = 2 + |P_{ij}| \ .$$

Since the graph (D, E_D) is simple, we can choose the cheapest ij-edge in K_D and therefore

$$c(ij) = \lceil d(i,j) \rceil - 1 \leq 2 + |P_{i,j}| - 1 = 1 + |P_{i,j}| \ .$$

This implies

$$c(E_D) = \sum_{ij \in E_D} c(ij) \leq |E_D| + \sum_{ij \in E_D} |P_{i,j}| = k|D|/2 + \sum_{ij \in E_D} |P_{i,j}| \ .$$

It remains to show that $\sum_{ij \in E_D} |P_{i,j}| \leq k(k+2)|C|/2$. Note that $|P_{i,j}| = \sum_{m=i}^{j-1} |P_{m,m+1}|$ and that $\sum_{m=1}^{q-1} |P_{m,m+1}| = 2(|C|-1) \leq 2|C|$. By the construction of E_D, the number of times each $P_{m,m+1}$ is shortcut by the edges in E_D equals to $2 \cdot \sum_{p=1}^{k/2} p = k(k+2)/4$. This implies

$$\sum_{ij \in E_D} |P_{i,j}| = \sum_{ij \in E_D} \sum_{m=i}^{j-1} |P_{m,m+1}| \leq \frac{k(k+2)}{4} \sum_{m=1}^{q-1} |P_{m,m+1}| \leq \frac{k(k+2)|C|}{2} \ .$$

The proof of the lemma is complete. $\square$

We finish the proof of Lemma 3.3 for the case $|\Gamma_G(C)| \geq k+2$. Note that C is a connected dominating set in the subgraphs of G induced by $C \cup D$ for any $D \subseteq \Gamma_G(C)$. By Lemma 2.5, there is a k-dominating set D in $G[\Gamma_G(C)]$ of size $|D| \leq \Delta k|C|$. By Lemma 3.5, there is $E_D \subseteq K_D$ such that (D, E_D) is $2\lceil k/2\rceil$-connected, and

$$c(E_D) \leq \frac{2\lceil k/2\rceil |D|}{2} + \frac{2\lceil k/2\rceil(2\lceil k/2\rceil+2)|C|}{2} = 2\lceil k/2\rceil \left(|D|/2 + (\lceil k/2\rceil+1)|C|\right) .$$

Combining we get

$$c(E_D) \leq 2\lceil k/2\rceil \left(|D|/2 + (\lceil k/2\rceil+1)|C|\right) \leq 2\lceil k/2\rceil \left(\Delta k|C|/2 + (\lceil k/2\rceil+1)|C|\right) = \rho(k)|C| .$$

Let $E_0 = \{uv : u, v \in \Gamma_G(C), d(u,v) \leq 1\}$ be the set of edges of cost 0 in $K_{\Gamma_G(C)}$. Let $J_C = E_D \cup E_0$. Clearly, $J_C \subseteq K_{\Gamma_G(C)}$ and $c(J_C) = c(E_D) \leq \rho(k)|C|$. Since (D, E_D) is k-connected, and since D is a k-dominating set in $G[\Gamma_G(C)]$, the graph $(\Gamma_G(C), J_C)$ is k-connected, by Lemma 2.2.

This finishes the proof of Lemma 3.3, and thus also the proof of Lemma 3.1 is complete.

A tight example: The following example shows that our analysis is tight (up to constants). Given k points in a ball of radius $1/2$ with uniform requirements as an instance for SN-MSP, an optimal solution size is 1 – add one Steiner point in the ball. An optimal solution for the SNDP instance has cost $\binom{k}{2}$, as it is a union of two cliques on V: in one clique every edge uv has cost $\lceil d(u,v)\rceil - 1 = 0$, while in the other every edge uv has cost $\lceil d(u,v)\rceil = 1$.

References

[1] A. Agrawal, P. N. Klein, and R. Ravi. When trees collide: An approximation algorithm for the generalized Steiner problem on networks. *SIAM J. on Computing*, 24(3):440–456, 1995.

[2] E. Althaus, G. Calinescu, I. Mandoiu, K Prasad, N. Tchervenski, and A. Zelikovsky. Power efficient range assignment for symmetric connectivity in static ad hoc wireless networks. *Wireless Networks*, 12(3):287–299, 2006.

[3] J. Bredin, E. Demaine, M. Hajiaghayi, and D. Rus. Deploying sensor networks with guaranteed capacity and fault tolerance. In *MobiHoc*, pages 309–319, 2005.

[4] J. Cheriyan, S. Vempala, and A. Vetta. An approximation algorithm for the minimum-cost k-vertex connected subgraph. *SIAM J. on Computing*, 32(4):1050–1055, 2003.

[5] J. Chuzhoy and S. Khanna. Algorithms for single-source vertex-connectivity. In *FOCS*, pages 105–114, 2008.

[6] J. Chuzhoy and S. Khanna. An $O(k^3 \log n)$-approximation algorithm for vertex-connectivity survivable network design. In *FOCS*, pages 437–441, 2009.

[7] D-Z. Du, L. Wang, and B. Xu. The Euclidean bottleneck Steiner tree and Steiner tree with minimum number of Steiner points. In *COCOON*, pages 509–518, 2001.

[8] D. Estrin, L. Girod, G. Pottie, and M. Srivastava. Instrumenting the world with wireless sensor networks. In *ICASSP*, pages 2033–2036, 2001.

[9] J. Fackharoenphol and B. Laekhanukit. An $O(\log^2 k)$-approximation algorithm for the k-vertex connected subgraph problem. In *STOC*, pages 153–158, 2008.

[10] A. Frank. Connectivity and network flows. In R. L. Graham, M. Grötschel, and L. Lovász, editors, *Handbook of Combinatorics*, chapter 2, pages 111–177. Elsevier Science, 1995.

[11] A. Frank and É. Tardos. An application of submodular flows. *Linear Algebra and its Applications*, 114/115:329–348, 1989.

[12] F. Harary. The maximum connectivity of a graph. In *Natl. Acad. Sci.*, pages 1142–1146, 1962.

[13] K. Jain. A factor 2 approximation algorithm for the generalized Steiner network problem. *Combinatorica*, 21(1):39–60, 2001.

[14] G. Kabatjansky and V. Levenstein. Bounds for packing of the sphere and in space. *Prob. Information Trans.*, 14:1–17, 1978.

[15] C. Klein and R. Ravi. A nearly best-possible approximation algorithm for node-weighted Steiner trees. *J. of Algorithms*, 19(1):104–115, 1995.

[16] G. Kortsarz, V. S. Mirrokni, Z. Nutov, and E. Tsanko. Approximating minimum-power degree and connectivity problems. In *LATIN*, pages 423–435, 2008.

[17] G. Kortsarz and Z. Nutov. Approximating k-node connected subgraphs via critical graphs. *SIAM J. on Computing*, 35(1):247–257, 2005.

[18] Z. Nutov. Approximating minimum cost connectivity problems via uncrossable bifamilies. Manuscript, 2010. Preliminary version in *FOCS* 2009, pages 417-426.

[19] Z. Nutov. An almost $O(\log k)$-approximation for k-connected subgraphs. In *SODA*, pages 912–921, 2009.

[20] Z. Nutov. Approximating minimum-power k-connectivity. *Ad Hoc & Sensor Wireless Networks*, 9(1-2):129–137, 2010.

[21] Z. Nutov. Approximating Steiner networks with node weights. *SIAM J. on Computing*, 39(7):3001–3022, 2010.

[22] Z. Nutov and A. Yaroshevitch. Wireless network design via 3-decompositions. *Informayion Processing Letters*, 109(19):1136–1140, 2009.

[23] G. Robins and J. Salowe. Low-degree minimum spanning trees. *Discrete & Computational Geometry*, 14(2):151–165, 1995.