

An $O(k^3 \log n)$ -Approximation Algorithm for Vertex-Connectivity Survivable Network Design

Julia Chuzhoy*
 Toyota Technological Institute
 Chicago, IL 60637
 cjulia@tti-c.org

Sanjeev Khanna†
 University of Pennsylvania
 Philadelphia PA 19104
 sanjeev@cis.upenn.edu

Abstract— In the Survivable Network Design problem (SNDP), we are given an undirected graph $G(V, E)$ with costs on edges, along with a connectivity requirement $r(u, v)$ for each pair u, v of vertices. The goal is to find a minimum-cost subset E^* of edges, that satisfies the given set of pairwise connectivity requirements. In the *edge-connectivity version* we need to ensure that there are $r(u, v)$ edge-disjoint paths for every pair u, v of vertices, while in the *vertex-connectivity version* the paths are required to be vertex-disjoint. The edge-connectivity version of SNDP is known to have a 2-approximation. However, no non-trivial approximation algorithm has been known so far for the vertex version of SNDP, except for special cases of the problem.

We present an extremely simple algorithm to achieve an $O(k^3 \log |T|)$ -approximation for this problem, where k denotes the maximum connectivity requirement, and T is the set of vertices that participate in one or more pairs with non-zero connectivity requirements. We also give a simple proof of the recently discovered $O(k^2 \log |T|)$ -approximation algorithm for the single-source version of vertex-connectivity SNDP. Our results establish a natural connection between vertex-connectivity and a well-understood generalization of edge-connectivity, namely, *element-connectivity*, in that, any instance of vertex-connectivity can be expressed by a small number of instances of the element-connectivity problem.

Keywords-survivable network design; vertex-connectivity.

1. INTRODUCTION

In the Survivable Network Design problem (SNDP), we are given an undirected graph $G(V, E)$ with costs on edges, and a connectivity requirement $r(u, v)$ for each pair u, v of vertices. The goal is to find a minimum cost subset E^* of edges, such that each pair (u, v) of vertices is connected by $r(u, v)$ paths. In the edge-connectivity version (EC-SNDP), these paths are required to be edge-disjoint, while in the vertex-connectivity version (VC-SNDP), they need to be vertex-disjoint. It is not hard to show that EC-SNDP can be cast as a special case of VC-SNDP. We denote by n the number of vertices in the graph and by k the maximum pairwise connectivity requirement, that is, $k = \max_{u, v \in V} \{r(u, v)\}$. We also define a subset $T \subseteq V$ of vertices called *terminals*: a vertex $u \in T$ iff $r(u, v) > 0$ for some vertex $v \in V$.

*Research supported in part by NSF CAREER award CCF-0844872.

†Research supported in part by a Guggenheim Fellowship, an IBM Faculty Award, and by NSF Award CCF-0635084.

General VC-SNDP: While a celebrated result of Jain [16] gives a 2-approximation algorithm for EC-SNDP, no non-trivial approximation algorithms are known for VC-SNDP, except for restricted special cases. Agrawal et al. [1] showed a 2-approximation algorithm for the special case when maximum connectivity requirement $k = 1$. For $k = 2$, a 2-approximation algorithm was given by Fleischer [11]. The k -vertex connected spanning subgraph problem, a special case of VC-SNDP where for all $u, v \in V$ $r_{u,v} = k$, has been studied extensively. Cheriyan et al. [2], [3] gave an $O(\log k)$ -approximation algorithm for this case when $k \leq \sqrt{n/6}$, and an $O(\sqrt{n/\epsilon})$ -approximation algorithm for $k \leq (1-\epsilon)n$. For large k , Kortsarz and Nutov [20] improved the preceding bound to an $O(\ln k \cdot \min\{\sqrt{k}, \frac{n}{n-k} \ln k\})$ -approximation. Fakcharoenphol and Laekhanukit [10] improved it to an $O(\log n \log k)$ -approximation, and further obtained an $O(\log^2 k)$ -approximation for $k < n/2$. Very recently, Nutov [25] improved this to $O(\log k \cdot \log \frac{n}{n-k})$ -approximation.

Kortsarz et al. [18] showed that VC-SNDP is hard to approximate to within a factor of $2^{\log^{1-\epsilon} n}$ for any $\epsilon > 0$, when k is polynomially large in n . This result was subsequently strengthened by Chakraborty et al. [4] to a k^ϵ -hardness for all $k > k_0$, where k_0 and ϵ are fixed positive constants. However, the existence of good approximation algorithms for small values of k has remained an open problem, even for $k \geq 3$. In particular, when each connectivity requirement $r_{u,v} \in \{0, 3\}$, the best known approximation factor is polynomially large ($\tilde{O}(n)$ to best of our knowledge) while only an APX-hardness is known. The main result of our paper is an $O(k^3 \log |T|)$ -approximation algorithm for VC-SNDP.

Single-Source VC-SNDP: A special case of VC-SNDP that has received much attention recently is the single-source version. In this problem there is a special vertex s called the *source*, and all non-zero connectivity requirements involve s , that is, if $u \neq s$ and $v \neq s$, then $r(u, v) = 0$. Kortsarz et al [18] showed that even this restricted special case of VC-SNDP is hard to approximate up to factor $\Omega(\log n)$, and recently Lando and Nutov [22] improved this to $(\log n)^{2-\epsilon}$.

hardness of approximation for any constant $\epsilon > 0$. We note that both results only hold when k is polynomially large in n . On the algorithmic side, Chakraborty et. al. [4] gave an $2^{O(k^2)} \log^4 n$ -approximation for the problem. This result was later independently improved to an $O(k^{O(k)} \log n)$ -approximation by Chekuri and Korula [5], and to an $O(k^2 \log n)$ -approximation by Chuzhoy and Khanna [8], and by Nutov [23]. Recently, Chekuri and Korula [6] simplified the analysis of the algorithm of [8]. We note that for the uniform case, where all non-zero connectivity requirements are k , Chuzhoy and Khanna [8] show a slightly better $O(k \log n)$ -approximation algorithm, and the results of [6] extend to this special case. In this paper we give a simple $O(k^2 \log |T|)$ -approximation algorithm for single-source VC-SNDP.

Element-Connectivity SNDP: A closely related problem to EC-SNDP and VC-SNDP is the element-connectivity SNDP. The input to the element-connectivity SNDP is the same as for EC-SNDP and VC-SNDP. As before, we define the set $T \subseteq V$ of terminals to be vertices that participate in one or more pairs with a positive connectivity requirement. Given a problem instance, an *element* is any edge or any *non-terminal* vertex in the graph. We say that a pair (s, t) of vertices is *k-element connected* iff for every subset X of at most $(k - 1)$ elements, s and t remain connected by a path when X is removed from the graph. In other words, there are k element-disjoint paths connecting s to t ; these paths are allowed to share terminals. Observe that if a pair (s, t) is k -vertex connected, then it is also k -element connected, and similarly, if a pair (s, t) is k -element connected, then it is also k -edge connected. But the converse relationships do not hold, that is, if a pair (s, t) is k -edge connected, then it need not be k -element connected, and similarly, if a pair (s, t) is k -element connected, then it need not be k -vertex connected. Thus the notion of element-connectivity resides in between edge-connectivity and vertex-connectivity. The goal in the element-connectivity SNDP is to select a minimum-cost subset E^* of edges, such that in the graph induced by E^* , each pair (u, v) of vertices is $r(u, v)$ -element connected. The element-connectivity SNDP was introduced in [17] as a problem of intermediate difficulty between edge-connectivity and vertex-connectivity, and the authors gave a primal-dual $O(\log k)$ -approximation for this problem. Subsequently, Fleischer et al. [12] gave a 2-approximation algorithm for element-connectivity SNDP via the iterative rounding technique, matching the 2-approximation guarantee of Jain [16] for EC-SNDP. We use this result as a building block for our algorithm.

Our Results: Our main result is as follows.

Theorem 1: There is a polynomial-time randomized $O(k^3 \log |T|)$ -approximation algorithm for VC-SNDP, where k is the largest pairwise connectivity requirement.

The proof of this result is based on a randomized reduc-

tion that maps a given instance of VC-SNDP to a family of instances of element-connectivity SNDP. The reduction creates $O(k^3 \log |T|)$ instances, and has the property that any collection of edges that is feasible for *each one* of the element-connectivity SNDP instances generated above, is a feasible solution for the given VC-SNDP instance. We can thus use the known 2-approximation algorithm for element-connectivity SNDP to obtain the desired result.

We use these ideas to also give an alternative simple proof of the $O(k^2 \log |T|)$ -approximation algorithm for the single-source VC-SNDP problem.

As noted earlier, the notion of element-connectivity is trivially subsumed by vertex-connectivity. Our result shows that in a weak sense, the converse also holds in that any set of pairwise vertex-connectivity requirements can be captured by a collection of element-connectivity instances.

Remark 1: We note that very recently, subsequent to our work, Nutov [24] has shown an $O(k^2)$ -approximation algorithm for single-source VC-SNDP. He also studied the more general version of VC-SNDP, where the costs are on vertices (instead of edges), and has given an $O(k^4 \log^2 |T|)$ -approximation algorithm for the general problem, and an $O(k^2 \log |T|)$ -approximation for the single-source version. The latter result improves upon the recent $O(k^8 \log^2 n)$ -approximation [8].

Organization: We present the proof of Theorem 1 in Section 2. Section 3 presents an alternative proof of the $O(k^2 \log |T|)$ -approximation result for single-source VC-SNDP. In Section 4 we show a connection between our techniques and a well-studied notion of cover-free families. Using this connection we show that our algorithms are essentially tight, and that similar techniques cannot give significantly better approximation guarantees.

2. THE ALGORITHM FOR VC-SNDP

Recall that in VC-SNDP we are given an undirected graph $G(V, E)$ with costs on edges, and a connectivity requirement $r(u, v) \leq k$ for all $u, v \in V$. Additionally, we have a subset $T \subseteq V$ of terminals, and $r(u, v) > 0$ only if $u, v \in T$. The pairs of terminals with non-zero connectivity requirements are called *source-sink pairs*. We will use OPT to denote the cost of an optimal solution to the given VC-SNDP instance.

Our algorithm is as follows. We create p identical copies of our input graph G , say G_1, G_2, \dots, G_p , where p is a parameter to be determined later. For each copy G_i we define a subset $T_i \subseteq T$ of terminals. We then view G_i as an instance of element-connectivity SNDP, where the connectivity requirements are induced by the set T_i of terminals as follows. For each $s, t \in T_i$ the new connectivity requirement is the same as the original one. For all other pairs the connectivity requirements are 0. Observe that for each G_i the cost of an optimal solution for the induced element-connectivity SNDP instance is at most OPT . We

then apply the 2-approximation algorithm of [12] to each one of the p instances of the k -element connectivity problem. Let E_i denote the set of edges output by the 2-approximation algorithm on the instance defined on the G_i . Our final solution is $E^* = E_1 \cup E_2 \cup \dots \cup E_p$. Since any solution to the original VC-SNDP instance is also a feasible solution for each one of the p element-connectivity instances created above, the cost of the solution above is bounded by $2p \cdot \text{OPT}$.

We now show that for $p = O(k^3 \log |T|)$, there exist subsets T_1, T_2, \dots, T_p such that the solution E^* produced above is a feasible solution for VC-SNDP. Moreover, we show a simple randomized algorithm to create the sets T_1, T_2, \dots, T_p .

Definition 2.1: Let \mathcal{M} be the input collection of source-sink pairs, and let T be the corresponding set of terminals. We say that a family $\{T_1, \dots, T_p\}$ of subsets of T is k -resilient iff for each source-sink pair $(s, t) \in \mathcal{M}$, for each subset $X \subseteq T \setminus \{s, t\}$ of size at most $(k-1)$, there is a subset T_i , $1 \leq i \leq p$, such that $s, t \in T_i$ and $X \cap T_i = \emptyset$.

We show below that a k -resilient family of subsets exists for $p = O(k^3 \log |T|)$, and give a poly-time randomized algorithm to find such a family with high probability. We start by proving that such a family guarantees that the algorithm produces a feasible solution.

Lemma 2: Let $\{T_1, \dots, T_p\}$ be a k -resilient family of subsets. Then the output E^* of the above algorithm is a feasible solution to the VC-SNDP instance.

Proof: Let $(s, t) \in \mathcal{M}$ be any source-sink pair, and let $X \subseteq V \setminus \{s, t\}$ be any collection of at most $(r(s, t) - 1) \leq (k-1)$ vertices. It is enough to show that the removal of X from the graph induced by E^* does not separate s from t . Let $X' = X \cap T$. Since $\{T_1, \dots, T_p\}$ is a k -resilient family of subsets, there is some T_i such that $s, t \in T_i$ while $T_i \cap X' = \emptyset$. Recall that set E_i of edges defines a feasible solution to the element-connectivity SNDP instance corresponding to T_i . Then X is a set of non-terminal vertices with respect to T_i . Since s is $r(s, t)$ -element connected to t in the graph induced by E_i , the removal of X from the graph does not disconnect s from t . ■

We now show how to construct a k -resilient family of subsets $\{T_1, \dots, T_p\}$. Let $p = 128k^3 \log |T|$, and set $q = p/(2k) = 64k^2 \log |T|$. Each terminal $t \in T$ selects q random indices uniformly and independently from the set $\{1, 2, \dots, p\}$ (repetitions are allowed). Let $\phi(t)$ denote the set of indices chosen by the terminal t . For each $1 \leq i \leq p$, we then define $T_i = \{t \mid i \in \phi(t)\}$.

Lemma 3: With high probability, the resulting family $\{T_1, \dots, T_p\}$ of subsets is k -resilient.

Proof: We extend the definition of $\phi()$ to an arbitrary subset Z of vertices by defining $\phi(Z) = \bigcup_{t \in Z \cap T} \phi(t)$. Fix any source-sink pair (s, t) . Let X be an arbitrary set of at most $(k-1)$ vertices that does not include s, t . Note that $|\phi(X)| \leq (k-1)q < p/2$. We say that the *bad event*

$\mathcal{E}_1(s, t, X)$ occurs if $|\phi(s) \cap \phi(X)| \geq \frac{3q}{4}$. The expected value of $|\phi(s) \cap \phi(X)|$ is at most $q/2$, so by Chernoff bounds,

$$\Pr[\mathcal{E}_1(s, t, X)] \leq e^{-\frac{q}{32}}.$$

We say that the *bad event* $\mathcal{E}_2(s, t, X)$ occurs if $\phi(s) \cap \phi(t) \subseteq \phi(X)$. We say that the set X is a *bad set* for a pair (s, t) if the event $\mathcal{E}_2(s, t, X)$ occurs. Note that if there is no bad set X of size at most $(k-1)$ for every pair $(s, t) \in \mathcal{M}$, then $\{T_1, \dots, T_p\}$ is a k -resilient family.

We observe that if event $\mathcal{E}_1(s, t, X)$ does not happen, then $|\phi(s) \setminus \phi(X)| \geq q/4$, so

$$\Pr[\mathcal{E}_2(s, t, X) \mid \overline{\mathcal{E}_1(s, t, X)}] \leq \left(1 - \frac{q/4}{p}\right)^q \leq e^{-\frac{q^2}{4p}} \leq e^{-\frac{q}{8k}}$$

Thus we can bound the probability of the event $\mathcal{E}_2(s, t, X)$ as follows:

$$\begin{aligned} \Pr[\mathcal{E}_2(s, t, X)] &= \Pr[\mathcal{E}_2(s, t, X) \mid \mathcal{E}_1(s, t, X)] \Pr[\mathcal{E}_1(s, t, X)] \\ &\quad + \Pr[\mathcal{E}_2(s, t, X) \mid \overline{\mathcal{E}_1(s, t, X)}] \Pr[\overline{\mathcal{E}_1(s, t, X)}] \\ &\leq \Pr[\mathcal{E}_1(s, t, X)] + \Pr[\mathcal{E}_2(s, t, X) \mid \overline{\mathcal{E}_1(s, t, X)}] \\ &\leq e^{-\frac{q}{32}} + e^{-\frac{q}{8k}} \\ &< |T|^{-4k}. \end{aligned}$$

Hence, using the union bound, the probability that some bad set X of size at most $(k-1)$ exists for any pair (s, t) can be bounded by $|T|^{-2k}$. ■

Combining Lemmas 2 and 3, we obtain the following corollary.

Corollary 1: There is a randomized $O(k^3 \log |T|)$ -approximation algorithm for VC-SNDP.

Remark 2: We note that this result implies that the standard set-pair relaxation for VC-SNDP [14] has an integrality gap of $O(k^3 \log |T|)$. This follows from the fact that the 2-approximation result of [12] also establishes an upper bound of 2 on the integrality gap of the set-pair relaxation for element-connectivity. We also note that a lower bound of $\tilde{\Omega}(k^{1/3})$ is known on the integrality gap of the set-pair relaxation for VC-SNDP [4].

Remark 3: We also note that our reduction carries over to the node-weighted version of VC-SNDP, and in particular an α -approximation algorithm for the node-weighted element-connectivity SNDP would imply an $O(\alpha k^3 \log |T|)$ -approximation for the node-weighted VC-SNDP.

3. THE ALGORITHM FOR SINGLE-SOURCE VC-SNDP

In this section we show that an $O(k^2 \log |T|)$ -approximation algorithm can be easily achieved using the above ideas for the single-source version of VC-SNDP. Several algorithms achieving similar approximation factors have been proposed recently [8], [6], [23]. While the algorithm

and the analysis proposed here are elementary, we make use of the (relatively involved) 2-approximation algorithm of [12] as a black box. The algorithms of [8], [6] have the advantage that they are presented “from scratch”, using only elementary tools, and when viewed as such they are rather simple.

Recall that the input to the single-source VC-SNDP is a graph $G(V, E)$ with a special vertex s called the source, and a subset T of terminals, where for each $t \in T$, we are given a connectivity requirement $r(s, t) \leq k$. The goal is to select a minimum-cost subset $E' \subseteq E$ of edges, such that in the graph induced by E' every terminal $t \in T$ is $r(s, t)$ -vertex connected to s . This is clearly a special case of VC-SNDP, where all source-sink pairs are of the form $\{(s, t)\}_{t \in T}$. As before, we create a family $\{T_1, \dots, T_p\}$ of subsets of terminals, $T_i \subseteq T$ for all $1 \leq i \leq p$. We also create p identical copies of our input graph G , say G_1, \dots, G_p . For each G_i we solve the single-source element-connectivity SNDP instance with connectivity requirements induced by terminals in T_i . Let E_i be the 2-approximate solution to instance G_i . Our final solution is $E^* = \bigcup_{i=1}^p E_i$. Clearly, the cost of the solution is at most $2p \cdot \text{OPT}$.

Definition 3.1: A family $\{T_1, \dots, T_p\}$ of subsets of terminals is *weakly k -resilient* iff for each terminal $t \in T$, for each subset $X \subseteq T \setminus \{t\}$ of at most $(k-1)$ terminals, there is $i : 1 \leq i \leq p$, such that $t \in T_i$ and $X \cap T_i = \emptyset$.

Lemma 4: If $\{T_1, \dots, T_p\}$ is a weakly k -resilient family of subsets then the above algorithm produces a feasible solution.

Proof: Let $t \in T$ and let $X \subseteq V \setminus \{s, t\}$ be any subset of at most $r(s, t) - 1 \leq (k-1)$ vertices excluding s and t . It is enough to prove that the removal of X from the graph induced by E^* does not disconnect s from t . Let $X' = X \cap T$. Since $\{T_1, \dots, T_p\}$ is a weakly k -resilient family, there is some $i : 1 \leq i \leq p$ such that $t \in T_i$ and $T_i \cap X' = \emptyset$. Consider the solution E_i to the corresponding k -element connectivity instance. Since vertices of X are non-terminal vertices for the instance G_i , their removal from the graph induced by E_i does not disconnect s from t . ■

Let $p = 4k^2 \log |T|$ and $q = p/(2k) = 2k \log |T|$. Each terminal $t \in T$ selects q indices from the set $\{1, 2, \dots, p\}$ uniformly at random with repetitions. Let $\phi(t)$ denote the set of indices chosen by the terminal t . For each $1 \leq i \leq p$, we then define $T_i = \{t \mid i \in \phi(t)\}$.

Lemma 5: With high probability, the resulting family of subsets $\{T_1, \dots, T_p\}$ is weakly k -resilient.

Proof: Let $t \in T$ be any terminal and let X be any subset of at most $r(s, t) - 1 \leq (k-1)$ terminals. As before, we extend the function ϕ to an arbitrary subset Z of vertices by defining $\phi(Z) = \bigcup_{t \in Z \cap T} \phi(t)$. We say that *bad event* $\mathcal{E}(t, X)$ occurs iff $\phi(t) \subseteq \phi(X)$.

The probability of $\mathcal{E}(t, X)$ is at most

$$\left(1 - \frac{kq}{p}\right)^q = \left(\frac{1}{2}\right)^q \leq |T|^{-2k}$$

Therefore, with high probability the event $\mathcal{E}(t, X)$ does not happen for any t, X and then $\{T_1, \dots, T_p\}$ is weakly k -resilient. ■

Combining Lemmas 4 and 5, we obtain the following corollary.

Corollary 2: There is a randomized $O(k^2 \log |T|)$ -approximation algorithm for single-source VC-SNDP.

4. RESILIENT VS. COVER-FREE FAMILIES

The notion of a k -resilient and weakly k -resilient families is closely related to a well-studied notion in coding theory and combinatorics, namely, *cover-free* families of sets. A family \mathcal{F} of sets over a universe $U = \{1, 2, \dots, p\}$ is said to be *r -cover-free* if for all distinct $A, S_1, \dots, S_r \in \mathcal{F}$, it satisfies the property that $A \not\subseteq \bigcup_{j=1}^r S_j$. This is precisely the property underlying our construction of a weakly k -resilient family. In particular, $\{T_1, T_2, \dots, T_p\}$ is weakly k -resilient iff $\mathcal{F} = \{\phi(t) \mid t \in T\}$ is a $(k-1)$ -cover-free family.

Let $N(r, \lambda)$ denote the smallest integer p such that there exists an r -cover-free family with λ sets over a universe of p elements. It is easy to see that the smaller the value $N(r, \lambda)$, the better the approximation guarantee achieved by the algorithm of Section 3. A classical result of Dyachkov and Rykov [9] (see the note by Füredi [15] for a simple proof of this lower bound result) shows that

$$N(r, \lambda) = \Omega\left(\frac{r^2 \log \lambda}{\log r}\right).$$

An immediate corollary of this result is that for any weakly k -resilient family for a set T of terminals, the parameter p must be $\Omega\left(\frac{k^2 \log |T|}{\log k}\right)$. Thus the bound achieved by the simple randomized construction given in Lemma 5 is tight to within a $O(\log k)$ factor.

Kumar, Rajagopalan, and Sahai [21] gave an elegant deterministic construction for cover-free families based on Reed-Solomon codes. The construction gives slightly weaker guarantees than the randomized construction. For sake of completeness, we briefly describe their construction. Let $\mathbb{F}_q = \{u_1, u_2, \dots, u_q\}$ be a finite field for some prime q . Moreover, let $F_{q,d}$ be the set of all polynomials over \mathbb{F}_q of degree at most d where $d = q/k$. Consider the family of sets $\mathcal{F} = \{S_f \mid f \in F_{q,d+1}\}$ defined over the universe $U = \mathbb{F}_q \times \mathbb{F}_q$ where $S_f = \{\langle u_1, f(u_1) \rangle, \dots, \langle u_q, f(u_q) \rangle\}$. Then \mathcal{F} is a $(k-1)$ -cover-free family since any two distinct polynomials in $F_{q,d}$ can agree on at most d points. Since the size of the underlying universe U is $p = q^2$ and $|\mathcal{F}| = \Omega(q^d)$, we get a deterministic construction for a weakly k -resilient family with $p = O\left(\frac{k^2 \log^2 |T|}{\log^2(k \log |T|)}\right)$.

A natural generalization of r -cover-free family is a (w, r) -cover-free family that is defined as follows. A family \mathcal{F} of sets over a universe $U = \{1, 2, \dots, p\}$ is said to be (w, r) -cover-free if for all any $A_1, A_2, \dots, A_w \in \mathcal{F}$ and any other $S_1, \dots, S_r \in \mathcal{F}$, it satisfies the property that $\bigcap_{i=1}^w A_i \not\subseteq \bigcup_{j=1}^r S_j$. It is easy to see that $\{T_1, T_2, \dots, T_p\}$ is k -resilient iff $\mathcal{F} = \{\phi(t) \mid t \in T\}$ is a $(2, k-1)$ -cover-free family. Let $N(w, r, \lambda)$ denote the smallest integer p such that there exists a (w, r) -cover-free family with λ sets over a universe of p elements. Stinson, Wei, and Zhu [27] showed that for any $r \geq 1$, there exists a λ_0 that depends only on r , such that for all $\lambda \geq \lambda_0$

$$N(2, r, \lambda) = \Omega\left(\frac{r^3 \log \lambda}{\log r}\right).$$

An immediate corollary of this result is that for any k -resilient family for a set T of terminals, the parameter p must be $\Omega\left(\frac{k^3 \log |T|}{\log k}\right)$. Thus the bound achieved by the simple randomized construction given in Lemma 3 is tight to within a $O(\log k)$ factor.

ACKNOWLEDGEMENTS

We thank Chandra Chekuri for his helpful comments on an earlier version of this paper.

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. of Computing*, 24(3):440–456, 1995.
- [2] J. Cheriyan, S. Vempala, and A. Vetta. An approximation algorithm for the minimum-cost k -vertex connected subgraph. *SIAM J. of Computing*, 32(4):1050–1055, 2003.
- [3] J. Cheriyan, S. Vempala, and A. Vetta. Network design via iterative rounding of setpair relaxations. *Combinatorica*, 26(3):255–275, 2006.
- [4] T. Chakraborty, J. Chuzhoy, and S. Khanna. Network Design for Vertex Connectivity. In *Proceedings of ACM Symp. on Theory of Computing (STOC)*, 2008.
- [5] C. Chekuri and N. Korula. Single-Sink Network Design with Vertex Connectivity Requirements. *FSTTCS*, 2008.
- [6] C. Chekuri and N. Korula. A Graph Reduction Step Preserving Element -Connectivity and Applications. *ICALP*, 2009.
- [7] J. Chuzhoy and S. Khanna. An $O(k^3 \log n)$ -Approximation Algorithm for Vertex-Connectivity Survivable Network Design. arXiv:0812.4442v1, 2008.
- [8] J. Chuzhoy and S. Khanna. Algorithms for single-source vertex connectivity. In *Proceedings of the IEEE Symp. on Foundations of Computer Science (FOCS)*, 2008.
- [9] A. G. Dyachkov and V. V. Rykov. Bounds on the length of disjunctive codes. *Problemy Peredachi Informatsii*, 18(3): 7–13, 1982 (in Russian).
- [10] J. Fakcharoenphol and B. Laekhanukit. An $O(\log^2 k)$ -approximation algorithm for the k -vertex connected subgraph problem. In *Proceedings of ACM Symposium on Theory of Computing (STOC)*, 2008.
- [11] L. Fleischer. A 2-Approximation for Minimum Cost $\{0, 1, 2\}$ Vertex Connectivity. *IPCO*, pp. 115-129, 2001.
- [12] L. Fleischer, K. Jain, and D. P. Williamson. An Iterative Rounding 2-Approximation Algorithm for the Element Connectivity Problem. In *Proc. of the IEEE Symp. on Foundations of Computer Science (FOCS)*, pp. 339-347, 2001.
- [13] L. Fleischer, K. Jain, and D. P. Williamson. Iterative rounding 2-approximation algorithms for minimum cost vertex connectivity problems. *Journal of Computer and System Sciences*, 72(5):838–867, 2006.
- [14] A. Frank and T. Jordan. Minimal edge-coverings of pairs of sets. *Journal of Combinatorial Theory, Series B*, 65(1):73–110, 1995.
- [15] Z. Füredi. On r -Cover-free Families. *J. Comb. Theory, Ser. A* 73(1):172–173, 1996.
- [16] K. Jain. Factor 2 approximation algorithm for the generalized steiner network problem. In *Proceedings of the thirty-ninth annual IEEE Foundations of Computer Science (FOCS)*, pages 448–457, 1998.
- [17] K. Jain, I. Mandoiu, V.V. Vazirani and D.P. Williamson. A primal-dual schema based approximation algorithm for the element connectivity problem. *J. Algorithms* 45(1), pp. 1-15.
- [18] G. Kortsarz, R. Krauthgamer, and J. R. Lee. Hardness of approximation for vertex-connectivity network design problems. *SIAM J. of Computing*, 33(3):704–720, 2004.
- [19] G. Kortsarz and Z. Nutov. Approximating node connectivity problems via set covers. *Algorithmica*, 37(2):75–92, 2003.
- [20] G. Kortsarz and Z. Nutov. Approximating k -node connected subgraphs via critical graphs. *SIAM J. of Computing*, 35(1):247–257, 2005.
- [21] R. Kumar, S. Rajagopalan, and A. Sahai. Coding Constructions for Blacklisting Problems without Computational Assumptions. In *CRYPTO*, pp. 609–623, 1999.
- [22] Y. Lando and Z. Nutov. Inapproximability of Survivable Networks APPROX 2008, *Lecture Notes in Computer Science*, 5171, pp. 146-152.
- [23] Z. Nutov. A note on Rooted Survivable Networks. Manuscript, 2008. <http://www.openu.ac.il/home/nutov/R-SND.pdf>.
- [24] Z. Nutov. Approximating minimum cost connectivity problems via uncrossable bifamilies and spider-cover decompositions. In *Proc. of the IEEE Symposium on Foundations of Computer Science (FOCS)*, 2009.
- [25] Z. Nutov. An almost $O(\log k)$ -approximation for k -connected subgraphs. In *Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 912-921, 2009.
- [26] R. Ravi and D. P. Williamson. An approximation algorithm for minimum-cost vertex-connectivity problems. *Algorithmica*, 18(1):21–43, 1997.
- [27] D. R. Stinson, R. Wei, and L. Zhu. Some New Bounds for Cover-Free Families. *J. Comb. Theory, Ser. A*, 90(1):224–234, 2000.