Approximation Algorithms and Hardness of the k-Route Cut Problem

Julia Chuzhoy[*]   Yury Makarychev[†]   Aravindan Vijayaraghavan[‡]   Yuan Zhou[§]

November 26, 2011

Abstract

We study the k-route cut problem: given an undirected edge-weighted graph $G = (V, E)$, a collection $\{(s_1, t_1), (s_2, t_2), \ldots, (s_r, t_r)\}$ of source-sink pairs, and an integer connectivity requirement k, the goal is to find a minimum-weight subset E' of edges to remove, such that the connectivity of every pair (s_i, t_i) falls below k. Specifically, in the edge-connectivity version, EC-kRC, the requirement is that there are at most $(k - 1)$ edge-disjoint paths connecting s_i to t_i in $G \setminus E'$, while in the vertex-connectivity version, VC-kRC, the same requirement is for vertex-disjoint paths. Prior to our work, poly-logarithmic approximation algorithms have been known for the special case where $k \leq 3$, but no nontrivial approximation algorithms were known for any value $k > 3$, except in the single-source setting. We show an $O(k \log^{3/2} r)$-approximation algorithm for EC-kRC with uniform edge weights, and several polylogarithmic bi-criteria approximation algorithms for EC-kRC and VC-kRC, where the connectivity requirement k is violated by a constant factor. We complement these upper bounds by proving that VC-kRC is hard to approximate to within a factor of k^{ϵ} for some fixed $\epsilon > 0$.

We then turn to study a simpler version of VC-kRC, where only one source-sink pair is present. We give a simple bi-criteria approximation algorithm for this case, and show evidence that even this restricted version of the problem may be hard to approximate. For example, we prove that the single source-sink pair version of VC-kRC has no constant-factor approximation, assuming Feige's Random κ-AND assumption.

1 Introduction

Multi-commodity flows and cuts are among the most extensively studied graph optimization problems. Due to their rich connections to many combinatorial optimization problems, algorithms for various versions of flow and cut problems are a powerful and a widely used toolkit. One of the central problems in this area is *minimum multicut*: given an n-vertex graph $G = (V, E)$ with non-negative weights w_e on edges $e \in E$ and a collection $\{(s_1, t_1), (s_2, t_2), \ldots, (s_r, t_r)\}$ of source-sink pairs, find a minimum-weight subset E' of edges to delete, so that each pair (s_i, t_i) is disconnected in the resulting graph $G \setminus E'$. The dual to minimum multicut is the *maximum multi-commodity flow* problem, where the goal is to find a maximum flow between the pairs (s_i, t_i), with the restriction that each edge e carries at most w_e flow units. It is easy to see that minimum multicut can be viewed as revealing a bottleneck in the routing capacity of G, as the value of any multi-commodity flow cannot exceed the value of the minimum multicut in G. A fundamental result, due to Leighton and Rao [LR99] and Garg, Vazirani and Yannakakis [GVY95] shows that the value of minimum multicut is within an $O(\log r)$ factor of that of maximum multicommodity flow in any graph, where r is the number of the source-sink pairs. This result can be seen as an extension of the famous min-cut max-flow theorem to the multicommodity setting, and it also gives an efficient $O(\log r)$-approximation algorithm for minimum multicut — the best currently known approximation guarantee for it.

In this paper we study a natural generalization of minimum multicut - the *minimum k-route cut* problem. In this problem, the input again consists of an n-vertex graph $G = (V, E)$ with non-negative weights w_e on edges $e \in E$, and a collection $\{(s_1, t_1), (s_2, t_2), \ldots, (s_r, t_r)\}$ of r source-sink pairs. Additionally, we are given an integral connectivity threshold $k > 0$. The goal is to find a minimum-weight subset $E' \subseteq E$ of edges to delete, such that the connectivity of each pair (s_i, t_i) falls below k in the resulting graph $G \setminus E'$. We study two versions of this problem: in the edge-connectivity version (EC-kRC), the requirement is

that for each $1 \leq i \leq r$, the number of **edge-disjoint** paths connecting s_i to t_i in graph $G \setminus E'$ is less than k. In the vertex-connectivity version (VC-kRC), the requirement is that the number of **vertex-disjoint** paths connecting s_i to t_i is less than k. It is not hard to see that VC-kRC captures EC-kRC as a special case (see the full version of the paper for details), and hence is more general. It is also easy to see that minimum multicut is a special case of both EC-kRC and VC-kRC, with the connectivity requirement $k = 1$. We also consider a special case of EC-kRC, where all edges have unit weight, and we refer to it as the *uniform* EC-kRC. We note that for VC-kRC, the uniform and the non-uniform edge-weight versions are equivalent up to a small loss in the approximation factor, and so we do not distinguish between them.

The primary motivation for studying k-route cuts comes from multi-commodity flows in fault tolerant settings, where the resilience to edge and node failures is important. An elementary k-route flow between a pair s and t of vertices is a set of k disjoint paths connecting s to t. A k-route (st)-flow is just a combination of such elementary k-route flows, where each elementary flow is assigned some fractional value. This is a natural generalization of the standard (st)–flows, which ensures that the flow is resilient to the failure of up to $(k-1)$ edges or vertices. Multi-route flows were first introduced by Kishimoto [Kis96], and have since been studied in the context of communication networks [BCSK07, BCK03, ACKN07]. In a series of papers, Kishimoto [Kis96], Kishimoto and Takeuchi [KT93] and Aggarwal and Orlin [AO02] have developed a number of efficient algorithms for computing maximum multi-route flows. As in the case of standard flows, we can extend k-route (st)-flows to the multi-commodity setting, where the goal is to maximize the total k-route flow among all source-destination pairs. It is easy to see that the minimum k-route cut is a natural upper bound on the maximum k-route flow – just like minimum multicut upper-bounds the value of the maximum multi-commodity flow. Hence, as in the case with the standard multicut, multi-route cuts can be seen as revealing the network bottleneck, and so the minimum k-route cut in a graph captures the robustness of real-life computer and transportation networks.

The first approximation algorithm for the EC-kRC problem, due to Chekuri and Khanna [CK08], achieved a factor $O(\log^2 n \log r)$-approximation for the special case where $k = 2$, by rounding a Linear Programming relaxation. This was improved by Barman and Chawla [BC10] to give an $O(\log^2 r)$-approximation algorithm for the same version, by generalizing the region-growing LP-rounding scheme of [LR99, GVY95]. They note that it seems unlikely that their algorithm can be extended to handle higher values of k using similar techniques. Very recently, Kolman and Scheideler [KS11] obtained a $O(\log^3 r)$ approximation to EC-3RC ($k = 3$ case) from the linear program of [BC10] by using a multi-level ball growing rounding. To the best of our knowledge, no approximation algorithms with sub-polynomial (in n) guarantees are known for any variant of the problem, for any value $k > 3$, except in the single-source setting that we discuss later. Our first result is an $O(k \log^{1.5} r)$-approximation algorithm for the uniform version of EC-kRC.

Since the problem appears to be computationally difficult, it is natural to turn to bi-criteria approximation, by slightly relaxing the connectivity requirement. Given parameters $\alpha, \beta > 1$, we say that an algorithm is an (α, β)-bi-criteria approximation for EC-kRC (or VC-kRC), if it is guaranteed to produce a valid k'-route cut of weight at most $\beta \cdot \mathsf{OPT}$, where $k' \leq \alpha k$, and OPT is the value of the optimal k-route cut. Indeed, we can do much better in the bi-criteria setting: we obtain a $(1 + \delta, O(\frac{1}{\delta} \log^{1.5} r))$-bi-criteria approximation for any constant $0 < \delta < 1$, for the uniform EC-kRC problem (notice that the factors do not depend on k). When edge weights are arbitrary, we obtain a $(2, \tilde{O}(\log^{2.5} r))$-bi-criteria approximation in $n^{O(k)}$ time, and an $(O(\log r), O(\log^3 r))$-bi-criteria approximation in time polynomial in n and k. We also show an $O(\log^{1.5} r)$-approximation for the special case where $k = 2$, thus slightly improving the result of [BC10]. The previously known upper bounds and our results for EC-kRC are summarized in Table 1.

	Previous results	Current paper
$k = 2$	$O(\log^2 r)$ [BC10]	$O(\log^{1.5} r)$
$k = 3$	$O(\log^3 r)$ [KS11]	
arbitrary k, uniform	-	$O(k \log^{1.5} r)$, $(1 + \delta, O\left(\frac{1}{\delta} \log^{1.5} r\right))$ for any constant $0 < \delta < 1$
arbitrary k, general	-	$(2, O(\log^{2.5} r \log \log r))$ in time $n^{O(k)}$; $(O(\log r), O(\log^3 r))$ in poly(n)-time

Table 1: Upper bounds for EC-kRC. Running time is polynomial in n and k unless stated otherwise.

We note that on the inapproximability side, it is easy to show that for any value of k, EC-kRC is at least as hard as minimum multicut, up to small

constant factors[1]. While multicut is known to be hard to approximate up to any constant factor assuming the Unique Games Conjecture [KV05, CKK$^+$06], it is only known to be NP-hard to approximate to within a small constant factor [DJP$^+$94]. In fact one of the motivations for studying k-route cuts is that inapproximability results may yield insights into approximation hardness of multicut.

We now turn to the more general VC-kRC problem. The $O(\log^2 n \log r)$-approximation of [CK08], and the $O(\log^2 r)$-approximation of [BC10] for 2-route cuts extend to the vertex-connectivity version as well, as does our $O(\log^{1.5} r)$-approximation algorithm. Prior to our work, no non-trivial approximation algorithms were known for any higher values of k. In this paper, we show a $\left(2, \tilde{O}(kd\log^{2.5} r)\right)$-bi-criteria approximation algorithm for VC-kRC, with running time $n^{O(k)}$, where d is the maximum number of demand pairs in which any terminal participates. We note that, as in the case of EC-kRC, for any value of k, VC-kRC is at least as hard to approximate as minimum multicut (up to small constant factors), and to the best of our knowledge, no other inapproximability results have been known for this problem. We show that VC-kRC is hard to approximate up to any factor better than $\Omega(k^\epsilon)$, for some constant $\epsilon > 0$. Our results for VC-kRC are summarized in Table 2.

	Previous results	Current paper
$k = 2$	$O(\log^2 r)$ [BC10]	$O(\log^{1.5} r)$
arbitrary k	APX-hard [DJP$^+$94] no constant factor approximation under UGC [KV05, CKK$^+$06]	$\left(2, O(dk\log^{2.5} r \log\log r)\right)$-approximation algorithm, running time $n^{O(k)}$, where d is the maximum number of demand pairs in which any terminal participates $\Omega(k^\epsilon)$-hardness for some constant $\epsilon > 0$

Table 2: Results for VC-kRC.

In order to better understand the multi-route cut problem computationally, it is instructive to consider a simpler special case, where we are only given a single source-sink pair (s, t). We refer to this special case of VC-kRC and EC-kRC as (st)-VC-kRC and (st)-EC-kRC, respectively. As in the general case, it is easy to see that (st)-EC-kRC can be cast as a special case of (st)-VC-kRC. When the connectivity requirement k is constant, both

problems can be solved efficiently as follows: guess a set E' of $(k - 1)$ edges, and compute the minimum edge (st)–cut in graph $G \setminus E'$. The algorithm for (st)-VC-kRC is similar except that we guess a set V' of $(k - 1)$ vertices, and compute the minimum edge (st)–cut in graph $G \setminus V'$. However, for larger values of k, only a $2(k-1)$-approximation is known for (st)-EC-kRC, for the special case where the edge weights are uniform, due to Bruhn et al [BČH$^+$08][2]. Barman and Chawla [BC10] show that a generalization of (st)-EC-kRC where edges are allowed to have capacities is NP-hard. As no good approximation guarantees are known for the problem, it is natural to turn to bi-criteria approximation. For general values of k, Barman and Chawla [BC10] have shown a $(4, 4)$-bi-criteria approximation algorithm for (st)-EC-kRC, and a $(2, 2)$-bi-criteria approximation for uniform (st)-EC-kRC. In fact all these algorithms extend to a single-source multiple-sink scenario, except that the factor $(4, 4)$-approximation requires that the number of terminals is constant. In this paper we focus on the more general node-connectivity version of the problem. We start by showing a simple factor $(k + 1)$-approximation algorithm for (st)-VC-kRC, and a factor $\left(1+\frac{1}{c}, 1+c\right)$-bi-criteria approximation for any constant c. We then complement these upper bounds by providing evidence that the problem is hard to approximate. Specifically, we show that for any constant C, there is no $(1 + \gamma, C)$-bi-criteria approximation for (st)-VC-kRC, assuming Feige's Random κ-AND Hypothesis, where γ is some small constant depending on C. We also show that a factor ρ approximation algorithm for (st)-VC-kRC would lead to a factor $2\rho^2$-approximation for the Densest κ-Subgraph problem. These inapproximability results are inspired by the recent work of Arora et al. [AAM$^+$11], who have ruled out a constant factor approximation for Densest κ-Subgraph assuming Feige's Random κ-AND hypothesis.

Recall that the Densest κ-Subgraph problem takes as input a graph $G(V, E)$ on n vertices and a parameter κ, and asks for a subgraph of G on at most κ vertices containing the maximum number of edges. While it is a fundamental graph optimization problem, there is a huge gap between the best known approximation algorithm and the known inapproximability results. The current best approximation algorithm due to [BCC$^+$10] gives an $O(n^{1/4+\epsilon})$-factor approximation algorithm which runs in time $n^{O(1/\epsilon)}$ for any constant $\epsilon > 0$. On the inapproximability side, Feige [Fei02] initially showed a small constant factor inapproximability using the random 3-SAT assumption, and later

---

[1]A simple reduction replaces every vertex v of the multicut instance by a set S_v of M vertices, where $M \gg k$, and every edge (u, v) by a set of M^2 edges connecting every vertex of S_v to every vertex of S_u.

[2]This result also extends to the single-source multiple-sinks setting.

Khot [Kho04] used quasi-random PCPs to rule out a PTAS, assuming $\mathsf{NP} \not\subseteq \bigcap_{\epsilon>0} \mathsf{BPTIME}(2^{n^\epsilon})$. Raghavendra and Steurer [RS10] and Alon et al. [AAM+11] ruled out constant factor approximation algorithms for Densest κ-Subgraph under other less standard complexity assumptions. The Densest κ-Subgraph problem can also be generalized to λ-uniform hypergraphs, where the goal is again to find a subset of κ vertices containing maximum possible number of hyperedges. We show that for any constant $\lambda \geq 2$, a factor ρ approximation algorithm for (st)-VC-kRC would lead to a factor $(2\rho^\lambda)$-approximation for the λ-uniform Densest κ-subgraph. We note that Applebaum [App11] has recently shown that for $\lambda \geq 3$, the λ-uniform Densest κ-subgraph problem is hard to approximate to within n^ϵ-factor for some constant $\epsilon > 0$ assuming the existence of a certain family of one-way functions.

All our inapproximability results for (st)-VC-kRC are proved using a "proxy" problem, Small Set Vertex Expansion (SSVE). In this problem, we are given a bipartite graph $G = (U, V, E)$ and a parameter $0 \leq \alpha \leq 1$. The goal is to find a subset $S \subseteq U$ of $\alpha \cdot |U|$ vertices, while minimizing the number of its neighbors $|\Gamma(S)|$. We show an approximation-preserving reduction from SSVE to (st)-VC-kRC, and then prove inapproximability results for the SSVE problem. In particular, we show that approximating SSVE is almost as hard as approximating Densest κ-subgraph problem (that is, if there is a ρ approximation algorithm for SSVE then there is a $(2\rho^2)$-approximation algorithm for the Densest κ-subgraph problem). This result suggests that although the SSVE problem looks similar to the Small Set Expansion (SSE) problem [RS10], it might be much harder than SSE. On the other hand, the SSVE problem is of independent interest – besides its application to the (st)-VC-kRC problem, Applebaum et al. [ABW10] used a "planted" version of SSVE as a hardness assumption to construct a public key encryption scheme.

Other Related Work Another version of the EC-kRC problem that has received a significant amount of attention recently is the single-source setting. In this setting we are given a single source s and a set T of r terminals. The source-sink pairs are then set to be $\{(s, t)\}_{t \in T}$. Bruhn et al. [BČH+08] have shown a factor $2(k - 1)$-approximation for the uniform version of this problem, and Barman and Chawla [BC10] have shown a factor $(6, O(\sqrt{r} \ln r))$-bi-criteria approximation for the general version, a factor $(4, 4)$-bi-criteria approximation for the general version where r is a constant, and a factor $(2, 4)$-approximation for the uniform version and arbitrary r.

The (st)-EC-kRC and (st)-VC-kRC problems cap-

ture two natural budgetted cut minimization problems. The first is the Minimum Unbalanced cut problem [HKPS05], in which we are given a graph G with a source vertex s and a budget B. The goal is to find a cut $(S, \bar{S})$ in G with $s \in S$ and $|E(S, \bar{S})| \leq B$, while minimizing $|S|$. Hayrapetyan et al. [HKPS05] obtain a $(1 + 1/\lambda, \lambda)$-bi-criteria approximation algorithm for any $\lambda > 1$, by rounding a Lagrangean relaxation for the problem. Given an instance $G = (V, E)$ of the Minimum Unbalanced cut problem, we can transform it into an instance of (st)-EC-kRC, by setting the weights of all edges in E to ∞, adding a sink t, that connects to every vertex in V with a unit-weight edge, and setting $k = B$. The other problem is the Minimum k-size (st)-cut problem, where we are given a graph $G = (V, E)$ with a special source vertex s and a parameter k, and the goal is to find a cut $(S, \bar{S})$ in G with $s \in S$ and $|S| \leq k$, minimizing the size of the cut $|E(S, \bar{S})|$. Li and Zhang [LZ10] give an $O(\log n)$-approximation to this problem using Räcke's graph decomposition [Räc08]. This problem can be reduced to (st)-EC-kRC by assigning unit weights to the edges of E, and adding a sink t with infinity-weight edges (v, t) for each $v \in V$; the parameter k remains unchanged.

Our results and techniques The following two theorems summarize our results for the EC-kRC problem.

THEOREM 1.1. *There is an efficient $O(k \log^{1.5} r)$-approximation algorithm, and a $\left(1 + \delta, O\left(\frac{1}{\delta} \log^{1.5} r\right)\right)$-bi-criteria approximation algorithm for any constant $\delta \in (0, 1)$, for the uniform EC-kRC problem.*

THEOREM 1.2. *There is a $\left(2, O(\log^{2.5} r \log \log r)\right)$-bi-criteria approximation algorithm with running time $n^{O(k)}$ and an $\left(O(\log r), O(\log^3 r)\right)$-bi-criteria approximation algorithm with running time $\mathrm{poly}(n)$ for the EC-kRC problem.*

We now proceed to discuss our techniques. Our algorithms are based on a simple iterative approach: find a "sparse" cut that separates some demand pairs, remove all cut edges except for the $(k-1)$ most expensive ones from the graph, also remove all demand pairs that are no longer k-connected, and then recursively solve the obtained instance. The main challenge in this approach is to ensure that the cost of the removed edges is bounded by the cost of the optimal solution. In fact, in the first step of the algorithm, we use a modified notion of sparsity — we use the k-route sparsity of a cut, which is the cost of all but $(k-1)$ most expensive edges of the cut divided by the number of separated terminals (see below for formal definitions). This is necessary since

the standard sparsest cut can be prohibitively expensive; its cost cannot be bounded in terms of the cost of the optimal solution. We prove however that the cost of the k-route sparsest cut can be bounded in terms of the cost of the optimal solution and thus obtain guarantees on the performance of our algorithms. This is the most technically challenging part of the analysis of our algorithms.

We extend our bi-criteria approximation for EC-kRC to the more general VC-kRC problem in the following theorem.

THEOREM 1.3. *There is a $\left(2, O(dk \log^{2.5} r \log \log r)\right)$-bi-criteria approximation algorithm for VC-kRC with running time $n^{O(k)}$, where d is the maximum number of demand pairs in which any terminal participates.*

We also prove the following hardness of approximation result for VC-kRC, whose proof uses ideas similar to those used by Kortsarz et al. [KKL04] and Chakraborty et al. [CCK08] to prove hardness of vertex-connectivity network design:

THEOREM 1.4. *There are constants $0 < \epsilon < 1$, $k_0 > 1$, such that for any constant η, for any $k = O\left(2^{(\log n)^{1-\eta}}\right)$, where $k > k_0$, there is no k^ϵ-approximation algorithm for VC-kRC, under the assumption that $\mathsf{P} \neq \mathsf{NP}$ for constant k, and under the assumption that $\mathsf{NP} \not\subseteq \mathsf{DTIME}(n^{\mathrm{poly} \log n})$ for super-constant k.*

Finally, for the special case of $k = 2$, we obtain a slightly improved approximation algorithm:

THEOREM 1.5. *There is an efficient factor $O(\log^{1.5} r)$-approximation algorithm for both VC-kRC and EC-kRC, when $k = 2$.*

We now turn to the single (st)–pair version of the problems. We start with a simple approximation algorithm, summarized in the next theorem.

THEOREM 1.6. *There is an efficient factor $(k + 1)$-approximation algorithm, and for every constant $c > 0$, there is an efficient $\left(1 + \frac{1}{c}, 1 + c\right)$-bi-criteria approximation algorithm for both (st)-VC-kRC and (st)-EC-kRC.*

We then proceed to show inapproximability results for the single (st)–pair version of the problem. Our first inapproximability result uses Feige's random κ-AND assumption [Fei02]. Given parameters n and Δ, a random κ-AND instance is defined to be a κ-AND formula on n variables and $m = \Delta n$ clauses, where each clause chooses κ literals uniformly at random from the set of $2n$ available literals. We say that a formula Φ is α-satisfiable iff there is an assignment to the variables that satisfies an α-fraction of the clauses. Notice that a random assignment satisfies a $1/2^\kappa$-fraction of the clauses in expectation, and we expect that this is a typical number of simultaneously satisfiable clauses for a random κ-AND formula. We next state the Random κ-AND conjecture of Feige [Fei02] and our inapproximability result for (st)-VC-kRC.

HYPOTHESIS 1.1.
(Random κ-AND assumption: Hypothesis 3 in [FEI02]). *For some constant $c_0 > 0$, for every κ, there is a value of Δ_0, such that for every $\Delta > \Delta_0$, there is no polynomial time algorithm that for random κ-AND formulas with n variables and $m = \Delta n$ clauses, outputs 'typical' with probability $1/2$, but never outputs 'typical' on instances with $m/2^{c_0 \sqrt{\kappa}}$ simultaneously satisfiable clauses.*

THEOREM 1.7. *For every constant $C > 0$, there exists a small constant $0 < \gamma < 1$ which depends only on C, such that assuming Hypothesis 1.1, there is no polynomial time algorithm which obtains a $(1 + \gamma, C)$-bi-criteria approximation for the (st)–VC-kRC problem.*

We also prove a slightly different inapproximability result based on the slightly weaker Random 3-SAT assumption of Feige. Given parameters n and Δ, a random 3-SAT formula on n variables and $m = \Delta n$ clauses is constructed as follows: each clause chooses 3 literals uniformly at random among all available literals. Notice that a random assignment satisfies a $7/8$-fraction of clauses in expectation. Below is Feige's 3-SAT assumption and our inapproximability result for (st)-VC-kRC.

HYPOTHESIS 1.2.
(Random 3-SAT assumption: Hypothesis 2 from [FEI02]). *For every fixed $\epsilon > 0$, for Δ a sufficiently large constant independent of n, there is no polynomial time algorithm that on a random 3CNF formula with n variables and $m = \Delta n$ clauses, outputs 'typical' with probability at least $1/2$, but never outputs 'typical' when the formula is $(1 - \epsilon)$-satisfiable (i.e. there is an assignment satisfying simultaneously $(1 - \epsilon)m$ clauses).*

THEOREM 1.8. *Assuming Hypothesis 1.2, for any constant $\epsilon > 0$, no polynomial-time algorithm achieves a $\left(\frac{25}{24} - \epsilon, 1.1 - \epsilon\right)$-bi-criteria approximation for (st)–VC-kRC.*

Finally, we show that an existence of a good approximation algorithm for (st)-VC-kRC would imply a good approximation for the λ-uniform Hypergraph Densest κ-subgraph problem. Recall that in the λ-uniform Hypergraph Densest κ-subgraph problem, we are given a graph $G(V, E)$ where E is the set of λ-uniform hyperedges, and a parameter κ. The goal is to find a subset $S \subseteq V(G)$ of κ vertices, maximizing the number of hyper-edges $e \subseteq S$. Notice that for $\lambda = 2$, this is the standard Densest κ-subgraph problem.

THEOREM 1.9. *For any constant $\lambda \geq 2$, and for any approximation factor ρ (that may depend on n), if there is an efficient factor ρ approximation algorithm for the (st)–VC-kRC problem, then there is an efficient factor $(2\rho^\lambda)$-approximation algorithm for the λ-uniform Hypergraph Densest κ-subgraph problem.*

We note that Theorem 1.9, combined with the recent result of [AAM+11] immediately implies superconstant inapproximability for (st)-VC-kRC, under Hypothesis 1.1. However, our proof of Theorem 1.7 is conceptually simpler, and also leads to a bi-criteria inapproximability.

Organization We present notation and definitions and prove some results that we use throughout the paper in Section 2. We study the uniform case of EC-kRC in Section 3, and the non-uniform case in Section 4. We describe our results for VC-kRC in Section 5. Then we present an algorithm for 2-route cuts in Section 6. We prove n^ϵ hardness of VC-kRC in Section 7. Finally, we study the single source-sink case in Section 8, where we present an approximation algorithm and prove several hardness results for the problem. Due to lack of space, some details are deferred to the full version of the paper, available from the authors' web pages.

2 Preliminaries

In all our problems, the input is an undirected n-vertex graph $G = (V, E)$ with non-negative weights $w(e)$ on edges $e \in E$ and a parameter k. Additionally, we are given a set $D = \{(s_1, t_1), \ldots, (s_r, t_r)\}$ of source-sink pairs, that we also refer to as demand pairs. We let $T \subseteq V$ be the subset of vertices that participate in any demand pairs, and we refer to the vertices in T as terminals. For every vertex $v \in V$, let D_v be the number of demand pairs in which v participates. Given a subset $S \subseteq V$ of vertices, let $D(S) = \sum_{v \in S} D_v$ be the total number of terminals in S, counting multiplicities. We also denote by $D(S, \bar{S})$ the number of demand pairs (s_i, t_i) with $s_i \in S$, $t_i \in \bar{S}$, or $s_i \in \bar{S}$ and $t_i \in S$. Given any subset $E' \subseteq E$ of edges, we denote by $w(E') = \sum_{e \in E'} w(e)$ its weight. Throughout the paper,

we denote by E^* the optimal solution to the given EC-kRC or VC-kRC problem instance, and by OPT its value.

One of the main ideas in our algorithms is to relate the value of the appropriately defined sparsest cut in graph G to the value of the optimal solution to the k-route cut problem. We now define the different variations of the sparsest cut problem that we use.

The Sparsest Cut Problem. Given any cut $(S, \bar{S})$ in graph G, its *uniform sparsity* is defined to be

$$\Phi(S) = \frac{w(E(S, \bar{S}))}{\min\left\{D(S), D(\bar{S})\right\}}.$$

The uniform sparsity $\Phi(G)$ of the graph G is the minimum sparsity of any cut in G,

$$\Phi(G) = \min_{\substack{S \subset V: \\ D(S), D(\bar{S}) > 0}} \{\Phi(S)\}.$$

We use the $O(\sqrt{\log r})$-approximation algorithm for the uniform sparsest cut problem due to Arora, Rao and Vazirani [ARV04]. Let $\mathcal{A}_{\mathrm{ARV}}$ denote this algorithm, and let $\alpha_{\mathrm{ARV}}(r) = O(\sqrt{\log r})$ denote its approximation factor. Given an edge-weighted graph G and a set D of r demand pairs, algorithm $\mathcal{A}_{\mathrm{ARV}}$ finds a subset $S \subseteq V$ of vertices with $\Phi(S) \leq \alpha_{\mathrm{ARV}}(r) \cdot \Phi(G)$.

Given any cut $(S, \bar{S})$ in graph G, its *non-uniform sparsity* is defined to be

$$\tilde{\Phi}(S) = \frac{w(E(S, \bar{S}))}{D(S, \bar{S})}.$$

The non-uniform sparsity $\tilde{\Phi}(G)$ of the graph G is:

$$\tilde{\Phi}(G) = \min_{\substack{S \subset V: \\ D(S, \bar{S}) > 0}} \left\{\tilde{\Phi}(S)\right\}.$$

We also use the $O(\sqrt{\log r} \cdot \log \log r)$-approximation algorithm for the non-uniform sparsest cut problem of Arora, Lee and Naor [ALN05]. Let $\mathcal{A}_{\mathrm{ALN}}$ denote this algorithm, and let $\alpha_{\mathrm{ALN}}(r) = O(\sqrt{\log r} \cdot \log \log r)$ denote its approximation factor. Given an edge-weighted graph G with a set D of r demand pairs, algorithm $\mathcal{A}_{\mathrm{ALN}}$ finds a subset $S \subseteq V$ of vertices with $\tilde{\Phi}(S) \leq \alpha_{\mathrm{ALN}}(r) \cdot \tilde{\Phi}(G)$.

We next generalize the notion of the sparsest cut to the multi-route setting. Given a subset $S \subseteq V$ of vertices, let F denote the set of $(k - 1)$ most expensive edges of $E(S, \bar{S})$, breaking ties arbitrarily, and we refer to F as the set of *free edges* for cut $(S, \bar{S})$. We then define $w^{(k)}(S, \bar{S}) = \sum_{e \in E(S, \bar{S}) \setminus F} w_e$.

The *uniform k-route sparsity* of set S is defined to be:

$$\Phi^{(k)}(S) = \frac{w^{(k)}(S, \bar{S})}{\min\left\{D(S), D(\bar{S})\right\}},$$

and the uniform k-route sparsity of the graph G is:

$$\Phi^{(k)}(G) = \min_{\substack{S \subset V: \\ D(S), D(\bar{S}) > 0}} \left\{ \Phi^{(k)}(S) \right\}.$$

Similarly, the *non-uniform k-route sparsity* of S is:

$$\tilde{\Phi}^{(k)}(S) = \frac{w^{(k)}(S, \bar{S})}{D(S, \bar{S})},$$

and the non-uniform k-route sparsity of the graph G is:

$$\tilde{\Phi}^{(k)}(G) = \min_{\substack{S \subset V: \\ D(S, \bar{S}) > 0}} \left\{ \tilde{\Phi}^{(k)}(S) \right\}.$$

Note that $\Phi^{(1)}(G) = \Phi(G)$ and $\tilde{\Phi}^{(1)}(G) = \tilde{\Phi}(G)$ are the standard uniform and non-uniform sparsest cut values, respectively. We now show that there is an efficient algorithm to find an approximate k-route sparsest cut when k is a constant.

THEOREM 2.1. *There is an algorithm that, given a graph $G = (V, E)$ with r source-sink pairs and an integer k, computes in time $n^{O(k)}$ a cut $S \subseteq V$, with $\Phi^{(k)}(S) \leq \alpha_{\mathrm{ARV}}(r) \cdot \Phi^{(k)}(G)$. Similarly, there is an algorithm that computes in time $n^{O(k)}$ a cut S, with $\tilde{\Phi}^{(k)}(S) \leq \alpha_{\mathrm{ALN}}(r) \cdot \tilde{\Phi}^{(k)}(G)$.*

Proof. We start with the uniform k-route sparsest cut. We go over all subsets $F \subseteq E$ of $k-1$ edges. For each such subset F, we compute the $\alpha_{\mathrm{ARV}}(r)$-approximate sparsest cut in the graph $G \setminus F$ using the algorithm $\mathcal{A}_{\mathrm{ARV}}$, and output the best cut over all such subsets F. The algorithm for the non-uniform sparsest k-route cut is similar, except that we use the algorithm $\mathcal{A}_{\mathrm{ALN}}$ for the non-uniform sparsest cut. $\qquad\square$

The above theorem works well for constant values of k. However, when k is super-constant, the running time of the algorithm is no longer polynomial. For such cases, we use a bi-criteria approximation algorithm for the k-route sparsest cut problem, summarized in the next theorem.

THEOREM 2.2. *There is an efficient algorithm that, given an edge-weighted graph $G = (V, E)$, an integer $k > 1$, and a set $D = \{(s_i, r_i)\}_{i=1}^r$ of r demand pairs, finds a cut $S \subseteq V$ with $\tilde{\Phi}^{(k')}(S) = O(\log r) \cdot \tilde{\Phi}^{(k)}(G)$, where $k' = Ck \log r$ for some absolute constant C.*

Proof. We use as a sub-routine the approximation algorithm of Englert et al. [EGK+10] for the ℓ-Multicut problem. In the ℓ-Multicut problem, we are given a

graph $G = (V, E)$ with weights on edges, a set D of r demand pairs, and an integer ℓ. The goal is to find a minimum-weight subset $E' \subseteq E$ of edges, such that at least ℓ of the demand pairs are disconnected in the graph $G \setminus E'$. Engert et al. [EGK+10] give an efficient $O(\log r)$-approximation algorithm for this problem. We denote their algorithm by $\mathcal{A}_{\mathrm{EGK+}}$, and the approximation factor it achieves by $\alpha_{\mathrm{EGK+}} = O(\log r)$.

Let $(S^*, \bar{S}^*)$ be the optimal non-uniform k-route cut in G, and let $F^* \subseteq E_G(S^*, \bar{S}^*)$ be the subset of the $(k-1)$ most expensive edges in this cut. Then $w(E(S^*, \bar{S}^*) \setminus F^*) = \tilde{\Phi}^{(k)}(G) \cdot D(S^*, \bar{S}^*)$. Let $W^* = w(E(S^*, \bar{S}^*) \setminus F^*)$ and let $r^* = D(S^*, \bar{S}^*)$.

Assume first that our algorithm is given the values of W^* and r^*. We define new edge weights as follows: for each edge $e \in E$, $\tilde{w}_e = \min\{w_e, W^*/(k-1)\}$. We use the algorithm $\mathcal{A}_{\mathrm{EGK+}}$ on the resulting instance of the ℓ-Multicut problem, with $\ell = r^*$. Let $(S, \bar{S})$ be the output of this algorithm, and let F be the set of $2\alpha_{\mathrm{EGK+}}(r)(k-1)$ most expensive edges of $E(S, \bar{S})$, with respect to the original weights w_e, breaking ties arbitrarily.

The output of our algorithm is the cut $(S, \bar{S})$. In order to complete the proof, it is enough to show that $w(E(S, \bar{S}) \setminus F) \leq O(\log r) \cdot \tilde{\Phi}^{(k)}(G) \cdot D(S, \bar{S})$.

Note that the value of the optimal solution to the ℓ-Multicut problem instance is at most

$$\tilde{w}(E(S^*, \bar{S}^*)) \leq \tilde{w}(E(S^*, \bar{S}^*) \setminus F^*) + |F^*| \cdot \frac{W^*}{k-1}$$
$$\leq w(E(S^*, \bar{S}^*) \setminus F^*) + W^* = 2W^*.$$

Therefore, $\tilde{w}(E(S, \bar{S})) \leq 2\alpha_{\mathrm{EGK+}}(r)W^*$. In particular, $E(S, \bar{S})$ may contain at most $2\alpha_{\mathrm{EGK+}}(r)(k-1)$ edges e with $\tilde{w}_e = W^*/(k-1)$, and so all such edges lie in F. For edges $e \notin F$, $\tilde{w}_e < W^*/(k-1)$ must hold, and therefore, $\tilde{w}_e = w_e$. We conclude that

$$w(E(S, \bar{S}) \setminus F) = \tilde{w}(E(S, \bar{S}) \setminus F) \leq \tilde{w}(E(S, \bar{S}))$$
$$\leq 2\alpha_{\mathrm{EGK+}}(r)W^* = 2\alpha_{\mathrm{EGK+}}(r)\tilde{\Phi}^{(k)}(G) \cdot r^*$$
$$\leq O(\log r)\tilde{\Phi}^{(k)}(G)D(S, \bar{S})$$

as required.

Of course, our algorithm does not know the values of W^* and r^*. Instead, we perform the procedure described above for all possible values of $r' \in \{1, \ldots, r\}$ and (say) all values of W' in $\{\tau w_e : e \in E, 1 \leq \tau \leq |E|\}$, and then output the best of the cuts found. One of the values of r' will be equal to r^*, and one the values of W' will be within a factor of 2 of W^*: if e is the most expensive edge in $E(S^*, \bar{S}^*) \setminus F^*$, and $\tau = \lceil W^*/w_e \rceil$, then $W^* \leq \tau w_e \leq (2\lfloor W^*/w_e \rfloor)w_e \leq 2W^*$. For these

values of r' and W', the algorithm will find a cut that satisfies the conditions of the lemma. □

Laminar Families of Minimum Cuts Our main tool in establishing the connection between the values of the k-route sparsest cut and the cost of the optimal solution to the k-route cut problem is the following theorem, which shows that there is a laminar family of minimum cuts disconnecting the source-sink pairs in the graph G.

LEMMA 2.1. *There is an efficient algorithm, that, given any edge-weighted graph $G = (V, E)$ with a set $D = \{(s_i, t_i)\}_{i=1}^{r}$ of r source-sink pairs, finds a laminar family $\mathcal{S} = \{S_1, \cdots, S_r\}$ of vertex subsets, such that for all $1 \leq i \leq r$:*

- *$(S_i, V \setminus S_i)$ is a minimum cut separating s_i from t_i in G, and*

- *$D(S_i) \leq r$ (so S_i contains at most half the terminals, counting multiplicities).*

Proof. We use a Gomory–Hu tree T_{GH} for the graph G. Recall that it is a weighted tree, whose vertex set is V. Let c_e denote the costs of the edges $e \in E(T_{GH})$. Tree T_{GH} has the following key property: for every pair $(u, v) \in V$ of vertices, the value of the minimum cut separating u from v in graph G equals the value of the minimum cut separating u from v in T_{GH}. Note that the latter cut contains only one edge – the minimum-cost edge on the unique path connecting u to v in the tree.

We start with a Gomory–Hu tree T_{GH} for the graph G. For each $1 \leq i \leq r$, let (L_i, R_i) be a minimum cut separating s_i from t_i in T_{GH}. If $D(L_i) < D(R_i)$, then we set $S_i = L_i$. If $D(R_i) < D(L_i)$, we set $S_i = R_i$. Otherwise, if $D(R_i) = D(L_i)$, we let S_i to be the side containing the vertex s_1. We use this tie-breaking rule that enforces consistency across different source-sink pairs later.

This finishes the definition of the family $\mathcal{S} = \{S_1, \ldots S_r\}$ of vertex subsets. It is immediate to see that for each $1 \leq i \leq r$, $(S_i, V \setminus S_i)$ is a minimum cut separating s_i from t_i in G, and that $D(S_i) \leq r$. It now only remains to show that $S_1, \ldots, S_r$ form a laminar family.

Assume for contradiction that for some $i \neq j$, $S_i \cap S_j \neq \emptyset$, but $S_i \setminus S_j \neq \emptyset$, and $S_j \setminus S_i \neq \emptyset$. Let e_i be the unique edge of T_{GH} lying in the cut $(S_i, V \setminus S_i)$ in tree T_{GH}, and let e_j be the unique edge of T_{GH} lying in the cut $(S_j, V \setminus S_j)$. Observe that $T_{GH} \setminus \{e_i, e_j\}$ consists of three non-empty connected components. Let C_1 denote

the component that is incident on both e_i and e_j, C_2 the component incident on e_j only, and C_3 the remaining component. We claim that $S_i = C_1 \cup C_2$. Otherwise, since edge e_i separates S_i from $V \setminus S_i$ in T_{GH}, $S_i = C_3$ must hold. But then either $S_j = C_2$ and so $S_i \cap S_j = \emptyset$, or $S_j = C_1 \cup C_3$ and then $S_i \subseteq S_j$, a contradiction. Therefore, $S_i = C_1 \cup C_2$ and $V \setminus S_i = C_3$. Similarly, $S_j = C_1 \cup C_3$ and $V \setminus S_j = C_2$.

From the definition of S_i, either $D(S_i) < D(V \setminus S_i)$, or $D(S_i) = D(V \setminus S_i)$ and $s_1 \in S_i$. Assume first that $D(S_i) < D(V \setminus S_i)$. Then $V \setminus S_j = C_2 \subseteq S_i$, and so $D(V \setminus S_j) < D(S_j)$, contradicting the definition of S_j. We reach a similar contradiction if $D(S_j) < D(V \setminus S_j)$. Therefore, $D(S_i) = D(V \setminus S_i)$ and $D(S_j) = D(V \setminus S_j)$ must hold. In other words, $D(V \setminus S_i) = D(C_3) = r$, and $D(V \setminus S_j) = D(C_2) = r$. Since C_2 and C_3 are disjoint, this means that $D(C_1) = 0$. But from the definitions of S_i and S_j, $s_1 \in S_i \cap S_j$ must hold, a contradiction. □

First Algorithmic Framework Most our algorithms belong to one of two simple algorithmic frameworks. The first framework uses a divide-and-conquer paradigm: We start with the graph $G = (V, E)$ and a set D of $r \geq 1$ demand pairs, and then find a cut $(S, \bar{S})$ in G, with $D(S), D(\bar{S}) \geq 1$. We then select a subset $E_0 \subseteq E(S, \bar{S})$ of edges to delete, and apply the algorithm recursively to the sub-instances induced by $G[S]$ and $G[\bar{S}]$. Here, the sub-instance induced by $G[S]$ consists of the graph $G[S]$ and the collection of the original demand pairs (s_i, t_i), with both $s_i, t_i \in S$. The sub-instance induced by $G[\bar{S}]$ is defined similarly. Let E_1 and E_2 be the solutions returned by the two recursive calls, respectively. The final solution is $E' = E_0 \cup E_1 \cup E_2$. The specific cut $(S, \bar{S})$, and the subsets $E_0 \subseteq E(S, \bar{S})$ of edges computed will differ from algorithm to algorithm, and we will need to select them in a way that ensures the feasibility of the final solution. However, the analysis of the solution cost is similar in all these algorithms, and is summarized in the following theorem.

THEOREM 2.3. *Let $\mathcal{A}$ be any algorithm in the above framework, and assume that the algorithm guarantees that $w(E_0) \leq \frac{\alpha \cdot \mathsf{OPT}}{r} \cdot \min \left\{ D(S), D(\bar{S}) \right\}$, for some factor α. Then $w(E') \leq 4\alpha \ln(1 + r) \cdot \mathsf{OPT}$.*

Proof. The proof is by induction on r. If $r = 1$ then $E' = E_0$, and the statement trivially holds. Assume now that the statement holds for instances with fewer than r demand pairs, for some $r > 1$. Consider the cut $(S, \bar{S})$ computed by the algorithm $\mathcal{A}$ on the current instance. Let a be the number of demand pairs (s_i, t_i) with $s_i, t_i \in S$, let b be the number of

demand pairs (s_i, t_i) with $s_i, t_i \in \bar{S}$, and assume w.l.o.g. that $a \geq b$. Then $D(S) \geq 2a$, $D(\bar{S}) \geq 2b$, and so $D(S) \leq 2r - D(\bar{S}) \leq 2(r - b)$ and $D(\bar{S}) \leq 2(r - a)$. Therefore,

$$w(E_0) \leq \frac{\alpha \cdot \mathsf{OPT}}{r} \cdot \min\left\{D(S), D(\bar{S})\right\} \leq \frac{2\alpha \cdot \mathsf{OPT}}{r}(r - a).$$

The optimal solutions to the EC-kRC instances on graphs $G[S]$ and $G[\bar{S}]$ have costs at most $w(E^* \cap E(S))$ and $w(E^* \cap E(\bar{S}))$, respectively. By the induction hypothesis, the total cost of solutions E_1 and E_2 on graphs $G[S]$ and $G[\bar{S}]$ is at most

$$4\alpha w(E^* \cap E(S))\ln(1 + a) + 4\alpha w(E^* \cap E(\bar{S}))\ln(1 + b)$$
$$\leq 4\alpha\Big(w(E^* \cap E(S)) + w(E^* \cap E(\bar{S}))\Big)\ln(1 + a)$$
$$\leq 4\alpha \cdot \mathsf{OPT} \cdot \ln(1 + a).$$

The total solution cost is then bounded by:

$$w(E') \leq 4\alpha \cdot \mathsf{OPT} \cdot \ln(1 + a) + \frac{2\alpha \cdot \mathsf{OPT}}{r}(r - a)$$
$$\leq 4\alpha \cdot \mathsf{OPT}\left(\ln(1 + a) + \frac{r - a}{2r}\right).$$

The theorem follows from the following inequality:

$$\ln(1 + a) + \frac{r - a}{2r} = \ln(1 + r) + \ln\left(\frac{1 + a}{1 + r}\right) + \frac{r - a}{2r}$$
$$\leq \ln(1 + r) - \frac{r - a}{1 + r} + \frac{r - a}{2r}$$
$$\leq \ln(1 + r),$$

where we have used the fact that $\ln\left(\frac{1+a}{1+r}\right) = \ln\left(1 - \frac{r-a}{1+r}\right) \leq -\frac{r-a}{1+r}$, since $\ln(1 + x) \leq x$ for all $x > -1$. $\square$

Second Algorithmic Framework The algorithmic framework presented above has some limitations. Specifically, we can only use it in scenarios where there is a cheap collection E' of edges (with cost roughly comparable to OPT), whose removal decomposes our instance G into two disjoint sub-instances, $G[S], G[\bar{S}]$, which can then be solved separately. This is the case for the uniform EC-kRC, and the non-uniform EC-kRC and VC-kRC when $k = 2$. For higher values of k in the non-uniform setting, such a decomposition may not exist. Instead, we use the following framework. Given a graph G and a set D of $r \geq 1$ demand pairs, we find a collection E_0 of edges to delete, together with a subset D_0 of demand pairs to remove, where $|D_0| \geq 1$. We then solve the problem recursively on the graph $G' = G \setminus E_0$, and the

set $D \setminus D_0$ of the remaining demand pairs. Let E_1 be the subset of edges returned by the recursive call. Then the solution computed by the algorithm is $E' = E_0 \cup E_1$. The specific subset E_0 of edges to remove and the subset D_0 of demands will again be computed differently by each algorithm, in a way ensures that the final solution is feasible. The analysis of the solution cost of such algorithms is summarized in the next theorem.

THEOREM 2.4. *Let $\mathcal{A}$ be any algorithm in the above framework, and assume that we are guaranteed that $w(E_0) \leq \alpha \cdot \mathsf{OPT} \cdot \frac{|D_0|}{r}$, for some factor α. Then $w(E') \leq 2\alpha \ln(1 + r) \cdot \mathsf{OPT}$.*

Proof. The proof is by induction on r. If $r = 1$ then $E' = E_0$, and the statement trivially holds. Assume now that the statement holds for instances with fewer than r demand pairs, for some $r > 1$. We prove the theorem for instances with r demand pairs. Let $a = |D_0|$. Then by the induction hypothesis, $w(E_1) \leq 2\alpha \cdot \mathsf{OPT} \cdot \ln(1 + r - a)$. Therefore,

$$w(E') \leq 2\alpha \cdot \mathsf{OPT} \cdot \ln(1 + r - a) + \alpha \cdot \mathsf{OPT} \cdot \frac{a}{r}$$
$$= 2\alpha \cdot \mathsf{OPT}\left(\ln(r + 1) + \ln\left(\frac{1 + r - a}{r + 1}\right) + \frac{a}{2r}\right)$$
$$\leq 2\alpha \cdot \mathsf{OPT}\left(\ln(r + 1) - \frac{a}{r + 1} + \frac{a}{2r}\right)$$
$$\leq 2\alpha \ln(r + 1)\mathsf{OPT}.$$

$\square$

3 Uniform EC-kRC

This section is dedicated to proving Theorem 1.1. We first show an $O(k \log^{1.5} r)$-approximation algorithm, and provide a bi-criteria algorithm later. Recall that we are given an unweighted graph $G = (V, E)$, a set $\{(s_i, t_i)\}_{i=1}^r$ of demand pairs, and an integer k. Our goal is to find a collection E' of $O(k \log^{3/2} r) \cdot \mathsf{OPT}$ edges, such that for each demand pair (s_i, t_i), there are at most $(k - 1)$ edge-disjoint paths in graph $G \setminus E'$ connecting them.

We assume w.l.o.g. that each source-sink pair (s_i, t_i) is k-edge connected in the current graph G. Our algorithm views the graph G as an instance of the uniform sparsest cut problem. We use the algorithm $\mathcal{A}_{\mathrm{ARV}}$ to find a partition $(S, \bar{S})$ of V with $\Phi(S) \leq \alpha_{\mathrm{ARV}}(r) \cdot \Phi(G)$, add the edges in $E(S, \bar{S})$ to the solution E', and delete the demand pairs (s_i, t_i) that are no longer k-edge connected from the list of source-sink pairs. Notice that each remaining source-sink pair must be contained either in S or in $\bar{S}$. We then recursively

solve the EC-kRC problem on the sub-instances induced by $G[S]$ and $G[\bar{S}]$. The algorithm is summarized in Figure 1.

Input: An unweighted graph $G = (V, E)$ with demand pairs $\{(s_i, t_i)\}_{1 \leq i \leq r}$
Output: A set E' of edges, such that each pair (s_i, t_i) has at most $(k-1)$ edge-disjoint paths connecting s_i to t_i in $G \setminus E'$.

1. If $r = 0$ return $E' = \emptyset$.

2. Find a partition $(S, \bar{S})$ of V using the algorithm $\mathcal{A}_{\mathrm{ARV}}$, with $\Phi(S) \leq \alpha_{\mathrm{ARV}}(r) \cdot \Phi(G)$.

3. Let $E_0 = E(S, \bar{S})$, $G' = G \setminus E_0$.

4. Remove all demand pairs (s_i, t_i) that are no longer k-edge connected in G'.

5. Solve the instances induced by $G[S]$ and $G[V \setminus S]$ recursively, to obtain solutions E_1' and E_2', respectively.

6. Return $E' = E_1' \cup E_2' \cup E_0$.

Figure 1: Approximation algorithm for uniform EC-kRC.

The heart of the analysis of the algorithm is the following theorem, that relates the value of the uniform sparsest cut in graph G to the value OPT of the optimal solution for EC-kRC.

THEOREM 3.1. *Suppose that we are given an unweighted graph $G = (V, E)$ with r source-sink pairs $\{(s_i, t_i)\}_{i=1}^r$, such that for each pair (s_i, t_i), there are at least k edge-disjoint paths connecting s_i to t_i in G, and let OPT be the cost of the optimal solution of EC-kRC on G. Then $\Phi(G) \leq \frac{2k}{r} \cdot$ OPT.*

Proof. Consider the graph $H = G \setminus E^*$. We use Lemma 2.1 with edge weights $w_e = 1$ on graph H to obtain the laminar family $\mathcal{S} = \{S_i\}_{i=1}^r$ of vertex subsets. Consider all maximal sets in the laminar family, that is, sets S_i that are not contained in other sets. Assume w.l.o.g. that these sets are $S_1, \ldots, S_q$, for some $q \leq r$. Then $\sum_{i=1}^q D(S_i) \geq r$ must hold. Note that for each i, $|E_H(S_i, V \setminus S_i)| \leq k - 1$ since s_i and t_i are not k-edge connected in H, and $(S_i, V \setminus S_i)$ is a minimum cut separating s_i from t_i in H. On the other hand, $|E_G(S_i, V \setminus S_i)| \geq k$ since s_i and t_i are k-edge connected

in G. Therefore,

$$
\begin{aligned}
|E_G(S_i, V \setminus S_i)| &= |E_H(S_i, V \setminus S_i)| \\
&\quad + |E_G(S_i, V \setminus S_i) \cap E^*| \\
&\leq (k-1) + |E_G(S_i, V \setminus S_i) \cap E^*| \\
(3.1) &\leq k|E_G(S_i, V \setminus S_i) \cap E^*|.
\end{aligned}
$$

Note that every edge $e \in E^*$ belongs to at most two cuts $E_G(S_i, V \setminus S_i)$ and $E_G(S_j, V \setminus S_j)$. Therefore,

$$
\begin{aligned}
\sum_{i=1}^q |E_G(S_i, V \setminus S_i)| &\leq \sum_{i=1}^q k|E_G(S_i, V \setminus S_i) \cap E^*| \\
&\leq 2k \cdot \mathsf{OPT}
\end{aligned}
$$

On the other hand,

$$
\begin{aligned}
\sum_{i=1}^q |E_G(S_i, V \setminus S_i)| &= \sum_{i=1}^q \Phi(S_i) \cdot D(S_i) \\
&\geq \sum_{i=1}^q \Phi(G) \cdot D(S_i) \quad \geq r \cdot \Phi(G).
\end{aligned}
$$

We conclude that $\Phi(G) \leq 2k \cdot \mathsf{OPT}/r$. $\qquad\square$

We now analyze the algorithm. Since the algorithm removes a demand pair (s_i, t_i) only when s_i and t_i are no longer k-edge connected, and terminates when all demand pairs are removed, the algorithm is guaranteed to find a feasible solution to the problem. In order to bound the solution cost, note that

$$
\begin{aligned}
|E_0| &= \Phi(S) \cdot \min\{D(S), D(\bar{S})\} \\
&\leq \alpha_{\mathrm{ARV}}(r) \cdot \Phi(G) \cdot \min\{D(S), D(\bar{S})\} \\
&\leq \frac{2k\alpha_{\mathrm{ARV}}(r)}{r} \cdot \mathsf{OPT} \cdot \min\{D(S), D(\bar{S})\}.
\end{aligned}
$$

We can now use Theorem 2.3 with $\alpha = 2k\alpha_{\mathrm{ARV}}(r)$ to conclude that $|E'| = O(k\alpha_{\mathrm{ARV}}(r) \log r)\mathsf{OPT} = O(k \log^{3/2} r)\mathsf{OPT}$.

Bi-criteria approximation algorithm We now slightly modify the algorithm from Figure 1, to obtain a $\left(1 + \delta, O(\frac{1}{\delta} \log^{1.5} r)\right)$-bi-criteria approximation algorithm for any constant $0 < \delta < 1$. The algorithm works exactly as before, except that it removes a demand pair (s_i, t_i) in step 4 iff s_i and t_i are no longer $(1 + \delta)k$ edge-connected. We also assume w.l.o.g. that in the original instance G, every demand pair (s_i, t_i) has at least $(1 + \delta)k$ edge-disjoint paths connecting s_i to t_i. As before, it is straightforward to verify that if E' is the final solution produced by the algorithm, then each demand pair (s_i, t_i) pair has fewer than $(1 + \delta)k$ edge-disjoint paths connecting them in $G \setminus E'$. In order to bound the solution cost, we prove the following analogue of Theorem 3.1.

THEOREM 3.2. *Suppose that we are given an un-weighted graph G with r demand pairs $\{(s_i, t_i)\}_{i=1}^r$, where for each pair (s_i, t_i), there are at least $(1+\delta)k$ edge-disjoint paths connecting s_i to t_i in G. Then $\Phi(G) \leq \frac{2\mathsf{OPT}}{r} \cdot (1 + 1/\delta)$.*

Proof. As before, we compute the laminar family of minimum cuts in graph $H = G \setminus E^*$, using Lemma 2.1, and we consider the collection of all maximal cuts in this family. Assume w.l.o.g. that it is $\{S_1, \ldots, S_q\}$, for $q \leq r$, and recall that $\sum_{i=1}^q D(S_i) \geq r$. As before, for each $1 \leq i \leq q$, $|E_G(S_i, V \setminus S_i)| \leq (k-1) + |E_G(S_i, V \setminus S_i) \cap E^*|$. Since $|E_G(S_i, V \setminus S_i)| \geq (1+\delta)k$, we get that $|E_G(S_i, V \setminus S_i) \cap E^*| \geq \delta k$, and so $(k-1) \leq |E_G(S_i, V \setminus S_i) \cap E^*|/\delta$. We get that:

$$
\begin{aligned}
|E_G(S_i, V \setminus S_i)| &\leq (k-1) + |E_G(S_i, V \setminus S_i) \cap E^*| \\
&\leq (1 + 1/\delta)|E_G(S_i, V \setminus S_i) \cap E^*|.
\end{aligned}
$$

On the other hand, $|E_G(S_i, V \setminus S_i)| \geq \Phi(G) \cdot D(S_i)$. Summing up over all $1 \leq i \leq q$, we get that:

$$
\begin{aligned}
2\mathsf{OPT} &\geq \sum_{i=1}^q |E_G(S_i, V \setminus S_i) \cap E^*| \\
&\geq \frac{\delta}{\delta+1} \sum_{i=1}^q |E_G(S_i, V \setminus S_i)| \\
&\geq \frac{\delta}{\delta+1} \Phi(G) \sum_{i=1}^q D(S_i) \geq \frac{\delta}{\delta+1} \Phi(G) \cdot r.
\end{aligned}
$$

We conclude that $\Phi(G) \leq \frac{2\mathsf{OPT}}{r}(1 + 1/\delta)$. $\qquad\square$

In order to bound the final solution cost, observe that

$$
\begin{aligned}
|E_0| &= \Phi(S) \cdot \min\left\{D(S), D(\bar{S})\right\} \\
&\leq \alpha_{\mathrm{ARV}}(r) \cdot \Phi(G) \cdot \min\left\{D(S), D(\bar{S})\right\} \\
&\leq \frac{2\mathsf{OPT}\alpha_{\mathrm{ARV}}(r)}{r} \cdot (1 + 1/\delta) \cdot \min\left\{D(S), D(\bar{S})\right\}.
\end{aligned}
$$

We now use Theorem 2.3 with $\alpha = 2\alpha_{\mathrm{ARV}}(r)(1 + 1/\delta)$ to conclude that $|E'| = O(\alpha_{\mathrm{ARV}}(r) \log r/\delta)\mathsf{OPT} = O(\log^{1.5} r/\delta)\mathsf{OPT}$, when $0 < \delta < 1$.

This concludes the proof of Theorem 1.1.

4 Non-uniform EC-kRC

In this section we prove Theorem 1.2. We start with a $(2, \tilde{O}(\log^{2.5} r))$-bi-criteria algorithm with running time $n^{O(k)}$, and we show an algorithm whose running time is polynomial in n and k later.

Abusing the notation, for each cut $(S, \bar{S})$ in graph G, we denote by $D(S, \bar{S})$ both the set of demand pairs (s_i, t_i) with $|\{s_i, t_i\} \cap S| = 1$, and the number of such pairs.

4.1 A $(2, \tilde{O}(\log^{5/2} r))$ bi-criteria approximation in time $n^{O(k)}$

We cannot employ the first algorithmic framework for EC-kRC on weighted graphs. A natural approach in using it would be to find an appropriately defined sparse cut $(S, \bar{S})$, remove all but $k-1$ most expensive edges of this cut, and then recursively solve the problem on instances $G[S]$ and $G[\bar{S}]$. Let E_0 be the subset of edges removed, and let $G' = G \setminus E_0$ be the remaining graph. This approach does not work because it is possible that a demand pair (s_i, t_i) with both $s_i, t_i \in S$ is connected by a path that visits $G[\bar{S}]$ in graph G'. So if we solve the problem recursively on $G[S]$ and $G[\bar{S}]$, then the combined solution is not necessarily a feasible solution to the problem. On the other hand, if, instead, we remove all or almost all edges of $E(S, \bar{S})$, then the resulting solution cost may be too high. Therefore, we employ the second algorithmic framework instead.

We assume w.l.o.g. that in the input graph G, each demand pair (s_i, t_i) has at least $(2k-1)$ edge-disjoint paths connecting them. Our algorithm, summarized in Figure 2, starts by finding an approximate non-uniform $(2k-1)$-route sparse cut $(S, \bar{S})$ in G, using Theorem 2.1. That is, $\tilde{\Phi}^{(2k-1)}(S) \leq \alpha_{\mathrm{ALN}}(r)\tilde{\Phi}^{(2k-1)}(G)$. Let F be the set of the $(2k-2)$ most expensive edges of $E(S, \bar{S})$, let $E_0 = E(S, \bar{S}) \setminus F$, and let $G' = G \setminus E_0$. We remove all demand pairs that are no longer $(2k-1)$ connected in G' (we denote the set of these demand pairs by D_0), and then recursively solve the resulting instance.

It is immediate to verify that the algorithm returns a feasible solution. The running time of the algorithm is dominated by computing an approximate $(2k-1)$-route sparsest cut, and is therefore bounded by $n^{O(k)}$. In order to bound the solution cost, we use the following lemma that relates the value of $\tilde{\Phi}^{(2k-1)}(G)$ to OPT.

THEOREM 4.1. *Suppose that we are given a graph $G(V, E)$ with edge weights w_e, and a set $D = \{(s_i, t_i)\}_{i \in [r]}$ of r demand pairs, where every pair (s_i, t_i) has at least $(2k-1)$ edge-disjoint paths connecting s_i to t_i in G. Then $\mathsf{OPT} = \Omega\left(\frac{r}{\log r}\right)\tilde{\Phi}^{(2k-1)}(G)$.*

Proof. Consider the graph $H = G \setminus E^*$. Let $\mathcal{S} = \{S_1, \ldots, S_r\}$ be the laminar family of minimum cuts in H, guaranteed by Lemma 2.1. Recall that for all $1 \leq i \leq r$, $|E_H(S_i, V \setminus S_i)| \leq k - 1$. We need the following lemma.

Input: A weighted graph $G(V, E)$ with a set $D = \{(s_i, t_i)\}_{1 \leq i \leq r}$ of demand pairs, and edge weights $\{w_e\}_{e \in E}$.
Output: A set E' of edges, such that each demand pair (s_i, t_i) is no longer $(2k-1)$-edge connected in $G \setminus E'$.

1. If $r = 0$ return $E' = \emptyset$.

2. Find an approximate non-uniform $(2k-1)$-route sparsest cut $(S, \bar{S})$ with $\tilde{\Phi}^{(2k-1)}(S) \leq \alpha_{\mathrm{ALN}}(r) \cdot \tilde{\Phi}^{(2k-1)}(G)$, using Theorem 2.1. Let F be the set of the $(2k-2)$ most expensive edges in $E(S, \bar{S})$, breaking ties arbitrarily.

3. Let $E_0 = E(S, \bar{S}) \setminus F$, $G' = G \setminus E_0$, and let D_0 be the set of all demand pairs that are no longer $(2k-1)$-connected in G'.

4. Recursively solve the problem on G' with the demand set $D \setminus D_0$, to obtain a solution E_1.

5. Return $E' = E_0 \cup E_1$.

Figure 2: A bi-criteria approximation algorithm for non-uniform EC-kRC in time $n^{O(k)}$.

LEMMA 4.1. *We can efficiently find a collection $\mathcal{P}$ of mutually disjoint vertex subsets, such that:*

- *For each $U \in \mathcal{P}$, $D(U) \leq r$;*

- *For each $U \in \mathcal{P}$, $|E_H(U, \overline{U})| \leq 2(k-1)$, and*

- $\sum_{U \in \mathcal{P}} D(U, \overline{U}) \geq \dfrac{r}{8 \log r}$.

Proof. We will define each set $U \in \mathcal{P}$ to be either some set $S \in \mathcal{S}$, or a difference of two sets, $U = S \setminus S'$, for $S, S' \in \mathcal{S}$. Since for each set $S \in \mathcal{S}$, $D(S) \leq r$, this will ensure the first condition. Since $|E_H(U, \overline{U})| \leq |E_H(S, \bar{S})| + |E_H(S', \bar{S}')| \leq 2(k-1)$, this will also ensure the second condition.

We now turn to define the sets $U \in \mathcal{P}$ so that the third condition is also satisfied. For simplicity, if collection $\mathcal{S}$ contains identical sets, we discard them, keeping at most one copy of each set in $\mathcal{S}$. Recall that for each set $S \in \mathcal{S}$, $D(S, \bar{S})$ is the set of all demand pairs (s_j, t_j) with $|\{s_j, t_j\} \cap S| = 1$. Let $D'(S, \bar{S})$ be the union of $D(S', \bar{S}')$ for all sets $S' \in \mathcal{S}$ where $S' \subset S$, and let $q(S) = |D(S, \bar{S}) \setminus D'(S, \bar{S})|$.

We partition the family $\mathcal{S}$ into subsets $\mathcal{S}_x$, for $1 \leq x \leq \lfloor \log_2 r \rfloor + 1$, as follows: Collection $\mathcal{S}_x$ contains all sets $S \in \mathcal{S}$ with $2^{x-1} \leq q(S) < 2^x$. Since

$\sum_{S \in \mathcal{S}} q(S) = r$, there is at least one index x^*, for which $\sum_{S \in \mathcal{S}_{x^*}} q(S) \geq \frac{r}{2 \log r}$. Fix any such index x^*.

Consider the decomposition forest $\mathcal{F}$ for the sets in $\mathcal{S}_{x^*}$. The nodes of the forest are the sets in $\mathcal{S}_{x^*}$, and for a pair $S, S' \in \mathcal{S}_{x^*}$, set S is the parent of S' iff $S' \subset S$, and there is no other set $S'' \in \mathcal{S}_{x^*}$ with $S' \subset S'' \subset S$. Let $\mathcal{S}' \subseteq \mathcal{S}_{x^*}$ be the collection of sets that have at most one child in this forest. We are now ready to define the collection $\mathcal{P}$ of vertex subsets. If $S \in \mathcal{S}'$ is a leaf in $\mathcal{F}$, then we add S to $\mathcal{P}$. Otherwise, if S is a non-leaf set in $\mathcal{S}'$, and S' is the unique child of S in $\mathcal{F}$, then we add $S \setminus S'$ to $\mathcal{P}$.

It now only remains to prove that $\sum_{U \in \mathcal{P}} D(U, \overline{U}) \geq \frac{r}{8 \log r}$. In order to do so, observe that $|\mathcal{S}'| \geq |\mathcal{S}_{x^*}|/2$, and recall that for each $S \in \mathcal{S}_{x^*}$, $2^{x^*-1} \leq q(S) < 2^{x^*}$. Therefore, $\sum_{U \in \mathcal{P}} D(U, \overline{U}) \geq \sum_{S \in \mathcal{S}'} q(S) \geq \frac{|\mathcal{S}_{x^*}|}{2} \cdot 2^{x^*-1} = |\mathcal{S}_{x^*}| \cdot 2^{x^*-2}$. On the other hand, $\sum_{S \in \mathcal{S}_{x^*}} q(S) \geq \frac{r}{2 \log r}$, and so $|\mathcal{S}_{x^*}| \geq \frac{r}{2^{x^*+1} \log r}$. We conclude that $\sum_{U \in \mathcal{P}} D(U) \geq \frac{r}{8 \log r}$. □

Let $\mathcal{P}$ be the set family computed by Lemma 4.1. Clearly, for each $U \in \mathcal{P}$,

$$(4.2) \qquad w^{(2k-1)}(U, \overline{U}) = \tilde{\Phi}^{(2k-1)}(U) \cdot D(U, \overline{U})$$
$$\geq \tilde{\Phi}^{(2k-1)}(G) \cdot D(U, \overline{U}).$$

On the other hand, since $|E_H(U, \overline{U})| \leq 2k - 2$, $w(E^* \cap E_G(U, \overline{U})) \geq w^{(2k-1)}(U, \overline{U})$ must hold. Therefore,

$$(4.3)$$
$$\sum_{U \in \mathcal{P}} w^{(2k-1)}(U, \overline{U}) \leq \sum_{U \in \mathcal{P}} w(E^* \cap E_G(U, \overline{U})) \leq 2\mathsf{OPT}.$$

Combining Equations (4.2) and (4.3), we get that:

$$2\mathsf{OPT} \geq \sum_{U \in \mathcal{P}} w^{(2k-1)}(U, \overline{U}) \geq \tilde{\Phi}^{(2k-1)}(G) \sum_{U \in \mathcal{P}} D(U, \overline{U})$$
$$\geq \tilde{\Phi}^{(2k-1)}(G) \cdot \frac{r}{8 \log r}$$

Therefore, $\mathsf{OPT} = \Omega\left(\frac{r}{\log r}\right) \tilde{\Phi}^{(2k-1)}(G)$. □

In order to bound the cost $w(E')$ of the solution, we note that $D(S, \bar{S}) \subseteq D_0$, and so

$$w(E_0) = w^{(2k-1)}(S, \bar{S}) = \tilde{\Phi}^{(2k-1)}(S) \cdot D(S, \bar{S})$$
$$\leq \alpha_{\mathrm{ALN}}(r) \cdot \tilde{\Phi}^{(2k-1)}(G) \cdot |D_0|$$
$$= O(\alpha_{\mathrm{ALN}}(r) \log r) \cdot \frac{|D_0|}{r} \cdot \mathsf{OPT}.$$

We can now use Theorem 2.4 with $\alpha = O(\alpha_{\text{ALN}}(r) \log r)$ to conclude that $w(E') = O(\alpha_{\text{ALN}}(r) \log^2 r) = O(\log^{2.5} r \log \log r)$.

4.2 A polynomial-time bi-criteria approximation algorithm

In this section, we extend the algorithm from Section 4.1 to handle higher values of k in polynomial time. Notice that the bottleneck in the algorithm from Section 4.1 is computing the approximate multi-route sparsest cut via Theorem 2.1, which is done in time $n^{O(k)}$. We use Theorem 2.2 instead, that gives an efficient algorithm for computing the k-route sparsest cut, albeit with somewhat weaker guarantees.

Our algorithm is identical to the algorithm in Figure 2, except for the following changes. First, in step 2, we use Theorem 2.2 to find an approximate $(2k - 1)$-route non-uniform sparsest cut S. That is, $\tilde{\Phi}^{(k')}(S) = O(\log r)\tilde{\Phi}^{(2k-1)}(G)$, where $k' = C(2k - 1) \log r$, and C is the constant from Theorem 2.2. Note that $\tilde{\Phi}^{(C \log r(2k-1))}(G) \leq \tilde{\Phi}^{(2k-1)}(G) \leq O\left(\frac{\log r}{r}\right) \mathsf{OPT}$ by Theorem 4.1. Therefore, $w(E_0) \leq \tilde{\Phi}^{(C \log r(2k-1))}(S) \cdot D(S, \bar{S}) \leq O(\log r)\tilde{\Phi}^{(C \log r(2k-1))}(G) \cdot D_0 \leq O\left(\frac{\log^2 r}{r}\right) \mathsf{OPT} \cdot D_0$. Using Theorem 2.4 with $\alpha = O(\log^2 r)$, we get that the algorithm returns a bi-criteria $(O(\log r), O(\log^3 r))$-approximate solution to the problem.

5 Vertex Connectivity

In this section, we extend our approximation algorithms for EC-kRC to handle vertex-connectivity and prove Theorem 1.3. We start by extending some of our technical definitions and results to the vertex-connectivity setting.

Let (s, t) be any pair of vertices in graph G, and let $\Delta \subseteq V$ be any subset of vertices. We say that Δ is a *separator* for s and t, or that Δ *separates* s and t, iff $s, t \notin \Delta$, and s and t belong to two distinct connected components of $V \setminus \Delta$. We say that Δ is a minimum cost separator for (s, t) iff for each subset Δ' separating s from t, $|\Delta| \leq |\Delta'|$. Given any pair $S, T \subseteq V$ of vertex subsets, let $E(S, T)$ be the set of edges with one endpoint in S and the other endpoint in T. Similarly, we say that Δ separates S from T iff $S \cap \Delta = \emptyset$, $T \cap \Delta = \emptyset$, and $E(S, T) = \emptyset$. Notice that in general $G \setminus \Delta$ may contain more than two connected components. A *vertex cut* in graph $G = (V, E)$ is a tri-partition (S, Δ, T) of V, where $E(S, T) = \emptyset$. For any subset $\Delta \subseteq V$, we will sometimes refer to $|\Delta|$ as the cost of Δ.

We start with the following lemma, which is an analogue of Lemma 2.1 for vertex cuts. For technical reasons, it is more convenient to state it for graphs with costs on vertices. Given a graph $G = (V, E)$ with costs c_v on vertices $v \in V$, a cost of a subset $\Delta \subseteq V$ of vertices is $\sum_{v \in \Delta} c_v$.

LEMMA 5.1. LAMINAR FAMILY OF MIN. VERTEX CUTS). *Suppose we are given a graph $G = (V, E)$ with costs c_v on vertices $v \in V$, and a set $\{(s_1, t_1), (s_2, t_2), \cdots, (s_r, t_r)\}$ of r demand pairs. Let $\mathcal{T}$ be the set of all vertices participating in the demand pairs. Suppose additionally that for every demand pair (s_i, t_i), for every minimum-cost separator Δ for (s_i, t_i), $\Delta \cap \mathcal{T} = \emptyset$. Then there exists a family of vertex cuts (S_i, Δ_i, T_i) such that:*

1. *For every $i \in \{1, \cdots, r\}$, Δ_i is a minimum cost separator for (s_i, t_i) (note that s_i may belong either to S_i or T_i); and*

2. *Sets $\{S_i\}_{i=1}^r$ form a laminar family.*

For the edge-connectivity case, we used the existence of Gomory–Hu trees to prove the corresponding laminar decomposition (Lemma 2.1). For the vertex-connectivity case, we need a more intensive case analysis, which we provide in the full version of the paper.

Proof of Theorem 1.3 In this section, we prove Theorem 1.3, by showing a $(2, O(dk \log^{5/2} r \log \log r))$ bi-criteria approximation algorithm VC-kRC, where d is the maximum number of demand pairs in which any terminal participates. The running time of the algorithm is $n^{O(k)}$.

We start by extending the definition of the k-route sparsest cut to the vertex connectivity version. Given two disjoint subsets S, T of vertices, let $D(S, T)$ be the set of all demand pairs (s_i, t_i) with exactly one of the vertices s_i, t_i lying in S, and the other one lying in T. Given any pair (S, Δ) of disjoint subsets, let $\Upsilon^{(\Delta)}(S) = \sum_{e \in E(S, V \setminus (S \cup \Delta))} w_e$, where $E(S, V \setminus (S \cup \Delta))$ is the subset of all edges with one endpoint in S and the other endpoint in $V \setminus (S \cup \Delta)$.

The k-route vertex sparsity of the set S is then defined to be:

$$\tilde{\Psi}^{(k)}(S) = \min_{\substack{\Delta \subseteq V \setminus S: \\ |\Delta| \leq k-1}} \left\{ \frac{\Upsilon^{(\Delta)}(S)}{|D(S, V \setminus (S \cup \Delta))|} \right\},$$

and the k-route vertex sparsity of the graph G is:

$$\tilde{\Psi}^{(k)}(G) = \min_{S \subset V} \left\{ \tilde{\Psi}^{(k)}(S) \right\}$$

It is easy to see that, similarly to the edge version of k-route sparsest cut, the k-route vertex sparsest cut can be approximated in time $n^{O(k)}$ to within a factor of $\alpha_{\text{ALN}}(r)$, as we show in the next theorem.

THEOREM 5.1. *There is an algorithm that finds, in time $n^{O(k)}$, disjoint subsets $S, \Delta \subset V$ of vertices, with $|\Delta| \leq k - 1$ such that*

$$\Upsilon^{(\Delta)}(S) \leq \alpha_{\text{ALN}}(r) \cdot \tilde{\Psi}^{(k)}(G) \cdot |D(S, V \setminus (S \cup \Delta))|.$$

Proof. For every subset $\Delta \subset V$ of at most $k-1$ vertices, we use the algorithm $\mathcal{A}_{\text{ALN}}$ to find an approximate sparsest cut in the graph $G \setminus \Delta$, and output the sparsest cut among all such cuts. $\square$

Our algorithm for VC-kRC is very similar to the algorithm for EC-kRC from Section 4. The only difference is that we use Theorem 5.1 to find an approximate k-route vertex sparsest cut. The algorithm is summarized in Figure 3.

Input: A weighted graph $G(V, E)$ with a set $D = \{(s_i, t_i)\}_{1 \leq i \leq r}$ of demand pairs, and edge weights $\{w_e\}_{e \in E}$.
Output: A subset E' of edges, such that no demand pair s_i and t_i is $(2k - 1)$-vertex connected in $G \setminus E'$.

1. If $r = 0$ return $E' = \emptyset$.

2. Find sets U and Δ with $|\Delta| \leq 2k - 1$ and $\Upsilon^{(\Delta)}(U) \leq \alpha_{\text{ALN}}(r) \cdot \tilde{\Psi}^{(2k-1)}(G) \cdot |D(U, V \setminus (U \cup \Delta))|$ using Theorem 5.1.

3. Let $E_0 = E(U, V \setminus (U \cup \Delta))$, and let $G' = G \setminus E_0$.

4. Let D_0 be the set of all demand pairs (s_i, t_i) that are no longer $(2k - 1)$-vertex connected in G'.

5. Solve the problem recursively on G' with the set $D \setminus D_0$ of demand pairs to obtain a solution E_1.

6. Return $E' = E_0 \cup E_1$.

Figure 3: Bi-criteria approximation algorithm for VC-kRC in time $n^{O(k)}$.

It is easy to verify that if E' is the solution computed by the algorithm, then for each demand pair (s_i, t_i) there are at most $(2k - 1)$ vertex-disjoint paths connecting them in $G \setminus E'$. This is since the algorithm only removes a demand pair (s_i, t_i) when the terminals s_i and t_i are no longer $(2k - 1)$-vertex connected, and it terminates, since it removes at least one demand pair in each iteration.

In order to analyze the performance of the algorithm, we use the following theorem, that relates the value $\tilde{\Psi}^{(k)}(G)$ of the k-route vertex sparsest cut in graph G to the value OPT of the optimal solution to VC − kRC.

THEOREM 5.2.

$$\tilde{\Psi}^{(2k-1)}(G) \leq O\left(\frac{dk \log r}{r}\right) \cdot \text{OPT}.$$

Proof. Let $H = G \setminus E^*$. The proof roughly follows the proof of Theorem 3.1, except that we need one additional step, that is summarized in the following lemma.

LEMMA 5.2. *There exists a subset $D' \subseteq D$ of $r' = \Omega(r/(dk))$ demand pairs, and a collection of vertex cuts $\{(S_i, \Delta_i, T_i)\}_{(s_i, t_i) \in D'}$, such that:*

- *For all $(s_i, t_i) \in D'$, Δ_i is a separator for (s_i, t_i) in H, $|\Delta_i| < k$, and $\Delta_i \cap T' = \emptyset$, where T' is the set of all terminals participating in demand pairs in D'.*

- *$\{S_i\}_{(s_i, t_i) \in D'}$ is a laminar family of vertex subsets.*

Proof. For each $1 \leq i \leq r$, let Δ'_i be a minimum vertex separator for s_i and t_i in H. Since s_i and t_i are not k-vertex connected in H, $|\Delta'_i| < k$. We construct an auxiliary graph Z, whose vertex set is $[r]$, that is, each vertex i of Z represents the demand pair (s_i, t_i). We say that a demand i *blocks* another demand j iff Δ'_i contains either s_j or t_j (or both). We connect i and j with an edge in Z iff one of them blocks the other. Since $|\Delta'_i| \leq k - 1$ and each vertex in Δ'_i participates in at most d demand pairs, demand i blocks at most $(d-1)k$ demands. Therefore, there are at most $(d - 1)kr$ edges in Z. By Turan's theorem, there is an independent set I of size $\Omega(r/(dk))$ in Z. Let $r' = |I|$, and let $D' = \{(s_i, t_i) \mid i \in I\}$.

Next, we apply Lemma 5.1 to graph G with the set D' of demand pairs, where we define the cost c_u of every vertex u as follows: $c_u = |V|$ if $u = s_i$ or $u = t_i$ for some $(s_i, t_i) \in D'$, and $c(v) = 1$ otherwise. Since demand pairs in D' do not block each other, the minimum *cost* vertex cut for each of them has cost at most $k - 1 < |V|$. Let $\{(S_i, \Delta_i, T_i)\}_{(s_i, t_i) \in D'}$ be the collection of cuts returned by Lemma 5.1. It is easy to see that these cuts satisfy the conditions of the lemma. $\square$

We apply Lemma 5.2 and find the subset D' of demand pairs and vertex cuts (S_i, Δ_i, T_i). We assume w.l.o.g. that $D' = \{(s_1, t_1), \ldots, (s_{r'}, t_{r'})\}$. Now we need a counterpart of Lemma 4.1.

LEMMA 5.3. *There is a family $\mathcal{P} = \{U_1, \ldots, U_p\}$ of disjoint vertex subsets, and a collection $\{(U_j, \Lambda_j, R_j)\}_{j=1}^p$ of vertex cuts in graph H, such that:*

- *for each $1 \leq j \leq p$, $|\Lambda_j| < 2k - 1$,*

- $\sum_{j=1}^p |D(U_j, R_j)| \geq \frac{r'}{8 \log r'}$.

The proof closely follows the proof of Lemma 4.1 and can be found in the full version. Consider the family $\mathcal{P} = \{U_1, \ldots, U_p\}$ and the corresponding cuts (U_i, Λ_i, R_i) as in Lemma 5.3. Since all sets in $\mathcal{P}$ are mutually disjoint, and for each such set $U_i \in \mathcal{P}$, $|E_H(U_i, R_i)| \leq 2k - 1$,

$$\sum_{j=1}^p \Upsilon^{(\Lambda_j)}(U_j) \leq 2\mathsf{OPT},$$

and so

$$\frac{\sum_{j=1}^p \Upsilon^{(\Lambda_j)}(U_j)}{\sum_{j=1}^p |D(U_j, R_j)|} \quad \leq O\left(\frac{\log r'}{r'}\right) \cdot \mathsf{OPT}$$
$$\leq O\left(\frac{dk \log r}{r}\right) \cdot \mathsf{OPT}.$$

Therefore, there is an index $1 \leq j \leq p$, such that

$$\frac{\Upsilon^{(\Lambda_j)}(U_j)}{|D(U_j, R_j)|} \leq O\left(\frac{dk \log r}{r}\right) \cdot \mathsf{OPT}.$$

The left hand side of this inequality is at least $\tilde{\Psi}^{(2k-1)}(G)$ since $|\Lambda_j| \leq 2k - 2$. We conclude that $\tilde{\Psi}^{(2k-1)}(G) \leq O\left(\frac{dk \log r}{r}\right) \cdot \mathsf{OPT}$. □

In order to complete the proof of Theorem 1.3, observe that $w(E_0) = \Upsilon^{(\Delta)}(U)$, and by Theorem 5.2,

$$w(E_0) \leq \alpha_{\text{ALN}}(r)\Psi^{(2k-1)}(G)|D(U, V \setminus (U \cup \Delta))|$$
$$\leq O\left(\frac{dk \log r}{r}\right) \cdot \mathsf{OPT} \cdot |D(U, V \setminus (U \cup \Delta))|.$$

Note that we remove all demand pairs in $D(U, V \setminus (U \cup \Delta))$ in step 4 of the algorithm. We can now use Theorem 2.4 with $\alpha = O(dk \log r \cdot \alpha_{\text{ALN}}(r))$ to conclude that the cost of the solution returned by the algorithm is bounded by $O(dk \log^{5/2} r \log \log r) \cdot \mathsf{OPT}$.

6 Algorithms for 2-route cuts

In this section we prove Theorem 1.5. Since we prove in the full version that EC-kRC can be cast as a special case of VC-kRC, and the connectivity value k remains unchanged in this reduction, it is enough to prove the theorem for VC-kRC, where $k = 2$. In the rest of this section we show an efficient $O(\log^{3/2} r)$-approximation algorithm for VC-kRC with $k = 2$.

Given a subset S of vertices in graph G, the uniform vertex 2-route sparsity of S is:

$$\Psi^{(2)}(S) = \min_{\substack{\Delta \subseteq V \setminus S: \\ |\Delta| \leq 1}} \left\{ \frac{\Upsilon^{(\Delta)}(S)}{\min\{D(S), D(V \setminus (S \cup \Delta))\}} \right\},$$

and the uniform vertex 2-route sparsity of the graph G is:

$$\Psi^{(2)}(G) = \min_{S \subset V} \left\{ \Psi^{(2)}(S) \right\}$$

As before, we can efficiently approximate the uniform vertex 2-route sparsest cut in any graph, as shown in the next theorem.

THEOREM 6.1. *There is a polynomial time algorithm that finds disjoint subsets $S \subset V$ and $\Delta \subset V$ of vertices, with $|\Delta| \leq 1$ and $0 < D(S) \leq r$, such that*

$$\Upsilon^{(\Delta)}(S) \leq \alpha_{\text{ARV}}(r) \cdot \Psi^{(2)}(G) \cdot D(S).$$

Proof. For every subset $\Delta \subset V$ of size at most 1, we use the algorithm $\mathcal{A}_{\text{ARV}}$ to find the $\alpha_{\text{ARV}}(r)$-approximate uniform sparsest cut in graph $G \setminus \Delta$, and output the cut with the smallest sparsity. □

The approximation algorithm for VC-kRC with $k = 2$ is shown in Figure 4.

In order to analyze the algorithm, we start by showing that it is guaranteed to produce a feasible solution.

CLAIM 6.1. *The algorithm outputs a feasible solution to the problem.*

Proof. The proof is by induction on the number of vertices in G. Assume that the algorithm outputs a feasible solution for all graphs containing fewer than n vertices, and consider a graph G containing n vertices. Let (s_i, t_i) be any demand pair, and assume for contradiction that there are at least two vertex-disjoint simple paths P_1, P_2 connecting s_i to t_i in $G \setminus E'$. Observe first that either $s_i, t_i \in S \cup \Delta$ or $s_i, t_i \in T \cup \Delta$ must hold. Otherwise, one

Input: A weighted graph $G = (V, E)$ with demand pairs $\{(s_i, t_i)\}_{1 \leq i \leq r}$, and edge weights $\{w_e\}_{e \in E}$, such that each demand pair has at least 2 vertex-disjoint paths connecting them in G

Output: A set E' of edges such that each pair (s_i, t_i) is no longer 2-vertex connected in $G \setminus E'$.

1. If $r = 0$ return $E' = \emptyset$.

2. Find disjoint subsets S, Δ of vertices with $|\Delta| = 1$, $0 < D(S) \leq r$, such that $\Upsilon^{(\Delta)}(S) \leq \alpha_{\text{ARV}}(r) \cdot \Psi^{(2)}(G) \cdot D(S)$, using Theorem 6.1. Let $T = V \setminus (S \cup \Delta)$.

3. Let $E_0 = E(S, T)$; $G' = G \setminus E_0$.

4. Remove all demand pairs (s_i, t_i) that are no longer 2-vertex connected in G'.

5. Recursively solve the sub-instances induced by $G[S \cup \Delta]$ and $G[T \cup \Delta]$ to obtain solutions E_1 and E_2. The set of demand pairs for the instance induced by $G[S \cup \Delta]$ is defined to be the subset of all remaining demand pairs contained in $S \cup \Delta$. The set of demand pairs for the instance induced by $G[T \cup \Delta]$ is defined similarly.

6. Return $E' = E_0 \cup E_1 \cup E_2$.

Figure 4: Approximation algorithm for VC-kRC, $k = 2$ (weighted case).

of the two vertices must belong to S and the other to T. But Δ is a separator for S and T in graph G', and since $|\Delta| = 1$, the paths P_1 and P_2 cannot be vertex-disjoint. Assume w.l.o.g. that $s_i, t_i \in S \cup \Delta$. By the induction hypothesis, E_1 is a feasible solution to the instance induced by $G[S \cup \Delta]$, and in particular $G[S \cup \Delta] \setminus E_1$ cannot contain two vertex-disjoint paths connecting s_i to t_i. Therefore, at least one of the two paths, say P_1, must contain a vertex of T. But since Δ is a separator for S and T, $|\Delta| = 1$, and both $s_i, t_i \notin T$, path P_1 cannot be a simple path, a contradiction. $\square$

It now remains to bound the cost of the solution produced by the algorithm. As before, we do so by relating the value of the 2-route vertex sparsest cut to the value OPT of the optimal solution to the VC-kRC problem.

THEOREM 6.2. *Suppose that we are given an undirected graph $G = (V, E)$ with edge weights w_e, and r demand pairs $(s_1, t_1), \ldots, (s_r, t_r)$. Let OPT be cost of the optimal solution to the corresponding VC-kRC problem instance,*

and assume that $k = 2$. Then

$$\Psi^{(2)}(G) \leq \frac{4\text{OPT}}{r}.$$

The proof proceeds by considering a block decomposition of the graph $G \setminus OPT$, and is deferred to the full version.

Let a be the number of demand pairs contained in $S \cup \Delta$ and b be the number of demand pairs contained in $(V \setminus S) \cup \Delta$ in graph G'. From Theorem 6.2,

$$\begin{aligned}
w(E_0) &\leq \alpha_{\text{ARV}}(r) \cdot \Psi^{(2)}(G) \cdot D(S) \\
&\leq 4\alpha_{\text{ARV}}(r)\text{OPT} \min\{a, b\}/r \\
&\leq 4\alpha_{\text{ARV}}(r)\text{OPT} \min\{r - a, r - b\}/r
\end{aligned}$$

Therefore, by setting $\alpha = 4\alpha_{\text{ARV}}(r)$, we get the same recurrence as in the proof of Theorem 2.3:

$$w(E') \leq w(E'_1) + w(E'_2) + 2\alpha \cdot \frac{\min\{r - a, r - b\} \text{OPT}}{r}$$

Solving this recurrence as in Theorem 2.3, we get that $w(E') \leq O(\log^{3/2} r)\text{OPT}$.

7 A factor k^ϵ-hardness for k-VC-kRC

We prove Theorem 1.4 through a gadget reduction from the 3SAT(5) problem (each variable occurs in at most five clauses). Given a 3SAT(5) formula ϕ, we say that it is a Yes-Instance if it is satisfiable, and it is a No-Instance iff the maximum fraction of simultaneously satisfiable clauses is δ, for some constant parameter δ. From the PCP theorem, there is some constant δ, for which it is NP-hard to distinguish whether an input 3SAT(5) formula ϕ is a Yes-Instance or a No-Instance. We use the Raz verifier for 3SAT(5) with ℓ parallel repetitions. Given the 3SAT(5) instance φ, the verifier chooses, independently at random, ℓ clauses $C_1, \ldots, C_\ell$, and for each $i : 1 \leq i \leq \ell$, a variable x_i participating in clause C_i is chosen at random. The verifier then sends one query to each one of the two provers, while the query to the first prover consists of the indices of the variables $x_1 \ldots, x_\ell$, and the query to the second prover contains the indices of the clauses $C_1, \ldots, C_\ell$. The first prover returns an assignment to variables $x_1, \ldots, x_\ell$. The second prover is expected to return an assignment to all the variables in clauses $C_1, \ldots, C_\ell$, which must satisfy the clauses. Finally, the verifier checks that the answers of the two provers are consistent.

The following theorem is obtained by combining the PCP theorem[AS98, ALM+98] with the parallel repetition theorem [Raz98].

THEOREM 7.1. ([AS98, ALM$^+$98, Raz98]) *There exists a constant $\gamma > 0$, such that:*

- *If φ is a Yes-Instance, then there is a strategy of the provers, for which the acceptance probability is 1.*

- *If φ is a No-Instance, then for any strategy of the provers, the acceptance probability is at most $2^{-\gamma\ell}$.*

We denote the set of all the random strings of the verifier by R, $|R| = (5n)^\ell$, and the sets of all the possible queries of the first and the second prover by Q_1 and Q_2 respectively, $|Q_1| = n^\ell$, $|Q_2| = (5n/3)^\ell$. For each query $q \in Q$, let $A(q)$ be the collection of all the possible answers to q (if q is a query of the second prover, then $A(q)$ only contains answers that satisfy all the clauses of the query). Let $A = 2^\ell$, $A' = 7^\ell$. Then for each $q \in Q_1$, $|A(q)| = A$, and for each $q' \in Q_2$, $|A(q')| = A'$. Given a random string $r \in R$, let $q_1(r), q_2(r)$ be the queries sent to the first and the second prover respectively, when the verifier chooses r. For each $q \in Q_1$, let $N(q) = \{q' \in Q_2 \mid \exists r \in R : q_1(r) = q, q_2(r) = q'\}$, and for each $q' \in Q_2$, let $N(q') = \{q \in Q_1 \mid \exists r \in R : q_1(r) = q, q_2(r) = q'\}$. Notice that for all $q \in Q_1$, $|N(q)| = 5^\ell$, and for all $q' \in Q_2$, $|N(q')| = 3^\ell$.

Construction: We now turn to describe our reduction. For each query $q \in Q_1$ of the first prover, for each answer $a \in A(q)$, we have an edge $e(q, a)$, whose endpoints are denoted by $v(q, a)$, $u(q, a)$, and whose cost is $(5/3)^\ell$. We will think of $v(q, a)$ as the first endpoint of $e(q, a)$ and of $u(q, a)$ as its second endpoint, even though the graph is undirected. Similarly, for each query $q \in Q_2$ of the second prover, for each answer $a \in A(q)$, there is an edge $e(q, a) = (v(q, a), u(q, a))$, of cost 1. As before, $v(q, a)$ is called the first endpoint and $u(q, a)$ the second endpoint of $e(q, a)$. Let E_0 be the set of all resulting edges. For each $q \in Q$, let $V(q) = \{v(q, a), u(q, a) \mid a \in A(q)\}$.

For each random string $r \in R$ of the verifier, we introduce a source-sink pair $s(r), t(r)$, and two collections of edges $E_1(r), E_2(r)$, whose costs are ∞. Let $E_1 = \bigcup_{r \in R} E_1(r)$ and $E_2 = \bigcup_{r \in R} E_2(r)$. The set of edges in the final graph is $E_0 \cup E_1 \cup E_2$.

We now fix some random string $r \in R$, and define the set $E_1(r)$ of edges. Let $q = q_1(r)$, $q' = q_2(r)$. Let $(a_1, a_2, \ldots, a_A)$ be any ordering of the set $A(q)$ of answers to q_1. For each $1 \le i \le A$, let $b_1(a_i), b_2(a_i), \ldots, b_{z_i}(a_i)$ be the set of all answers to q' that are consistent with the answer a_i to q. We start by connecting the edges corresponding to $b_1(a_i), b_2(a_i), \ldots, b_{z_i}(a_i)$ into a single path P_i as follows: for $1 \le j < z_i$, we connect the second endpoint

of the edge $e(q', b_j(a_i))$ to the first endpoint of edge $e(q', b_{j+1}(a_i))$. We will refer to $v(q', b_1(a_i))$ as the first vertex on path P_i, and to $u(q', b_{z_i}(a_i))$ as the last vertex. Next, we connect the source $s(r)$ to the first vertex of $e(q, a_1)$ and the first vertex of P_1. We also connect the second vertex of $e(q, a_A)$ and the last vertex of P_A to the sink $t(r)$. Finally, for all $1 \le i < A$, we connect the last vertex of P_i to the first vertices of $e(q, a_{i+1})$ and P_{i+1}, and the second vertex of $e(q, a_i)$ to the first vertices of $e(q, a_{i+1})$ and P_{i+1}. This finishes the definition of the set $E_1(r)$ of edges. Let $G(r)$ be the graph whose vertex set is $V(q) \cup V(q') \cup \{s(r), t(r)\}$, and the edge set consists of $E_1(r)$ and the edges of E_0 representing the answers to q and q', that is: $\{e(q, a) \mid a \in A(q)\} \cup \{e(q', a') \mid a' \in A(q')\}$. Then $G(r)$ is an "almost layered" graph, where for each $1 \le i \le A$, layer i consists of the edge $e(q_1(r), a_i)$ and of the path P_i (see Figure 5). Notice that the only way to disconnect $s(r)$ from $t(r)$ in graph $G(r)$, without deleting edges of $E_1(r)$ (whose cost is ∞), is to delete a pair $e(q, a)$, $e(q', a')$ of edges, where a and a' are matching answers to queries q and q', respectively.
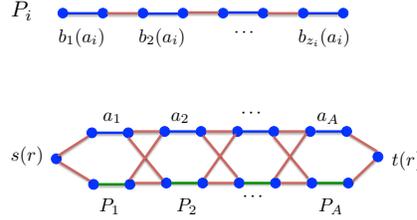

Figure 5: Graph $G(r)$. Red edges belong to $E_1(r)$ and have cost ∞.

Finally, we define the sets $E_2(r)$ of edges for all $r \in R$. Given a random string $r \in R$, let $N(r) = N(q_1(r)) \cup N(q_2(r))$, and let $U(r) = \bigcup_{q \in N(r)} V(q)$. Notice that $|U(r)| = |N(q_1(r))| \cdot 7^\ell + |N(q_2(r))| \cdot 2^\ell = 35^\ell + 6^\ell$. We connect $s(r)$ to every vertex in $U(r)$, and we connect every vertex in $U(r)$ to $t(r)$. We denote the resulting set of edges by $E_2(r)$, and we set the costs of these edges to be ∞. Finally, we set the parameter k to be $|U(r)| + 1 = 35^\ell + 6^\ell + 1$ (this value is identical for all $r \in R$). Let G be the final instance of the VC-kRC problem. The completeness and soundness analysis are included in the full version.

8   Single Source-Sink Pair

In this section we study the single source-sink pair version of EC-kRC and VC-kRC, denoted by (st)–EC-kRC and (st)–VC-kRC, respectively. We start with algorithmic results in Section 8.1, and complement them

with inapproximability results in Section 8.2.

8.1 Algorithms for the Single (s, t)-pair Version

This section is devoted to proving Theorem 1.6. Since we show in the full version that VC-kRC captures EC-kRC as a special case, and this reduction remains valid for the single source-sink pair version, it is enough to prove the theorem for VC-kRC. We start with describing the bi-critera approximation algorithm. We reduce the (st)-VC-kRC problem to the problem of finding a minimum-weight vertex (st)-cut in a new graph G'. Recall that in this problem, we are given a graph G', with non-negative weights $w(v)$ on vertices $v \in V(G')$, and two special vertices s and t. The goal is to find a minumum-weight subset $S \subseteq V(G') \setminus \{s, t\}$ of vertices, whose removal disconnects s from t in G'. This problem can be solved efficiently by standard techniques.

Given an instance G of the (st)-VC-kRC problem, let OPT denote the value of the optimal solution (that we guess). We produce an instance G' of the minimum-weight vertex (st)-cut problem, as follows. Graph G' is obtained from graph G, after sub-dividing every edge $e \in E(G)$ by a vertex v_e. The weight of this new vertex is set to be $w(v_e) = w_e$, and for each original vertex $v \in V(G)$, we set its weight to $\frac{c}{k-1} \cdot$ OPT.

For the correct guess of OPT, it is easy to see that there is vertex (st)-cut of value at most $(1 + c) \cdot$ OPT. Hence, the minimum value cut returns a $((1 + 1/c), (1 + c))$-bi-criteria approximate solution. In order to obtain a factor k-approximation algorithm, we use the above algorithm, setting the parameter $c = k$.

8.2 Inapproximability of (st)–VC-kRC

In this section we complement our upper bounds from Section 8.1 by inapproximability results, and prove Theorems 1.7, 1.8 and 1.9. The starting point for all three reductions is similar. We define the Small Set Vertex Expansion (SSVE) problem, and show an approximation preserving reduction from SSVE to (st)–VC-kRC, in Section 8.2.1. We then show inapproximability results for SSVE in subsequent sections, which are used to establish the lower bounds on the approximability of (st)–VC-kRC.

8.2.1 Small Set Vertex Expansion

DEFINITION 8.1.
(SMALL SET VERTEX EXPANSION PROBLEM (SSVE)).
Given a bipartite graph $G = (U, V, E)$ and a parameter $0 < \alpha \leq 1$, the aim is to find a subset $S \subseteq U$ of vertices, $|S| \geq \alpha|U|$, minimizing the number of its neighbors, $|\Gamma(S)|$.

We present a gap-preserving reduction from SSVE to (st)-VC-kRC, that will allow us to later focus on proving inapproximability of SSVE.

THEOREM 8.1. *Let $G = (U, V, E)$ be any bipartite graph with $|U| = m, |V| = n$, and let $N = 2mn + 1$. We can efficiently construct an edge-weighted graph G' with two special vertices $s, t \in V(G')$, such that for any $0 < \alpha < 1$, and any integer $0 \leq C \leq |V|$, the following property holds: there is a subset $S \subseteq U$ in graph G with $|S| \geq \alpha|U|$ and $|\Gamma(S)| \leq C$ iff there is a solution of cost at most $C \cdot N$ to the (st)-VC-kRC problem on graph G', where the parameter k is set to be $k = |U|(1 - \alpha) + 1$.*

Proof. Given an SSVE instance $G = (U, V, E)$, with $|U| = m, |V| = n$, let $N = 2nm + 1$. In order to construct the graph G', we start with the bipartite graph $G = (U, V, E)$, and then replace every vertex $v \in V$ with a clique $K(v)$ on N new vertices. All edges of the clique $K(v)$ have cost ∞. Let $V' = \bigcup_{v \in V} V(K(v))$ be the set of all vertices in all such cliques. We add an edge of cost ∞ between $u \in U$ and every vertex in $K(v)$ if $(u, v) \in E(G)$. We also add two additional vertices s and t. For every vertex $u \in U$, add an ∞-cost edge (s, u), and for every $v' \in V'$, add a cost-1 edge (v', t) to G'. This completes the description of graph G'. Given a parameter $0 < \alpha < 1$, we set $k = |U|(1 - \alpha) + 1$.

Given a subset $S \subseteq U$ with $|S| = \alpha|U|$ and $|\Gamma(S)| \leq C$ in graph G, the corresponding k-route cut consists of the edges, joining the cliques corresponding to $\Gamma(S)$ to the special terminal t. The $k-1$ 'free' nodes correspond to the vertices $U \setminus S$. The completeness and soundness proofs of the reduction are included in the full version. $\square$

8.2.2 Inapproximability from the Random κ-AND Assumption

This section is devoted to proving Theorem 1.7. We prove the following inapproximability result for SSVE.

THEOREM 8.2. *For every large enough constant κ, there are parameters α, β that depend on κ only, such that, assuming Hypothesis 1.1, no polynomial-time algorithm, given a bipartite graph $G = (U, V, E)$, can distinguish between the following two cases:*

- *Completeness: there is a subset $S \subseteq U$, with $|S| = \alpha|U|$ and $|\Gamma(S)| \leq \beta|V|$.*

- *Soundness: for any subset $S \subseteq U$ with $|S| \geq \frac{\alpha}{2^{\kappa/2}}|U|$, $|\Gamma(S)| > \beta \cdot 2^{\sqrt{\kappa}/c}|V|$,*

where c is a constant independent of κ.

The proof of Theorem 8.2 proceeds in two steps. In the first step, we show a simple reduction that gives a weak inapproximability result for the SSVE problem. Next, we amplify the inapproximability factor, by using graph products. Combining Theorem 8.2 with Theorem 8.1, we get that there is no polynomial-time algorithm for (st)–VC-kRC, that distinguishes between the cases where there is a solution of cost $\beta|V| \cdot N$ for parameter $k = |U|(1 - \alpha) + 1$, and the cases where there is no solution of cost $\beta/2^{\sqrt{\kappa}/c}$ and parameter $k = |U|\left(1 - \frac{\alpha}{2^{\kappa/2}}\right) + 1$. Since α and κ are constants, this will complete the proof of Theorem 1.7.

We refer the reader to the full version for the proof of Theorem 8.2, and the details regarding the reduction from Random 3-SAT assumption (Theorem 1.8).

8.2.3 Reduction from the Densest κ-Subgraph Problem
In this section we prove Theorem 1.9. As before, we do so by proving a similar result for the SSVE problem.

THEOREM 8.3. *For any constant $\lambda \geq 2$, and for any approximation factor ρ (that may depend on n), if there is an efficient factor ρ approximation algorithm for the* SSVE *problem, then there is an efficient factor $(2\rho^\lambda)$-approximation algorithm for the λ-uniform Hypergraph Densest κ-subgraph problem.*

Observe that combining Theorem 8.3 with Theorem 8.1 immediately implies Theorem 1.9. We now focus on proving Theorem 8.3.

Proof. Given a λ-uniform Hypergraph Densest κ-subgraph instance $G = (V, E)$, we construct an instance $G' = (U', V', E')$ of SSVE as follows. For each hyper-edge $e \in E$, we add a vertex u_e to U'. The set V' of vertices is $V' = V$. We add an edge between $u_e \in U'$ and $v \in V'$ iff vertex v belongs to the hyper-edge e. Since λ is a constant, we can assume that $\kappa \gg \lambda^2$ (otherwise, the optimal solution to the Densest κ-subgraph instance can be found efficiently by exhaustive search).

Let $\mathcal{A}$ be a factor ρ approximation algorithm for the SSVE problem. We now show a factor $2\rho^\lambda$-approximation algorithm for the Densest κ-subgraph problem. The algorithm will guess the value m' of the optimal solution to the Densest k-subgraph instance G. It will then apply algorithm $\mathcal{A}$ to instance G', with value $\alpha = m'/|U'|$. If value m' was guessed correctly, then there is a subset $S \subseteq U'$ of vertices, with $|S| = m'$, and $|\Gamma(S)| = \kappa$. Therefore, algorithm $\mathcal{A}$ must return a subset $S' \subseteq U'$ of vertices, with $|S'| = m'$, and $|\Gamma(S')| \leq \kappa \cdot \rho$. Let $V' = \Gamma(S')$. Set V' is also a subset of vertices in the initial instance G, and we are now guaranteed that $|V'| \leq \kappa \cdot \rho$, while the number of edges contained in V' is at least m'. Let V'' be a random subset of κ vertices from V'. Observe that for a hyper-edge $e \subseteq V'$, the probability that e is contained in V'' is at least $\left(\frac{\kappa - \lambda}{\rho\kappa}\right)^\lambda \geq \frac{2}{3} \cdot \left(\frac{1}{\rho}\right)^\lambda$ (since $\kappa \gg \lambda^2$), and so the expected number of hyper-edges contained in V'' is at least $2m'/3\rho^\lambda$. $\qquad\square$

References

[AAM+11] Noga Alon, Sanjeev Arora, Rajsekar Manokaran, Dana Moshkovitz, and Omri Weinstein. Manuscript. 2011.

[ABW10] Benny Applebaum, Boaz Barak, and Avi Wigderson. Public-key cryptography from different assumptions. In Leonard J. Schulman, editor, *STOC*, pages 171–180. ACM, 2010.

[ACKN07] Yash P. Aneja, Ramaswamy Chandrasekaran, Santosh N. Kabadi, and K. P. K. Nair. Flows over edge-disjoint mixed multipaths and applications. *Discrete Applied Mathematics*, 155:1979–2000, 2007.

[ALM+98] Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy. Proof verification and the hardness of approximation problems. *J. ACM*, 45:501–555, May 1998.

[ALN05] Sanjeev Arora, James R. Lee, and Assaf Naor. Euclidean distortion and the sparsest cut. In *STOC '05: Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, pages 553–562, New York, NY, USA, 2005. ACM.

[AO02] Charu C. Aggarwal and James B. Orlin. On multi-route maximum flows in networks. *Networks*, 39:43–52, 2002.

[App11] Benny Applebaum. Pseudorandom generators with long stretch and low locality from random local one-way functions. *Electronic Colloquium on Computational Complexity (ECCC)*, 18:7, 2011.

[ARV04] Sanjeev Arora, Satish Rao, and Umesh Vazirani. Expander flows, geometric embeddings and graph partitioning. In *STOC '04: Proceedings of the thirty-sixth annual ACM symposium on Theory of computing*, pages 222–231, New York, NY, USA, 2004. ACM.

[AS98] Sanjeev Arora and Shmuel Safra. Probabilistic checking of proofs: a new characterization of np. *J. ACM*, 45:70–122, January 1998.

[BC10] Siddharth Barman and Shuchi Chawla. Region growing for multi-route cuts. In *Proceedings of the*

Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '10, pages 404–418, Philadelphia, PA, USA, 2010. Society for Industrial and Applied Mathematics.

[BCC⁺10] Aditya Bhaskara, Moses Charikar, Eden Chlamtac, Uriel Feige, and Aravindan Vijayaraghavan. Detecting high log-densities: an $o(n^{1/4})$ approximation for densest k-subgraph. In *STOC '10: Proceedings of the 42nd ACM symposium on Theory of computing*, pages 201–210, New York, NY, USA, 2010. ACM.

[BČH⁺08] Henning Bruhn, Jakub Černý, Alexander Hall, Petr Kolman, and Jiří Sgall. Single source multiroute flows and cuts on uniform capacity networks. *Theory of Computing*, 4(1):1–20, 2008.

[BCK03] Amitabha Bagchi, Amitabh Chaudhary, and Petr Kolman. Short length menger's theorem and reliable optical routing. In *Theoretical Computer Science*, pages 246–247, 2003.

[BCSK07] Amitabha Bagchi, Amitabh Chaudhary, Christian Scheideler, and Petr Kolman. Algorithms for fault-tolerant routing in circuit-switched networks. *SIAM J. Discret. Math.*, 21:141–157, February 2007.

[CCK08] T. Chakraborty, J. Chuzhoy, and S. Khanna. Network design for vertex connectivity. In *Proceedings of ACM Symposium on Theory of Computing (STOC)*, pages 167–176, 2008.

[CK08] Chandra Chekuri and Sanjeev Khanna. Algorithms for 2-route cut problems. In *Automata, Languages and Programming*, volume 5125 of *Lecture Notes in Computer Science*, pages 472–484. Springer Berlin / Heidelberg, 2008.

[CKK⁺06] Shuchi Chawla, Robert Krauthgamer, Ravi Kumar, Yuval Rabani, and D. Sivakumar. On the hardness of approximating multicut and sparsest-cut. *Comput. Complex.*, 15:94–114, June 2006.

[DJP⁺94] E. Dahlhaus, D. S. Johnson, C. H. Papadimitriou, P. D. Seymour, and M. Yannakakis. The complexity of multiterminal cuts. *SIAM Journal on Computing*, 23:864–894, 1994.

[EGK⁺10] Matthias Englert, Anupam Gupta, Robert Krauthgamer, Harald Räcke, Inbal Talgam, and Kunal Talwar. Vertex sparsifiers: New results from old techniques. In *Approximation Algorithms for Combinatorial Optimization*, pages 152–165, 2010.

[Fei02] Uriel Feige. Relations between average case complexity and approximation complexity. In *Proceedings of the 34th annual ACM Symposium on Theory of Computing (STOC'02)*, pages 534–543. ACM Press, 2002.

[GVY95] N. Garg, V.V. Vazirani, and M. Yannakakis. Approximate max-flow min-(multi)-cut theorems and their applications. *SIAM Journal on Computing*, 25:235–251, 1995.

[HKPS05] Ara Hayrapetyan, David Kempe, Martin Päl, and Zoya Svitkina. Unbalanced graph cuts. In *In Proc. 13th European Symp. on Algorithms*, 2005.

[Kho04] Subhash Khot. Ruling out PTAS for graph min-bisection, densest subgraph and bipartite clique. In

Proceedings of the 44th Annual IEEE Symposium on the Foundations of Computer Science (FOCS'04), pages 136–145, 2004.

[Kis96] W. Kishimoto. A method for obtaining the maximum multiroute flows in a network. *Networks*, 27:279–291, 1996.

[KKL04] G. Kortsarz, R. Krauthgamer, and J. R. Lee. Hardness of approximation for vertex-connectivity network design problems. *SIAM Journal of Computing*, 33(3):704–720, 2004.

[KS11] Petr Kolman and Christian Scheideler. Towards Duality of Multicommodity Multiroute Cuts and Flows: Multilevel Ball-Growing. In *28th International Symposium on Theoretical Aspects of Computer Science (STACS 2011)*, volume 9, pages 129–140, Dagstuhl, Germany, 2011. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.

[KT93] W. Kishimoto and M. Takeuchi. On m-route flows in a network. *Ieice Transactions*, 1993.

[KV05] Subhash Khot and Nisheeth K. Vishnoi. The unique games conjecture, integrality gap for cut problems and embeddability of negative type metrics into ℓ_1. In *FOCS*, pages 53–62, 2005.

[LR99] F. T. Leighton and S. Rao. Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms. *Journal of the ACM*, 46:787–832, 1999.

[LZ10] Angsheng Li and Peng Zhang. Unbalanced graph partitioning. In *ISAAC (1)'10*, pages 218–229, 2010.

[Räc08] Harald Räcke. Optimal hierarchical decompositions for congestion minimization in networks. In *STOC '08: Proceedings of the 40th annual ACM symposium on Theory of computing*, pages 255–264, New York, NY, USA, 2008. ACM.

[Raz98] Ran Raz. A parallel repetition theorem. *SIAM J. Comput.*, 27(3):763–803, 1998.

[RS10] Prasad Raghavendra and David Steurer. Graph expansion and the unique games conjecture. In *STOC '10: Proceedings of the 42nd ACM symposium on Theory of computing*, pages 755–764, New York, NY, USA, 2010. ACM.