

Approximate minimum cuts and their enumeration*

Calvin Beideman

Karthekeyan Chandrasekaran

Weihang Wang

Abstract

We show that every α -approximate minimum cut in a connected graph is the *unique* minimum (S, T) -terminal cut for some subsets S and T of vertices each of size at most $\lfloor 2\alpha \rfloor + 1$. This leads to an alternative proof that the number of α -approximate minimum cuts in a n -vertex connected graph is $n^{O(\alpha)}$ and they can all be enumerated in deterministic polynomial time for constant α .

1 Introduction

Let $G = (V, E)$ be a graph with positive edge costs $c : E \rightarrow \mathbb{R}_+$. A partitioning of the vertex set into 2 non-empty parts is known as a *cut*. For a non-empty proper subset U of vertices, we will use \bar{U} to denote $V \setminus U$, (U, \bar{U}) to denote the cut associated with U , $\delta(U)$ to denote the set of edges crossing the cut (U, \bar{U}) , and $d(U) := \sum_{e \in \delta(U)} c(e)$ to denote the cut value of U . We will use λ to denote the value of a minimum cut in G , i.e.,

$$\lambda := \min\{d(U) : \emptyset \neq U \subsetneq V\}.$$

For $\alpha \geq 1$, a cut (U, \bar{U}) is an α -approximate minimum cut if $d(U) \leq \alpha\lambda$. A fundamental graph structural result states that the number of α -approximate minimum cuts in a n -vertex connected graph is $O(n^{2\alpha})$. This structural result forms the backbone of several algorithmic and representation results in graphs—e.g., fast randomized construction of skeletons [12], existence of cut sparsifiers which in turn finds applications in streaming and sketching [1–3, 14], approximation algorithms for TSP [4–6, 17], computing reliability of probabilistic networks [11], and polynomial time algorithms for connectivity augmentation [16]. The structural result along with a randomized polynomial-time algorithm to enumerate all constant-approximate minimum cuts was first shown via Karger’s random contraction technique [13]. Subsequently, the splitting-off technique and the tree packing technique have been used to show the structural result along with a *deterministic* polynomial-time algorithm to enumerate all constant-approximate minimum cuts [8, 15].

In this work, we give a fourth technique to bound the number of α -approximate minimum cuts by $n^{O(\alpha)}$ along with a deterministic polynomial-time algorithm to enumerate all constant-approximate minimum cuts. Let S, T be disjoint non-empty subsets of vertices. A 2-partition (U, \bar{U}) is an (S, T) -terminal cut if $S \subseteq U \subseteq V \setminus T$. Here, the set U is known as the *source set* and the set \bar{U} is known as the *sink set*. An (S, T) -terminal cut with minimum cut value will be denoted as a *minimum (S, T) -terminal cut*. The following is our main result.

Theorem 1.1. *Let $G = (V, E)$ be a connected graph with positive edge costs and (U, \bar{U}) be an α -approximate minimum cut for some $\alpha \geq 1$. Then, there exist subsets $S, T \subseteq V$ with $|S|, |T| \leq \lfloor 2\alpha \rfloor + 1$ such that (U, \bar{U}) is the unique minimum (S, T) -terminal cut.*

*University of Illinois, Urbana-Champaign. Email: {calvinb2, karthe, weihang3}@illinois.edu. Supported in part by NSF grants CCF-1814613 and CCF-1907937.

In other words, every α -approximate minimum cut (U, \bar{U}) in a connected graph can be recovered as the *unique* minimum (S, T) -terminal cut for some subsets S and T of sizes at most $\lfloor 2\alpha \rfloor + 1$. A few remarks are in order.

Remark 1.1. *An immediate consequence of Theorem 1.1 is that the number of distinct α -approximate minimum cuts in a n -vertex connected graph is at most $n^{4\alpha+2}$ and they can all be enumerated using $n^{4\alpha+2}$ minimum (S, T) -terminal cut computations. We recall that minimum (S, T) -terminal cut in a given graph with given subsets S and T of vertices can be computed in deterministic polynomial time.*

Remark 1.2. *The size bound for subsets S and T in Theorem 1.1 is tight up to an additive factor of one: Consider the n -vertex cycle graph G with α being a positive integer where $n \geq 4\alpha$ and with all edge costs being 1. Let (U, \bar{U}) be a cut in G with $d(U) = 2\alpha$ and moreover, each connected component in $G - \delta(U)$ has at least 2 vertices (see Figure 1 for such an example where $\alpha = 4$). For such a cut (U, \bar{U}) , the only choice of S and T for which (U, \bar{U}) is the unique minimum (S, T) -terminal cut is given by*

$$S := \{v \in U : v \text{ is an end vertex of an edge in } \delta(U)\} \text{ and}$$

$$T := \{v \in \bar{U} : v \text{ is an end vertex of an edge in } \delta(U)\}.$$

Hence, this example shows that both S and T need to have at least 2α vertices in order to recover (U, \bar{U}) as the unique minimum (S, T) -terminal cut.




Figure 1: An example showing the tightness of Theorem 1.1: The set U consists of vertices within all gray boxes and the set \bar{U} consists of all vertices outside gray boxes; the dotted edges correspond to the edges in the cut set $\delta(U)$. The only choice of S and T for which (U, \bar{U}) is the unique minimum (S, T) -terminal cut is as follows: S consists of all vertices that are shaded black and T consists of all vertices that are shaded gray.

Theorem 1.1 is inspired by a result of Goemans and Ramakrishnan [9]. They showed the following structural theorem for $(4/3 - \epsilon)$ -approximate minimum cuts: fix an arbitrary vertex t in a connected graph $G = (V, E)$ and let (U, \bar{U}) be a $(4/3 - \epsilon)$ -approximate minimum cut with $t \in \bar{U}$ for some $\epsilon > 0$. Then, there exists a subset S of size at most 2 such that (U, \bar{U}) is the unique minimum $(S, \{t\})$ -terminal cut. Their result leads to the following natural question: for every $\alpha \geq 1$, can

every α -approximate minimum cut (U, \overline{U}) be obtained as a minimum (S, T) -terminal cut for some subsets S and T of vertices each of size at most some function of α ? Our Theorem 1.1 answers this question affirmatively. However, our proof of Theorem 1.1 deviates significantly from the proof approach of Goemans-Ramakrishnan’s structural result for $(4/3 - \epsilon)$ -approximate minimum cuts. Their proof proceeds via contradiction and relies on the *submodular triple inequality* which holds for the graph cut function. The submodular triple inequality shows that 3 non-empty sets satisfying certain crossing properties can be uncrossed to obtain 4 disjoint non-empty sets without increasing the sum of their cut values; thus, if all 3 sets have cut value less than $(4/3)\lambda$, then one of the 4 sets should have cut value strictly smaller than λ , a contradiction to the definition of λ . However, there does not seem to be a generalization of the submodular triple inequality to larger number of sets. Our proof of Theorem 1.1 circumvents the submodular triple inequality but achieves its intended purpose in this context—we show that a sufficiently large number of α -approximate minimum cuts with certain crossing properties can be uncrossed to obtain a cut of value cheaper than λ , thus leading to a contradiction again.

1.1 Preliminaries

Let $G = (V, E)$ be a graph with positive edge costs $c : E \rightarrow \mathbb{R}_+$. For disjoint non-empty subsets $S, T \subseteq V$, there can be multiple minimum (S, T) -terminal cuts. We will be interested in *source minimal* minimum (S, T) -terminal cuts. A minimum (S, T) -terminal cut (U, \overline{U}) is a source minimal minimum (S, T) -terminal cut if for all minimum (S, T) -terminal cut (X, \overline{X}) , we have $U \subseteq X$. For every pair of disjoint non-empty subsets S and T of vertices, there exists a unique source minimal minimum (S, T) -terminal cut (this is due to submodularity—e.g., see [9]).

We recall that the graph cut function $d : 2^V \rightarrow \mathbb{R}$ is symmetric and submodular, i.e., $d(A) = d(V \setminus A)$ and $d(A) + d(B) \geq d(A \cap B) + d(A \cup B)$ for all $A, B \subseteq V$. We will need an uncrossing result that relies on more careful counting of edges than simply employing the submodularity inequality. We begin with some notation that will help in such careful counting. Let (Y_1, \dots, Y_p, W, Z) be a partition of V into $p + 2$ non-empty parts. We define

$$\sigma(Y_1, \dots, Y_p, W, Z) := 2 \left(\sum_{uv \in E: u \in Y_i, v \in Y_j \text{ for distinct } i, j \in [p]} c(uv) + \sum_{uv \in E: u \in W, v \in Z} c(uv) \right) + \sum_{uv \in E: u \in \cup_{i \in [p]} Y_i, v \in W \cup Z} c(uv)$$

We note that

$$\sum_{i \in [p]} d(Y_i) \leq \sigma(Y_1, \dots, Y_p, W, Z). \quad (1)$$

The following lemma uncrosses a collection of p sets to obtain a partition of the vertex set into $p + 2$ non-empty parts with small σ -value. The lemma was shown by Kamidoi, Yoshida, and Nagamochi in the context of minimum k -cuts [10]. It was also extended to hypergraphs by Chandrasekaran and Chekuri [7]. See Figure 2 for an illustration of the sets that appear in the statement of the lemma.

Lemma 1.1. [7, 10] *Let $G = (V, E)$ be a graph with positive edge costs and $\emptyset \neq R \subsetneq U \subsetneq V$. Let $S = \{u_1, \dots, u_p\} \subseteq U \setminus R$ for $p \geq 2$. For each $i \in [p]$, let (\overline{A}_i, A_i) be a minimum $((S \cup R) \setminus \{u_i\}, \overline{U})$ -terminal cut. Suppose that $u_i \in A_i \setminus (\cup_{j \in [p] \setminus \{i\}} A_j)$ for every $i \in [p]$. Let*

$$Z := \cap_{i=1}^p \overline{A}_i, \quad W := \cup_{1 \leq i < j \leq p} (A_i \cap A_j), \quad \text{and } Y_i := A_i - W \quad \forall i \in [p].$$

Then, (Y_1, \dots, Y_p, W, Z) is a partition of V into $p + 2$ non-empty parts with

$$\sigma(Y_1, \dots, Y_p, W, Z) \leq \min\{d(A_i) + d(A_j) : i, j \in [p], i \neq j\}.$$




Figure 2: Illustration of the sets that appear in the statement of Lemma 1.1.

We briefly remark on the lemma: known proofs of the lemma are by induction on p with an uncrossing argument that involves careful edge counting. The use of $\sigma(Y_1, \dots, Y_p, W, Z)$ instead of $\sum_{i \in [p]} d(Y_i)$ on the LHS is crucial for the induction argument—this is why we prefer to state the lemma as giving an upper bound on $\sigma(Y_1, \dots, Y_p, W, Z)$ although we will use it only to conclude an upper bound on $\sum_{i \in [p]} d(Y_i)$ via inequality (1).

2 Proof of Theorem 1.1

We will derive Theorem 1.1 from the following theorem.

Theorem 2.1. *Let $G = (V, E)$ be a connected graph with positive edge costs and (U, \bar{U}) be an α -approximate minimum cut for some $\alpha \geq 1$. Then, there exists a subset $S \subseteq U$ with $|S| \leq \lfloor 2\alpha \rfloor + 1$ such that (U, \bar{U}) is the unique minimum (S, \bar{U}) -terminal cut.*

Proof. Consider the collection $\mathcal{C} := \{Q \subsetneq V : \bar{U} \subsetneq Q, d(Q) \leq d(U)\}$. Suppose \mathcal{C} is empty. Let $S = \{x\}$ for some arbitrary $x \in U$ and let (X, \bar{X}) be a minimum (S, \bar{U}) -terminal cut. Then, $\bar{U} \subseteq \bar{X} \subsetneq V$ and $d(\bar{X}) \leq d(U)$. Since $\bar{X} \notin \mathcal{C}$ (because \mathcal{C} is empty), we must have $\bar{X} = \bar{U}$. Therefore, (U, \bar{U}) is the unique minimum (S, \bar{U}) -terminal cut.

Next, suppose $\mathcal{C} \neq \emptyset$. Let $S \subseteq U$ be a minimal transversal of the collection \mathcal{C} , i.e., $S \subseteq U$ is a minimal set with $S \cap Q \neq \emptyset$ for all $Q \in \mathcal{C}$. Proposition 2.1 and Lemma 2.1 complete the proof of Theorem 2.1. \square

Proposition 2.1. *(U, \bar{U}) is the unique minimum (S, \bar{U}) -terminal cut.*

Proof. Suppose (X, \bar{X}) be the source minimal minimum (S, \bar{U}) -terminal cut with $X \neq U$. Since (U, \bar{U}) is a (S, \bar{U}) -terminal cut, we have $d(\bar{X}) = d(X) \leq d(U)$. Moreover, we have $\emptyset \neq X \subsetneq U$, which implies that $\bar{U} \subsetneq \bar{X} \subsetneq V$. Hence, we have $\bar{X} \in \mathcal{C}$, and thus $S \cap \bar{X} \neq \emptyset$. This contradicts with the assumption that (X, \bar{X}) is a (S, \bar{U}) -terminal cut. Therefore, we must have $X = U$. \square

Lemma 2.1. $|S| \leq \lfloor 2\alpha \rfloor + 1$.

Proof. For the sake of contradiction, suppose that $|S| \geq \lfloor 2\alpha \rfloor + 2$. Let λ denote the value of a minimum cut in G . Our proof strategy is to arrive at a cut with value cheaper than λ , thus contradicting the definition of λ . For convenience, we will denote $p := |S|$ and write $S = \{u_1, \dots, u_p\}$. For each $i \in [p]$, let (\overline{A}_i, A_i) be the source minimal minimum $(S - u_i, \overline{U})$ -terminal cut.

Claim 2.1. *For every $i \in [p]$, we have $d(A_i) \leq d(U)$ and $u_i \in A_i$.*

Proof. Let $i \in [p]$. Since S is a minimal transversal of the collection \mathcal{C} , there exists a set $B_i \in \mathcal{C}$ such that $B_i \cap S = \{u_i\}$. Hence, (\overline{B}_i, B_i) is a $(S - u_i, \overline{U})$ -terminal cut. Therefore,

$$d(A_i) \leq d(B_i) \leq d(U).$$

We will show that A_i is in the collection \mathcal{C} . By definition, $A_i \subseteq V \setminus (S - u_i) \subsetneq V$ and $\overline{U} \subseteq A_i$. If $A_i = \overline{U}$, then the above inequalities are equations (since $d(A_i) = d(\overline{U}) = d(U)$) implying that (\overline{B}_i, B_i) is a minimum $(S - u_i, \overline{U})$ -terminal cut, and consequently, (\overline{B}_i, B_i) contradicts source minimality of the minimum $(S - u_i, \overline{U})$ -terminal cut (\overline{A}_i, A_i) . Therefore, $\overline{U} \subsetneq A_i$. Hence, A_i is in the collection \mathcal{C} .

We recall that the set S is a transversal for the collection \mathcal{C} and moreover, none of the elements of $S - u_i$ are in A_i by definition of A_i . Therefore, the vertex u_i must be in A_i . \square

We recall our assumption that $p \geq \lfloor 2\alpha \rfloor + 2$. We note that $p - 1 > 2\alpha \geq 2$. Using Claim 2.1, we observe that the sets U , $R := \{u_p\}$, $S = \{u_1, \dots, u_{p-1}\}$, and partitions (\overline{A}_i, A_i) for $i \in [p-1]$ satisfy the conditions of Lemma 1.1. Therefore, applying Lemma 1.1 shows the existence of a partition $(Y_1, \dots, Y_{p-1}, W, Z)$ of V into $p + 1$ non-empty parts such that

$$\sigma(Y_1, \dots, Y_{p-1}, W, Z) \leq \min\{d(A_i) + d(A_j) : i, j \in [p-1], i \neq j\} \leq 2d(U) \leq 2\alpha\lambda.$$

By inequality (1), this implies that

$$\sum_{i \in [p-1]} d(Y_i) \leq 2\alpha\lambda.$$

Therefore, there exists $i \in [p-1]$ such that

$$d(Y_i) \leq \frac{2\alpha\lambda}{p-1} < \lambda.$$

The last inequality is because $p - 1 > 2\alpha$ and $\lambda > 0$ since the graph is connected. Thus, we have a cut (Y_i, \overline{Y}_i) with cut value smaller than λ , a contradiction to the definition of λ . \square

Applying Theorem 2.1 to (\overline{U}, U) yields the following corollary.

Corollary 2.1. *Let $G = (V, E)$ be a connected graph with positive edge costs and (U, \overline{U}) be an α -approximate minimum cut for some $\alpha \geq 1$. Then, there exists a subset $T \subseteq \overline{U}$ with $|T| \leq \lfloor 2\alpha \rfloor + 1$ such that (U, \overline{U}) is the unique minimum (U, T) -terminal cut.*

We now restate Theorem 1.1 and prove it using Theorem 2.1 and Corollary 2.1.

Theorem 1.1. *Let $G = (V, E)$ be a connected graph with positive edge costs and (U, \overline{U}) be an α -approximate minimum cut for some $\alpha \geq 1$. Then, there exist subsets $S, T \subseteq V$ with $|S|, |T| \leq \lfloor 2\alpha \rfloor + 1$ such that (U, \overline{U}) is the unique minimum (S, T) -terminal cut.*

Proof. By Theorem 2.1, there exists a subset $S \subseteq U$ with $|S| \leq \lfloor 2\alpha \rfloor + 1$ such that (U, \overline{U}) is the unique minimum (S, \overline{U}) -terminal cut. By Corollary 2.1, there exists a subset $T \subseteq \overline{U}$ with $|T| \leq \lfloor 2\alpha \rfloor + 1$ such that (U, \overline{U}) is the unique minimum (U, T) -terminal cut. For these choices of subsets S and T , we now show that (U, \overline{U}) is the unique minimum (S, T) -terminal cut.

Let (Y, \overline{Y}) be an arbitrary minimum (S, T) -terminal cut. It suffices to show that $Y = U$. Since (U, \overline{U}) is a (S, T) -terminal cut, we have that

$$d(Y) \leq d(U).$$

Since $(Y \cap U, \overline{Y \cap U})$ is a (S, \overline{U}) -terminal cut and (U, \overline{U}) is a minimum (S, \overline{U}) -terminal cut, we have that

$$d(U) \leq d(Y \cap U).$$

Since $(Y \cup U, \overline{Y \cup U})$ is a (U, T) -terminal cut and (U, \overline{U}) is a minimum (U, T) -terminal cut, we have that

$$d(U) \leq d(Y \cup U).$$

Using the above inequalities in conjunction with the submodularity of the cut function, we obtain that

$$2d(U) \leq d(U \cap Y) + d(U \cup Y) \leq d(U) + d(Y) \leq 2d(U).$$

Hence, all of the above inequalities should be equations. Consequently, $(Y \cap U, \overline{Y \cap U})$ is a minimum (S, \overline{U}) -terminal cut, and $(Y \cup U, \overline{Y \cup U})$ is a minimum (U, T) -terminal cut. This implies that $Y \cap U = U$ since (U, \overline{U}) is the unique minimum (S, \overline{U}) -terminal cut and $Y \cup U = U$ since (U, \overline{U}) is the unique minimum (U, T) -terminal cut. Thus, we have that $Y = U$. \square

3 An Open Problem

The fundamental structural result that the number of α -approximate minimum cuts in a n -vertex connected graph is at most $n^{O(\alpha)}$ has been extended to r -rank hypergraphs (the *rank* of a hypergraph is the size of the largest hyperedge): Kogan and Krauthgamer [14] used the random contraction technique to show that the number of α -approximate minimum cuts in a r -rank n -vertex connected hypergraph is $O(2^{\alpha r} n^{2\alpha})$ and they can all be enumerated in randomized polynomial time for constant α and r . It would be interesting to give a *deterministic* proof of this bound.

References

- [1] K. Ahn and S. Guha, *Graph sparsification in the semi-streaming model*, Proceedings of the 36th International Colloquium on Automata, Languages and Programming: Part II, ICALP, 2009, pp. 328–338.
- [2] K. Ahn, S. Guha, and A. McGregor, *Analyzing graph structure via linear measurements*, Proceedings of the 23rd Annual ACM-SIAM Symposium on Discrete Algorithms, SODA, 2012, pp. 459–467.
- [3] ———, *Graph sketches: Sparsification, spanners, and subgraphs*, Proceedings of the 31st Symposium on Principles of Database Systems, PODS, 2012, pp. 5–14.
- [4] A. Benczur, *A representation of cuts within 6/5 times the edge connectivity with applications*, Proceedings of the 36th Annual IEEE Foundations of Computer Science, FOCS, 1995, pp. 92–102.

- [5] ———, *Cut structures and randomized algorithms in edge-connectivity problems*, Ph.D. thesis, MIT, 1997.
- [6] A. Benczur and M. Goemans, *Deformable polygon representation and near-mincuts*, Building Bridges: Between Mathematics and Computer Science, M. Groetschel and G.O.H. Katona, Eds., Bolyai Society Mathematical Studies **19** (2008), 103–135.
- [7] K. Chandrasekaran and C. Chekuri, *Hypergraph k -cut for fixed k in deterministic polynomial time*, Mathematics of Operations Research (2022), Prelim. version in Proceedings of the 61st Annual Symposium on Foundations of Computer Science (FOCS 2020) pages 810–821.
- [8] C. Chekuri, K. Quanrud, and C. Xu, *LP relaxation and tree packing for minimum k -cut*, SIAM Journal on Discrete Mathematics **34** (2020), no. 2, 1334–1353.
- [9] M. X. Goemans and V. S. Ramakrishnan, *Minimizing submodular functions over families of sets*, Combinatorica **15** (1995), 499–513.
- [10] Y. Kamidoi, N. Yoshida, and H. Nagamochi, *A Deterministic Algorithm for Finding All Minimum k -Way Cuts*, SIAM Journal on Computing **36** (2007), no. 5, 1329–1341.
- [11] D. Karger, *A randomized fully polynomial time approximation scheme for the all-terminal network reliability problem*, SIAM Journal on Computing **29** (1999), no. 2, 492–514.
- [12] ———, *Minimum cuts in near-linear time*, Journal of the ACM **47** (2000), no. 1, 46–76.
- [13] D. Karger and C. Stein, *A new approach to the minimum cut problem*, Journal of the ACM **43** (1996), no. 4, 601–640.
- [14] D. Kogan and R. Krauthgamer, *Sketching cuts in graphs and hypergraphs*, Proceedings of the 2015 Conference on Innovations in Theoretical Computer Science, ITCS, 2015, pp. 367–376.
- [15] H. Nagamochi, K. Nishimura, and T. Ibaraki, *Computing all small cuts in an undirected network*, SIAM Journal on Discrete Mathematics **10** (1997), no. 3, 469–481.
- [16] D. Naor, D. Gusfield, and C. Martel, *A fast algorithm for optimally increasing the edge connectivity*, SIAM Journal on Computing **26** (1997), no. 4, 1139–1165.
- [17] R. Zenklusen, *A 1.5-Approximation for Path TSP*, Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA, 2019, pp. 1539–1549.