

Counting and enumerating optimum cut sets for hypergraph k -partitioning problems for fixed k^*

Calvin Beideman Karthekeyan Chandrasekaran Weihang Wang

Abstract

We consider the problem of enumerating optimal solutions for two hypergraph k -partitioning problems—namely, HYPERGRAPH- k -CUT and MINMAX-HYPERGRAPH- k -PARTITION. The input in hypergraph k -partitioning problems is a hypergraph $G = (V, E)$ with positive hyperedge costs along with a fixed positive integer k . The goal is to find a partition of V into k non-empty parts (V_1, V_2, \dots, V_k) —known as a k -partition—so as to minimize an objective of interest.

1. If the objective of interest is the maximum cut value of the parts, then the problem is known as MINMAX-HYPERGRAPH- k -PARTITION. A subset of hyperedges is a MINMAX- k -CUT-SET if it is the subset of hyperedges crossing an optimum k -partition for MINMAX-HYPERGRAPH- k -PARTITION.
2. If the objective of interest is the total cost of hyperedges crossing the k -partition, then the problem is known as HYPERGRAPH- k -CUT. A subset of hyperedges is a MIN- k -CUT-SET if it is the subset of hyperedges crossing an optimum k -partition for HYPERGRAPH- k -CUT.

We give the first polynomial bound on the number of MINMAX- k -CUT-SETS and a polynomial-time algorithm to enumerate all of them in hypergraphs for every fixed k . Our technique is strong enough to also enable an $n^{O(k)}p$ -time deterministic algorithm to enumerate all MIN- k -CUT-SETS in hypergraphs, thus improving on the previously known $n^{O(k^2)}p$ -time deterministic algorithm, where n is the number of vertices and p is the size of the hypergraph. The correctness analysis of our enumeration approach relies on a structural result that is a strong and unifying generalization of known structural results for HYPERGRAPH- k -CUT and MINMAX-HYPERGRAPH- k -PARTITION. We believe that our structural result is likely to be of independent interest in the theory of hypergraphs (and graphs).

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

1 Introduction

In hypergraph k -partitioning problems, the input consists of a hypergraph $G = (V, E)$ with positive hyperedge-costs $c : E \rightarrow \mathbb{R}_+$ and a fixed positive integer k (e.g., $k = 2, 3, 4, \dots$). The goal is to find a partition of the vertex set into k *non-empty* parts V_1, V_2, \dots, V_k so as to minimize an objective of interest. There are several natural objectives of interest in hypergraph k -partitioning problems. In this work, we focus on two particular objectives: MINMAX-HYPERGRAPH- k -PARTITION and HYPERGRAPH- k -CUT:

1. In MINMAX-HYPERGRAPH- k -PARTITION, the objective is to minimize the maximum cut value of the parts of the k -partition—i.e., minimize $\max_{i=1}^k c(\delta(V_i))$; here $\delta(V_i)$ is the set of hyperedges intersecting both V_i and $V \setminus V_i$ and $c(\delta(V_i)) = \sum_{e \in \delta(V_i)} c(e)$ is the total cost of hyperedges in $\delta(V_i)$.
2. In HYPERGRAPH- k -CUT¹, the objective is to minimize the cost of hyperedges crossing the k -partition—i.e., minimize $c(\delta(V_1, \dots, V_k))$; here $\delta(V_1, \dots, V_k)$ is the set of hyperedges that intersect at least two sets in $\{V_1, \dots, V_k\}$ and $c(\delta(V_1, \dots, V_k)) = \sum_{e \in \delta(V_1, \dots, V_k)} c(e)$ is the total cost of hyperedges in $\delta(V_1, \dots, V_k)$.

If the input G is a graph, then we will refer to these problems as MINMAX-GRAPH- k -PARTITION and GRAPH- k -CUT respectively. We note that the case of $k = 2$ corresponds to global minimum cut in both objectives. In this work, we focus on the problem of enumerating all optimum solutions to MINMAX-HYPERGRAPH- k -PARTITION and HYPERGRAPH- k -CUT.

Motivations and Related Problems. We consider the problem of counting and enumerating optimum solutions for partitioning problems over hypergraphs for three reasons. Firstly, hyperedges provide more powerful modeling capabilities than edges and consequently, several problems in hypergraphs become non-trivial in comparison to graphs. Although hypergraphs and partitioning problems over hypergraphs (including MINMAX-HYPERGRAPH- k -PARTITION) were discussed as early as 1973 by Lawler [43], most of these problems still remain open. The powerful modeling capability of hyperedges has been useful in a variety of modern applications, which in turn, has led to a resurgence in the study of hypergraphs with recent works focusing on min-cuts, cut-sparsifiers, spectral-sparsifiers, etc. [6, 7, 11, 14, 18, 19, 23, 24, 35, 42, 56]. Our work adds to this rich and emerging theory of hypergraphs.

Secondly, hypergraph k -partitioning problems are special cases of submodular k -partitioning problems. In submodular k -partitioning problems, the input is a finite ground set V , a submodular function² $f : 2^V \rightarrow \mathbb{R}$ provided by an evaluation oracle³ and a positive integer k (e.g., $k = 2, 3, 4, \dots$). The goal is to partition the ground set V into k non-empty parts V_1, V_2, \dots, V_k so as to minimize an objective of interest. Two natural objectives are of interest: (1) In MINMAX-SUBMOD- k -PARTITION, the objective is to minimize $\max_{i=1}^k f(V_i)$ and (2) In MINSUM-SUBMOD- k -PARTITION, the objective is to minimize $\sum_{i=1}^k f(V_i)$. If the given submodular function is symmetric⁴, then we denote the resulting problems as MINMAX-SYMSUBMOD- k -PARTITION and MINSUM-SYMSUBMOD- k -PARTITION respectively. Since the hypergraph cut function is symmetric submodular, it follows that MINMAX-HYPERGRAPH- k -PARTITION is a special case of MINMAX-SYMSUBMOD- k -PARTITION. Moreover, HYPERGRAPH- k -CUT is a special case of MINSUM-SUBMOD- k -PARTITION (this reduction is slightly non-trivial with the submodular function in the reduction being asymmetric—e.g., see [50] for the

¹We emphasize that the objective of HYPERGRAPH- k -CUT is not equivalent to minimizing $\sum_{i=1}^k c(\delta(V_i))$.

²A real-valued set function $f : 2^V \rightarrow \mathbb{R}$ is submodular if $f(A) + f(B) \geq f(A \cap B) + f(A \cup B) \forall A, B \subseteq V$.

³An evaluation oracle for a set function f over a ground set V returns the value of $f(S)$ given $S \subseteq V$.

⁴A real-valued set function $f : 2^V \rightarrow \mathbb{R}$ is symmetric if $f(A) = f(V \setminus A) \forall A \subseteq V$.

reduction). Queyranne claimed, in 1999, a polynomial-time algorithm for MINSUM-SYMSUBMOD- k -PARTITION for every fixed k [52], however the claim was retracted subsequently (see [27]). The complexity status of submodular k -partitioning problems (for fixed $k \geq 4$) are open, so recent works have focused on hypergraph k -partitioning problems as a stepping stone towards submodular k -partitioning [7, 11, 12, 27, 50, 60, 61]. Our work contributes to this stepping stone by advancing the state of the art in hypergraph k -partitioning problems. We emphasize that the complexity status of two other variants of hypergraph k -partitioning problems which are also special cases of MINSUM-SUBMOD- k -PARTITION are still open (see [50, 60, 61] for these variants).

Thirdly, counting and enumeration of optimum solutions for *graph* k -partitioning problems are fundamental to graph theory and extremal combinatorics. They have found farther reaching applications than initially envisioned. We discuss some of the results and applications for $k = 2$ and $k > 2$ now. For $k = 2$ in connected graphs, it is well-known that the number of min-cuts and the number of α -approximate min-cuts are at most $\binom{n}{2}$ and $O(n^{2\alpha})$ respectively, and they can all be enumerated in polynomial time for constant α . These combinatorial results have been the crucial ingredients of several algorithmic and representation results in graphs. On the algorithmic front, these results enable fast randomized construction of graph skeletons which, in turn, plays a crucial role in fast algorithms to solve graph min-cut [37]. On the representation front, counting results form the backbone of cut sparsifiers which in turn have found applications in sketching and streaming [2–4, 42]. A polygon representation of the family of $6/5$ -approximate min-cuts in graphs was given by Benczur and Goemans in 1997 (see [8–10])—this representation was used in the recent groundbreaking $(3/2 - \epsilon)$ -approximation for metric TSP [39]. On the approximation front, in addition to the $(3/2 - \epsilon)$ -approximation for metric TSP [39], counting results also led to the recent 1.5-approximation for path TSP [59]. For $k > 2$, we note that fast algorithms for GRAPH- k -CUT have been of interest since they help in generating cutting planes while solving TSP [5, 20]. A recent series of works aimed towards improving the bounds on the number of optimum solutions for GRAPH- k -CUT culminated in a drastic improvement in the run-time to solve GRAPH- k -CUT [28, 31, 32]. Given the status of counting and enumeration results for k -partitioning in graphs and their algorithmic and representation implications that were discovered subsequently, we believe that a similar understanding in hypergraphs could serve as an important ingredient in the algorithmic and representation theory of hypergraphs.

The Enumeration Problem. There is a fundamental structural distinction between hypergraphs and graphs that becomes apparent while enumerating optimum solutions to k -partitioning problems. In connected graphs, the number of optimum k -partitions for GRAPH- k -CUT and for MINMAX-GRAPH- k -PARTITION are $n^{O(k)}$ and $n^{O(k^2)}$ respectively and they can all be enumerated in polynomial time, where n is the number of vertices in the input graph [13, 17, 28, 32, 38, 57]. In contrast, a connected hypergraph could have exponentially many optimum k -partitions for both MINMAX-HYPERGRAPH- k -PARTITION and HYPERGRAPH- k -CUT even for $k = 2$ —e.g., consider the hypergraph with a single hyperedge containing all vertices; we will denote this as the spanning-hyperedge-example. Hence, enumerating all optimum k -partitions for hypergraph k -partitioning problems in polynomial time is impossible. Instead, our goal in the enumeration problems is to enumerate k -cut-sets corresponding to optimum k -partitions. We will call a subset $F \subseteq E$ of hyperedges to be a k -cut-set if there exists a k -partition (V_1, \dots, V_k) such that $F = \delta(V_1, \dots, V_k)$; we will call a 2-cut-set as a cut-set. In the enumeration problems that we will consider, the input consists of a hypergraph $G = (V, E)$ with positive hyperedge-costs $c : E \rightarrow \mathbb{R}_+$ and a fixed positive integer k (e.g., $k = 2, 3, 4, \dots$).

1. For an optimum k -partition (V_1, \dots, V_k) for MINMAX-HYPERGRAPH- k -PARTITION in (G, c) , we will denote $\delta(V_1, \dots, V_k)$ as a MINMAX- k -CUT-SET. In ENUM-MINMAX-HYPERGRAPH- k -

PARTITION, the goal is to enumerate all MINMAX- k -CUT-SETS.

2. For an optimum k -partition (V_1, \dots, V_k) for HYPERGRAPH- k -CUT in (G, c) , we will denote $\delta(V_1, \dots, V_k)$ as a MIN- k -CUT-SET. In ENUM-HYPERGRAPH- k -CUT, the goal is to enumerate all MIN- k -CUT-SETS.

We observe that in the spanning-hyperedge-example, although the number of optimum k -partitions for MINMAX-HYPERGRAPH- k -PARTITION (as well as HYPERGRAPH- k -CUT) is exponential, the number of MINMAX- k -CUT-SETS (as well as MIN- k -CUT-SETS) is only one.

1.1 Results

In contrast to graphs, whose representation size is the number of edges, the representation size of a hypergraph $G = (V, E)$ is $p := \sum_{e \in E} |e|$. Throughout, our algorithmic discussion will focus on the case of fixed k (e.g., $k = 2, 3, 4, \dots$).

There are no prior results regarding ENUM-MINMAX-HYPERGRAPH- k -PARTITION in the literature. We recall the status of MINMAX-HYPERGRAPH- k -PARTITION. As mentioned earlier, MINMAX-HYPERGRAPH- k -PARTITION was discussed as early as 1973 by Lawler [43] with its complexity status being open until recently. We note that the objective here could be viewed as aiming to find a *fair* k -partition, i.e., a k -partition where no part pays too much in cut value. Motivated by this connection to fairness, Chandrasekaran and Chekuri (2021) [12] studied the more general problem of MINMAX-SYMSUBMOD- k -PARTITION. They gave the first (deterministic) polynomial-time algorithm to solve MINMAX-SYMSUBMOD- k -PARTITION and as a consequence, obtained the first polynomial-time algorithm to solve MINMAX-HYPERGRAPH- k -PARTITION. Their algorithm does not show any bound on the number of MINMAX- k -CUT-SETS since it solves the more general problem of MINMAX-SYMSUBMOD- k -PARTITION for which the number of optimum k -partitions can indeed be exponential (recall the spanning-hyperedge-example). Focusing on hypergraphs raises the question of whether all k -cut-sets corresponding to optimum solutions can be enumerated efficiently for every fixed k . We answer this question affirmatively by giving the first polynomial-time algorithm for ENUM-MINMAX-HYPERGRAPH- k -PARTITION.

Theorem 1.1. *There exists a deterministic algorithm to solve ENUM-MINMAX-HYPERGRAPH- k -PARTITION that runs in time $O(kn^{4k^2-2k+1}p)$, where n is the number of vertices and p is the size of the input hypergraph. Moreover, the number of MINMAX- k -CUT-SETS in a n -vertex hypergraph is $O(n^{4k^2-2k})$.*

We emphasize that our result shows the first polynomial bound on the number of MINMAX- k -CUT-SETS in hypergraphs for every fixed k (in addition to a polynomial-time algorithm to enumerate all of them for every fixed k). Our upper bound of $n^{O(k^2)}$ on the number of MINMAX- k -CUT-SETS is tight—there exist n -vertex connected *graphs* for which the number of MINMAX- k -CUT-SETS is $n^{\Theta(k^2)}$ (see Section 6).

Next, we briefly recall the status of HYPERGRAPH- k -CUT and ENUM-HYPERGRAPH- k -CUT. HYPERGRAPH- k -CUT was shown to be solvable in randomized polynomial time only recently [14, 23]; the randomized algorithms also showed that the number of MIN- k -CUT-SETS is $O(n^{2k-2})$ and they can all be enumerated in randomized polynomial time. A subsequent deterministic algorithm was designed to solve HYPERGRAPH- k -CUT in time $n^{O(k)}p$ by Chandrasekaran and Chekuri [11]. Chandrasekaran and Chekuri’s techniques were extended to design the first deterministic polynomial-time algorithm to solve ENUM-HYPERGRAPH- k -CUT in [7]. The algorithm for ENUM-HYPERGRAPH- k -CUT given in [7] runs in time $n^{O(k^2)}p$. We note that this run-time has a quadratic dependence on

k in the exponent of n although the number of MIN- k -CUT-SETS has only linear dependence on k in the exponent of n (since it is $O(n^{2k-2})$). So, an open question that remained from [7] is whether one can obtain an $n^{O(k)}p$ -time deterministic algorithm for ENUM-HYPERGRAPH- k -CUT. We resolve this question affirmatively.

Theorem 1.2. *There exists a deterministic algorithm to solve ENUM-HYPERGRAPH- k -CUT that runs in time $O(n^{16k-25}p)$, where n is the number of vertices and p is the size of the input hypergraph.*

Our algorithms for both ENUM-MINMAX-HYPERGRAPH- k -PARTITION and ENUM-HYPERGRAPH- k -CUT are based on a structural theorem that allows for efficient recovery of optimum k -cut-sets via minimum (s, t) -terminal cuts (see Theorem 1.4). Our structural theorem builds on structural theorems that have appeared in previous works on MINMAX-HYPERGRAPH- k -PARTITION and HYPERGRAPH- k -CUT [7, 11, 12]. Our structural theorem may appear to be natural/incremental in comparison to ones that appeared in previous works, but formalizing the theorem and proving it is a significant part of our contribution. Moreover, our single structural theorem is strong enough to enable efficient algorithms for both ENUM-HYPERGRAPH- k -CUT as well as ENUM-MINMAX-HYPERGRAPH- k -PARTITION in contrast to previously known structural theorems. In this sense, our structural theorem can be viewed as a strong and unifying generalization of structural theorems that have appeared in previous works. We believe that our structural theorem will be of independent interest in the theory of cuts and partitioning in hypergraphs (as well as graphs).

1.2 Technical overview and main structural result

We focus on the unit-cost variant of ENUM-HYPERGRAPH- k -CUT and ENUM-MINMAX-HYPERGRAPH- k -PARTITION in the rest of this work for the sake of notational simplicity—i.e., the cost of every hyperedge is 1. Throughout, we will allow multigraphs and hence, this is without loss of generality. Our algorithms extend in a straightforward manner to arbitrary hyperedge costs. They rely only on minimum (s, t) -terminal cut computations and hence, they are strongly polynomial-time algorithms.

Notation and background. Let $G = (V, E)$ be a hypergraph. Throughout this work, n will denote the number of vertices in G , m will denote the number of hyperedges in G , and $p := \sum_{e \in E} |e|$ will denote the representation size of G . We will denote a partition of the vertex set into h non-empty parts by an ordered tuple (V_1, \dots, V_h) and call such an ordered tuple as an h -partition. For a partition $\mathcal{P} = (V_1, V_2, \dots, V_h)$, we will say that a hyperedge e crosses the partition \mathcal{P} if it intersects at least two parts of the partition. We will refer to a 2-partition as a cut. For a non-empty proper subset U of vertices, we will use \bar{U} to denote $V \setminus U$, $\delta(U)$ to denote the set of hyperedges crossing the cut (U, \bar{U}) , and $d(U) := |\delta(U)|$ to denote the cut value of U . We observe that $\delta(U) = \delta(\bar{U})$, so we will use $d(U)$ to denote the value of the cut (U, \bar{U}) . More generally, given a partition $\mathcal{P} = (V_1, V_2, \dots, V_h)$, we denote the set of hyperedges crossing the partition by $\delta(V_1, V_2, \dots, V_h)$ (also by $\delta(\mathcal{P})$ for brevity) and the number of hyperedges crossing the partition by $|\delta(V_1, V_2, \dots, V_h)|$. We will denote the optimum value of MINMAX-HYPERGRAPH- k -PARTITION and HYPERGRAPH- k -CUT respectively by

$$OPT_{\text{minmax-}k\text{-partition}} := \min \left\{ \max_{i \in [k]} |\delta(V_i)| : (V_1, \dots, V_k) \text{ is a } k\text{-partition of } V \right\} \text{ and}$$

$$OPT_{k\text{-cut}} := \min \{ |\delta(V_1, \dots, V_k)| : (V_1, \dots, V_k) \text{ is a } k\text{-partition of } V \}.$$

A key algorithmic tool will be the use of fixed-terminal 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. A minimum-valued (S, T) -terminal

cut is known as a *minimum* (S, T) -terminal cut. Since there could be multiple minimum (S, T) -terminal cuts, we will be interested in *source minimal* minimum (S, T) -terminal cuts. For every pair of disjoint non-empty subsets S and T of vertices, there exists a unique source minimal minimum (S, T) -terminal cut and it can be found in deterministic polynomial time via standard maxflow algorithms. In particular, the source minimal minimum (S, T) -terminal cut can be found in time $O(np)$ [18].

Our technique to enumerate all MINMAX- k -CUT-SETS and all MIN- k -CUT-SETS will build on the approaches of Chandrasekaran and Chekuri for HYPERGRAPH- k -CUT and MINMAX-SYMSUBMOD- k -PARTITION [7, 11, 12]. We need the following structural theorem that was shown in [7].

Theorem 1.3. [7] *Let $G = (V, E)$ be a hypergraph and let $OPT_{k\text{-cut}}$ be the optimum value of HYPERGRAPH- k -CUT in G for some integer $k \geq 2$. Suppose (U, \bar{U}) is a 2-partition of V with $d(U) < OPT_{k\text{-cut}}$. Then, for every pair of vertices $s \in U$ and $t \in \bar{U}$, there exist subsets $S \subseteq U \setminus \{s\}$ and $T \subseteq \bar{U} \setminus \{t\}$ with $|S| \leq 2k - 3$ and $|T| \leq 2k - 3$ such that (U, \bar{U}) is the unique minimum $(S \cup \{s\}, T \cup \{t\})$ -terminal cut in G .*

ENUM-HYPERGRAPH- k -CUT. We first focus on ENUM-HYPERGRAPH- k -CUT. We note that Theorem 1.3 will allow us to recover those parts V_i of an optimum k -partition (V_1, \dots, V_k) for which $d(V_i) < OPT_{k\text{-cut}}$. However, recall that our goal is *not* to recover all optimum k -partitions for HYPERGRAPH- k -CUT, but rather to recover all MIN- k -CUT-SETS (i.e., not to recover the parts of every optimum k -partition, but rather only to recover the k -cut-set of every optimum k -partition). The previous work [7] that designed an $n^{O(k^2)}p$ -time deterministic enumeration algorithm achieved this by proving the following structural result: suppose (V_1, \dots, V_k) is an optimum k -partition for HYPERGRAPH- k -CUT for which $d(V_1) = OPT_{k\text{-cut}}$. Then, they showed that for every subset $T \subseteq \bar{V}_1$ satisfying $T \cap V_j \neq \emptyset$ for all $j \in \{2, \dots, k\}$, there exists a subset $S \subseteq V_1$ with $|S| \leq 2k$ such that the source minimal minimum (S, T) -terminal cut (A, \bar{A}) satisfies $\delta(A) = \delta(V_1)$. This structural theorem in conjunction with Theorem 1.3 allows one to enumerate a candidate family \mathcal{F} of $n^{O(k^2)}$ subsets of hyperedges such that every MIN- k -CUT-SET is present in the family. The drawback of their structural theorem is that it is driven towards recovering the cut-set $\delta(V_i)$ of every part V_i of every optimum k -partition (V_1, \dots, V_k) . Hence, their algorithmic approach ends up with a run-time of $n^{O(k^2)}p$. In order to improve the run-time, we prove a stronger result: we show that for an *arbitrary* cut (U, \bar{U}) with cut value $OPT_{k\text{-cut}}$ (as opposed to only those sets V_i of an optimum k -partition (V_1, \dots, V_k)), its cut-set $\delta(U)$ can be recovered as the cut-set of *any* minimum (S, T) -terminal cut for some S and T of small size. The following is the main structural theorem of this work.

Theorem 1.4. *Let $G = (V, E)$ be a hypergraph and let $OPT_{k\text{-cut}}$ be the optimum value of HYPERGRAPH- k -CUT in G for some integer $k \geq 2$. Suppose (U, \bar{U}) is a 2-partition of V with $d(U) = OPT_{k\text{-cut}}$. Then, there exist sets $S \subseteq U$, $T \subseteq \bar{U}$ with $|S| \leq 2k - 1$ and $|T| \leq 2k - 1$ such that every minimum (S, T) -terminal cut (A, \bar{A}) satisfies $\delta(A) = \delta(U)$.*

We encourage the reader to compare and contrast Theorems 1.3 and 1.4. The former helps to recover cuts whose cut value is strictly smaller than $OPT_{k\text{-cut}}$ while the latter helps to recover *cut-sets* whose size is equal to $OPT_{k\text{-cut}}$. So, the latter theorem is weaker since it only recovers cut-sets, but we emphasize that this is the best possible that one can hope to do (as seen from the spanning-hyperedge-example). However, proving the latter theorem requires us to work with cut-sets (as opposed to cuts) which is a technical barrier to overcome. Indeed, our proof of Theorem 1.4 deviates significantly from the proof of Theorem 1.3 since we have to work with cut-sets. Our proof also deviates from the structural result in [7] that was mentioned in the paragraph above Theorem 1.4 since our result is stronger than their result—our result helps to recover the cut-set $\delta(U)$ of an

arbitrary cut (U, \bar{U}) whose cut value is $d(U) = OPT_{k\text{-cut}}$ while their result helps only to recover the cut-set $\delta(V_i)$ of a part V_i of an optimum k -partition (V_1, \dots, V_k) for HYPERGRAPH- k -CUT whose cut value is $d(V_i) = OPT_{k\text{-cut}}$; moreover, their proof technique crucially relies on a containment property with respect to the part V_i , whereas under the hypothesis of our structural theorem, the containment property fails with respect to the set U and consequently, our proof technique differs from theirs.

Theorems 1.3 and 1.4 lead to a deterministic $n^{O(k)}$ -time algorithm to enumerate all MIN- k -CUT-SETS via a divide-and-conquer approach. We describe this algorithm now: For each pair (S, T) of disjoint subsets of vertices S and T with $|S|, |T| \leq 2k - 1$, compute the source minimal minimum (S, T) -terminal cut (A, \bar{A}) ; (i) if $G - \delta(A)$ has at least k connected components, then add $\delta(A)$ to the candidate family \mathcal{F} ; (ii) otherwise, add the set A to a collection \mathcal{C} . We note that the sizes of the family \mathcal{F} and the collection \mathcal{C} are $O(n^{4k-2})$. Next, for each subset A in the collection \mathcal{C} , recursively enumerate all MIN- $k/2$ -CUT-SETS in the subhypergraphs induced by A and \bar{A} respectively⁵—denoted $G[A]$ and $G[\bar{A}]$ respectively—and add $\delta(A) \cup F_1 \cup F_2$ to the family \mathcal{F} for each F_1 and F_2 being MIN- $k/2$ -CUT-SET in $G[A]$ and $G[\bar{A}]$ respectively. Finally, return the subfamily of k -cut-sets from the family \mathcal{F} that are of smallest size.

We sketch the correctness analysis of the above approach: let $F = \delta(V_1, \dots, V_k)$ be a MIN- k -CUT-SET with (V_1, \dots, V_k) being an optimum k -partition for HYPERGRAPH- k -CUT. We will show that the family \mathcal{F} contains F . Let $U := \cup_{i=1}^{k/2} V_i$. We note that $\delta(U) \subseteq F$. We have two possibilities: (1) Say $d(U) = |F|$. Then, $d(U) = OPT_{k\text{-cut}}$. Consequently, by Theorem 1.4, the MIN- k -CUT-SET F will be added to the family \mathcal{F} by step (i). (2) Say $d(U) < |F|$. Then, by Theorem 1.3, the set $U = \cup_{i=1}^{k/2} V_i$ will be added to the collection \mathcal{C} by step (ii); moreover, $F_1 := F \cap E(G[U])$ and $F_2 := F \cap E(G[\bar{U}])$ are MIN- $k/2$ -CUT-SETS in $G[U]$ and $G[\bar{U}]$ respectively and they would have been enumerated by recursion, and hence, the set $\delta(U) \cup F_1 \cup F_2 = F$ will be added to the family \mathcal{F} . The size of the family \mathcal{F} can be shown to be $n^{O(k \log k)}$ and the run-time is $n^{O(k \log k)} p$ (see Theorem 4.1). Using the known fact that the number of MIN- k -CUT-SETS in a n -vertex hypergraph is $O(n^{2k-2})$, we can improve the run-time analysis of this approach to $n^{O(k)} p$ (see Lemma 4.1).

ENUM-MINMAX-HYPERGRAPH- k -PARTITION. Next, we focus on ENUM-MINMAX-HYPERGRAPH- k -PARTITION. There is a fundamental technical issue in enumerating MINMAX- k -CUT-SETS as opposed to MIN- k -CUT-SETS. We highlight this technical issue now. Suppose we find an optimum k -partition (V_1, \dots, V_k) for MINMAX-HYPERGRAPH- k -PARTITION (say via Chandrasekaran and Chekuri’s algorithm [12]) and store only the MINMAX- k -CUT-SET $F = \delta(V_1, \dots, V_k)$ but forget to store the partition (V_1, \dots, V_k) ; now, by knowing a MINMAX- k -CUT-SET F , can we recover *some* optimum k -partition for MINMAX-HYPERGRAPH- k -PARTITION (not necessarily (V_1, \dots, V_k))? Or by knowing a MINMAX- k -CUT-SET F , is it even possible to find the value $OPT_{\text{minmax-}k\text{-partition}}$ without solving MINMAX-HYPERGRAPH- k -PARTITION from scratch again—i.e., is there an advantage to knowing a MINMAX- k -CUT-SET in order to solve MINMAX-HYPERGRAPH- k -PARTITION? We are not aware of such an advantage. This is in stark contrast to HYPERGRAPH- k -CUT where knowing a MIN- k -CUT-SET enables a linear-time solution to HYPERGRAPH- k -CUT⁶.

Why is this issue significant while solving ENUM-MINMAX-HYPERGRAPH- k -PARTITION? We recall that in our approach for ENUM-HYPERGRAPH- k -CUT, the algorithm computed a polynomial-sized family \mathcal{F} containing all MIN- k -CUT-SETS and returned the ones with smallest size—the smallest size ones will exactly be MIN- k -CUT-SETS. It is unclear if a similar approach could work for enu-

⁵Subhypergraph $G[A]$ has vertex set A and contains all hyperedges of G which are entirely contained within A .

⁶Suppose we know a MIN- k -CUT-SET F . Then consider the connected components Q_1, \dots, Q_i in $G - F$ and create a partition (P_1, \dots, P_k) by taking $P_i = Q_i$ for every $i \in [k - 1]$ and $P_k = \cup_{j=k}^i Q_j$; such a k -partition (P_1, \dots, P_k) will be an optimum k -partition for HYPERGRAPH- k -CUT.

merating MINMAX- k -CUT-SETS: suppose we do have an algorithm to enumerate a polynomial-sized family \mathcal{F} containing all MINMAX- k -CUT-SETS; now, in order to return all MINMAX- k -CUT-SETS (which is a subfamily of \mathcal{F}), note that we need to identify them among the ones in the family \mathcal{F} —i.e., we need to verify if a given subset $F \in \mathcal{F}$ of hyperedges is a MINMAX- k -CUT-SET; this verification problem is closely related to the question mentioned in the previous paragraph. We do not know how to address this verification problem directly. So, our algorithmic approach for ENUM-MINMAX-HYPERGRAPH- k -PARTITION has to overcome this technical issue.

Our ingredient to overcome this technical issue is to enumerate *representatives* for MINMAX- k -CUT-SETS. For a k -partition (V_1, \dots, V_k) and disjoint subsets $U_1, \dots, U_k \subseteq V$, we will call the k -tuple (U_1, \dots, U_k) to be a *k -cut-set representative* of (V_1, \dots, V_k) if $U_i \subseteq V_i$ and $\delta(U_i) = \delta(V_i)$ for all $i \in [k]$. We note that a fixed k -partition (V_1, \dots, V_k) could have several k -cut-set representatives and a fixed k -tuple (U_1, \dots, U_k) could be the k -cut-set representative of several k -partitions. Yet, it is possible to efficiently verify if a given k -tuple (U_1, \dots, U_k) is a k -cut-set representative (Theorem 5.1). Moreover, knowing a k -cut-set representative (U_1, \dots, U_k) of a k -partition (V_1, V_2, \dots, V_k) allows one to recover the k -cut-set $F := \delta(V_1, \dots, V_k)$ since $F = \cup_{i=1}^k \delta(U_i)$. Thus, in order to enumerate all MINMAX- k -CUT-SETS, it suffices to enumerate k -cut-set representatives of *all* optimum k -partitions for MINMAX-HYPERGRAPH- k -PARTITION. At this point, the astute reader may wonder if there exists a polynomial-sized family of k -cut-set representatives of all optimum k -partitions for MINMAX-HYPERGRAPH- k -PARTITION given that the number of optimum k -partitions for MINMAX-HYPERGRAPH- k -PARTITION could be exponential. For example, is there a polynomial-sized family of k -cut-set representatives of all optimum k -partitions for MINMAX-HYPERGRAPH- k -PARTITION in the spanning-hyperedge-example? Indeed, in the spanning-hyperedge-example, even though the number of optimum k -partitions for MINMAX-HYPERGRAPH- k -PARTITION is exponential, there exists a $(k! \binom{n}{k})$ -sized family of k -cut-set representatives of all optimum k -partitions: consider the family $\{(\{v_1\}, \dots, \{v_k\}) : v_1, \dots, v_k \in V, v_i \neq v_j \forall \text{ distinct } i, j \in [k]\}$.

It turns out that Theorems 1.3 and 1.4 are strong enough to enable efficient enumeration of k -cut-set representatives of all optimum k -partitions for MINMAX-HYPERGRAPH- k -PARTITION. We describe the algorithm to achieve this: For each pair (S, T) of disjoint subsets of vertices with $|S|, |T| \leq 2k - 1$, compute the source minimal minimum (S, T) -terminal cut (U, \bar{U}) and add U to a candidate collection \mathcal{C} . We note that the size of the collection \mathcal{C} is $O(n^{4k-2})$. Next, for each k -tuple $(U_1, \dots, U_k) \in \mathcal{C}^k$, verify if (U_1, \dots, U_k) is a k -cut-set representative (using Theorem 5.1) and if so, then add the k -tuple to the candidate family \mathcal{D} . Finally, return $\arg \min\{\max_{i=1}^k d(U_i) : (U_1, \dots, U_k) \in \mathcal{D}\}$, i.e., prune and return the subfamily of k -cut-set representatives (U_1, \dots, U_k) from the family \mathcal{D} that have minimum $\max_{i=1}^k d(U_i)$.

We note that the size of the family \mathcal{D} is $n^{O(k^2)}$ and consequently, the run-time is $n^{O(k^2)}p$. We sketch the correctness analysis of the above approach: let (V_1, \dots, V_k) be an optimum k -partition for MINMAX-HYPERGRAPH- k -PARTITION. We will show that the family \mathcal{D} contains a k -cut-set representative of (V_1, \dots, V_k) . By noting that $OPT_{\text{minmax-}k\text{-partition}} \leq OPT_{k\text{-cut}}$ and by Theorems 1.3 and 1.4, for every $i \in [k]$, we have a set U_i in the collection \mathcal{C} with $U_i \subseteq V_i$ and $\delta(U_i) = \delta(V_i)$. Hence, the k -tuple $(U_1, \dots, U_k) \in \mathcal{C}^k$ is a k -cut-set representative and it will be added to the family \mathcal{D} . The final pruning step will not remove (U_1, \dots, U_k) from the family \mathcal{D} and hence, it will be in the subfamily returned by the algorithm.

Significance of our technique. As mentioned earlier, our techniques build on the structural theorems that appeared in previous works [7, 11, 12]. The main technical novelty of our contribution lies in Theorem 1.4 which can be viewed as the culmination of structural theorems developed in those previous works. We also emphasize that using minimum (s, t) -terminal cuts to solve global partitioning problems is not a new technique per se (e.g., minimum (s, t) -terminal cut is the first

and most natural approach to solve global minimum cut). This technique of using minimum (s, t) -terminal cuts to solve global partitioning problems has a rich variety of applications in combinatorial optimization: e.g., (1) it was used to design the first efficient algorithm for GRAPH- k -CUT for fixed k [26], (2) it was used to design efficient algorithms for certain constrained submodular minimization problems [25, 49], and (3) more recently, it was used to design fast algorithms for global minimum cut in graphs as well as to obtain fast Gomory-Hu trees in unweighted graphs [1, 44]. The applicability of this technique relies on identifying and proving appropriate structural results. Our Theorem 1.4 is such a structural result. The merit of the structural result lies in its ability to solve two different enumeration problems in hypergraph k -partitioning which was not possible via structural theorems that were developed before. Moreover, it leads to the first polynomial bound on the number of MINMAX- k -CUT-SETS in hypergraphs for every fixed k .

Organization. We discuss related work in Section 1.3. In Section 1.4, we recall properties of the hypergraph cut function. In Section 2, we prove a special case of Theorem 1.4. In Section 3, we use this special case to prove Theorem 1.4. In Section 4, we design an efficient algorithm for ENUM-HYPERGRAPH- k -CUT and prove Theorem 1.2. In Section 5, we design an efficient algorithm for ENUM-MINMAX-HYPERGRAPH- k -PARTITION and prove Theorem 1.1. We present a lower bound example in Section 6. We conclude with an open question in Section 7.

1.3 Related work

We briefly discuss the status of the enumeration problems for $k = 2$ followed by the status for $k \geq 2$ in graphs and hypergraphs.

Enumeration problems for $k = 2$. For $k = 2$, both ENUM-MINMAX-HYPERGRAPH- k -PARTITION and ENUM-HYPERGRAPH- k -CUT are equivalent to enumerating optimum solutions for global minimum cut in hypergraphs. For graphs that are connected, it is well-known that the number of minimum cuts and the number of α -approximate minimum cuts are at most $\binom{n}{2}$ and $O(n^{2\alpha})$ respectively and they can all be enumerated in polynomial time for constant α [17, 21, 25, 33, 36, 48]. For connected hypergraphs, the number of minimum cuts can be exponential as seen from the spanning-hyperedge-example. However, recent results have shown that the number of minimum cut-sets in a hypergraph is at most $\binom{n}{2}$ via several different techniques and they can all be enumerated in polynomial time [7, 14, 18, 23, 24]. On the other hand, there exist hypergraphs with exponential number of 2-approximate minimum cut-sets.⁷

GRAPH- k -CUT and ENUM-GRAPH- k -CUT. When k is part of input, GRAPH- k -CUT is NP-hard [26] and admits a $2(1 - 1/k)$ -approximation [30, 51, 54, 55, 61]. Manurangsi [47] showed that there is no polynomial-time $(2 - \epsilon)$ -approximation for any constant $\epsilon > 0$ assuming the *Small Set Expansion Hypothesis* [53]. We note that GRAPH- k -CUT is $W[1]$ -hard when parameterized by k [22] and admits a fixed-parameter approximation scheme when parameterized by k [29, 30, 41, 45], and is fixed-parameter tractable when parameterized by k and the solution size [40].

GRAPH- k -CUT for fixed k was shown to be polynomial-time solvable by Goldschmidt and Hochbaum [26]. Subsequently, Karger and Stein [38] gave a randomized polynomial-time algorithm whose analysis also showed that the number of optimum k -partitions in a connected graph is $O(n^{2k-2})$ and they can all be enumerated in polynomial time for every fixed k (also see [17, 34, 57, 58]). The number of optimum k -partitions has recently been improved to $O(n^k)$ for fixed k thereby leading to a faster algorithm for GRAPH- k -CUT for fixed k [28, 31, 32].

⁷Consider the n -vertex hypergraph $G = (V, E)$ obtained from the complete bipartite graph $K_{n/2, n/2} = (L \cup R, E')$ by adding $n^2/4$ copies of two hyperedges e_1 and e_2 , where $e_1 := L$ and $e_2 := R$.

HYPERGRAPH- k -CUT and ENUM-HYPERGRAPH- k -CUT. When k is part of input, HYPERGRAPH- k -CUT is at least as hard as the *densest k -subgraph* problem [16]. Combined with results in [46], this implies that HYPERGRAPH- k -CUT is unlikely to have a sub-polynomial factor approximation ratio. Moreover, HYPERGRAPH- k -CUT is $W[1]$ -hard even when parameterized by k and the solution size (see [11]). These two hardness results illustrate that HYPERGRAPH- k -CUT differs significantly from GRAPH- k -CUT in complexity.

HYPERGRAPH- k -CUT for fixed k was recently shown to be polynomial-time solvable via a randomized algorithm [14, 23]. The analysis of the randomized algorithm also showed that the number of MIN- k -CUT-SETS is $O(n^{2k-2})$ and they can all be enumerated in randomized polynomial time. A deterministic polynomial-time algorithm was given by Chandrasekaran and Chekuri [11]. Subsequently, a deterministic polynomial-time algorithm for ENUM-HYPERGRAPH- k -CUT for fixed k was given by Beideman, Chandrasekaran, and Wang [7].

MINMAX-GRAPH- k -PARTITION and ENUM-MINMAX-GRAPH- k -PARTITION. When k is part of input, MINMAX-GRAPH- k -PARTITION is NP-hard [13] while its approximability is open—we do not yet know if it admits a constant factor approximation. When parameterized by k , it is $W[1]$ -hard and admits a fixed-parameter approximation scheme [13].

MINMAX-GRAPH- k -PARTITION for fixed k is polynomial-time solvable via the following observation (see [12, 13]): in connected graphs, an optimum k -partition for MINMAX-GRAPH- k -PARTITION is a k -approximate solution to GRAPH- k -CUT. The randomized algorithm of Karger and Stein implies that the number of k -approximate solutions to GRAPH- k -CUT is $n^{O(k^2)}$ and they can all be enumerated in polynomial time [17, 28, 32, 38]. These two facts together imply that MINMAX-GRAPH- k -PARTITION can be solved in time $n^{O(k^2)}$ and moreover, the number of optimum k -partitions for MINMAX-GRAPH- k -PARTITION in a connected graph is $n^{O(k^2)}$ and they can all be enumerated in polynomial time for constant k .

MINMAX-HYPERGRAPH- k -PARTITION and ENUM-MINMAX-HYPERGRAPH- k -PARTITION. MINMAX-HYPERGRAPH- k -PARTITION was discussed as early as 1973 by Lawler [43]. When k is part of input, MINMAX-HYPERGRAPH- k -PARTITION is at least as hard as the densest k -subgraph problem (this follows from the reduction in [16] and was observed by [15]). For fixed k , the approach for MINMAX-GRAPH- k -PARTITION described above does not extend to MINMAX-HYPERGRAPH- k -PARTITION. This is because, the number of k -approximate solutions to HYPERGRAPH- k -CUT can be exponential and hence, they cannot be enumerated efficiently (e.g., we have already seen that the number of 2-approximate minimum cut-sets in a hypergraph can be exponential). Chandrasekaran and Chekuri [12] gave a deterministic polynomial-time algorithm for the more general problem of MINMAX-SYMSUBMOD- k -PARTITION for fixed k which in turn, implies that MINMAX-HYPERGRAPH- k -PARTITION is also solvable efficiently for fixed k . Their algorithm finds an optimum k -partition for MINMAX-HYPERGRAPH- k -PARTITION and is not conducive to enumerate all MINMAX- k -CUT-SETS. We emphasize that no polynomial bound on the number of MINMAX- k -CUT-SETS for fixed k was known prior to our work.

For a detailed discussion on other hypergraph k -partitioning problems that are special cases of MINSUM-SUBMOD- k -PARTITION, we refer the reader to [50, 60, 61].

1.4 Preliminaries

Let $G = (V, E)$ be a hypergraph. Throughout, we will follow the notation mentioned in the second paragraph of Section 1.2. For disjoint $A, B \subseteq V$, we define $E(A, B) := \{e \in E : e \subseteq A \cup B, e \cap A \neq \emptyset, e \cap B \neq \emptyset\}$, and $E[A] := \{e \in E : e \subseteq A\}$. We will repeatedly rely on the fact that the hypergraph cut function $d : 2^V \rightarrow \mathbb{R}_+$ is symmetric and submodular. We recall that a

set function $f : 2^V \rightarrow \mathbb{R}$ is *symmetric* if $f(U) = f(\bar{U})$ for all subsets $U \subseteq V$ and is *submodular* if $f(A) + f(B) \geq f(A \cap B) + f(A \cup B)$ for all subsets $A, B \subseteq V$.

We will need the following partition uncrossing theorem that was proved in previous works on HYPERGRAPH- k -CUT and ENUM-HYPERGRAPH- k -CUT (see Figure 1 for an illustration of the sets that appear in the statement of Theorem 1.5):

Theorem 1.5. [7, 11] *Let $G = (V, E)$ be a hypergraph, $k \geq 2$ be an integer and $\emptyset \neq R \subsetneq U \subsetneq V$. Let $S = \{u_1, \dots, u_p\} \subseteq U \setminus R$ for $p \geq 2k - 2$. Let (\bar{A}_i, A_i) be a minimum $((S \cup R) \setminus \{u_i\}, \bar{U})$ -terminal cut. Suppose that $u_i \in A_i \setminus (\cup_{j \in [p] \setminus \{i\}} A_j)$ for every $i \in [p]$. Then, the following two hold:*

1. *There exists a k -partition (P_1, \dots, P_k) of V with $\bar{U} \subsetneq P_k$ such that*

$$|\delta(P_1, \dots, P_k)| \leq \frac{1}{2} \min\{d(A_i) + d(A_j) : i, j \in [p], i \neq j\}.$$

2. *Moreover, if there exists a hyperedge $e \in E$ such that e intersects $W := \cup_{1 \leq i < j \leq p} (A_i \cap A_j)$, e intersects $Z := \cap_{i \in [p]} \bar{A}_i$, and e is contained in $W \cup Z$, then the inequality in the previous conclusion is strict.*

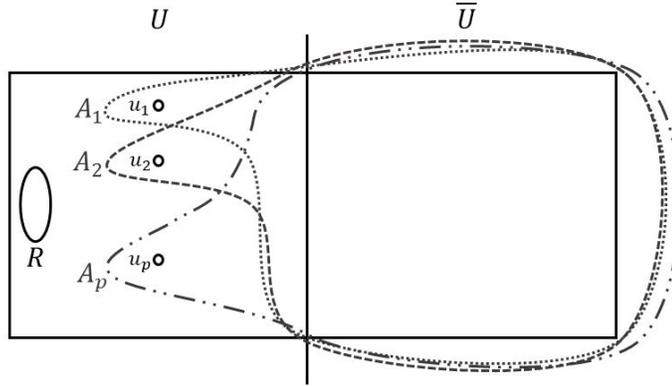


Figure 1: Illustration of the sets that appear in the statement of Theorem 1.5.

2 A special case of Theorem 1.4

The following is the main theorem of this section. Theorem 2.1 implies Theorem 1.4 in the special case where the 2-partition (U, \bar{U}) of interest to Theorem 1.4 is such that $|\bar{U}| \leq 2k - 1$.

Theorem 2.1. *Let $G = (V, E)$ be a hypergraph and let $OPT_{k\text{-cut}}$ be the optimum value of HYPERGRAPH- k -CUT in G for some integer $k \geq 2$. Suppose (U, \bar{U}) is a 2-partition of V with $d(U) = OPT_{k\text{-cut}}$. Then, there exists a set $S \subseteq U$ with $|S| \leq 2k - 1$ such that every minimum (S, \bar{U}) -terminal cut (A, \bar{A}) satisfies $\delta(A) = \delta(U)$.*

Proof. Consider the collection

$$\mathcal{C} := \{Q \subseteq V : \bar{U} \subsetneq Q, d(Q) \leq d(U), \text{ and } \delta(Q) \neq \delta(U)\}.$$

Let S be an inclusion-wise minimal subset of U such that $S \cap Q \neq \emptyset$ for all $Q \in \mathcal{C}$, i.e., the set S is completely contained in U and is a minimal transversal of the collection \mathcal{C} . Proposition 2.1 and Lemma 2.1 complete the proof of Theorem 2.1 for this choice of S . \square

Proposition 2.1. *Every minimum (S, \overline{U}) -terminal cut (A, \overline{A}) has $\delta(A) = \delta(U)$.*

Proof. Let (A, \overline{A}) be a minimum (S, \overline{U}) -terminal cut. If $A = U$, then we are done, so we may assume that $A \neq U$. This implies that $S \subseteq A$ and $\overline{U} \subsetneq \overline{A}$. Since (U, \overline{U}) is a (S, \overline{U}) -terminal cut, we have that $d(\overline{A}) = d(A) \leq d(U)$. Since S intersects every set in the collection \mathcal{C} , we have that $\overline{A} \notin \mathcal{C}$. Hence, $\delta(\overline{A}) = \delta(U)$, and by symmetry of cut-sets, $\delta(A) = \delta(U)$. \square

Lemma 2.1. *The size of the subset S is at most $2k - 1$.*

Proof. For the sake of contradiction, suppose $|S| \geq 2k$. Our proof strategy is to show the existence of a k -partition with fewer crossing hyperedges than $OPT_{k\text{-cut}}$, thus contradicting the definition of $OPT_{k\text{-cut}}$. Let $S := \{u_1, u_2, \dots, u_p\}$ for some $p \geq 2k$. For notational convenience, we will use $S - u_i$ to denote $S \setminus \{u_i\}$ and $S - u_i - u_j$ to denote $S \setminus \{u_i, u_j\}$. For a subset $X \subseteq U$, we denote the source minimal minimum (X, \overline{U}) -terminal cut by $(H_X, \overline{H_X})$.

Our strategy to arrive at a k -partition with fewer crossing hyperedges than $OPT_{k\text{-cut}}$ is to apply the second conclusion of Theorem 1.5. The next few claims will set us up to obtain sets that satisfy the hypothesis of Theorem 1.5.

Claim 2.1. *For every $i \in [p]$, we have $\overline{H_{S-u_i}} \in \mathcal{C}$.*

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,

$$d(\overline{H_{S-u_i}}) \leq d(B_i) \leq d(U).$$

Since $(H_{S-u_i}, \overline{H_{S-u_i}})$ is a $(S - u_i, \overline{U})$ -terminal cut, we have that $\overline{U} \subseteq \overline{H_{S-u_i}}$. If $d(\overline{H_{S-u_i}}) < d(U)$, then $\delta(\overline{H_{S-u_i}}) \neq \delta(U)$ and $\overline{U} \subsetneq \overline{H_{S-u_i}}$, and consequently, $\overline{H_{S-u_i}} \in \mathcal{C}$. So, we will assume henceforth that $d(\overline{H_{S-u_i}}) = d(U)$.

Since $(H_{S-u_i} \cap \overline{B_i}, \overline{H_{S-u_i} \cap \overline{B_i}})$ is a $(S - u_i, \overline{U})$ -terminal cut, we have that

$$d(H_{S-u_i} \cap \overline{B_i}) \geq d(H_{S-u_i}).$$

Since $(H_{S-u_i} \cup \overline{B_i}, \overline{H_{S-u_i} \cup \overline{B_i}})$ is a $(S - u_i, \overline{U})$ -terminal cut, we have that

$$d(H_{S-u_i} \cup \overline{B_i}) \geq d(H_{S-u_i}).$$

Therefore, by submodularity of the hypergraph cut function, we have that

$$2d(U) \geq d(H_{S-u_i}) + d(B_i) \geq d(H_{S-u_i} \cap \overline{B_i}) + d(H_{S-u_i} \cup \overline{B_i}) \geq 2d(H_{S-u_i}) = 2d(U). \quad (1)$$

Therefore, all inequalities above should be equations. In particular, we have that $d(H_{S-u_i} \cap \overline{B_i}) = d(U) = d(B_i) = d(H_{S-u_i})$ and hence, $(H_{S-u_i} \cap \overline{B_i}, \overline{H_{S-u_i} \cap \overline{B_i}})$ is a minimum $(S - u_i, \overline{U})$ -terminal cut. Since $(H_{S-u_i}, \overline{H_{S-u_i}})$ is a source minimal minimum $(S - u_i, \overline{U})$ -terminal cut, we must have $H_{S-u_i} \cap \overline{B_i} = H_{S-u_i}$, and thus $H_{S-u_i} \subseteq \overline{B_i}$. Therefore, $B_i \subseteq \overline{H_{S-u_i}}$. Since $B_i \in \mathcal{C}$, we have $\delta(B_i) \neq \delta(U)$. However, $d(B_i) = d(U)$. Therefore $\delta(U) \setminus \delta(B_i) \neq \emptyset$. Let $e \in \delta(U) \setminus \delta(B_i)$. Since $e \in \delta(\overline{U})$, but $e \notin \delta(B_i)$, and $\overline{U} \subseteq B_i$, we have that $e \subseteq B_i$, and thus $e \subseteq \overline{H_{S-u_i}}$. Thus, we conclude that $\delta(U) \setminus \delta(\overline{H_{S-u_i}}) \neq \emptyset$, and so $\delta(\overline{H_{S-u_i}}) \neq \delta(U)$. This also implies that $\overline{U} \subsetneq \overline{H_{S-u_i}}$. Thus, $\overline{H_{S-u_i}} \in \mathcal{C}$. \square

Claim 2.1 implies the following Corollary.

Corollary 2.1. *For every $i \in [p]$, we have $u_i \in \overline{H_{S-u_i}}$.*

Proof. By definition, $S - u_i \subseteq H_{S-u_i}$, so $S \cap \overline{H_{S-u_i}} \subseteq \{u_i\}$. By Claim 2.1 we have that $\overline{H_{S-u_i}} \in \mathcal{C}$. Since S is a transversal of the collection \mathcal{C} , we have that $S \cap \overline{H_{S-u_i}} \neq \emptyset$. So, the vertex u_i must be in $\overline{H_{S-u_i}}$. \square

Having obtained Corollary 2.1, the next few claims (Claims 2.2, 2.3, 2.4, and 2.5) are similar to the claims appearing in the proof of a structural theorem that appeared in [7]. Since the hypothesis of the structural theorem that we are proving here is different from theirs, we present the complete proofs of these claims here. The way in which we use the claims will also be different from [7].

The following claim will help in showing that $u_i, u_j \notin H_{S-u_i-u_j}$, which in turn, will be used to show that the hypothesis of Theorem 1.5 is satisfied by suitably chosen sets.

Claim 2.2. *For every $i, j \in [p]$, we have $H_{S-u_i-u_j} \subseteq H_{S-u_i}$.*

Proof. We may assume that $i \neq j$. We note that $(H_{S-u_i-u_j} \cap H_{S-u_i}, \overline{H_{S-u_i-u_j} \cap H_{S-u_i}})$ is a $(S - u_i - u_j, \overline{U})$ -terminal cut. Therefore,

$$d(H_{S-u_i-u_j} \cap H_{S-u_i}) \geq d(H_{S-u_i-u_j}). \quad (2)$$

Also, $(H_{S-u_i-u_j} \cup H_{S-u_i}, \overline{H_{S-u_i-u_j} \cup H_{S-u_i}})$ is a $(S - u_i, \overline{U})$ -terminal cut. Therefore,

$$d(H_{S-u_i-u_j} \cup H_{S-u_i}) \geq d(H_{S-u_i}). \quad (3)$$

By submodularity of the hypergraph cut function and inequalities (2) and (3), we have that

$$\begin{aligned} d(H_{S-u_i-u_j}) + d(H_{S-u_i}) &\geq d(H_{S-u_i-u_j} \cap H_{S-u_i}) + d(H_{S-u_i-u_j} \cup H_{S-u_i}) \\ &\geq d(H_{S-u_i-u_j}) + d(H_{S-u_i}). \end{aligned}$$

Therefore, inequality (2) is an equation, and consequently, $(H_{S-u_i-u_j} \cap H_{S-u_i}, \overline{H_{S-u_i-u_j} \cap H_{S-u_i}})$ is a minimum $(S - u_i - u_j, \overline{U})$ -terminal cut. If $H_{S-u_i-u_j} \setminus H_{S-u_i} \neq \emptyset$, then $(H_{S-u_i-u_j} \cap H_{S-u_i}, \overline{H_{S-u_i-u_j} \cap H_{S-u_i}})$ contradicts source minimality of the minimum $(S - u_i - u_j, \overline{U})$ -terminal cut $(H_{S-u_i-u_j}, \overline{H_{S-u_i-u_j}})$. Hence, $H_{S-u_i-u_j} \setminus H_{S-u_i} = \emptyset$ and consequently, $H_{S-u_i-u_j} \subseteq H_{S-u_i}$. \square

Claim 2.2 implies the following Corollary.

Corollary 2.2. *For every $i, j \in [p]$, we have $u_i, u_j \notin H_{S-u_i-u_j}$.*

Proof. By Corollary 2.1, we have that $u_i \notin H_{S-u_i}$. Therefore, $u_i, u_j \notin H_{S-u_i} \cap H_{S-u_j}$. By Claim 2.2, $H_{S-u_i-u_j} \subseteq H_{S-u_i}$ and $H_{S-u_i-u_j} \subseteq H_{S-u_j}$. Therefore, $H_{S-u_i-u_j} \subseteq H_{S-u_i} \cap H_{S-u_j}$, and thus, $u_i, u_j \notin H_{S-u_i-u_j}$. \square

The next claim will help in controlling the cost of the k -partition that we will obtain by applying Theorem 1.5.

Claim 2.3. *For every $i, j \in [p]$, we have $d(H_{S-u_i}) = d(U) = d(H_{S-u_i-u_j})$.*

Proof. Let $a, b \in [p]$. We will show that $d(H_{S-u_a}) = d(U) = d(H_{S-u_a-u_b})$. Since (U, \overline{U}) is a $(S-u_a, \overline{U})$ -terminal cut, we have that $d(H_{S-u_a}) \leq d(U)$. Since $(H_{S-u_a}, \overline{H_{S-u_a}})$ is a $(S-u_a-u_b, \overline{U})$ -terminal cut, we have that $d(H_{S-u_a-u_b}) \leq d(H_{S-u_a}) \leq d(U)$. Thus, in order to prove the claim, it suffices to show that $d(H_{S-u_a-u_b}) \geq d(U)$.

Suppose for contradiction that $d(H_{S-u_a-u_b}) < d(U)$. Let $\ell \in [p] \setminus \{a, b\}$ be an arbitrary element (which exists since we have assumed that $p \geq 2k$ and $k \geq 2$). Let $R := \{u_\ell\}$, $S' := S - u_a - u_\ell$, and $A_i := \overline{H_{S-u_a-u_i}}$ for every $i \in [p] \setminus \{a, \ell\}$. We note that $|S'| = p - 2 \geq 2k - 2$. By definition, $(\overline{A_i}, A_i)$ is a minimum $(S - u_a - u_i, \overline{U})$ -terminal cut for every $i \in [p] \setminus \{a, \ell\}$. Moreover, by Corollary 2.2, we have that $u_i \in A_i \setminus (\cup_{j \in [p] \setminus \{a, i, \ell\}} A_j)$ for every $i \in [p] \setminus \{a, \ell\}$. Hence, the sets U , R , and S' , and the cuts $(\overline{A_i}, A_i)$ for $i \in [p] \setminus \{a, \ell\}$ satisfy the conditions of Theorem 1.5. Therefore, by the first conclusion of Theorem 1.5, there exists a k -partition \mathcal{P}' with

$$|\delta(\mathcal{P}')| \leq \frac{1}{2} \min\{d(H_{S-u_a-u_i}) + d(H_{S-u_a-u_j}) : i, j \in [p] \setminus \{a, \ell\}\}.$$

By assumption, $d(H_{S-u_a-u_b}) < d(U)$ and $b \in [p] \setminus \{a, \ell\}$, so $\min\{d(H_{S-u_a-u_i}) : i \in [p] \setminus \{a, \ell\}\} < d(U)$. Since (U, \overline{U}) is a $(S - u_a - u_i, \overline{U})$ -terminal cut, we have that $d(H_{S-u_a-u_i}) \leq d(U)$ for every $i \in [p] \setminus \{a, \ell\}$. Therefore,

$$\frac{1}{2} \min\{d(H_{S-u_a-u_i}) + d(H_{S-u_a-u_j}) : i, j \in [p] \setminus \{a, \ell\}\} < d(U) = OPT_{k\text{-cut}}.$$

Thus, we have that $|\delta(\mathcal{P}')| < OPT_{k\text{-cut}}$, which is a contradiction. \square

The next two claims will help in arguing the existence of a hyperedge satisfying the conditions of the second conclusion of Theorem 1.5. In particular, we will need Claim 2.5. The following claim will help in proving Claim 2.5.

Claim 2.4. *For every $i, j \in [p]$, we have*

$$d(H_{S-u_i} \cap H_{S-u_j}) = d(U) = d(H_{S-u_i} \cup H_{S-u_j}).$$

Proof. Since $(H_{S-u_i} \cap H_{S-u_j}, \overline{H_{S-u_i} \cap H_{S-u_j}})$ is a $(S - u_i - u_j, \overline{U})$ -terminal cut, we have that $d(H_{S-u_i} \cap H_{S-u_j}) \geq d(H_{S-u_i-u_j})$. By Claim 2.3, we have that $d(H_{S-u_i-u_j}) = d(U) = d(H_{S-u_i})$. Therefore,

$$d(H_{S-u_i} \cap H_{S-u_j}) \geq d(H_{S-u_i}). \quad (4)$$

Since $(H_{S-u_i} \cup H_{S-u_j}, \overline{H_{S-u_i} \cup H_{S-u_j}})$ is a $(S - u_j, \overline{U})$ -terminal cut, we have that

$$d(H_{S-u_i} \cup H_{S-u_j}) \geq d(H_{S-u_j}). \quad (5)$$

By submodularity of the hypergraph cut function and inequalities (4) and (5), we have that

$$d(H_{S-u_i}) + d(H_{S-u_j}) \geq d(H_{S-u_i} \cap H_{S-u_j}) + d(H_{S-u_i} \cup H_{S-u_j}) \geq d(H_{S-u_i}) + d(H_{S-u_j}).$$

Therefore, inequalities (4) and (5) are equations. Thus, by Claim 2.3, we have that

$$d(H_{S-u_i} \cap H_{S-u_j}) = d(H_{S-u_i}) = d(U),$$

and

$$d(H_{S-u_i} \cup H_{S-u_j}) = d(H_{S-u_j}) = d(U).$$

\square

Claim 2.5. For every $i, j, \ell \in [p]$ with $i \neq j$, we have $H_{S-u_\ell} \subseteq H_{S-u_i} \cup H_{S-u_j}$.

Proof. If $\ell = i$ or $\ell = j$ the claim is immediate. Thus, we assume that $\ell \notin \{i, j\}$. Let $Q := H_{S-u_\ell} \setminus (H_{S-u_i} \cup H_{S-u_j})$. We need to show that $Q = \emptyset$. We will show that $(H_{S-u_\ell} \setminus Q, \overline{H_{S-u_\ell} \setminus Q})$ is a minimum $(S - u_\ell, \overline{U})$ -terminal cut. Consequently, Q must be empty (otherwise, $H_{S-u_\ell} \setminus Q \subsetneq H_{S-u_\ell}$ and hence, $(H_{S-u_\ell} \setminus Q, \overline{H_{S-u_\ell} \setminus Q})$ contradicts source minimality of the minimum $(S - u_\ell, \overline{U})$ -terminal cut $(H_{S-u_\ell}, \overline{H_{S-u_\ell}})$).

We now show that $(H_{S-u_\ell} \setminus Q, \overline{H_{S-u_\ell} \setminus Q})$ is a minimum $(S - u_\ell, \overline{U})$ -terminal cut. Since $H_{S-u_\ell} \setminus Q = H_{S-u_\ell} \cap (H_{S-u_i} \cup H_{S-u_j})$, we have that $S - u_i - u_j - u_\ell \subseteq H_{S-u_\ell} \setminus Q$. We also know that u_i and u_j are contained in both H_{S-u_ℓ} and $H_{S-u_i} \cup H_{S-u_j}$. Therefore, $S - u_\ell \subseteq H_{S-u_\ell} \setminus Q$. Thus, $(H_{S-u_\ell} \setminus Q, \overline{H_{S-u_\ell} \setminus Q})$ is a $(S - u_\ell, \overline{U})$ -terminal cut. Therefore,

$$d(H_{S-u_\ell} \cap (H_{S-u_i} \cup H_{S-u_j})) = d(H_{S-u_\ell} \setminus Q) \geq d(H_{S-u_\ell}). \quad (6)$$

We also have that $(H_{S-u_\ell} \cup (H_{S-u_i} \cup H_{S-u_j}), \overline{H_{S-u_\ell} \cup (H_{S-u_i} \cup H_{S-u_j})})$ is a $(S - u_i, \overline{U})$ -terminal cut. Therefore, $d(H_{S-u_\ell} \cup (H_{S-u_i} \cup H_{S-u_j})) \geq d(H_{S-u_i})$. By Claims 2.3 and 2.4, we have that $d(H_{S-u_i}) = d(V_1) = d(H_{S-u_i} \cup H_{S-u_j})$. Therefore,

$$d(H_{S-u_\ell} \cup (H_{S-u_i} \cup H_{S-u_j})) \geq d(H_{S-u_i} \cup H_{S-u_j}). \quad (7)$$

By submodularity of the hypergraph cut function and inequalities (6) and (7), we have that

$$\begin{aligned} d(H_{S-u_\ell}) + d(H_{S-u_i} \cup H_{S-u_j}) &\geq d(H_{S-u_\ell} \cap (H_{S-u_i} \cup H_{S-u_j})) + d(H_{S-u_\ell} \cup (H_{S-u_i} \cup H_{S-u_j})) \\ &\geq d(H_{S-u_\ell}) + d(H_{S-u_i} \cup H_{S-u_j}). \end{aligned}$$

Therefore, inequalities (6) and (7) are equations, so $(H_{S-u_\ell} \setminus Q, \overline{H_{S-u_\ell} \setminus Q})$ is a minimum $(S - u_\ell, \overline{U})$ -terminal cut. \square

Let $R := \{u_p\}$, $S' := S - u_p$, and $(\overline{A_i}, A_i) := (H_{S-u_i}, \overline{H_{S-u_i}})$ for every $i \in [p-1]$. By definition, $(\overline{A_i}, A_i)$ is a minimum $(S - u_i, \overline{U})$ -terminal cut for every $i \in [p-1]$. Moreover, by Corollary 2.1, we have that $u_i \in A_i \setminus (\cup_{j \in [p-1] \setminus \{i\}} A_j)$. Hence, the sets U , R , and S' , and the cuts $(\overline{A_i}, A_i)$ for $i \in [p-1]$ satisfy the conditions of Theorem 1.5. We will use the second conclusion of Theorem 1.5. We now show that there exists a hyperedge satisfying the conditions mentioned in the second conclusion of Theorem 1.5. We will use Claim 2.6 below to prove this. Let $W := \cup_{1 \leq i < j \leq p-1} (A_i \cap A_j)$ and $Z := \cap_{i \in [p-1]} \overline{A_i}$ as in the statement of Theorem 1.5.

Claim 2.6. There exists a hyperedge $e \in E$ such that $e \cap W \neq \emptyset$, $e \cap Z \neq \emptyset$, and $e \subseteq W \cup Z$.

Proof. We note that $S \subseteq (S - u_i) \cup (S - u_j) \subseteq H_{S-u_i} \cup H_{S-u_j}$ for every distinct $i, j \in [p-1]$. Therefore, $S \cap (A_i \cap A_j) = \emptyset$ for every distinct $i, j \in [p-1]$, and thus $S \cap W = \emptyset$. Since S is a transversal of the collection \mathcal{C} , it follows that the set W is not in the collection \mathcal{C} .

By definition, $\overline{U} \subseteq A_i$ for every $i \in [p-1]$, and thus $\overline{U} \subseteq W$. Since $W \notin \mathcal{C}$, either $d(W) > d(U)$ or $\delta(W) = \delta(U)$. By Claim 2.1, we have that $\overline{H_{S-u_p}} \in \mathcal{C}$, and thus, $d(\overline{H_{S-u_p}}) \leq d(U)$ and $\delta(\overline{H_{S-u_p}}) \neq \delta(U)$. Consequently, $d(W) \geq d(\overline{H_{S-u_p}})$, and $\delta(W) \neq \delta(\overline{H_{S-u_p}})$, and thus, $\delta(W) \setminus \delta(\overline{H_{S-u_p}}) \neq \emptyset$. Let $e \in \delta(W) \setminus \delta(\overline{H_{S-u_p}})$. We will show that this choice of e achieves the desired properties.

For each $i \in [p]$, let $Y_i := \overline{H_{S-u_i}} \setminus W$. By Claim 2.5, for every $i, j, \ell \in [p]$ with $i \neq j$ we have that $H_{S-u_\ell} \subseteq H_{S-u_i} \cup H_{S-u_j}$. Therefore $\overline{H_{S-u_i}} \cap \overline{H_{S-u_j}} \subseteq \overline{H_{S-u_\ell}}$ for every such $i, j, \ell \in [p]$, and hence $W \subseteq \overline{H_{S-u_\ell}}$ for every $\ell \in [p]$. Thus, $W \subseteq \overline{H_{S-u_p}}$. Since $e \in \delta(W) \setminus \delta(\overline{H_{S-u_p}})$, we have that $e \subseteq W \cup Y_p$, $e \cap W \neq \emptyset$ and $e \cap Y_p \neq \emptyset$. Therefore, in order to show that e has the three desired properties as in the claim, it suffices to show that $Y_p \subseteq Z$. We prove this next.

By definition, $Y_p \cap W = \emptyset$. By Claim 2.5, for every $i \in [p-1]$, we have that $\overline{H_{S-u_p}} \cap \overline{H_{S-u_i}} \subseteq \overline{H_{S-u_1}}$ and $\overline{H_{S-u_p}} \cap \overline{H_{S-u_i}} \subseteq \overline{H_{S-u_2}}$, so $\overline{H_{S-u_p}} \cap \overline{H_{S-u_i}} \subseteq \overline{H_{S-u_1}} \cap \overline{H_{S-u_2}} \subseteq W$. Thus, for every $i \in [p-1]$, $Y_p \cap Y_i \subseteq \overline{H_{S-u_p}} \cap \overline{H_{S-u_i}} \subseteq W$, so since $Y_p \cap W = \emptyset$, we have that $Y_p \cap Y_i = \emptyset$ for every $i \in [p-1]$. Therefore,

$$Y_p \subseteq W \cup \overline{\left(\bigcup_{i=1}^{p-1} Y_i \right)} = \bigcup_{i=1}^{p-1} \overline{H_{S-u_i}} = \bigcap_{i=1}^{p-1} H_{S-u_i} = Z.$$

□

By Claim 2.6, there is a hyperedge e satisfying the conditions of the second conclusion of Theorem 1.5. Therefore, by Theorem 1.5, there exists a k -partition \mathcal{P}' with

$$\begin{aligned} |\delta(\mathcal{P}')| &< \frac{1}{2} \min\{d(A_i) + d(A_j) : i, j \in [p-1], i \neq j\} \\ &= d(U) && \text{(By Claim 2.3)} \\ &= OPT_{k\text{-cut}}. && \text{(By assumption of the theorem)} \end{aligned}$$

Thus, we have obtained a k -partition \mathcal{P}' with $|\delta(\mathcal{P}')| < OPT_{k\text{-cut}}$, which is a contradiction.

□

3 Proof of Theorem 1.4

We prove Theorem 1.4 in this section. Applying Theorem 2.1 to (\overline{U}, U) yields the following corollary.

Corollary 3.1. *Let $G = (V, E)$ be a hypergraph and let $OPT_{k\text{-cut}}$ be the optimum value of HYPERGRAPH- k -CUT in G for some integer $k \geq 2$. Suppose (U, \overline{U}) is a 2-partition of V with $d(U) = OPT_{k\text{-cut}}$. Then, there exists a set $T \subseteq \overline{U}$ with $|T| \leq 2k - 1$ such that every minimum (U, T) -terminal cut (A, \overline{A}) satisfies $\delta(A) = \delta(U)$.*

We now restate Theorem 1.4 and prove it using Theorem 2.1 and Corollary 3.1.

Theorem 1.4. *Let $G = (V, E)$ be a hypergraph and let $OPT_{k\text{-cut}}$ be the optimum value of HYPERGRAPH- k -CUT in G for some integer $k \geq 2$. Suppose (U, \overline{U}) is a 2-partition of V with $d(U) = OPT_{k\text{-cut}}$. Then, there exist sets $S \subseteq U$, $T \subseteq \overline{U}$ with $|S| \leq 2k - 1$ and $|T| \leq 2k - 1$ such that every minimum (S, T) -terminal cut (A, \overline{A}) satisfies $\delta(A) = \delta(U)$.*

Proof. By Theorem 2.1, there exists a subset $S \subseteq U$ with $|S| \leq 2k - 1$ such that every minimum (S, \overline{U}) -terminal cut (A, \overline{A}) has $\delta(A) = \delta(U)$. By Corollary 3.1, there exists a subset $T \subseteq \overline{U}$ with $|T| \leq 2k - 1$ such that every minimum (U, T) -terminal cut (A, \overline{A}) has $\delta(A) = \delta(U)$. We will show that every minimum (S, T) -terminal cut (A, \overline{A}) has $\delta(A) = \delta(U)$. We will need the following claim.

Claim 3.1. *Let (Y, \overline{Y}) be the source minimal minimum (S, T) -terminal cut. Then $\delta(Y) = \delta(U)$.*

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

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

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

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

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

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

Thus, by the submodularity of the hypergraph cut function we have that

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

Therefore, we have that $d(U \cap Y) = d(U)$, so $(U \cap Y, \overline{U \cap Y})$ is a minimum (S, T) -terminal cut. Since (Y, \overline{Y}) is the source minimal (S, T) -terminal cut, we have that $U \cap Y = Y$, and hence $Y \subseteq U$. Therefore, (Y, \overline{Y}) is a minimum (S, \overline{U}) -terminal cut. By the choice of S , we have that $\delta(Y) = \delta(U)$. \square

Applying Claim 3.1 to both sides of the partition (U, \overline{U}) , we have that the source minimal minimum (S, T) -terminal cut (Y, \overline{Y}) has $\delta(Y) = \delta(U)$, and the source minimal minimum (T, S) -terminal cut (Z, \overline{Z}) has $\delta(Z) = \delta(U)$. Therefore, for every $e \in \delta(U)$, we have that $e \cap Y \neq \emptyset$ and $e \cap Z \neq \emptyset$.

Let (A, \overline{A}) be a minimum (S, T) -terminal cut. Since (Y, \overline{Y}) is the source minimal minimum (S, T) -terminal cut, we have that $Y \subseteq A$. Since (Z, \overline{Z}) is the source minimal minimum (T, S) -terminal cut, we have that $Z \subseteq \overline{A}$. Since every $e \in \delta(U)$ intersects both Y and Z , it follows that every $e \in \delta(U)$ intersects both A and \overline{A} , and hence, $\delta(U) \subseteq \delta(A)$. Since (A, \overline{A}) is a minimum (S, T) -terminal cut, $d(A) \leq d(U)$, and thus we have that $\delta(A) = \delta(U)$. \square

4 Algorithm for ENUM-HYPERGRAPH- k -CUT

In this section, we design a deterministic algorithm for ENUM-HYPERGRAPH- k -CUT that is based on divide and conquer and has a run-time of $n^{O(k)}$ source minimal minimum (s, t) -terminal cut computations, where n is the number of vertices in the input hypergraph. The high-level idea is to use minimum (S, T) -terminal cuts to enumerate a collection of candidate cuts such that for every optimum k -partition for HYPERGRAPH- k -CUT, either the union of some $k/2$ parts of the optimum k -partition is contained in the candidate collection or we find the set of hyperedges crossing this optimum k -partition. This helps in cutting the recursion depth to $\log k$ which saves on overall run-time. We describe the algorithm in Figure 2 and its guarantees in Theorem 4.1. We recall that for a hypergraph $G = (V, E)$ and a subset $A \subseteq V$, the subgraph $G[A]$ induced by A is given by $G[A] = (A, E')$, where $E' := \{e \in E : e \subseteq A\}$.

Theorem 4.1 is a self-contained proof that the number of MIN- k -CUT-SETS in a n -vertex hypergraph is $O(n^{8k \log k})$ and the run-time of the algorithm in Figure 2 is $O(n^{8k \log k})$ source minimal minimum (s, t) -terminal cut computations. In Lemma 4.1, we improve the run-time analysis of the same algorithm to $O(n^{16k})$ source minimal minimum (s, t) -terminal cut computations. For this, we exploit the known fact that the number of MIN- k -CUT-SETS in a n -vertex hypergraph is $O(n^{2k-2})$ (via the randomized algorithm in [14]).

Theorem 4.1 and Lemma 4.1 together imply Theorem 1.2 since the source minimal minimum (s, t) -terminal cut in a n -vertex hypergraph of size p can be computed in time $O(np)$ [18].

Theorem 4.1. *Let $G = (V, E)$ be a n -vertex hypergraph of size p and let k be a positive integer. Then, Algorithm Enum-Cut-Sets(G, k) in Figure 2 returns the family of all MIN- k -CUT-SETS in G and it can be implemented to run in time $O(n^{(8k-6) \log k})T(n, p)$, where $T(n, p)$ denotes the time complexity for computing the source minimal minimum (s, t) -terminal cut in a n -vertex hypergraph of size p . Moreover, the cardinality of the family returned by the algorithm is $O(n^{(8k-6) \log k})$.*

<p>Algorithm Enum-Cut-Sets($G = (V, E), k$)</p> <p>Input: Hypergraph $G = (V, E)$ and an integer $k \geq 1$</p> <p>Output: Family of all MIN-k-CUT-SETS in G</p> <p>If $k = 1$</p> <p style="padding-left: 20px;">Return $\{\emptyset\}$</p> <p>Else</p> <p style="padding-left: 20px;">Initialize $\mathcal{C} \leftarrow \emptyset, \mathcal{F} \leftarrow \emptyset$</p> <p style="padding-left: 20px;">For each pair (S, T) such that $S, T \subseteq V$ with $S \cap T = \emptyset$ and $S , T \leq 2k - 1$</p> <p style="padding-left: 40px;">Compute the source minimal minimum (S, T)-terminal cut (A, \bar{A})</p> <p style="padding-left: 40px;">If $G - \delta(A)$ has at least k connected components</p> <p style="padding-left: 60px;">$\mathcal{F} \leftarrow \mathcal{F} \cup \{\delta(A)\}$</p> <p style="padding-left: 40px;">Else</p> <p style="padding-left: 60px;">$\mathcal{C} \leftarrow \mathcal{C} \cup \{A\}$</p> <p style="padding-left: 20px;">For each $A \in \mathcal{C}$ such that $A \geq \lfloor k/2 \rfloor$ and $\bar{A} \geq k - \lfloor k/2 \rfloor$</p> <p style="padding-left: 40px;">$\mathcal{F}_A \leftarrow \text{DC-Enum-Cuts}(G[A], \lfloor k/2 \rfloor)$</p> <p style="padding-left: 40px;">$\mathcal{F}'_A \leftarrow \text{DC-Enum-Cuts}(G[\bar{A}], k - \lfloor k/2 \rfloor)$</p> <p style="padding-left: 40px;">$\mathcal{F} \leftarrow \mathcal{F} \cup \{\delta(A) \cup F \cup F' : F \in \mathcal{F}_A, F' \in \mathcal{F}'_A\}$</p> <p style="padding-left: 20px;">Among all k-cut-sets in the family \mathcal{F}, return the subfamily that are of smallest size</p>

Figure 2: Divide-and-conquer algorithm to enumerate hypergraph minimum k -cut-sets

Proof. We begin by showing correctness. The last step of the algorithm considers only k -cut-sets in the family \mathcal{F} , so the algorithm returns a subfamily of k -cut-sets. We only have to show that every MIN- k -CUT-SET is in the family \mathcal{F} ; this will also guarantee that every k -cut-set in the returned subfamily is indeed a MIN- k -CUT-SET. We show this by induction on k .

For the base case of $k = 1$, the only MIN- k -CUT-SET is the empty set which is contained in the returned family. We now show the induction step. Assume that $k \geq 2$. Let $F \subseteq E$ be a MIN- k -CUT-SET in G and let (V_1, \dots, V_k) be an optimum k -partition for HYPERGRAPH- k -CUT such that $F = \delta(V_1, \dots, V_k)$. We will show that F is in the family returned by the algorithm. Let $U := \cup_{i=1}^{\lfloor k/2 \rfloor} V_i$. We distinguish between the following two cases:

1. Suppose $d(U) < OPT_{k\text{-cut}}$.

By Theorem 1.3, there exist disjoint subsets $S, T \subseteq V$ with $|S|, |T| \leq 2k - 2$ such that (U, \bar{U}) is the unique minimum (S, T) -terminal cut. Hence, the set U is in the collection \mathcal{C} . Moreover, U contains $\lfloor k/2 \rfloor$ non-empty sets $V_1, V_2, \dots, V_{\lfloor k/2 \rfloor}$, so we have $|U| \geq \lfloor k/2 \rfloor$. Similarly, we have $|\bar{U}| \geq k - \lfloor k/2 \rfloor$. Since (V_1, \dots, V_k) is an optimum k -partition for HYPERGRAPH- k -CUT, the set $\{e \in F \setminus \delta(U) : e \subseteq U\}$ is a MIN- $\lfloor k/2 \rfloor$ -CUT-SET in $G[U]$. Similarly, the set $\{e \in F \setminus \delta(U) : e \subseteq \bar{U}\}$ is a MIN- $(k - \lfloor k/2 \rfloor)$ -CUT-SET in $G[\bar{U}]$. Since $d(U) < OPT_{k\text{-cut}}$, we know that $G - \delta(U)$ has less than k connected components. Therefore, the set U is in the collection \mathcal{C} . By induction hypothesis, we know that the set $\{e \in F \setminus \delta(U) : e \subseteq U\}$ is contained in the family \mathcal{F}_U and the set $\{e \in F \setminus \delta(U) : e \subseteq \bar{U}\}$ is contained in the family \mathcal{F}'_U . Therefore, the set F is added to the family \mathcal{F} in the second for-loop.

2. Suppose $d(U) = OPT_{k\text{-cut}}$.

By Theorem 1.4, there exist sets $S \subseteq U$ and $T \subseteq \bar{U}$ with $|S|, |T| \leq 2k - 1$ such that the source minimal minimum (S, T) -terminal cut (A, \bar{A}) satisfies $\delta(A) = \delta(U) = F$. Therefore, the set A is in the collection \mathcal{C} . Since $F = \delta(A)$, the hypergraph $G - \delta(A)$ contains at least k connected components. Therefore, the set $F = \delta(A)$ is added to the family \mathcal{F} in the first for-loop.

Thus, in both cases, we have shown that the set F is contained in the family \mathcal{F} . Since the algorithm

returns the subfamily of hyperedge sets in \mathcal{F} that are MIN- k -CUT-SETS, the set F is in the family returned by the algorithm.

Next, we bound the run time of the algorithm. Let $N(k, n)$ denote the run time of the algorithm for a n -vertex hypergraph. Then, we have $N(1, n) = O(1)$. For $k \geq 2$, there are $O(n^{4k-2})$ pairs of subsets $S, T \subseteq V$ with $|S|, |T| \leq 2k - 1$ and $S \cap T = \emptyset$. Hence, the first for-loop performs $O(n^{4k-2})$ source minimal minimum (S, T) -terminal cut computations. The collection \mathcal{C} and the family \mathcal{F} at the end of the first for-loop each have $O(n^{4k-2})$ sets. This implies that the first for-loop can be implemented to run in $O(n^{4k-2})T(n, p)$ time. For each $A \in \mathcal{C}$, the computation of $\text{DC-Enum-Cuts}(G[A], \lfloor k/2 \rfloor)$ in the second for-loop runs in $N(\lfloor k/2 \rfloor, n)$ time. The computation of $\text{DC-Enum-Cuts}(G[\bar{A}], k - \lfloor k/2 \rfloor)$ in the second for-loop runs in $N(k - \lfloor k/2 \rfloor, n)$ time. Hence, the second for-loop can be implemented to run in $O(n^{4k-2})N(\lfloor k/2 \rfloor, n)N(k - \lfloor k/2 \rfloor, n)$ time. The last step to prune the family \mathcal{F} can be implemented to run in time that is linear in the time to implement the first and second for-loops: this is because for each member of \mathcal{F} , we can decide whether it is a minimum k -cut-set in time linear in the time to write this member in \mathcal{F} . Therefore, we have

$$N(k, n) = O\left(n^{4k-2}\right)T(n, p) + O\left(n^{4k-2}\right)N\left(\left\lfloor \frac{k}{2} \right\rfloor, n\right)N\left(k - \left\lfloor \frac{k}{2} \right\rfloor, n\right).$$

Since $N(1, n) = O(1)$, we have that $N(k, n) = O(n^{(8k-6)\log k})T(n, p)$.

Finally, we bound the cardinality of the family returned by the algorithm. Let $f(k, n)$ be the cardinality of the family returned by the algorithm for a n -vertex hypergraph. We note that $f(k, n)$ is at most the cardinality of the family \mathcal{F} computed by the algorithm. There are $O(n^{4k-2})$ pairs of subsets $S, T \subseteq V$ with $|S|, |T| \leq 2k - 1$ and $S \cap T = \emptyset$. Hence, the total cardinality of the collection \mathcal{C} and the family \mathcal{F} at the end of the first for-loop is $O(n^{4k-2})$. Consequently, for $k \geq 2$, by the recursion, we have that

$$f(k, n) = O(n^{4k-2})f\left(\left\lfloor \frac{k}{2} \right\rfloor, n\right)f\left(k - \left\lfloor \frac{k}{2} \right\rfloor, n\right)$$

and $f(1, n) = 1$. So, $f(k, n) = O(n^{(8k-6)\log k})$. □

We recall that the number of MIN- k -CUT-SETS in a n -vertex hypergraph is $O(n^{2k-2})$ [14]. Assuming this bound improves the run-time of Algorithm $\text{DC-Enum-Cuts}(G, k)$ in Figure 2.

Lemma 4.1. *Algorithm $\text{DC-Enum-Cuts}(G, k)$ in Figure 2 can be implemented to run in time $O(n^{16k-26})T(n, p)$, where n is the number of vertices, p is the size of the input hypergraph G , and $T(n, p)$ denotes the time complexity for computing the source minimal minimum (s, t) -terminal cut in a n -vertex hypergraph of size p .*

Proof. Let $N(k, n)$ denote the run time of the algorithm for a n -vertex hypergraph. Then, we have $N(1, n) = O(1)$. For $k \geq 2$, there are $O(n^{4k-2})$ pairs of subsets $S, T \subseteq V$ with $|S|, |T| \leq 2k - 1$ and $S \cap T = \emptyset$. Hence, the first for-loop performs $O(n^{4k-2})$ source minimal minimum (S, T) -terminal cut computations. The collection \mathcal{C} and the family \mathcal{F} at the end of the first for-loop each have $O(n^{4k-2})$ sets. This implies that the first for-loop can be implemented to run in $O(n^{4k-2})T(n, p)$ time. For each $A \in \mathcal{C}$, the computation of $\text{DC-Enum-Cuts}(G[A], \lfloor k/2 \rfloor)$ in the second for-loop runs in $N(\lfloor k/2 \rfloor, n)$ time. The computation of $\text{DC-Enum-Cuts}(G[\bar{A}], k - \lfloor k/2 \rfloor)$ in the second for-loop runs in $N(k - \lfloor k/2 \rfloor, n)$ time. We recall that \mathcal{F}_A consists of all minimum $\lfloor k/2 \rfloor$ -cut-sets in a n -vertex

graph and hence, has size $O(n^{2\lfloor k/2 \rfloor - 2})$. Similarly, \mathcal{F}'_A has size $O(n^{2(k - \lfloor k/2 \rfloor) - 2})$. Hence, the second for-loop can be implemented to run in time

$$\begin{aligned} & O\left(n^{4k-2}\right) \left(N\left(\left\lfloor \frac{k}{2} \right\rfloor, n\right) + N\left(k - \left\lfloor \frac{k}{2} \right\rfloor, n\right) + O(n^{2\lfloor k/2 \rfloor - 2})O(n^{2(k - \lfloor k/2 \rfloor) - 2}) \right) \\ &= O\left(n^{4k-2}\right) \left(N\left(\left\lfloor \frac{k}{2} \right\rfloor, n\right) + N\left(k - \left\lfloor \frac{k}{2} \right\rfloor, n\right) + O(n^{2k-4}) \right). \end{aligned}$$

Moreover, the size of the family \mathcal{F} at the end of the second for-loop is $O(n^{4k-2}) + O(n^{4k-2})|\mathcal{F}_A| \cdot |\mathcal{F}'_A| = O(n^{-2})O(n^{2k-4}) = O(n^{6k-6})$. Hence, the last step to prune the family \mathcal{F} can be implemented to run in time $O(n^{6k-6})$. Hence, the second for-loop and the last step can together be implemented to run in time

$$O\left(n^{4k-2}\right) \left(N\left(\left\lfloor \frac{k}{2} \right\rfloor, n\right) + N\left(k - \left\lfloor \frac{k}{2} \right\rfloor, n\right) + O(n^{2k-4}) \right).$$

Therefore, we have

$$N(k, n) = O\left(n^{4k-2}\right) \left(T(n, p) + N\left(\left\lfloor \frac{k}{2} \right\rfloor, n\right) + N\left(k - \left\lfloor \frac{k}{2} \right\rfloor, n\right) + O(n^{2k-4}) \right).$$

Solving the recursive relation gives $N(k, n) = O(n^{16k-26})T(n, p)$. \square

5 Algorithm for ENUM-MINMAX-HYPERGRAPH- k -PARTITION

In this section, we design a deterministic algorithm for ENUM-MINMAX-HYPERGRAPH- k -PARTITION that runs in time $n^{O(k^2)}p$, where n is the number of vertices and p is the size of the input hypergraph. For this, we rely on the notion of k -cut-set representatives.

We recall that for a k -partition (V_1, \dots, V_k) and disjoint subsets $U_1, \dots, U_k \subseteq V$, the k -tuple (U_1, \dots, U_k) is defined to be a k -cut-set representative of (V_1, \dots, V_k) if $U_i \subseteq V_i$ and $\delta(U_i) = \delta(V_i)$ for all $i \in [k]$. We first show that there exists a polynomial-time algorithm to verify whether a given k -tuple (U_1, \dots, U_k) is a k -cut-set representative.

Theorem 5.1. *Let $G = (V, E)$ be a n -vertex hypergraph of size p and let k be a positive integer. Then, there exists an algorithm that takes as input the hypergraph G and disjoint subsets $U_1, \dots, U_k \subseteq V$ and runs in time $O(knp)$ to decide if (U_1, \dots, U_k) is a k -cut-set representative of some k -partition (V_1, \dots, V_k) and if so, then return such a k -partition.*

Proof. We will use Algorithm Recover-Partition(G, U_1, \dots, U_k) in Figure 3.

We begin by showing correctness. Since Algorithm 3 maintains $U_i \subseteq P_i$ and $\delta(U_i) = \delta(P_i)$ for all $i \in [k]$, if it returns a k -partition, then the k -partition necessarily satisfies the required conditions. Next, we show that if (U_1, \dots, U_k) is a k -cut-set representative of a k -partition (V_1, \dots, V_k) , then the algorithm will indeed return a k -partition (P_1, \dots, P_k) with $U_i \subseteq P_i$ and $\delta(U_i) = \delta(P_i)$ for all $i \in [k]$ (however, (P_1, \dots, P_k) may not necessarily be the same as (V_1, \dots, V_k)).

Let (V_1, \dots, V_k) be a k -partition such that $U_i \subseteq V_i$ and $\delta(U_i) = \delta(V_i)$ for all $i \in [k]$. Let (P_1, \dots, P_k) be the sequence of subsets at the end of the for-loop. Moreover, for each $j \in [t]$, let (P_1^j, \dots, P_k^j) be the sequence of subsets at the end of the j th iteration of the for-loop. For notational convenience, for $i \in [k]$, we will define $P_i^0 := U_i$. We note that $(P_1, \dots, P_k) = (P_1^t, \dots, P_k^t)$ and that $P_i^0 \subseteq P_i^1 \subseteq \dots \subseteq P_i^t$ for every $i \in [k]$. We observe that $U_i \subseteq P_i^j$ and $\delta(U_i) = \delta(P_i^j)$ for all $j \in \{0, 1, 2, \dots, t\}$. Moreover, the subsets P_1^j, \dots, P_k^j are pairwise disjoint for each $j \in \{0, 1, 2, \dots, t\}$. Therefore, it suffices to show that $\cup_{i=1}^k P_i = V$.

<p>Algorithm Recover-Partition($G = (V, E), U_1, \dots, U_k$)</p> <p>Input: Hypergraph $G = (V, E)$ and disjoint subsets $U_1, \dots, U_k \subseteq V$</p> <p>Output: Decide if there exists a k-partition (V_1, \dots, V_k) of V with $U_i \subseteq V_i$ and $\delta(U_i) = \delta(V_i) \forall i \in [k]$, and return one if it exists</p> <p>Initialize $P_i \leftarrow U_i$ for all $i \in [k]$</p> <p>Let $C_1, \dots, C_t \subseteq V$ be the components of $G - \cup_{i=1}^k \delta(U_i)$ that are disjoint from $\cup_{i=1}^k U_i$</p> <p>For $j = 1, \dots, t$</p> <p style="padding-left: 20px;">If $\exists i \in [k]$ such that $\delta(P_i \cup C_j) = \delta(P_i)$</p> <p style="padding-left: 40px;">$P_i \leftarrow P_i \cup C_j$</p> <p>If (P_1, \dots, P_k) is a k-partition of V</p> <p style="padding-left: 20px;">Return (P_1, \dots, P_k)</p> <p>Else</p> <p style="padding-left: 20px;">Return NO</p>
--

Figure 3: Algorithm in Theorem 5.1

We claim that $C_1 \cup \dots \cup C_j \subseteq \cup_{i=1}^k P_i^j$ for each $j \in \{0, 1, \dots, t\}$. Applying this claim for $j = t$ gives that $\cup_{i=1}^k P_i = V$ as desired. We now show the claim by induction on j . The base case of $j = 0$ holds by definition. We now prove the induction step. By induction hypothesis, we have that $C_1 \cup \dots \cup C_{j-1} \subseteq \cup_{i=1}^k P_i^{j-1}$. We will show that there exists $i \in [k]$ such that $\delta(P_i^{j-1} \cup C_j) = \delta(P_i^{j-1})$. We know that C_j is contained in one of the sets in $\{V_1, \dots, V_k\}$, say $C_j \subseteq V_\ell$ for some $\ell \in [k]$. We will prove that $\delta(P_\ell^{j-1} \cup C_j) = \delta(P_\ell^{j-1})$ to complete the proof of the claim. Since C_j is a component of $G - \cup_{i=1}^k \delta(U_i) = G - \cup_{i=1}^k \delta(V_i)$, we know that each hyperedge in $\delta(C_j)$ crosses the k -partition (V_1, \dots, V_k) . Moreover, each hyperedge in $\delta(C_j)$ intersects $C_j \subseteq V_\ell$, and hence, $\delta(C_j) \subseteq \delta(V_\ell)$. Therefore,

$$\delta(P_\ell^{j-1} \cup C_j) - \delta(P_\ell^{j-1}) \subseteq \delta(C_j) \subseteq \delta(V_\ell) = \delta(U_\ell) = \delta(P_\ell^{j-1}).$$

This is possible only if $\delta(P_\ell^{j-1} \cup C_j) - \delta(P_\ell^{j-1}) = \emptyset$, i.e., $\delta(P_\ell^{j-1} \cup C_j) \subseteq \delta(P_\ell^{j-1})$. We now show the reverse inclusion. We have that

$$\begin{aligned} \delta(P_\ell^{j-1}) - \delta(P_\ell^{j-1} \cup C_j) &= \delta(U_\ell) - \delta(P_\ell^{j-1} \cup C_j) \\ &= \delta(V_\ell) - \delta(P_\ell^{j-1} \cup C_j) \\ &= E(V_\ell - P_\ell^{j-1} - C_j, V - V_\ell - P_\ell^{j-1}) \cup E(V_\ell \cap P_\ell^{j-1}, P_\ell^{j-1} - V_\ell) \\ &\subseteq E[V - P_\ell^{j-1}] \cup E[P_\ell^{j-1}]. \end{aligned}$$

We note that the LHS is a subset of $\delta(P_\ell^{j-1})$ while the RHS is disjoint from $\delta(P_\ell^{j-1})$ since $E[V - P_\ell^{j-1}] \cap \delta(P_\ell^{j-1}) = \emptyset$ and $E[P_\ell^{j-1}] \cap \delta(P_\ell^{j-1}) = \emptyset$. Hence, the above containment is possible only if $\delta(P_\ell^{j-1}) - \delta(P_\ell^{j-1} \cup C_j) = \emptyset$ and hence, $\delta(P_\ell^{j-1}) \subseteq \delta(P_\ell^{j-1} \cup C_j)$. Consequently, $\delta(P_\ell^{j-1} \cup C_j) = \delta(P_\ell^{j-1})$.

We now bound the run-time. We can verify if there exists $i \in [k]$ such that $\delta(P_i \cup C_j) = \delta(P_i)$ in time $O(kp)$. The number of iterations of the for-loop is $t \leq n$. Hence, the total run-time is $O(knp)$. \square

Next, we address the problem of enumerating all MINMAX- k -CUT-SETS. For this, we define a sub-problem—namely ENUM-MINMAX-HYPERGRAPH- k -CUT-SET-REPS. The input here is a hypergraph $G = (V, E)$ and a fixed positive integer k (e.g., $k = 2, 3, 4, \dots$). The goal is to enumerate a family \mathcal{F} of k -cut-set representatives satisfying the following two properties:

- (1) every k -tuple (U_1, \dots, U_k) in the family \mathcal{F} is a k -cut-set representative of some optimum k -partition (V_1, \dots, V_k) for MINMAX-HYPERGRAPH- k -PARTITION and

(2) for every optimum k -partition (V_1, \dots, V_k) for MINMAX-HYPERGRAPH- k -PARTITION, the family \mathcal{F} contains a k -cut-set representative (U_1, \dots, U_k) of (V_1, \dots, V_k) .

We note that if a family \mathcal{F} is a solution to ENUM-MINMAX-HYPERGRAPH- k -CUT-SET-REPS, then returning $\{\cup_{i=1}^k \delta(U_i) : (U_1, \dots, U_k) \in \mathcal{F}\}$ solves ENUM-MINMAX-HYPERGRAPH- k -PARTITION. Hence, it suffices to solve ENUM-MINMAX-HYPERGRAPH- k -CUT-SET-REPS in order to solve ENUM-MINMAX-HYPERGRAPH- k -PARTITION. We describe our algorithm for ENUM-MINMAX-HYPERGRAPH- k -CUT-SET-REPS in Figure 4 and its guarantees in Theorem 5.2. Theorem 1.1 follows from Theorem 5.2.

Theorem 5.2. *Let $G = (V, E)$ be a n -vertex hypergraph of size p and let k be a positive integer. Then, Algorithm Enum-MinMax-Reps(G, k) in Figure 4 solves ENUM-MINMAX-HYPERGRAPH- k -CUT-SET-REPS and it can be implemented to run in time $O(kn^{4k^2-2k+1}p)$. Moreover, the cardinality of the family returned by the algorithm is $O(n^{4k^2-2k})$.*

Algorithm Enum-MinMax-Reps($G = (V, E), k$)
Input: Hypergraph $G = (V, E)$ and an integer $k \geq 2$
Output: Family \mathcal{F} of k -cut-set representatives of all optimum k -partitions for MINMAX-HYPERGRAPH- k -PARTITION
Initialize $\mathcal{C} \leftarrow \emptyset$, $\mathcal{D} \leftarrow \emptyset$, and $\mathcal{F} \leftarrow \emptyset$
For each pair (S, T) such that $S, T \subseteq V$ with $S \cap T = \emptyset$ and $|S|, |T| \leq 2k - 1$
 Compute the source minimal minimum (S, T) -terminal cut (U, \bar{U})
 $\mathcal{C} \leftarrow \mathcal{C} \cup \{U\}$
For all $(U_1, \dots, U_k) \in \mathcal{C}^k$ such that U_1, \dots, U_k are pairwise disjoint
 If Recover-Partition(G, U_1, \dots, U_k) returns a k -partition
 $\mathcal{D} \leftarrow \mathcal{D} \cup \{(U_1, \dots, U_k)\}$
 $\lambda \leftarrow \min\{\max_{i \in [k]} d(U_i) : (U_1, \dots, U_k) \in \mathcal{D}\}$
For all $(U_1, \dots, U_k) \in \mathcal{D}$ such that $\max_{i \in [k]} d(U_i) = \lambda$:
 $\mathcal{F} \leftarrow \mathcal{F} \cup \{(U_1, \dots, U_k)\}$
Return \mathcal{F}

Figure 4: Algorithm in Theorem 5.2

Proof. We begin by showing correctness—i.e., the family \mathcal{F} returned by the algorithm satisfies properties (1) and (2) mentioned in the definition of ENUM-MINMAX-HYPERGRAPH- k -CUT-SET-REPS. By the second for-loop, each k -tuple added to the collection \mathcal{D} is a k -cut-set representative of some k -partition (it need not necessarily be a k -cut-set representative of an optimum k -partition for MINMAX-HYPERGRAPH- k -PARTITION). The algorithm returns a subfamily of \mathcal{D} and hence, it returns a subfamily of k -cut-set representatives. We only have to show that a k -cut-set representative of an arbitrary optimum k -partition for MINMAX-HYPERGRAPH- k -PARTITION is present in the family \mathcal{D} ; this will guarantee that the value λ computed by the algorithm will exactly be $OPT_{\text{minmax-}k\text{-partition}}$ and owing to the way in which the algorithm constructs the family \mathcal{F} from the family \mathcal{D} , it follows that the family \mathcal{F} satisfies properties (1) and (2).

Let $OPT_{\text{minmax-}k\text{-partition}}$ denote the optimum value of a minmax k -partition in G and let $OPT_{k\text{-cut}}$ denote the optimum value of a minimum k -cut in G . We note that $OPT_{\text{minmax-}k\text{-partition}} \leq OPT_{k\text{-cut}}$. This is because, if (P_1, \dots, P_k) is a k -partition with minimum $|\delta(P_1, \dots, P_k)|$ (i.e., an optimum k -partition for HYPERGRAPH- k -CUT), then

$$OPT_{\text{minmax-}k\text{-partition}} \leq \max_{i \in [k]} |\delta(P_i)| \leq |\delta(P_1, \dots, P_k)| = OPT_{k\text{-cut}}.$$

Let (V_1, \dots, V_k) be an arbitrary optimum k -partition for MINMAX-HYPERGRAPH- k -PARTITION. We will show that the family \mathcal{F} returned by the algorithm contains a k -cut-set representative of (V_1, \dots, V_k) . We have that $d(V_i) \leq OPT_{\text{minmax-}k\text{-partition}} \leq OPT_{k\text{-cut}}$ for all $i \in [k]$. Hence, by Theorems 1.3 and 1.4, there exist subsets $S_i \subseteq V_i, T_i \subseteq V - V_i$ with $|S_i|, |T_i| \leq 2k - 1$ such that the source minimal minimum (S_i, T_i) -terminal cut (U_i, \bar{U}_i) satisfies $\delta(U_i) = \delta(V_i)$ for all $i \in [k]$. Source minimality of the cut (U_i, \bar{U}_i) also guarantees that $U_i \subseteq V_i$ for all $i \in [k]$. Hence, the k -tuple (U_1, \dots, U_k) is a k -cut-set representative of (V_1, \dots, V_k) . It remains to show that this k -tuple is indeed present in the families \mathcal{D} and \mathcal{F} . We note that the sets U_1, \dots, U_k are added to the collection \mathcal{C} in the first for-loop. Since the k -tuple (U_1, \dots, U_k) is a k -cut-set representative of the k -partition (V_1, \dots, V_k) , the k -tuple (U_1, \dots, U_k) will be added to the family \mathcal{D} in the second for-loop. Since the family \mathcal{D} contains only k -cut-set representatives of k -partitions, it follows that $\lambda = OPT_{\text{minmax-}k\text{-partition}}$ and (U_1, \dots, U_k) will be added to the family \mathcal{F} in the third for-loop. Hence, the k -cut-set representative (U_1, \dots, U_k) of the optimum k -partition (V_1, \dots, V_k) for MINMAX-HYPERGRAPH- k -PARTITION is present in the family \mathcal{F} returned by the algorithm.

The bound on the size of the family \mathcal{F} returned by the algorithm is

$$|\mathcal{F}| \leq |\mathcal{D}| \leq |\mathcal{C}|^k = O(n^{k(4k-2)}).$$

Next, we bound the run time of the algorithm. The first for-loop can be implemented to run in time $O(n^{4k-2})T(n, p)$. The second for-loop executes the algorithm from Theorem 5.1 $O(n^{4k^2-2k})$ times and hence, the second for-loop can be implemented to run in time $O(kn^{4k^2-2k+1}p)$. The computation of λ and the third for-loop can be implemented to run in time $O(|\mathcal{D}|) = O(n^{4k^2-2k})$. Hence, the total run-time is $O(n^{4k-2})T(n, p) + O(kn^{4k^2-2k+1}p)$. We recall that $T(n, p) = O(np)$ and hence, the total run-time is $O(kn^{4k^2-2k+1}p)$. □

6 A lower bound on the number of MINMAX- k -CUT-SETS

In this section, we show that there exist n -vertex connected graphs for which the number of MINMAX- k -CUT-SETS is $n^{\Omega(k^2)}$. In particular, we show the following result.

Lemma 6.1. *For every positive integer $k \geq 2$, there exists a positive integer n such that the number of optimum k -partitions for MINMAX-GRAPH- k -PARTITION in the n -vertex complete graph is $n^{\Omega(k^2)}$.*

Proof. Let $k \geq 2$ be fixed, and let $G = (V, E)$ be the complete graph on $n = k(k-1)$ vertices (with all edge weights being uniformly 1). We will show that $OPT_{\text{minmax-}k\text{-partition}} = (k-1)^3$ and every partition of V into k parts of equal size is an optimum k -partition for MINMAX-GRAPH- k -PARTITION. Since the number of partitions of V into k parts of equal size is $\Omega(k^n) = \Omega(n^{k^2/2})$, the lemma follows.

First, we show that $OPT_{\text{minmax-}k\text{-partition}} \geq (k-1)^3$. For every partition of V into k non-empty parts, the largest part has at least $k-1$ vertices by pigeonhole principle, and at most $k(k-1) - (k-1) = (k-1)^2$ vertices since each of the remaining $k-1$ parts contain at least one vertex. Therefore, the cut value of the largest part is at least

$$\min_{x \in \{k-1, \dots, (k-1)^2\}} x(n-x) = (k-1)^3.$$

The equality follows since $n = k(k-1)$ and the function $f(x) = x(n-x)$ is convex and is minimized at the boundaries. This implies that $OPT_{\text{minmax-}k\text{-partition}} \geq (k-1)^3$.

Next, we show that $OPT_{\text{minmax-}k\text{-partition}} \leq (k-1)^3$ and every partition of V into k parts of equal size is an optimum k -partition for MINMAX-GRAPH- k -PARTITION. Let (V_1, \dots, V_k) be an arbitrary k -partition of V such that $|V_i| = k-1$ for all $i \in [k]$. The MINMAX-GRAPH- k -PARTITION objective value of this k -partition is $(k-1)(n-(k-1)) = (k-1)^3$. Thus, (V_1, \dots, V_k) is an optimum k -partition for MINMAX-GRAPH- k -PARTITION in G . \square

We note that our example exhibiting $n^{\Omega(k^2)}$ optimum k -partitions for MINMAX-GRAPH- k -PARTITION has the number of vertices n upper bounded by a function of k . We are not aware of examples that exhibit $n^{\Omega(k^2)}$ optimum k -partitions for MINMAX-GRAPH- k -PARTITION for fixed k but arbitrary n (e.g., $k = 2, 3, 4, \dots$ but n is arbitrary).

7 Conclusion

We showed the first polynomial bound on the number of MINMAX- k -CUT-SETS in hypergraphs for every fixed k and gave a polynomial-time algorithm to enumerate all MINMAX- k -CUT-SETS as well as all MIN- k -CUT-SETS in hypergraphs for every fixed k . Our main contribution is a structural theorem that is the backbone of the correctness analysis of our enumeration algorithms. In order to enumerate MINMAX- k -CUT-SETS in hypergraphs, we introduced the notion of k -cut-set representatives and enumerated k -cut-set representatives of all optimum k -partitions for MINMAX-HYPERGRAPH- k -PARTITION. Our technique builds on known structural results for HYPERGRAPH- k -CUT and MINMAX-HYPERGRAPH- k -PARTITION [7, 11, 12].

The technique underlying our enumeration algorithms is not necessarily novel—we simply rely on minimum (s, t) -terminal cuts. Using fixed-terminal cuts to address global partitioning problems is not a novel technique by itself—it is common knowledge that minimum (s, t) -terminal cuts can be used to solve global minimum cut. However, there are several problems where naive use of this technique fails to lead to efficient algorithms: e.g., multiway cut does not help in solving GRAPH- k -CUT since multiway cut is NP-hard. Adapting this technique for specific partitioning problems requires careful identification of structural properties. In fact, beautiful structural properties have been shown for a rich variety of partitioning problems in combinatorial optimization in order to exploit this technique: for example, it was used (1) to design the first efficient algorithm for GRAPH- k -CUT [26], (2) to solve certain constrained submodular minimization problems [25, 49], and (3) more recently, to design fast algorithms for global minimum cut in graphs and for Gomory-Hu tree in unweighted graphs [1, 44]. Our use of this technique also relies on identifying and proving a suitable structural property, namely Theorem 1.4. The advantage of our structural property is that it simultaneously enables enumeration of MIN- k -CUT-SETS as well as MINMAX- k -CUT-SETS in hypergraphs which was not possible via structural theorems that were developed before. Furthermore, it helps in showing the first polynomial bound on the number of MINMAX- k -CUT-SETS in hypergraphs for every fixed k .

We also emphasize a limitation of our technique. Although it helps in solving ENUM-HYPERGRAPH- k -CUT and ENUM-MINMAX-HYPERGRAPH- k -PARTITION, it does not help in solving a seemingly related hypergraph k -partitioning problem—namely, given a hypergraph $G = (V, E)$ and a fixed integer k , find a k -partition (V_1, \dots, V_k) of the vertex set that minimizes $\sum_{i=1}^k |\delta(V_i)|$. Natural variants of our structural theorem fail to hold for this objective. Resolving the complexity of this variant of the hypergraph k -partitioning problem for $k \geq 5$ remains open.

We mention an open question concerning HYPERGRAPH- k -CUT and the enumeration of MIN- k -CUT-SETS in hypergraphs for fixed k . We recall the status in graphs: the number of minimum k -partitions in a connected graph was known to be $O(n^{2k-2})$ via Karger-Stein’s algorithm [38] and

$\Omega(n^k)$ via the cycle example, where n is the number of vertices; recent works have improved on the upper bound to match the lower bound for fixed k —this improvement in upper bound also led to the best possible $O(n^k)$ -time algorithm for GRAPH- k -CUT for fixed k [28, 31, 32]. For hypergraphs, the number of MIN- k -CUT-SETS is known to be $O(n^{2k-2})$ and $\Omega(n^k)$. Can we improve the upper/lower bound? Is it possible to design an algorithm for HYPERGRAPH- k -CUT that runs in time $O(n^k p)$?

References

- [1] A. Abboud, R. Krauthgamer, and O. Trabelsi, *Subcubic algorithms for gomory–hu tree in unweighted graphs*, Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing, STOC, 2021, p. 1725–1737.
- [2] 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.
- [3] 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.
- [4] ———, *Graph sketches: Sparsification, spanners, and subgraphs*, Proceedings of the 31st Symposium on Principles of Database Systems, PODS, 2012, pp. 5–14.
- [5] D. Applegate, R. Bixby, V. Chvátal, and W. Cook, *The Traveling Salesman Problem: A Computational Study*, Princeton University Press, 2006.
- [6] N. Bansal, O. Svensson, and L. Trevisan, *New notions and constructions of sparsification for graphs and hypergraphs*, Proceedings of the 60th Annual IEEE Symposium on Foundations of Computer Science, FOCS, 2019, pp. 910–928.
- [7] C. Beideman, K. Chandrasekaran, and W. Wang, *Deterministic enumeration of all minimum k -cut-sets in hypergraphs for fixed k* , Proceedings of the 33rd Annual ACM-SIAM Symposium on Discrete Algorithms, SODA, 2022.
- [8] 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.
- [9] ———, *Cut structures and randomized algorithms in edge-connectivity problems*, Ph.D. thesis, MIT, 1997.
- [10] 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.
- [11] K. Chandrasekaran and C. Chekuri, *Hypergraph k -cut for fixed k in deterministic polynomial time*, Proceedings of the 61st Annual Symposium on Foundations of Computer Science, FOCS, 2020, pp. 810–821.
- [12] ———, *Min-max partitioning of hypergraphs and symmetric submodular functions*, Proceedings of the 32nd annual ACM-SIAM Symposium on Discrete Algorithms, SODA, 2021, pp. 1026–1038.

- [13] K. Chandrasekaran and W. Wang, *Fixed Parameter Approximation Scheme for Min-max k -cut*, Integer Programming and Combinatorial Optimization, IPCO, 2021, pp. 354–367.
- [14] K. Chandrasekaran, C. Xu, and X. Yu, *Hypergraph k -cut in randomized polynomial time*, Mathematical Programming (Preliminary version in SODA 2018) **186** (2019), 85–113.
- [15] K. Chandrasekaran and X. Zhang, Private communication.
- [16] C. Chekuri and S. Li, *On the hardness of approximating the k -way hypergraph cut problem*, Theory of Computing **16** (2020), no. 14, 1–8.
- [17] 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.
- [18] C. Chekuri and C. Xu, *Minimum cuts and sparsification in hypergraphs*, SIAM Journal on Computing **47** (2018), no. 6, 2118–2156.
- [19] Y. Chen, S. Khanna, and A. Nagda, *Near-linear size hypergraph cut sparsifiers*, Proceedings of the 61st Annual IEEE Symposium on Foundations of Computer Science, FOCS, 2020, pp. 61–72.
- [20] W. Cook, W. Cunningham, W. Pulleyblank, and A. Schrijver, *Combinatorial Optimization*, Wiley, 1997.
- [21] E. A. Dinitz, A. V. Karzanov, and M. V. Lomonosov, *On the structure of a family of minimum weighted cuts in a graph*, Studies in Discrete Optimization (A. A. Fridman, ed.), Nauka Publishers, 1976.
- [22] R. Downey, V. Estivill-Castro, M. Fellows, E. Prieto, and F. Rosamund, *Cutting up is hard to do: The parameterised complexity of k -cut and related problems*, Electronic Notes in Theoretical Computer Science **78** (2003), 209–222.
- [23] K. Fox, D. Panigrahi, and F. Zhang, *Minimum cut and minimum k -cut in hypergraphs via branching contractions*, Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA, 2019, pp. 881–896.
- [24] M. Ghaffari, D. Karger, and D. Panigrahi, *Random contractions and sampling for hypergraph and hedge connectivity*, Proceedings of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA, 2017, p. 1101–1114.
- [25] M. X. Goemans and V. S. Ramakrishnan, *Minimizing submodular functions over families of sets*, Combinatorica **15** (1995), 499–513.
- [26] O. Goldschmidt and D. Hochbaum, *A Polynomial Algorithm for the k -cut Problem for Fixed k* , Mathematics of Operations Research (Preliminary version in FOCS 1988) **19** (1994), no. 1, 24–37.
- [27] F. Guinez and M. Queyranne, *The size-constrained submodular k -partition problem*, Unpublished manuscript. Available at <https://docs.google.com/viewer?a=v&pid=sites&srcid=ZGVmYXVsdGRvbWFpbnxmbGF2aW9ndWluZXpob21lcGFnZXxneDoONDVlMThkMDg4ZWRlOGI1>. See also <https://smartech.gatech.edu/bitstream/handle/1853/43309/Queyranne.pdf>, 2012.
- [28] A. Gupta, D. Harris, E. Lee, and J. Li, *Optimal Bounds for the k -cut Problem*, Preprint in arXiv: 2005.08301, 2020.

- [29] A. Gupta, E. Lee, and J. Li, *An FPT Algorithm Beating 2-Approximation for k -Cut*, Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA, 2018, pp. 2821–2837.
- [30] ———, *Faster exact and approximate algorithms for k -cut*, Proceedings of the 59th IEEE annual Symposium on Foundations of Computer Science, FOCS, 2018, pp. 113–123.
- [31] ———, *The number of minimum k -cuts: improving the Karger-Stein bound*, Proceedings of the 51st ACM Symposium on Theory of Computing, STOC, 2019, pp. 229–240.
- [32] ———, *The Karger-Stein algorithm is optimal for k -cut*, Proceedings of the 52nd Annual ACM Symposium on Theory of Computing, STOC, 2020, pp. 473–484.
- [33] M. Henzinger and D. Williamson, *On the number of small cuts in a graph*, Information Processing Letters **59** (1996), 41–44.
- [34] 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.
- [35] M. Kapralov, R. Krauthgamer, J. Tardos, and Y. Yoshida, *Towards tight bounds for spectral sparsification of hypergraphs*, Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing, STOC, 2021, pp. 598–611.
- [36] D. Karger, *Global min-cuts in RNC, and other ramifications of a simple min-cut algorithm*, Proceedings of the 4th annual ACM-SIAM Symposium on Discrete Algorithms, SODA, 1993, pp. 21–30.
- [37] ———, *Minimum cuts in near-linear time*, Journal of the ACM **47** (2000), no. 1, 46–76.
- [38] D. Karger and C. Stein, *A new approach to the minimum cut problem*, Journal of the ACM **43** (1996), no. 4, 601–640.
- [39] A. Karlin, N. Klein, and S-O. Gharan, *A (slightly) improved approximation algorithm for metric tsp*, Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing, STOC, 2021, p. 32–45.
- [40] K. Kawarabayashi and M. Thorup, *The minimum k -way cut of bounded size is fixed-parameter tractable*, Proceedings of the 52nd Annual Symposium on Foundations of Computer Science, FOCS, 2011, pp. 160–169.
- [41] K-i Kawarabayashi and B. Lin, *A nearly $5/3$ -approximation FPT Algorithm for Min- k -Cut*, Proceedings of the 31st ACM-SIAM Symposium on Discrete Algorithms, SODA, 2020, pp. 990–999.
- [42] D. Kogan and R. Krauthgamer, *Sketching cuts in graphs and hypergraphs*, Proceedings of the 2015 Conference on Innovations in Theoretical Computer Science, ITCS '15, 2015, pp. 367–376.
- [43] E. Lawler, *Cutsets and Partitions of Hypergraphs*, Networks **3** (1973), 275–285.
- [44] J. Li and D. Panigrahi, *Deterministic Min-cut in Poly-logarithmic Max-flows*, Proceedings of the 61st Annual Symposium on Foundations of Computer Science, FOCS, 2020, pp. 85–92.

- [45] D. Lokshtanov, S. Saurabh, and V. Surianarayanan, *A Parameterized Approximation Scheme for MIN k -CUT*, Proceedings of the 61st IEEE annual Symposium on Foundations of Computer Science, FOCS, 2020, pp. 798–809.
- [46] P. Manurangsi, *Almost-polynomial Ratio ETH-hardness of Approximating Densest k -subgraph*, Proceedings of the 49th Annual ACM Symposium on Theory of Computing, STOC, 2017, pp. 954–961.
- [47] ———, *Inapproximability of Maximum Biclique Problems, Minimum k -Cut and Densest At-Least- k -Subgraph from the Small Set Expansion Hypothesis*, Algorithms **11**(1) (2018), 10.
- [48] 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.
- [49] M. Nägele, B. Sudakov, and R. Zenklusen, *Submodular minimization under congruency constraints*, Combinatorica **39** (2019), 1351–1386.
- [50] K. Okumoto, T. Fukunaga, and H. Nagamochi, *Divide-and-conquer algorithms for partitioning hypergraphs and submodular systems*, Algorithmica **62** (2012), no. 3, 787–806.
- [51] K. Quanrud, *Fast and Deterministic Approximations for k -Cut*, Proceedings of Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX, 2019, pp. 23:1–23:20.
- [52] M. Queyranne, *On optimum size-constrained set partitions*, 1999, Talk at Aussiois workshop on Combinatorial Optimization.
- [53] P. Raghavendra and D. Steurer, *Graph Expansion and the Unique Games Conjecture*, Proceedings of the 42nd Annual ACM Symposium on Theory of Computing, STOC, 2010, pp. 755–764.
- [54] R. Ravi and A. Sinha, *Approximating k -cuts using network strength as a lagrangean relaxation*, European Journal of Operational Research **186** (2008), no. 1, 77–90.
- [55] H. Saran and V. Vazirani, *Finding k Cuts within Twice the Optimal*, SIAM Journal on Computing **24** (1995), no. 1, 101–108.
- [56] T. Soma and Y. Yoshida, *Spectral sparsification of hypergraphs*, Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA, 2019, pp. 2570–2581.
- [57] M. Thorup, *Minimum k -way Cuts via Deterministic Greedy Tree Packing*, Proceedings of the 40th Annual ACM Symposium on Theory of Computing, STOC, 2008, pp. 159–166.
- [58] M. Xiao, *An Improved Divide-and-Conquer Algorithm for Finding All Minimum k -Way Cuts*, Proceedings of 19th International Symposium on Algorithms and Computation, ISAAC, 2008, pp. 208–219.
- [59] 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.
- [60] L. Zhao, *Approximation algorithms for partition and design problems in networks*, Ph.D. thesis, Graduate School of Informatics, Kyoto University, Japan, 2002.
- [61] L. Zhao, H. Nagamochi, and T. Ibaraki, *Greedy splitting algorithms for approximating multiway partition problems*, Mathematical Programming **102** (2005), no. 1, 167–183.