

Strongly polynomial bounds for multiobjective and parametric global minimum cuts in graphs and hypergraphs

Hassene Aissi¹ · A. Ridha Mahjoub¹ ·
S. Thomas McCormick² · Maurice Queyranne^{2,3}

Received: 30 April 2014 / Accepted: 13 August 2015 / Published online: 19 September 2015
© Springer-Verlag Berlin Heidelberg and Mathematical Optimization Society 2015

Abstract We consider multiobjective and parametric versions of the global minimum cut problem in undirected graphs and bounded-rank hypergraphs with multiple edge cost functions. For a fixed number of edge cost functions, we show that the total number of supported non-dominated (SND) cuts is bounded by a polynomial in the numbers of nodes and edges, i.e., is strongly polynomial. This bound also applies to the combinatorial facet complexity of the problem, i.e., the maximum number of facets (linear pieces) of the parametric curve for the parametrized (linear combination) objective, over the set of all parameter vectors such that the parametrized edge costs are nonnegative and the parametrized cut costs are positive. We sharpen this bound in the case of two objectives (the bicriteria problem), for which we also derive a strongly polynomial upper bound on the total number of non-dominated (Pareto optimal) cuts. In particular, the bicriteria global minimum cut problem in an n -node graph admits $O(n^3 \log n)$ SND cuts and $O(n^5 \log n)$ Pareto optimal cuts. These results significantly

This paper derives from the conference paper [1], substantially extended, re-written, and corrected.

✉ Hassene Aissi
aissi@lamsade.dauphine.fr

A. Ridha Mahjoub
mahjoub@lamsade.dauphine.fr

S. Thomas McCormick
tom.mccormick@sauder.ubc.ca

Maurice Queyranne
maurice.queyranne@sauder.ubc.ca

¹ PSL, LAMSADE, Université Paris-Dauphine, Paris, France

² Sauder School of Business at the University of British Columbia, Vancouver, BC, Canada

³ CORE, Université Catholique de Louvain, Louvain, Belgium

improve on earlier graph cut results by Mulmuley (SIAM J Comput 28(4):1460–1509, 1999) and Armon and Zwick (Algorithmica 46(1):15–26, 2006). They also imply that the parametric curve and all SND cuts, and, for the bicriteria problems, all Pareto optimal cuts, can be computed in strongly polynomial time when the number of objectives is fixed.

Keywords Multiobjective optimization · Parametric optimization · Global minimum cut

Mathematics Subject Classification Multiobjective optimization 90C29 · Parametric optimization 90C31 · Global minimum cut 05C40, 05C85

1 Introduction

We consider the multicriteria version of the global minimum cut problem in undirected graphs. Global minimum cut is extensively studied in combinatorial optimization since many practical problems in, e.g., routing (subtour elimination), communications and electrical networks, contain it as a subproblem [2]. Let $G = (V, E)$ be an undirected graph, and $c^1, \dots, c^k : E \rightarrow \mathbb{R}$ be k cost functions, or *criteria*, defined on its edges. A *cut* X in G is a subset of nodes $X \subseteq V$ such that $\emptyset \neq X \neq V$, and it determines the set $\delta(X)$ of edges with exactly one end in X . The *cost* of cut X w.r.t. criterion j is $c^j(X) := c^j(\delta(X))$.

Many concepts and algorithms for global minimum cut generalize to hypergraphs.¹ We recall that a *hypergraph* is a finite set of vertices V , together with a family E of subsets of V . Thus each $e \in E$ is a vertex subset $e \subseteq V$, and is called a hyperedge, or simply, an edge. We say that hypergraph $G = (V, E)$ is *rank- ρ* if $|e| \leq \rho$ for all $e \in E$. Thus an undirected graph is a rank-2 hypergraph, and is rank- ρ for all $\rho \geq 2$. A *cut* X in hypergraph G is a non-trivial node subset, i.e., $\emptyset \neq X \subset V$. It cuts the set of edges $\delta(X) := \{e \in E : e \cap X \neq \emptyset \neq e \setminus X\}$. Notice that this matches the definition of a cut for a graph. A hypergraph is *connected* if all cuts X are non-empty, i.e., $\delta(X) \neq \emptyset$ for every cut X . Given edge costs $c^j(e)$ ($e \in E$, $j = 1, \dots, k$), the cost of a cut X w.r.t. c^j is the total cost $c^j(X) := c^j(\delta(X)) = \sum_{e \in \delta(X)} c^j(e)$ of all the edges crossed by cut X .

Ideally we would like a cut that simultaneously minimizes all criteria, but such a solution usually does not exist. Therefore, we focus on *Pareto optimal solutions*, i.e., solutions that cannot be improved upon in any criterion without degrading another criterion. Each cut X is associated with its criteria vector (or point) $c(X) := (c^1(X), \dots, c^k(X))$ in the criteria space \mathbb{R}^k . Let \mathcal{C} denote the set of all cuts, and $Y := \{c(X) : X \in \mathcal{C}\}$ denote the image in the criterion space of all the cuts (note that different cuts may give rise to the same criteria point). A point $c(X')$ *dominates*

¹ A reader only interested in graph cuts may skip the rest of this paragraph and the whole of Sect. 2; and in the sequel replace all occurrences of “(bounded-rank) hypergraph(s)” with “graph(s)”, all $O_{\rho, \cdot}(\cdot)$ with $O(\cdot)$, and ignore all mentions of rank $\rho \geq 3$.

$c(X)$, and by extension, cut X' dominates X , if $c^j(X') \leq c^j(X)$ for all $j = 1, \dots, k$, and $c^j(X') < c^j(X)$ for at least one j . If there is no $c(X') \in Y$ that dominates $c(X)$, then the point $c(X)$, and by extension, the cut X itself, is *non-dominated*, or *Pareto optimal*. In order to derive strongly polynomial bounds and algorithms, when we study the *multicriteria minimum cut problem* we assume that all edge costs are nonnegative, as negative edge costs may give rise to NP-hard minimization problems; and that all cuts have strictly positive costs, i.e., that $c^j(X) > 0$ for every objective j and cut X , for otherwise the hypergraph $(V, E^{>0,j})$ induced by the positive-cost edges (i.e., with edge set $E^{>0,j} = \{e \in E : c^j(e) > 0\}$) is not connected and may have an exponential number of minimum (zero-cost) cuts.

The computation of *Pareto optimal points* (i.e. non-dominated points) is related to the field of *parametric optimization*. For $j = 1, \dots, k - 1$ we put a non-negative parameter μ_j on criterion j , and we consider the parametric optimization problem

$$c^*(\mu) := \min_{X \in \mathcal{C}} c_\mu(X) \quad \text{where} \quad c_\mu(X) := \sum_{j=1}^{k-1} \mu_j c^j(X) + c^k(X). \tag{1}$$

Notice that a multicriteria problem with k objectives corresponds to a parametric problem with a single objective and $k - 1$ parameters. We call c_μ the *parametrized objective*. If there is some positive parameter vector $\mu \in \mathbb{R}_{>0}^{k-1}$ (i.e., with all components $\mu_j > 0$) such that cut X minimizes c_μ then $c(X)$ is Pareto optimal, we call it a *supported non-dominated* (SND) point, and the cut X an *SND cut*. The Pareto optimal points that are not SND points are called *unsupported non-dominated* (UND) points, see Fig. 1.

A natural object to define is the *dominant* \mathcal{D} of the convex hull of Y , i.e., the set of all points $x \in \mathbb{R}^k$ that satisfy $x \geq y$ for some y in the convex hull of Y . The boundary of \mathcal{D} (its lower convex hull) is the set of all points in \mathcal{D} “accessible” as optimal solutions for parametrized objectives c_μ with $\mu \geq 0$. Thus all SND points must lie on the boundary of \mathcal{D} , and every vertex of \mathcal{D} is actually an SND point (when the costs are in “general position”, the converse will also be true, namely, that every SND point is a vertex of \mathcal{D} . Otherwise, non-vertex SND points may arise as alternate optima for certain parameter vectors μ). The optimum cost function $c^* : \mathbb{R}_+^{k-1} \mapsto \mathbb{R}$ is piecewise linear and concave, and its graph is in a sense the convex dual of the boundary of \mathcal{D} (which is piecewise linear and convex, e.g., [8, 10]). This duality interchanges vertices and facets, and so the fact that the vertices of \mathcal{D} are SND points means that the facets of the graph of c^* also correspond to SND points. See Fig. 1 for an example of this for $k = 2$.

In fact, our results apply to a slightly more general parametric minimum cut problem than arises from the multicriteria version, wherein we do not require either the objectives or the parameters to be nonnegative. However, in order to derive strongly polynomial bounds and algorithms, when we study the *parametric minimum cut problem* we make two assumptions on the parametric cost values that are similar to the assumption on the cost function values for the multicriteria version: we assume that all parametrized edge costs $c_\mu(e) := \sum_{j=1}^{k-1} \mu_j c^j(e) + c^k(e)$ are nonnegative (though a single-criterion cost function $c^j(e)$ may be negative), and that all cuts have positive




Fig. 1 **a** is an example of SND points, UND points, and \mathcal{D} . Each point corresponds to a cut X with values $(c^1(X), c^2(X))$. SND points are *black* and are labeled C_1, C_2, \dots . UND points are *gray*, and dominated points are *white*. **b** Plots $c^*(\mu)$ for the same data. The extreme points and line segments of both graphs are labeled such that extreme point C_1 of **a** corresponds to line segment C_1 of the **b**, extreme point b_2 of the chart **a** is an example of b_2 of the chart **b**, etc.

parametrized costs, i.e., that $c^*(\mu) > 0$. Thus we restrict attention to the following *relevant region* in the parameter space:

$$M = \left\{ \mu \in \mathbb{R}^{k-1} : c_\mu(e) \geq 0 \forall e \in E, \text{ and } c_\mu(X) > 0 \forall X \subset V, X \neq \emptyset \right\}.$$

Note that M is a convex polyhedral set, defined by a finite system of strict and non-strict linear inequalities. Readers who find this slight generalization confusing can imagine, as in the multicriteria version, that every objective c^j is nonnegative and its minimum is positive (i.e., its support graph is connected), in which case M is just the nonnegative orthant \mathbb{R}_+^{k-1} .

Define \mathcal{S} as the set of cuts $X \in \mathcal{C}$ which minimize $c_\mu(X)$ for some $\mu \in M$; i.e., \mathcal{S} is the set of all SND cuts.

We define the *combinatorial facet complexity*² of the parametric minimum cut problem to be the maximum number of facets of the graph of c^* , which is equivalent to the maximum number of vertices of \mathcal{D} , and to the maximum number of SND points. Our main interest here is to study the combinatorial facet complexity of global minimum cut for hypergraphs, and thus also for graphs. We also derive sharper bounds on combinatorial facet complexity for the bicriterion case $k = 2$.

A natural subproblem of parametric minimum cut is solving single-criterion (ordinary) minimum cut, e.g., for some fixed value of μ . For graphs, the fastest deterministic algorithms for this problem run in $O(|E| \cdot |V| + |V|^2 \log |V|)$ time (Nagamochi and Ibaraki [30], and Stoer and Wagner [42]). The fastest randomized algorithm runs in $O(|E| \log^3 |V|)$ time (Karger [17]). These algorithms are faster than minimum s - t -cut algorithms that are based on network flows. See [31] for a detailed treatment of graph connectivity problems. For hypergraphs there exist polynomial time algorithms for finding a minimum cost cut in a hypergraph with nonnegative edge costs, see [20, 25, 37].

The multicriteria versions of several combinatorial optimization problems have been extensively studied, see Ehrgott [8] for a comprehensive survey. These problems are often *intractable* in the sense that the size of the set of (supported) Pareto optimal points grows exponentially in the input size. Furthermore, it is often hard even to verify if a given point is Pareto optimal. For example, Carstensen [3] shows that the combinatorial facet complexity of the minimum s - t -cut problem in a graph (namely, to find a minimum cost cut X that *separates* two given vertices s and t , in the sense that $|X \cap \{s, t\}| = 1$) is exponential, even for a single parameter (two cost functions). Mulmuley [28] gives a simpler proof of this result. In addition, it follows from Papadimitriou and Yannakakis [34, Theorem 6] that the *decision version* of the bicriteria minimum s - t -cut problem in a graph is strongly NP-hard.

² This is also called *parametric complexity* by some authors, such as Mulmuley [28] (see also [24, 27]) for the bicriterion case. We use a different terminology to avoid conflict with a different concept of parametric (or “parametrized”, or “parameterized”) complexity, e.g., [7], which deals with finding efficient algorithms for problems in which certain input or output “parameters” (or properties) are fixed. On the other hand, Fernández-Baca and Venkatachalam [10] use the term *combinatorial complexity* to refer to the total number of *faces* of all dimensions (here, 0 to $k - 1$) of the graph of a parametric function (such as c^* here), whereas Schrijver [39] uses *facet complexity* to refer to the maximum input size of a rational linear inequality in a system that defines a polyhedron (such as \mathcal{D} here).

Multicriteria *global* minimum cut in graphs are an exception to such intractability results. Mulmuley [28, Theorem 3.10] considers the combinatorial facet complexity of the parametric minimum cut problem with two objectives that may be negative, but where all parametric edge costs are positive (a slight restriction, for the bicriteria case, of our more general parametric framework above). His Theorem 3.10 implies that the combinatorial facet complexity is $O(|V|^{19} \log |V| \log c_{\max})$, where $c_{\max} := \max_e \max\{c^1(e), c^2(e)\}$ is the maximum edge cost. This is a weakly polynomial bound, but Mulmuley also claims (without detailed proof) that the bit sizes of the costs can be assumed to have polynomial size without affecting the combinatorial facet complexity, and this yields a strongly polynomial bound. In addition, Armon and Zwick [2] show that, when the number k of criteria is fixed, the decision version of the multicriteria minimum cut problem in graphs can be solved in strongly polynomial time. They also point out the existence, when k is fixed, of a pseudo-polynomial time algorithm to find all the Pareto optimal points, and of a fully-polynomial time approximation scheme (FPTAS) for finding an *approximate Pareto set*, a notion defined by Papadimitriou and Yannakakis [34]. All these results may be surprising since the single-objective global minimum cut problem may be solved by solving $|V| - 1$ minimum s - t -cut problems (e.g., by fixing s and letting t vary over the other nodes). Thus, for example, the parametric function c^* is the pointwise minimum of $|V| - 1$ parametric minimum s - t cut functions, each of which possibly has an exponential number of breakpoints.

All these positive results rely on the deep and far-reaching fact that the single-criterion global minimum cut problem on graphs has at most a strongly polynomial number of near-optimal solutions. Indeed, given $\alpha \geq 1$, call a cut α -*approximate* if its cost is less than α times the minimum. For global minimum cut in graphs, Dinitz et al. [6] showed that there are at most $\binom{|V|}{2} = \theta(|V|^2)$ minimum (i.e., 1-approximate) cuts, and that this bound is tight. Nagamochi et al. [33] and Henzinger and Williamson [14] show that this bound also applies to the number of α -approximate cuts for all $\alpha < 4/3$ and all $\alpha < 3/2$, respectively; the latter authors also show that $3/2$ is the largest approximation factor for which this result holds. More generally, Karger [16] states (see [19] for a detailed proof) that for every $\alpha > 1$ the number of α -approximate cuts is $O(|V|^{2\alpha})$. All these results are subsumed by Karger's result [17] that for every $\alpha \geq 1$ the number of α -approximate cuts in a graph is $O(|V|^{\lceil 2\alpha \rceil})$.

In this paper, we give a strongly polynomial upper bound on the combinatorial facet complexity of global minimum cut on rank- ρ hypergraphs. Our bound hides a constant that depends at least exponentially on ρ . For the case of two nonnegative objectives (with positive minimum cut values), we also derive sharper bounds that are much smaller, and follow from a simpler proof, than in [28].

As pointed out above, our bounds of the combinatorial facet complexity of global minimum cut depend on a strongly polynomial upper bound on the number of approximate optimal solutions to global minimum cut. Hence our first step, in Sect. 2, is to generalize the results on approximate global minimum cut from graphs to hypergraphs. In particular, we extend a recent result of Kogan and Krauthgamer [21]; a proof of our extended result is found in the Appendix. Given this, we derive in Sect. 3 a strongly polynomial bound on the combinatorial facet complexity of global minimum cut on hypergraphs. Then Sect. 4 studies in detail the bicriteria case, i.e., when $k = 2$. Finally,

Sect. 5 shows how to adapt existing parametric algorithms for computing the graph of the parametric function c^* , the boundary of \mathcal{D} (SND points and cuts), and for $k = 2$, the set of all Pareto optimal cuts.

2 Approximate minimum cuts for a single objective

In this section we consider a hypergraph $G = (V, E)$ with a single, nonnegative edge cost function c that has positive minimum cut value (without loss of generality we can discard all zero cost edges, and then G has positive minimum cut value if and only if it is connected). We assume that all edges $e \in E$ have size $2 \leq |e| \leq |V| - 1$, as edges with $|e| = 1$ or $|V|$ do not contribute to any cut. We also assume that distinct edges in E span different sets of vertices, and thus that $|E| = O(|V|^\rho)$, since if G contains “parallel” edges that span the same set of vertices, we may replace them by a single edge with cost equal to the sum of the costs of the parallel edges.

Theorem 1 (Adapted from Kogan and Krauthgamer [21]) *For every fixed scalar $\alpha \geq 1$ and integer $\rho \geq 2$, the number of α -approximate cuts in all rank- ρ hypergraphs $G = (V, E)$ with nonnegative edge costs c and positive minimum cut cost, is $O(|V|^{2\alpha})$.*

While Kogan and Krauthgamer only present this result for the case when α is half-integer,³ the proof in the Appendix shows that it actually holds for all $\alpha \geq 1$.

Karger’s $O(|V|^{\lfloor 2\alpha \rfloor})$ bound for graphs [17] is tighter than that in Theorem 1 when $\rho = 2$. When combined with a simple approximation of hypergraph cuts as graph cuts, it also leads to tighter bounds for $\rho = 3$, and for $\rho = 4$ and a small (but important in the sequel) range of values of α . To show this, in the next Proposition, we recall a general approximation method (see, e.g., Ihler et al. [15]) and derive consequences for bounding the number of α -approximate cuts.

Proposition 2 *For every fixed integer $\rho \geq 3$, fixed scalar $\alpha \geq 1$, and rank- ρ hypergraph $G = (V, E)$ with nonnegative edge costs c and positive minimum cut cost, the number of α -approximate cuts is $O(|V|^{\lfloor 2B(\rho)\alpha \rfloor})$ where $B(\rho) = 1$ if $\rho = 3$, and $B(\rho) = \lfloor \rho^2/4 \rfloor / (\rho - 1) \leq \frac{\rho}{4} + \frac{1}{3}$ if $\rho \geq 4$.*

Proof We prove the result by approximating the minimum cut problem in hypergraph $G = (V, E)$ with edge costs c , by the minimum cut problem in the complete graph $K(V) = (V, E_{K(V)})$ with edge costs \tilde{c} , whereby each edge $e \in E$ is replaced by a clique on the nodes in e and each edge in the clique has cost $c(e)/(|e| - 1)$, i.e., by letting

$$\tilde{c}(i, j) = \sum_{e \in E: \{i, j\} \subseteq e} \frac{c(e)}{|e| - 1} \quad \text{for all } \{i, j\} \in E_{K(V)}. \quad (2)$$

In particular, every cardinality-2 edge $e = \{i, j\} \in E$ contributes its full cost $c(e)$ to $\tilde{c}(i, j)$, and thus to the cost $\tilde{c}(X)$ of every cut X in $K(V)$ that crosses it. Note also that

³ Their approach also seems to require that $\alpha\rho$ be integer, but this requirement is not mentioned in [21].

every cardinality-3 edge $e \in E$ that is crossed by cut X has two of its nodes on one side of the cut and the other node on the other side, and thus also contributes its exact cost $2(c(e)/2) = c(e)$ to $\tilde{c}(X)$. Therefore [as it is well known, e.g., Ihler et al. [15]], when $\rho \leq 3$ this transformation preserves all cut values, i.e., $\tilde{c}(X) = c(X)$ for every cut X . The result then follows from Karger's bound when $\rho \leq 3$.

Now assume that $\rho \geq 4$. A cut X that crosses an edge $e \in E$ with cardinality $|e| \geq 4$ crosses at least $|e| - 1$ edges in the clique $K(e)$ (when exactly one node of e is on one side of the cut), and at most $\lfloor |e|^2/4 \rfloor$ such edges (when half the nodes of e are on either side). Thus every edge $e \in E$ crossed by X contributes at least $c(e)$ and at most $(\lfloor |e|^2/4 \rfloor) \frac{c(e)}{|e|-1} \leq \frac{\lfloor \rho^2/4 \rfloor}{\rho-1} c(e) = B(\rho)c(e)$ to the cost $\tilde{c}(X)$. Therefore,

$$c(X) \leq \tilde{c}(X) \leq B(\rho) c(X).$$

Note that

$$B(\rho) \leq \left(\frac{\rho}{4} + \frac{1}{4} \left(1 + \frac{1}{\rho-1} \right) \right) \leq \frac{\rho}{4} + \frac{1}{3},$$

where the last inequality follows from $\rho \geq 4$.

Let X^* denote a minimum cut for $(K(V), \tilde{c})$. If X is an α -optimal cut for (G, c) , i.e., $c(X) \leq \alpha \min_{Y \in \mathcal{C}} c(Y)$, then we have

$$\tilde{c}(X) \leq B(\rho) c(X) \leq B(\rho) \alpha \min_{Y \in \mathcal{C}} c(Y) \leq B(\rho) \alpha c(X^*) \leq B(\rho) \alpha \tilde{c}(X^*),$$

implying that X is a $(B(\rho)\alpha)$ -optimal cut for $(K(V), \tilde{c})$. The result then again follows from Karger's bound when $\rho \geq 4$. \square

Combining Theorem 1 and Proposition 2 (and noting that when $\rho = 4$ and $1 < \alpha < 9/8$, $\lfloor 2B(\rho)\alpha \rfloor = 2 = \lfloor 2\alpha \rfloor$), we obtain:

Corollary 3 *For every fixed scalar $\alpha \geq 1$ and fixed integer $\rho \geq 2$, the number of α -approximate cuts in all rank- ρ hypergraphs $G = (V, E)$ with nonnegative edge costs c and positive minimum cut cost, is*

$$\begin{aligned} & O\left(|V|^{\lfloor 2\alpha \rfloor}\right) \text{ for } \rho \leq 3, \text{ and} \\ & \quad \text{for } \rho = 4 \text{ when } 1 \leq \alpha < 9/8; \\ & O\left(|V|^{2\alpha}\right) \text{ for } \rho \geq 4 \text{ when } \alpha \geq 9/8, \text{ and} \\ & \quad \text{for all } \rho \geq 5. \end{aligned}$$

These results show that the dependence of the number of α -approximate cuts is polynomially bounded in the number $|V|$ of vertices when α and ρ are fixed. On the other hand, it is known (see, e.g., an example in [14] and its immediate generalization to larger half-integer values of α), that the dependence on α is exponential. Similarly,

Kogan and Krauthgamer present in [21] a family of instances showing that the dependence on ρ (through the “hidden” multiplicative constant in the $O(|V|^{2\alpha})$ expression) is at least exponential.

3 Minimum cuts with multiple edge costs

We are now given a hypergraph $G = (V, E)$, and k edge cost functions $c^1, \dots, c^k : E \rightarrow \mathbb{R}$, where $k \geq 2$ is fixed. Before considering the multiojective version of the problem, we start with its parametric version. Recall that for $\mu \in \mathbb{R}^{k-1}$ the parametrized cut cost is $c_\mu(X) = \sum_{e \in \delta(X)} c_\mu(e)$ for all cuts X in G .

Our strategy for proving the following theorem is to express the parameter space M first as a union of regions defined by edges, then to cut each region into subregions via a hyperplane dissection. The region for edge a is the part of M where some minimum cut has a as a maximum cost edge w.r.t. c_μ . The hyperplanes belong to a family of hyperplanes constructed from $c_\mu(a)$ and at successive relative distance β for some appropriate $\beta > 1$, and from $c_\mu(e)$ for every other edge e . The resulting subregions are small enough that we can apply Corollary 3 to get a small number of minimum cuts per region, and then we sum up across all subregions to get our bound.

In view of Corollary 3, we introduce the following notation: given a number $r \in \mathbb{R} \cup \{\pm\infty\}$, and real-valued functions f_ρ and g , where f_ρ (and possibly g) depends on a fixed parameter ρ , we say that

$$f_\rho(n) = O_{\rho,r}(g(n)) \text{ iff } \begin{aligned} f_\rho(n) &= O(g(n)) \text{ when } \rho < r; \text{ and} \\ f_\rho(n) &= O(g(n)n^\epsilon) \text{ for all } \epsilon > 0, \text{ when } \rho \geq r. \end{aligned}$$

In particular, $O_{\rho,+\infty}(g(n))$ is just $O(g(n))$.⁴ Corollary 3 thus implies that, when $\epsilon > 0$ approaches 0, the number of $(1 + \epsilon)$ -approximate cuts in all rank- ρ hypergraphs $G = (V, E)$ with nonnegative edge costs c and positive minimum cut cost, approaches $O_{\rho,5}(|V|^2)$, i.e., $O(|V|^2)$ if $\rho \leq 4$ and $|V|^{2+o(1)}$ if $\rho \geq 5$.

Theorem 4 *Given a fixed integer $\rho \geq 2$, a rank- ρ hypergraph $G = (V, E)$ and k edge cost functions c^1, \dots, c^k , the total number of cuts X minimizing $c_\mu(X)$ over all μ in M is $O_{\rho,5}(|E|^k |V|^2 \log^{k-1} |V|)$.*

Proof For every cut X in G let

$$M(X) = \left\{ \mu \in M : c_\mu(X) = \min_W \{c_\mu(W) : \emptyset \neq W \subset V\} \right\}$$

⁴ This is also a slight (and parametrized) extension of the more familiar \tilde{O} notation. Indeed, recall that $f(n) = \tilde{O}(g(n))$ if and only if $f(n) = O(g(n) \log^p g(n))$ for some $p \geq 1$. Then $\tilde{O}(g(n))$ is $O_{\rho,-\infty}(g(n))$, but the converse need not hold, e.g., when $f(n) = g(n)h(n)$ with $h(n)$ “super-polylog” [i.e., $h(n) = \Omega(\log^p n)$ for all fixed $p \geq 1$] and $h(n) = O_{\rho,-\infty}(1)$.

denote the subset of parameter vectors in M for which X is a minimum cut. Recall that \mathcal{S} is the set of all SND cuts. For every edge $a \in E$ let

$$\mathcal{S}_a = \left\{ X \in \mathcal{S} : a \in \arg \max_{e \in X} c_\mu(e) \text{ for some } \mu \in M(X) \right\}$$

denote the set of all cuts X such that, for some μ for which X is a minimum cut, a is an edge in X with largest cost $c_\mu(a)$. It suffices to show that $|\mathcal{S}_a| = O_{\rho,5}(|E|^{k-1}|V|^2 \log^{k-1}|V|)$ and the result will follow.

Thus, in the rest of this proof, fix edge $a \in E$ such that $\mathcal{S}_a \neq \emptyset$. For every $X \in \mathcal{S}_a$ let

$$M_a(X) = \left\{ \mu \in M(X) : a \in \arg \max_{e \in X} c_\mu(e) \right\}$$

so that $M_a(X) \neq \emptyset$, and let

$$M_a = \bigcup_{X \in \mathcal{S}_a} M_a(X)$$

denote the set of parameter vectors for which edge a has maximum cost in a minimum cut. We will cover the parameter region M_a with a number $L = O(|E|^{k-1} \log^{k-1}|V|)$ of regions R_l ($l = 1, \dots, L$). We will then derive, for each region R_l , an $O_{\rho,5}(|V|^2)$ bound on the total number of cuts W in \mathcal{S}_a that are minimum for some $\mu \in R_l \cap M_a$, i.e., such that $R_l \cap M_a(W) \neq \emptyset$. The covering condition $M_a = \bigcup_{l=1}^L R_l$ implies $\mathcal{S}_a = \bigcup_{l=1}^L \{W \in \mathcal{S}_a : R_l \cap M_a(W) \neq \emptyset\}$, and the result will follow.

First, note that for every $\mu \in M_a$ and cut W such that $a \in \arg \max_{e \in W} c_\mu(e)$ we have

$$c_\mu(a) \leq c_\mu(W) \leq |E|c_\mu(a). \tag{3}$$

Since $c_\mu(W) > 0$ for every $W \in \mathcal{S}_a \neq \emptyset$, this implies $c_\mu(a) > 0$.

Now fix a constant $\beta > 1$ whose value will be specified later, and let $\alpha = \frac{\beta-1}{|E|} > 0$. Compute $p = 1 + \lceil \log \frac{|E|^2}{\beta-1} / \log \beta \rceil$ so that $\alpha\beta^{p-1} > |E|$, and observe that $p = O(\log |V|)$ (since $p = O(\log |E|)$, and $|E| = O(|V|^\rho)$ due to no parallel edges, and ρ being constant).

Define the functions $g_i : \mathbb{R}^{k-1} \rightarrow \mathbb{R}$ by $g_0(\mu) = 0$ and

$$g_i(\mu) = \alpha \beta^{i-1} c_\mu(a) \quad \text{for } i = 1, \dots, p, \quad \text{and}$$

$$g_{p+1}(\mu) = +\infty.$$

Inequality (3) implies that, for every $\mu \in M_a$, every cut W such that $a \in \arg \max_{e \in W} c_\mu(e)$, and every $e \in W$,

$$c_\mu(a) \leq c_\mu(W) \leq |E|c_\mu(a) < g_p(\mu). \tag{4}$$

Since $c_\mu(a) > 0$,

$$0 = g_0(\mu) < g_1(\mu) < \dots < g_p(\mu) < g_{p+1}(\mu) = +\infty \tag{5}$$

for all $\mu \in M_a$.

For each $e \in E$ and each $i = 1, \dots, p$, the hyperplane $H_{i,e} = \{\mu \in \mathbb{R}^{k-1} : g_i(\mu) = c_\mu(e)\}$ divides \mathbb{R}^{k-1} into at most two regions, one where $g_i(\mu) \leq c_\mu(e)$, and the other where $g_i(\mu) \geq c_\mu(e)$.

Recall that, for fixed dimension d, N hyperplanes divide the d -dimensional space \mathbb{R}^d into $O(N^d)$ regions ([41], see also [45] for further references). Thus the $O(|E| \log(|V|))$ hyperplanes $H_{i,e}$ define regions R_1, \dots, R_L that cover the set $M_a \subset \mathbb{R}^{k-1}$, and $L = O(|E|^{k-1} \log^{k-1} |V|)$. By (5), for each $l \in \{1, \dots, L\}$ and each $e \in E$ there exists an index $i(e, l) \in \{0, 1, \dots, p\}$ such that

$$R_l \cap M_a = \{\mu \in M_a : g_{i(e,l)}(\mu) \leq c_\mu(e) \leq g_{i(e,l)+1}(\mu) \forall e \in E\}.$$

Consider any such region R_l and any cut $W \in \mathcal{S}_a$. Note that (4) implies that $i(e, l) \leq p - 1$ for all $e \in \delta(W)$. Let

$$R_l(W) = \left\{ \mu \in R_l : a \in \arg \max_{e \in W} c_\mu(e) \right\},$$

the subset of R_l such that our fixed arc a has the maximum $c_\mu(a)$ value among all edges in cut W . Then for every $\mu \in R_l(W)$, the definition of $i(e, l)$ implies that

$$\begin{aligned} c_\mu(W) &\geq c_\mu(a) + \sum_{e \in \delta(W) \setminus \{a\} : 1 \leq i(e,l) \leq p-1} c_\mu(e) \\ &\geq c_\mu(a) \left(1 + \sum_{e \in \delta(W) \setminus \{a\} : 1 \leq i(e,l) \leq p-1} \alpha \beta^{i(e,l)-1} \right). \end{aligned}$$

If we define

$$\gamma_l(W) = 1 + \sum_{e \in \delta(W) \setminus \{a\} : 1 \leq i(e,l) \leq p-1} \alpha \beta^{i(e,l)-1}$$

then we can re-write this as

$$c_\mu(W) \geq c_\mu(a) + \sum_{e \in \delta(W) \setminus \{a\} : 1 \leq i(e,l) \leq p-1} c_\mu(e) \geq c_\mu(a) \gamma_l(W) \tag{6}$$

for every $\mu \in R_l(W)$.

The definition of α also implies that, for every $\mu \in R_l(W)$,

$$\begin{aligned} c_\mu(W) &= c_\mu(a) + \sum_{e \in \delta(W) \setminus \{a\} : i(e,l)=0} c_\mu(e) + \sum_{e \in \delta(W) \setminus \{a\} : 1 \leq i(e,l) \leq p-1} c_\mu(e) \\ &\leq c_\mu(a) + |E| \alpha c_\mu(a) + c_\mu(a) \beta (\gamma_l(W) - 1) \\ &= c_\mu(a) \beta \gamma_l(W). \end{aligned} \tag{7}$$

Now fix a cut $X \in \mathcal{S}_a$ such that there exist $\mu \in R_l \cap M_a(X)$. Thus X is a minimum cut for the cost c_μ . For every cut $W \in \mathcal{S}_a$ such that there exists $\nu \in R_l \cap M_a(W)$, we have

$$\begin{aligned}
 c_\mu(W) &\leq c_\mu(a) \beta \gamma_l(W) && \text{[by (7)]} \\
 &= \frac{c_\mu(a)}{c_\nu(a)} \beta c_\nu(a) \gamma_l(W) \\
 &\leq \frac{c_\mu(a)}{c_\nu(a)} \beta c_\nu(W) && \text{[by (6)]} \\
 &\leq \frac{c_\mu(a)}{c_\nu(a)} \beta c_\nu(X) && \text{(by optimality of } W \text{ for } \nu) \\
 &\leq \frac{c_\mu(a)}{c_\nu(a)} \beta^2 c_\nu(a) \gamma_l(X) && \text{[by (7)]} \\
 &= \beta^2 c_\mu(a) \gamma_l(X) \\
 &\leq \beta^2 c_\mu(X) && \text{[by (6)].}
 \end{aligned}$$

Therefore, every cut W such that $R_l \cap M_a(W) \neq \emptyset$ is a β^2 -approximate cut for the parametrized cost c_μ .

If $\rho \leq 4$ then choose any β such that $1 < \beta < \sqrt{9/8}$, so $\beta^2 < 9/8$ and, by Corollary 3, the total number of such cuts W is $O(|V|^2)$. Otherwise (when $\rho \geq 5$), for every $\epsilon > 0$ we may choose $\beta = \beta(\epsilon)$ such that $1 < \beta(\epsilon) < \sqrt{1 + \epsilon/2}$, and the total number of such cuts W is $O(|V|^{2+\epsilon})$. Therefore it is $O_{\rho,5}(|V|^2)$, and the proof is complete. \square

Recall that the combinatorial facet complexity of the parametrized minimum cut problem over the relevant set M of parameters is the maximum number of facets of the graph of $c^*(\mu)$ when $\mu \in M$; equivalently, of vertices of the dominant \mathcal{D} , and also of SND points. Thus Theorem 4 implies:

Corollary 5 *Given a fixed integer $\rho \geq 2$, a rank- ρ hypergraph $G = (V, E)$, and a fixed number k of edge cost functions c^1, \dots, c^k , the combinatorial facet complexity of the parametrized minimum cut problem over set M is $O_{\rho,5}(|E|^k |V|^2 \log^{k-1} |V|)$.*

We now briefly comment on the multicriteria minimum cut problem. As explained earlier, we now assume that all edge costs are nonnegative, i.e., all $c^j \geq 0$, and that every cut cost $c^j(X) > 0$ ($j = 1, \dots, k, \emptyset \neq X \subset V$) (and therefore that G is connected). Under these assumptions, the relevant region in the parameter space is the nonnegative orthant, $M = \mathbb{R}_+^{k-1}$ and, by Theorem 4, the number of SND points is $O_{\rho,5}(|E|^k |V|^2 \log^{k-1} |V|)$. In Theorem 8 in the next Section, we derive a strongly polynomial bound on the total number of (supported and unsupported) Pareto optimal points when $k = 2$. However, for $k \geq 3$ we leave the following:

Open Problem In a graph or a bounded-rank hypergraph, with a fixed number $k \geq 3$ of nonnegative edge cost functions such that each cut cost is positive, is the number of UND points polynomially bounded?

4 Tighter bounds for minimum cuts with two nonnegative edge cost functions

When the bound on the number of SND cuts from Theorem 4 is specialized to $k = 2$ edge cost functions, it becomes $O_{\rho,5}(|E|^2 |V|^2 \log |V|)$. This section sharpens this bound to $O_{\rho,5}(|V|^3 \log |V|)$, an improvement by a factor $|E|^2/|V|$. For graphs (and rank-3 hypergraphs), this improvement is from $O(|E|^2|V|^2 \log |V|)$ to $O(|V|^3 \log |V|)$. We also derive a strongly polynomial bound on the total number of (supported and unsupported) nondominated points. Our proof strategy will rely on partitioning the criteria space, instead of the parameter space as in Sect. 3. For this reason, we now assume that both edge cost functions c^1 and c^2 are nonnegative. Note that partitioning the criteria space fails to work in the more general case discussed in Sect. 3 as the cost functions c^j there may be negative. The results below also hold in the case where the edge parametric costs are $c_\mu(e) = \mu c^1(e) + (1 - \mu)c^2(e)$ for $\mu \in [0, 1]$.

For general optimization problems with several objectives, Papadimitriou and Yannakakis [34] use a partitioning of the criteria space into rectangular regions of exponentially increasing sizes and obtain an approximate Pareto set of weakly polynomially bounded size. Here we partition the criteria space according to value of criterion c^1 , and we use two properties of global minimum cuts: First (in the proofs of Lemmas 6, 7), the polynomial bound on the number of approximate cuts from Corollary 3; and second (in the proof of Theorem 8), properties of edge contractions. We derive *strongly* polynomial bounds on the total number of *all exact* SND cuts.

The main argument used in the proof of Lemma 6 below is that every SND cut in a c^1 -cost interval of relative width $\beta > 1$ is a β -approximation for the objective c_μ defined by a particular facet of the dominant polyhedron \mathcal{D} . Then, by choosing a small enough constant β , the number of SND cuts is shown to be $O_{\rho,5}(|V|^2)$ per interval.

Lemma 6 *Given a fixed integer $\rho \geq 2$, a rank- ρ hypergraph $G = (V, E)$, two nonnegative edge cost functions c^1 and c^2 such that all cut costs are positive, and two reals $b > a > 0$, the total number of SND cuts X with c^1 -cost satisfying $a \leq c^1(X) \leq b$ is $O_{\rho,5}(|V|^2 \log(b/a))$.*

Proof First, if $\rho \leq 4$ then choose any β such that $1 < \beta < 9/8$; otherwise, i.e., when $\rho \geq 5$, for every $\epsilon > 0$ choose any β such that $1 < \beta < 1 + \epsilon/2$, so $2\beta < 2 + \epsilon$. Next, choose $n = \lfloor \log(b/a) / \log \beta \rfloor + 1 = O(\log(b/a))$ to ensure that $\beta^{n-1}a \leq b < \beta^n a$. Thus we can cover the “target” interval $[a, b]$ with half-open intervals $I_j = [\beta^{j-1}a, \beta^j a)$ for $j = 1, \dots, n$, such that $b \in I_n$.

Let \mathcal{X}_j denote the set of SND cuts X such that $c^1(X) \in I_j$. If $\mathcal{X}_j \neq \emptyset$ then let W be a cut in \mathcal{X}_j with largest c^2 -cost. Let $\mu > 0$ be a parameter value for which W maximizes the parametrized objective $c_\mu = \mu c^1 + c^2$. Since we only have two criteria, for every SND cut $X \in \mathcal{X}_j$ such that $(c^1(X), c^2(X)) \neq (c^1(W), c^2(W))$, we have $\beta^{j-1}a \leq c^1(W) < c^1(X) < \beta^j a$ and $c^2(X) < c^2(W)$. Thus

$$\begin{aligned} c_\mu(X) &< \mu c^1(X) + c^2(W) < \mu \beta^j a + c^2(W) \\ &< \beta \left(\mu \beta^{j-1} a + c^2(W) \right) \leq \beta \left(\mu c^1(W) + c^2(W) \right) \end{aligned}$$

$$= \beta c_\mu(W). \quad (8)$$

Therefore, every cut in \mathcal{X}_j is a β -approximate cut for the parametrized objective c_μ . By Corollary 3 and the above choice of β , the total number $|\mathcal{X}_j|$ of such cuts is $O_{\rho,5}(|V|^2)$, and the total number of SND cuts X with $a \leq c^1(X) \leq b$ is at most $\sum_{j=1}^n |\mathcal{X}_j| = O_{\rho,5}(|V|^2 \log(b/a))$, as claimed. \square

In a given width- β interval I , let $c^1(I)$ be the c^1 -cost of the SND cut in I with minimum c^1 -cost, with $c^1(I)$ equalling the left endpoint of I if it contains no SND cut. The proof of the next lemma uses the fact that there are two types of Pareto optimal cuts in I : Those with c^1 -cost greater than $c^1(I)$, which are again β -approximations; and the others (which must be UND cuts with c^1 -cost less than $c^1(I)$), which we show are 2-approximations for a suitable c_μ objective.

Lemma 7 *Given a fixed integer $\rho \geq 2$, a rank- ρ hypergraph $G = (V, E)$, two nonnegative edge cost functions c^1 and c^2 such that all cut costs are positive, and two reals $b > a > 0$, the total number of Pareto optimal cuts X with c^1 -cost satisfying $a \leq c^1(X) \leq b$ is $O_{\rho,5}(|V|^4 \log(b/a))$.*

Proof Let $\mathcal{N}[a, b]$ denote the set of all Pareto optimal cuts W with $c^1(W) \in [a, b]$, so we need to bound $|\mathcal{N}[a, b]|$. Our proof considers a number of cases and subcases. Cases 1 to 1.1.1 comprise the bulk of the proof; the other cases and subcases yield to simpler treatment.

Case 1. There is at least one SND cut Z with $c^1(Z) \in [a, b]$: Define $a'' = \min\{c^1(X) : X \in \mathcal{S}\}$, so $a'' \leq c^1(Z)$ and there is no Pareto optimal cut W with $c^1(W) < a''$. Similarly, define $b'' = \max\{c^1(X) : X \in \mathcal{S}\}$, so $b'' \geq c^1(Z)$ and there is no Pareto optimal cut W with $c^1(W) > b''$. Now we restrict our attention to the interval $[a', b']$, where $a' = \max\{a, a''\}$ and $b' = \min\{b, b''\}$. Note that $[a', b'] \subseteq [a, b]$ and $\mathcal{N}[a', b'] = \mathcal{N}[a, b]$.

Case 1.1. $a' < b'$: Let $x_1 < x_2 < \dots < x_N$ denote the distinct values of $c^1(X)$ for all SND cuts X with $a' < c^1(X) < b'$. It follows from Lemma 6 that $N = O(|V|^2 \log(b'/a'))$. Let X_1, \dots, X_N be corresponding SND cuts, i.e., with $c^1(X_i) = x_i$ (Since different cuts may give rise to the same point in the criterion space there may be several choices for X_i , but this does not affect our arguments). Define $x_0 = \max\{c^1(X) : X \in \mathcal{S} \text{ with } c^1(X) \leq a'\}$ with corresponding SND cut X_0 , and $x_{N+1} = \min\{c^1(X) : X \in \mathcal{S} \text{ with } c^1(X) \geq b'\}$ with corresponding SND cut X_{N+1} . For $i = 0, 1, \dots, N$ define $v(i)$ as the value of μ such that $c_\mu(X_i) = c_\mu(X_{i+1}) = \min\{c_\mu(X) : X \in \mathcal{C}\}$. Thus, $v(0) > v(1) > \dots > v(N)$ are the successive breakpoints of the parametric curve $c^*(\mu)$ over the interval of μ values for which SND cuts X that minimize c_μ have $c^1(X) \in [a', b']$.

- For example, suppose that in Fig. 1a we have $c^1(C_2) < a < c^1(C_3)$ and $b > c^1(C_5)$. Then $a' = a$ and $b' = c^1(C_5)$; $N = 2$; $x_0 = c^1(C_2)$, $x_1 = c^1(C_3)$, $x_2 = c^1(C_4)$ and $x_3 = c^1(C_5)$; $v(0) = b_2$, $v(1) = b_3$ and $v(2) = b_4$.

Every Pareto optimal cut $W \in \mathcal{N}[a', b']$ has $c^1(W) \in [x_i, x_{i+1}]$ for some $i \in \{0, 1, \dots, N\}$. Since W is Pareto optimal, its c^2 cost satisfies $c^2(W) \in [c^2(X_{i+1}), c^2(X_i)]$. Since X_i and X_{i+1} are both minimum cuts for $\mu = v(i)$, we have $c_{v(i)}(X_i) = c_{v(i)}(X_{i+1})$ and so

$$\begin{aligned} c_{v(i)}(W) &\leq v(i) c^1(X_{i+1}) + c^2(X_i) \\ &< v(i) c^1(X_{i+1}) + c^2(X_{i+1}) + v(i) c^1(X_i) + c^2(X_i) \\ &= 2 c_{v(i)}(X_i). \end{aligned} \tag{9}$$

That is, W is a 2-approximate cut for the parametrized objective $c_{v(i)}$. As in Lemma 6, we cover the interval $[a', b']$ by half open intervals $[\beta^{j-1}a', \beta^j a')$, where $\beta > 1$ is an appropriately chosen constant. Now fix one of the intervals $I'_j = [\beta^{j-1}a', \beta^j a')$ and let $\mathcal{N}(I'_j) := \{W \in \mathcal{N}[a', b'] : c^1(W) \in I'_j\}$.

Case 1.1.1. There is an SND cut with c^1 cost in I'_j : Let X_i be the SND cut with the least c^1 cost among all SND cuts with c^1 cost in I'_j . Every $W \in \mathcal{N}(I'_j)$ with $c^1(W) < c^1(X_i)$ satisfies $c^1(W) \in [x_{i-1}, x_i]$ and so by (9) is a 2-approximate cut for the parametrized objective $c_{v(i-1)}$; there are again $O_{\rho,5}(|V|^4)$ such Pareto optimal cuts. Every other $W \in \mathcal{N}(I'_j)$ satisfies $c^1(X_i) \leq c^1(W) < \beta^j a'$. Since W is non-dominated, we must have $c^2(W) < c^2(X_i)$. Using a similar argument as for (8), we have

$$\begin{aligned} c_\mu(W) &< \mu c^1(W) + c^2(X_i) < \mu \beta^j a + c^2(X_i) \\ &< \beta \left(\mu \beta^{j-1} a + c^2(X_i) \right) \leq \beta \left(\mu c^1(X_i) + c^2(X_i) \right) \\ &= \beta c_\mu(X_i). \end{aligned}$$

This implies that W is a β -approximate cut for the parametrized objective $c_{v(i)}$; thus, for β chosen as in the proof of Lemma 6, there are $O_{\rho,5}(|V|^2)$ such cuts, and so a total of $O_{\rho,5}(|V|^4)$ Pareto optimal cuts in I'_j . Since $|\mathcal{N}(I'_j)| = O_{\rho,5}(|V|^4)$ for each of the $O(\log(b'/a'))$ intervals I'_j , we get $|\mathcal{N}[a, b]| = |\mathcal{N}[a', b']| = O_{\rho,5}(|V|^4 \log(b'/a')) = O(|V|^4 \log(b/a))$, as claimed.

Case 1.1.2. I'_j does not contain the c^1 cost of any SND cut: Then every $W \in \mathcal{N}(I'_j)$ has $c^1(W) \in [x_i, x_{i+1}]$ for some i such that $x_i < \beta^{j-1}a'$ and $\beta^j a' \leq x_{i+1}$, hence is a 2-approximate cut for the parametrized objective $c_{v(i)}$. By Corollary 3 there are $O_{\rho,5}(|V|^4)$ such cuts.

Case 1.2 $a' = b'$: Then $\mathcal{N}[a', b']$ is just the set of all SND cuts X with $c^1(X) = a'$, and the proof of Lemma 6 shows that $|\mathcal{N}[a', a']| = O_{\rho,5}(|V|^2) = O(|V|^4 \log(b/a))$.

Case 2. There is no SND cut Z with $c^1(Z) \in [a, b]$: Then either (i) there is no SND cut Z at all with $c^1(Z) < a$, and thus $\mathcal{N}[a, b] = \emptyset$; or else (ii) let as above X_0 be an SND cut with largest cost $c^1(X_0) < a$ and $v(0)$ the corresponding parameter value, then (9) implies that every $W \in \mathcal{N}[a, b]$ is a 2-approximate cut for the para-

metrized objective $c_{v(0)}$, and thus $|\mathcal{N}[a, b]| = O_{\rho,5}(|V|^4) = O_{\rho,5}(|V|^4 \log(b/a))$, as claimed.

The proof is complete. □

The proof of the next theorem uses the following observation: every cut either has its c^1 cost in the interval $\left[\frac{c^1(E)}{|E|}, c^1(E)\right]$ between the average and total c^1 edge costs, that has relative width $|E| = O(|V|^\rho)$ (hence $\log |E| = O(\log |V|)$) and to which we apply the two preceding lemmas; or else the cut cannot contain any edge e with cost $c^1(e) \geq \frac{c^1(E)}{|E|}$, in which case we can contract all such costlier-than-average edges and recurse. Since there is at least one such edge to contract, the number of vertices decreases by at least one in each iteration and we are done after fewer than $|V|$ rounds of edge contractions.

Theorem 8 *Given a fixed integer $\rho \geq 2$, a rank- ρ hypergraph $G = (V, E)$, and two nonnegative edge cost functions such that all cut costs are positive, the total number of SND cuts is $O_{\rho,5}(|V|^3 \log |V|)$ and the total number of Pareto optimal cuts is $O_{\rho,5}(|V|^5 \log |V|)$.*

Proof Our strategy is to repeatedly apply Lemma 6 to a sequence of hypergraphs obtained by contracting edges. As in [20,25], by *contracting* an edge e in hypergraph $G = (V, E)$ with edge costs c^j to get $G' = (V', E')$ we mean replacing all nodes in e by a single “contracted node” v_e (i.e., $V' = (V \setminus e) \cup \{v_e\}$), and every edge $f \in E$ by an edge f' of same costs, $c^j(f') = c^j(f)$, wherein all nodes in $e \cap f$, if any, are replaced by the single node v_e (i.e., $f' = (f \setminus e) \cup \{v_e\}$ if $f \cap e \neq \emptyset$, and $f' = f$ otherwise). We then remove from E all edges f with $|f| = 1$ (loops, which no cut in G' can cross), and all nodes that do not belong to any edge in E' (isolated nodes). Conversely, every cut in the contracted hypergraph corresponds, after expanding back the contracted node v_e , to a cut with the same cost in the original hypergraph. Note that this generalizes edge contraction in graphs, e.g., [31].

Suppose that X is a cut in G , and we contract e to get G' . If e is not in $\delta(X)$ (“ X does not cross e ”), then $X' := X \setminus \{v_e\}$ is a cut in G' where f' is in $\delta(X')$ if and only if f is in $\delta(X)$. Thus $c^j(X) = c^j(X')$ for $j = 1, 2$. On the other hand, if X is a cut in G with $e \in \delta(X)$ (“ X crosses e ”), then there is no cut in G' corresponding to X .

We construct a sequence of m rank- ρ hypergraphs, $i = 1, \dots, m$. Hypergraph i is $G_i = (V_i, E_i)$ with edge costs $c_i = (c_i^1, c_i^2)$, such that $(G_0, c_0) = (G, (c^1, c^2))$. Hypergraph (G_i, c_i) is derived from (G_{i-1}, c_{i-1}) by contracting all edges in $e \in E_{i-1}$ with cost $c_{i-1}^1(e) \geq a_{i-1} := \frac{c_{i-1}^1(E_{i-1})}{|E_{i-1}|}$, so long as $|V_{i-1}| > \rho$. Since there is at least one edge $e \in E_{i-1}$ with c^1 -cost no less than the average a_{i-1} , there will be at least one edge to contract. Thus $|V_i| < |V_{i-1}|$ and so $m \leq |V| - \rho$.

We associate with each hypergraph G_i in the sequence an interval $[a_i, b_i]$ where, as defined above, $a_i := \frac{c_i^1(E_i)}{|E_i|}$ is the average edge c_i^1 -cost in G_i , and $b_i = c_i^1(E_i) = |E_i|a_i$ is the total c_i^1 -cost. Since $c_i^1 \geq 0$, the c_i^1 -cost of every cut X in G_i satisfies $c_i^1(X) \leq c_i^1(E_i) = b_i$. Every cut X in G_i either has cost $c_i^1(X) \geq a_i$, and therefore $c_i^1(X) \in [a_i, b_i]$, or else it cannot cross any edge e with cost $c_i^1(e) \geq a_i$. In the latter

case, if $i < m$ then X does not cross any edge which is contracted when defining G_{i+1} . Thus there is a cut X' in G_{i+1} corresponding to X , and its cost $c_i^1(\delta_i(X)) = c_{i+1}^1(\delta_{i+1}(X')) \leq b_{i+1}$.

Let s_i (resp. n_i) denote the number of SND (resp. nondominated) cuts in G_i with cost $c_i^1(X)$ in $[a_i, b_i]$. It follows that s_i (resp. n_i) is also the number of SND (resp. nondominated) cuts in the original hypergraph G with cost $c^1(X) \in [a_i, b_i]$. Then the total number of SND cuts in G is at most $\sum_{i=1}^m s_i$, and it suffices to prove that each $s_i = O_{\rho,5}(|V|^2 \log |V|)$. Similarly, it suffices to prove that each $n_i = O_{\rho,5}(|V|^4 \log |V|)$. These are certainly true for $i = m$ since $|V_m| \leq \rho$ and thus $s_m \leq n_m < 2^\rho = O(1)$. For $i < m$, $s_i = O_{\rho,5}(|V_i|^2 \log |E_i|)$ by Lemma 6, and $n_i = O_{\rho,5}(|V_i|^4 \log |E_i|)$ by Lemma 7. Since $|E_i| = O(|V_i|^\rho)$ and ρ is fixed, $\log |E_i| = O(\log |V_i|) = O(\log |V|)$. \square

Since distinct facets of the graph of c^* are associated with distinct SND cuts, we have:

Corollary 9 *Under the assumptions of Theorem 8, the combinatorial facet complexity of the bi-objective parametrized minimum cut problem is $O_{\rho,5}(|V|^3 \log |V|)$.*

Notice that in the case of graphs (where $\rho = 2$) the bounds in Theorem 8 specialize to there being $O(|V|^3 \log |V|)$ SND cuts and $O(|V|^5 \log |V|)$ UND cuts. The conference version of this paper [1] claimed a slightly better bound of $O(|V|^3)$ SND cuts, but there was an error in its analysis of its Algorithm 2 in the proof of its Claim 3. On the other hand, the present $O(|V|^5 \log |V|)$ bound on the number of UND cuts is better than the $O(|V|^7)$ bound of [1].

5 Algorithms

In this section we show how to combine the preceding results with several existing algorithms, and derive strongly polynomial time deterministic algorithms for computing the optimum cost function c^* ; the sets of supported and, for $k = 2$, unsupported Pareto optimal points; and the sets of all corresponding minimum cuts. For this we use three types of existing algorithms:

- (1) Algorithms for solving the ordinary (single-objective) minimum cut problems, i.e., to determine the value $c^*(\mu)$ and a corresponding minimum cut for any given $\mu \in M$, see the Introduction. Let $MC = MC(|V|, |E|)$ denote the running time of such an algorithm for finding a minimum cut in a hypergraph with $|V|$ vertices and $|E|$ edges. Recall that $MC = O(|E| \cdot |V| + |V|^2 \log |V|)$ ([30,42] for graphs; [20,25,37] for hypergraphs).
- (2) Algorithms for construction of the optimum cost function c^* for μ in a convex parameter region, see [11, Section 3.2] or [10, Section 30.4.2]. In particular, Theorem 30.2 in [10] implies that if the upper envelope of the piecewise linear concave function c^* of d parameters has Φ facets and Ψ vertices then, it can be obtained by evaluating $c^*(\mu)$ at $O(\Phi + d \Psi)$ points μ , and then constructing the full upper convex hull derived from the resulting hyperplanes. The latter task is dual, and

computationally equivalent, to the well-studied question of constructing the convex hull of a given set of points, e.g., [36, Section 3.4], [26, Section 7.3]. Chan's algorithm [4], the fastest convex hull algorithm known today, achieves this in $O(\Phi \log \Phi + (\Psi \Phi)^{1+1/(\lfloor d/2 \rfloor + 1)} \log^{O(1)} \Phi)$ time. For $k = 2$ (hence $d = 1$), the whole graph can be constructed in $O(\Phi T)$ total time [9], where T is the time needed for an evaluation of $c^*(\mu)$ (i.e., MC for our problem).

- (3) Algorithms for producing, for a single objective, *all* minimum cuts and all 2-approximate cuts. We collect and summarize a number of methods and results in the next lemma and its proof:

Lemma 10 *Let $AMC1$ and $AMC2$ denote the total time to produce all minimum cuts and all 2-approximate cuts, respectively, in rank- ρ hypergraphs $G = (V, E)$ with nonnegative edge costs c and positive minimum cut value.*

- (i) *For $\rho = 2$ (i.e., for graphs), $AMC1 = O(|E|^2|V| + |V|^2|E|)$, and $AMC2 = O(|E| |V|^4)$.*
 (ii) *For $\rho = 3$, $AMC1 = O(|V|^5)$ and $AMC2 = O(|V|^6)$.*
 (iii) *For $\rho \geq 4$, $AMC1 = \min \{O(|V|^{2B(\rho)+2}), \tilde{O}(|V|^3|E|^2)\}$ and $AMC2 = \min \{O(|V|^{4B(\rho)+2}), \tilde{O}(|V|^5|E|^2)\}$.*

Proof (i) For the case where G is a graph, Dinitz et al. [6] present a ‘‘cactus representation’’ of all minimum cuts, which is a of tree of cycles on node set V such that cut X is minimum if and only if it cuts exactly two arcs in one cycle. Nagamochi et al. [32] use a cactus representation to show that $AMC1 = O(|E|^2|V| + |V|^2|E|)$, while it follows from [31] that $AMC2 = O(|E| |V|^4)$ (see also [18] and [35, Chapter 2] for randomized graph cut enumeration algorithms).

(ii) If G is a rank-3 hypergraph then, as shown in the proof of Proposition 2, the cost of a cut X is the same in graph G with edge cost function c as in the complete graph $K(V)$ with edge cost function \tilde{c} defined in the proof of that proposition. Therefore one may apply the method of [32] to $(K(V), \tilde{c})$ and produce in $AMC1 = O(|E_{K(V)}|^2|V| + |V|^2|E_{K(V)}|) = O(|V|^5)$ time all minimum cuts for (G, c) , and in $AMC2 = O(|E_{K(V)}| |V|^4) = O(|V|^6)$ time all 2-approximate cuts for (G, c) .

(iii) If G is a rank- ρ hypergraph with $\rho \geq 4$, then we choose the faster of the following two methods. First, we may apply the approximation to graph cuts in $K(V)$ from the proof of Proposition 2. By (2) with $\alpha = 1$ and $\alpha = 2$, it suffices to produce all $\alpha B(\rho)$ -approximate cuts in $(K(V), \tilde{c})$, and then only retain the α -approximate cuts for (G, c) . Since ρ is fixed, it takes $O(|E| + |V|^2)$ time to compute the weights of all clique edges (storing them in a $|V| \times |V|$ matrix). It then follows from [31] that $AMC1 = O(|E| + |V|^2 + |E_{K(V)}| |V|^{2B(\rho)}) = O(|E| + |V|^{2B(\rho)+2})$ and $AMC2 = O(|E| + |V|^2 + |E_{K(V)}| |V|^{4B(\rho)}) = O(|E| + |V|^{4B(\rho)+2})$.

Alternately, we may apply the Lawler-Murty (LM) general approach for ranking solutions to combinatorial optimization problems (see, e.g., [12,22] for treatments generalizing [29], and [13,43] or [31, Section 4.1] for its application to ranking graph cuts). The LM method produces a sequence (X^1, X^2, \dots) of cuts, where we require (w.l.o.g.) that node 1 be in each X^i , in nondecreasing cost order, i.e., $c(X^1) \leq c(X^2) \leq \dots$, and such that if $c(X) < c(X^i)$ then X must be one of X^1, \dots, X^{i-1} . Thus X^1 is

a minimum cut; X^2 is a “second best cut”, i.e., a least cost cut (with $1 \in X^2$) distinct from X^1 ; X^3 is a “third best cut”; and so on. The LM method requires solving $O(|V|)$ modified minimum cut problems per cut X^i produced. It is based on the repeated application of a simple idea: X^2 contains at least one vertex $v \notin X^1$ or its complement $V \setminus X^2$ contains at least one vertex $v \in X^1 \setminus \{1\}$. Then X^2 is a best cut among $|V| - 1$ candidate sets Y_v^1 , one for each $v \in V \setminus \{1\}$, where Y_v^1 is a least cost cut with $1 \in Y_v^1$ and either (a) $v \in Y_v^1$ if $v \notin X^1$, or (b) $v \notin Y_v^1$ if $v \in X^1$. Case (a), $1 \in Y_v^1$ and $v \in Y_v^1$, is easy to enforce, for example by merging vertices 1 and v ; but case (b), $1 \in Y_v^1$ and $v \notin Y_v^1$, gives rise to a hypergraph s - t -cut problem, viz, to find a minimum cost cut Y_v^1 that separates terminals $s = 1$ and $t = v$. Lawler [23] (see also [44]) shows that a minimum s - t -cut in a hypergraph can be computed by finding a minimum cut in an auxiliary digraph with $|V| + 2|E|$ vertices and $2d_G + |E|$ edges, where $d_G = \sum_{e \in E} |e|$ is the sum of the cardinalities of all edges in hypergraph G . Since $d_G = O(\rho|E|)$, a minimum s - t -cut in a hypergraph can be found in $\tilde{O}(|E|^2)$ time. By Theorem 1, using the LM method to produce the $O(|V|^{2\alpha})$ α -approximate cuts for $\alpha = 1$ and 2, respectively, leads to $AMC1 = \tilde{O}(|V|^3|E|^2)$ and $AMC2 = \tilde{O}(|V|^5|E|^2)$, running times that are independent of ρ . \square

Combining these algorithms and our earlier results we obtain:

Theorem 11 *Given a hypergraph $G = (V, E)$, and a fixed number k of nonnegative edge cost functions c^1, \dots, c^k such that $c^j(X)$ is positive for all cuts X in G :*

- (i) *If $k = 2$, there exists an algorithm that constructs the optimum cost function in $O_{\rho,5}(|V|^3 \log |V| MC)$ time, and all cuts defining SND and Pareto optimal points in $O_{\rho,5}(|V|^3 \log |V| AMC1)$ and $O_{\rho,5}(|V|^3 \log |V| AMC2)$ time, respectively.*
- (ii) *If $k \geq 3$, the time needed to construct the optimum cost function and all cuts defining SND points is*

$$O_{\rho,5} \left(|E|^k \lfloor \frac{k-1}{2} \rfloor |V|^{2 \lfloor \frac{k-1}{2} \rfloor} \log^{(k-1) \lfloor \frac{k-1}{2} \rfloor + O(1)} |V| \right) \tag{10}$$

Proof (i) Follows from the use of the Eisner and Severance method [9] with $\Phi = O_{\rho,5}(|V|^3 \log |V|)$ from Theorem 8, and from the fact, established in the proof of Lemma 7, that every Pareto optimal point is defined by a 2-approximate cut for a facet-inducing parametrized objective.

(ii) Follows from Theorem 30.2 in [10] using $d = k - 1$, $\Phi = O_{\rho,5}(\hat{\Phi})$ where $\hat{\Phi} = |E|^k |V|^2 \log^{k-1} |V|$ by Theorem 4, $\Psi = O(\hat{\Psi})$ where $\hat{\Psi} = \hat{\Phi}^{\lfloor \frac{k-1}{2} \rfloor}$ by McMullen’s Upper Bound Theorem (e.g., [26, Section 7.3]; see also [40]). Indeed, using Chan’s convex hull algorithm [4], the time needed for constructing the lower convex hull from the Φ hyperplanes is bounded by (10). Since the time MC for each “probe” (evaluation of $c^*(\mu)$ for a given $\mu \in M$) satisfies $MC = O_{\rho,5}(\hat{\Phi})$ when $k \geq 3$, the lower hull construction time (10) dominates the total time for constructing the optimum cost function. Furthermore, for each of the Φ SND points one may produce in $AMC1$ time all corresponding cuts. Since $AMC1 = O(\hat{\Phi})$ when $k \geq 3$, the lower hull construction time (10) again dominates the total computation time. \square

Remark 12 The results in Lemma 10 and Theorem 11 are based on deterministic algorithms. The design of randomized algorithms and the study of the tradeoffs between their running time and success probability are beyond the scope of this paper and are left to future work.

6 Conclusion

We have extended Karger's bounds [16, 17] on the number of approximate global minimum cuts from graphs to hypergraphs. This then allowed us to derive strongly polynomial bounds on the number of Supported (and for the case of two objectives, Unsupported) Pareto optimal cuts w.r.t. multiple objectives (or multiple parameters) in both hypergraphs and graphs. These bounds, when combined with existing algorithms, lead to strongly polynomial algorithms for computing all such cuts.

In addition to the open problem at the end of Sect. 3 (namely, for a fixed number $k \geq 3$ of criteria and fixed rank $\rho \geq 2$, is there a strongly polynomial bound on the total number of optimal points?) another open problem is to find a non-trivial *lower bound* on the complexity of parametric global minimum cut. It is easy to produce classes of instances with $\Theta(|V|)$ different facets in the graph of c^* , but so far we do not have any classes of instances with superlinear complexity. It would be interesting to close this gap; or even for graphs and just two criteria, to reduce the gap between the present $\Omega(|V|)$ and $O(|V|^3 \log |V|)$ bounds.

Acknowledgments We thank Volker Kaibel and Martin Skutella for helpful conversations around the Upper Bound Theorem, and anonymous referees for detailed and perceptive comments, particularly for pointing out reference [21]. The work of the third author was supported by a Discovery Grant from the Natural Sciences and Engineering Research Council (NSERC) of Canada. The work of the last author was supported by a Discovery Grant and a Discovery Accelerator Supplement Grant from NSERC, and by the Center for Operations Research and Econometrics (CORE) of the Université Catholique de Louvain, Belgium.

Appendix

In this Appendix we present a proof of Theorem 1 for arbitrary $\alpha \geq 1$. We adapt and extend Kogan and Krauthgamer's [21] approach and some of their notations.

We start with a simple upper bound on the minimum cut cost in a rank- ρ hypergraph. For every $i = 2, \dots, \rho$ let $W_i = \sum\{c(e) : e \in E \text{ and } |e| = i\}$ denote the total cost of all size- i edges, and $c(E) = \sum_{i=2}^{\rho} W_i = \sum_{e \in E} c(e)$ the total weight of all edges.

Lemma 13 ([21]) *The minimum cost of a cut in a rank- ρ hypergraph $G = (V, E)$ with positive edge costs c , is at most $\frac{1}{|V|} \sum_{i=2}^{\rho} i W_i \leq \frac{\rho}{|V|} c(E)$.*

Proof An edge $e \in E$ crosses a singleton cut $C = (\{v\}, V \setminus \{v\})$ (i.e., $e \in \delta(\{v\})$) if and only if $v \in e$. Thus an edge e crosses exactly $|e|$ singleton cuts. If we choose a node \mathbf{v} uniformly at random in V and consider the resulting (random) singleton cut $\{\mathbf{v}\}$, then

$$\begin{aligned} \min_{X \in \mathcal{C}} c(\delta(X)) &\leq \min_{v \in V} c(\delta(\{v\})) \leq \mathbb{E}[c(\{\mathbf{v}\})] = \sum_{v \in V} \frac{1}{|V|} \sum_{e: v \in e \in E} c(e) \\ &= \frac{1}{|V|} \sum_{e \in E} |e| c(e) = \frac{1}{|V|} \sum_{i=2}^{\rho} i W_i \leq \frac{\rho}{|V|} \sum_{i=2}^{\rho} W_i = \frac{\rho}{|V|} c(E) \end{aligned}$$

and the Lemma follows. □

Kogan and Krauthgamer use a probabilistic argument which extends to hypergraphs by an approach introduced by Karger [16, 19] for graph cuts. The first step is to define a generalization to hypergraphs of Karger’s randomized edge contraction algorithm. Similar generalizations for finding and counting minimum hypergraph cuts were outlined by Chekuri and Korula [5] and by Queyranne and Guíñez [38]. The algorithm below, adapted from [21], is similar to these, except that (as in [19]), it stops early so as to ensure a sufficiently high probability of generating any fixed α -approximate cut.

Algorithm 1 CONTRACT HYPERGRAPH

Input: a rank- ρ hypergraph $G = (V, E)$ with edge costs c , and a parameter $\alpha \geq 1$
Output: a cut X in G
 1: Let $G' = (V', E') \leftarrow G$ and $c' \leftarrow c$
 2: **while** $|V'| > \lfloor \alpha \rho \rfloor$ **do**
 3: Choose an edge $e' \in E'$ with probability $c'(e')/c'(E')$
 4: Contract e' by merging all its vertices and removing self-loops
 5: **end while**
 6: Choose uniformly at random a cut X' in the final hypergraph G'
 7: **return** the cut X in G induced by the cut X'

Theorem 1 will be an easy consequence of the following result, which is adapted from and extends (with minor changes) Theorem 3.4 in [21]. Since we allow 2α to be noninteger, we use, as in Karger and Stein [19, proof of Corollary 8.3], generalized binomial coefficients

$$\binom{x}{y} = \frac{\Gamma(x + 1)}{\Gamma(y + 1)\Gamma(x - y + 1)}$$

where, for simplicity, we assume $x > -1$ and $-1 < y < x + 1$, and where the Gamma function is the Euler integral $\Gamma(x) = \int_0^{+\infty} t^{x-1} e^{-t} dt$. Since $\Gamma(t + 1) = t \Gamma(t)$ for all $t > 0$, these (generalized) binomial coefficients satisfy the recurrence

$$\binom{x}{y} = \frac{x}{x - y} \binom{x - 1}{y} = \left(1 + \frac{y}{x - y}\right) \binom{x - 1}{y} \tag{11}$$

when $x > 0$ and $-1 < y < x$.

Theorem 14 *For every real $\alpha \geq 1$ and every integer $\rho \geq 2$ there exists a constant $K(\alpha, \rho) > 0$ such that, for every rank- ρ hypergraph $G = (V, E)$ with $|V| > \alpha \rho$*

vertices, nonnegative edge costs c and positive minimum cut cost, and for every particular α -approximate cut X in G , the probability that Algorithm 1 outputs X or $V \setminus X$ is at least $K(\alpha, \rho) \binom{|V| - \alpha(\rho - 2)}{2\alpha}^{-1}$.

We prove Theorem 14 below, after first establishing another lemma. In the rest of this Appendix we assume that $\alpha \geq 1$ is a fixed real number, and $\rho \geq 2$ a fixed integer.

For every integer $t \geq 2$ let \mathcal{G}_t denote the set of all edge-weighted rank- ρ hypergraphs (G, c) with t vertices, nonnegative edge costs and positive minimum cut cost. Let $\mathcal{G} = \bigcup_{t \geq 2} \mathcal{G}_t$ denote the set of all such edge-weighted hypergraphs with any number of vertices. For $(G, c) \in \mathcal{G}_t$ and cut X in G let $p_t(X | G, c)$ denote the probability that Algorithm 1, when applied to (G, c) , outputs cut X or $V \setminus X$. Let

$$p_t = \inf \{ p_t(X | G, c) : (G, c) \in \mathcal{G}_t \text{ and } X \text{ is an } \alpha\text{-approximate cut in } (G, c) \}$$

Lemma 15 *If $t \geq \lfloor \alpha \rho \rfloor + 1$ and we have a positive lower bound $\tilde{p}_u \leq p_u$ for every $u = t - \rho + 1, \dots, t - 1$, then $p_t \geq (t - \alpha \rho) \min_{i=2, \dots, \rho} \tilde{p}_{t-i+1} / (t - \alpha(\rho - i))$.*

Proof Assume that, as stated, $t \geq \lfloor \alpha \rho \rfloor + 1$ and $0 < \tilde{p}_u \leq p_u$ for every $u = t - \rho + 1, \dots, t - 1$. For any $(G, c) \in \mathcal{G}_t$ let ε denote the (random) edge selected for contraction in step 3 of Algorithm 1. Let $(G, c)/e$ denote the node-weighted hypergraph resulting from the contraction of an edge $e \in E$. Thus $(G, c)/e \in \mathcal{G}_{t-|e|+1}$. Let X be any α -approximate cut in (G, c) and let $A(X, (G, c))$ represent the event that Algorithm 1, when applied to (G, c) , outputs cut X or $V \setminus X$. Cut X will survive the current contraction if $\varepsilon \notin \delta(X)$, otherwise it certainly cannot be output by Algorithm 1. Letting X/e denote the node subset X after contraction of edge $e \notin \delta(X)$, conditioning over the size $|e|$ of the contracted edge, and using the assumed lower bounds, we have

$$\begin{aligned} \text{Prob} \{ A(X, (G, c)) \} &= \sum_{e \in E \setminus \delta(X)} \text{Prob} \{ A(X/e, (G, c))/e \} \text{Prob} \{ \varepsilon = e \} \\ &= \sum_{i=2}^{\rho} \sum_{e \in E \setminus \delta(X) : |e|=i} \text{Prob} \{ A(X/e, (G, c))/e \} \text{Prob} \{ \varepsilon = e \} \\ &\geq \sum_{i=2}^{\rho} \tilde{p}_{t-i+1} \text{Prob} \{ \varepsilon \notin \delta(X) \text{ and } |\varepsilon| = i \} = \sum_{i=2}^{\rho} \tilde{p}_{t-i+1} (x_i - y_i). \end{aligned}$$

where $x_i = \text{Prob} \{ |\varepsilon| = i \}$ and $y_i = \text{Prob} \{ \varepsilon \in \delta(X) \text{ and } |\varepsilon| = i \}$. Using the notations in Lemma 13 above, we have $x_i = W_i / c(E)$. Let $\widehat{w}(G, c) = \min_{X \in \mathcal{C}} c(\delta(X))$ denote the minimum cost of a cut in (G, c) . By Lemma 13 we have (as in [21]):

$$\begin{aligned} \sum_{i=2}^{\rho} y_i &= \text{Prob} \{ \varepsilon \in \delta(X) \} = \frac{c(\delta(X))}{c(E)} \\ &\leq \frac{\alpha \widehat{w}(G, c)}{c(E)} \leq \frac{\alpha}{t} \sum_{i=2}^{\rho} \frac{i W_i}{c(E)} = \frac{\alpha}{t} \sum_{i=2}^{\rho} i x_i \end{aligned} \tag{12}$$

Let $\eta = t \min_{i=2, \dots, \rho} \tilde{p}_{t-i+1} / (t - \alpha(\rho - i))$, so $\eta > 0$ and for every $i = 2, \dots, \rho$

$$\tilde{p}_{t-i+1} \geq \frac{t - \alpha(\rho - i)}{t} \eta$$

Therefore,⁵

$$\text{Prob} \{A(X, (G, c))\} \geq \sum_{i=2}^{\rho} \frac{t - \alpha(\rho - i)}{t} \eta (x_i - y_i) \tag{13}$$

$$= \frac{t - \alpha\rho}{t} \eta \sum_{i=2}^{\rho} (x_i - y_i) + \eta \sum_{i=2}^{\rho} \frac{\alpha i}{t} (x_i - y_i)$$

$$\geq \frac{t - \alpha\rho}{t} \eta \sum_{i=2}^{\rho} x_i + \eta \sum_{i=2}^{\rho} \left(\frac{\alpha i}{t} x_i - y_i \right) \tag{14}$$

$$\geq \frac{t - \alpha\rho}{t} \eta = (t - \alpha\rho) \min_{i=2, \dots, \rho} \frac{\tilde{p}_{t-i+1}}{t - \alpha(\rho - i)} \tag{15}$$

where in (14) we use $y \geq 0$ and, since $i \leq r, \alpha i \leq \alpha r \leq t$; and in (15) we use $\sum_{i=2}^r x_i = 1$, (12), and then the definition of η .

Since (13)–(15) holds for every $(G, c) \in \mathcal{G}_t$ and every α -approximate cut X in (G, c) , the proof of Lemma 15 is complete. \square

Proof of Theorem 14 For $u = |V| = \lfloor \alpha\rho \rfloor - \rho + 2, \dots, \lfloor \alpha\rho \rfloor$, Algorithm 1 directly chooses in step 6, uniformly at random, one of the $2^u - 2$ cuts in \mathcal{C} , so the probability it outputs cut X or $V \setminus X$ is exactly $\tilde{p}_u = (2^{u-1} - 1)^{-1}$.

For $t = \lfloor \alpha\rho \rfloor + 1, \dots, \lfloor \alpha\rho \rfloor + \rho - 1$ recursively define

$$\tilde{p}_t = (t - \alpha\rho) \min_{i=2, \dots, \rho} \tilde{p}_{t-i+1} / (t - \alpha(\rho - i)).$$

By Lemma 15 we have $p_t \geq \tilde{p}_t$ for all such t . Define

$$K(\alpha, \rho) = \min_{t=\lfloor \alpha\rho \rfloor + 1, \dots, \lfloor \alpha\rho \rfloor + \rho - 1} \tilde{p}_t \binom{t - \alpha(\rho - 2)}{2\alpha}$$

and $\hat{p}_t = K(\alpha, \rho) \binom{t - \alpha(\rho - 2)}{2\alpha}^{-1}$, so $p_t \geq \hat{p}_t > 0$ for all such t .

For $t \geq \lfloor \alpha\rho \rfloor + \rho$, applying Lemma 15 using \hat{p}_u in lieu of \tilde{p}_u we get

$$p_t \geq (t - \alpha\rho) \min_{i=2, \dots, \rho} \frac{K(\alpha, \rho) \binom{t-i+1-\alpha(\rho-2)}{2\alpha}^{-1}}{t - \alpha(\rho - i)} \tag{16}$$

⁵ Readers familiar with [21] may recognize in (13)–(15) below the use of linear programming duality, arguably a simpler alternative than the analysis of basic solutions in [21]. Indeed η is the value of the dual variable associated with constraint (12) above.

By induction on $i = 2, \dots, \rho$, equation (11) implies (as in Claim 3.6 in [21]):

$$\begin{aligned} \frac{\binom{t-\alpha(\rho-2)}{2\alpha}}{\binom{t-i+1-\alpha(\rho-2)}{2\alpha}} &= \prod_{j=2}^i \left(1 + \frac{2\alpha}{t-j+2-\alpha\rho} \right) \\ &\geq \left(1 + \frac{2\alpha}{t-\alpha\rho} \right)^{i-1} \geq 1 + (i-1) \frac{2\alpha}{t-\alpha\rho} \\ &\geq \frac{t-\alpha(\rho-i)}{t-\alpha\rho}. \end{aligned}$$

Then (16) implies

$$p_t \geq K(\alpha, \rho) \binom{t-\alpha(\rho-2)}{2\alpha}^{-1}.$$

completing the proof of Theorem 14. \square

Proof of Theorem 1 Theorem 14 implies that the number of α -approximate cuts in any $(G, c) \in \mathcal{G}$, where $G = (V, E)$ is a rank- ρ hypergraph, is at most $K(\alpha, \rho)^{-1} \binom{|V|-\alpha(\rho-2)}{2\alpha}$, which is $O(|V|^{2\alpha})$ when $\alpha \geq 1$ is a fixed real number and $\rho \geq 2$ is fixed. This proves Theorem 1. \square

References

1. Aissi, H., Mahjoub, A.R., McCormick, S.T., Queyranne, M.: A strongly polynomial time algorithm for multicriteria global minimum cuts. In: Integer Programming and Combinatorial Optimization (IPCO 2014), Springer LNCS 8494, pp. 25–36 (2014)
2. Armon, A., Zwick, U.: Multicriteria global minimum cuts. *Algorithmica* **46**(1), 15–26 (2006)
3. Carstensen, P.: Complexity of some parametric integer and network programming problems. *Math. Program.* **26**, 64–75 (1983)
4. Chan, T.M.: Output-sensitive results on convex hulls, extreme points, and related problems. *Discrete Comput. Geom.* **16**(4), 369–387 (1996)
5. Chekuri, C., Korula, N.: Personal Communication (2010). Cited in: T. Fukunaga. Computing minimum multiway cuts in hypergraphs from hypertree packings. In: Integer Programming and Combinatorial Optimization (IPCO 2010), Springer LNCS 6080, pp. 15–28 (2010)
6. Dinitz, E.A., Karzanov, A.V., Lomonosov, M.V.: On the structure of the system of minimum edge cuts in a graph. In: Fridman, A.A. (ed.) *Issledovaniya po Diskretnoi Optimizatsii (Studies in Discrete Optimization)*, pp. 290–306. Nauka Publishers, Moscow (1976)
7. Downey, R.G., Fellows, M.R.: *Parameterized Complexity*, vol. 3. Springer, Berlin (1999)
8. Ehrgott, M.: *Multicriteria Optimization*. Springer, Berlin (2005)
9. Eisner, M.J., Severance, D.G.: Mathematical techniques for efficient record segmentation in large shared databases. *J. ACM* **23**, 619–635 (1976)
10. Fernández-Baca, D., Venkatachalam, B.: Sensitivity analysis in combinatorial optimization, chapter 30. In: Gonzalez, T. (ed.) *Handbook of Approximation Algorithms and Metaheuristics*. Chapman and Hall/CRC Press, Boca Raton (2007)
11. Fernández-Baca, D., Seppäläinen, T., Slutzki, G.: Parametric multiple sequence alignment and phylogeny construction. *J Discrete Algorithms* **2**(2), 271–287 (2004)
12. Hamacher, H.W., Queyranne, M.: K best solutions to combinatorial optimization problems. *Ann. Oper. Res.* **4**(1), 123–143 (1985)
13. Hamacher, H.W., Picard, J.-C., Queyranne, M.: Ranking the cuts and cut-sets of a network. *N.-holl. Math. Stud.* **95**, 183–200 (1984)

14. Henzinger, M., Williamson, D.P.: On the number of small cuts in a graph. *Inf. Process. Lett.* **59**(1), 41–44 (1996)
15. Ihler, E., Wagner, D., Wagner, F.: Modeling hypergraphs by graphs with the same mincut properties. *Infor. Process. Lett.* **45**(4), 171–175 (1993)
16. Karger, D.R.: Global min-cuts in RNC, and other ramifications of a simple min-cut algorithm. In: *Proceedings on Fourth Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 21–30 (1993)
17. Karger, D.R.: Minimum cuts in near-linear time. *J. ACM* **47**(1), 46–76 (2000)
18. Karger, D.R., Panigrahi, D.: A near-linear time algorithm for constructing a cactus representation of minimum cuts. In: *Proceedings on Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 246–255 (2009)
19. Karger, D.R., Stein, C.: A new approach to the minimum cut problem. *J. ACM* **43**(4), 601–640 (1996)
20. Klimmek, R., Wagner, F.: A simple hypergraph min cut algorithm. Technical Report b 96-02, Inst. of Computer Science, Freie Universität Berlin (1996)
21. Kogan, D., Krauthgamer, R.: Sketching cuts in graphs and hypergraphs. arXiv preprint [arXiv:1409.2391](https://arxiv.org/abs/1409.2391) (2014)
22. Lawler, E.L.: A procedure for computing the K best solutions to discrete optimization problems and its application to the shortest path problem. *Manag. Sci.* **18**, 401–405 (1972)
23. Lawler, E.L.: Cutsets and partitions of hypergraphs. *Networks* **3**(3), 275–285 (1973)
24. Loff Barreto, B.S.: A medley for computational complexity: with applications of information theory, learning theory, and ketan mulmuley’s parametric complexity technique. PhD thesis, University of Amsterdam (2014)
25. Mak, W.K., Wong, D.F.: A fast hypergraph min-cut algorithm for circuit partitioning. *Integr. VLSI J.* **30**(1), 1–11 (2000)
26. Mulmuley, K.: *Computational Geometry: An Introduction Through Randomized Algorithms*. Prentice-Hall, Upper Saddle River (1994)
27. Mulmuley, K.: Parallel vs. parametric complexity. In: *Algorithms and Complexity (WADS’97)*. Springer, pp. 282–283 (1997)
28. Mulmuley, K.: Lower bounds in a parallel model without bit operations. *SIAM J. Comput.* **28**(4), 1460–1509 (1999)
29. Murty, K.G.: An algorithm for ranking all the assignments in increasing order of cost. *Oper. Res.* **16**, 682–687 (1968)
30. Nagamochi, H., Ibaraki, T.: Computing edge-connectivity in multigraphs and capacitated graphs. *SIAM J. Discrete Math.* **5**(1), 54–66 (1992)
31. Nagamochi, H., Ibaraki, T.: *Algorithmic Aspects of Graph Connectivity*. Cambridge University Press, Cambridge (2008)
32. Nagamochi, H., Nakamura, S., Ishii, T.: Constructing a cactus for minimum cuts of a graph in $O(mn + n^2 \log n)$ space. *IEICE Trans. Inf. Syst.* **E86-D**, 179–185 (2003)
33. Nagamochi, H., Nishimura, K., Ibaraki, T.: Computing all small cuts in undirected networks. *SIAM J. Discrete Math.* **10**, 469–481 (1997)
34. Papadimitriou, C.H., Yannakakis, M.: On the approximability of trade-offs and optimal access of web sources. In: *Proceedings on 41st Annual Symposium Foundations of Computer Science (FOCS 2000)*, IEEE, pp. 86–92 (2000)
35. Panigrahi, D.: Optimization problems in network connectivity. Ph.D. thesis, MIT (2012)
36. Preparata, F.P., Shamos, M.I.: *Computational Geometry*. Springer, Berlin (1985)
37. Queyranne, M.: Minimizing symmetric submodular functions. *Math. Program.* **82**(1–2), 3–12 (1998)
38. Queyranne, M., Guíñez, F.: On optimum k-way partitions with submodular costs and minimum part-size constraints. In: *Talk given at the Workshop on Modern Aspects of Submodularity*, Georgia Institute of Technology, Atlanta (GA) USA, March 21 (2012)
39. Schrijver, A.: *Theory of Linear and Integer Programming*. Wiley, Hoboken (1998)
40. Seidel, R.: The upper bound theorem for polytopes: an easy proof of its asymptotic version. *Comput. Geom.* **5**(2), 115–116 (1995)
41. Steiner, J.: Einige Gesetze über die Theilung der Ebene und des Raumes. *J. Reine Angew. Math.* **1**, 349–364 (1826)
42. Stoer, M., Wagner, F.: A simple min-cut algorithm. *J. ACM* **44**(4), 585–591 (1997)
43. Vazirani, V.V., Yannakakis, M.: Suboptimal cuts: Their enumeration, weight and number. *Automata, Languages and Programming. Lecture Notes in Computer Science 623*. Springer, pp. 366–377 (1992)

-
44. Yang, H.H., Wong, D.F.: Efficient network flow based min-cut balanced partitioning. *IEEE Trans. Comput. Aided Design* **15**(12), 1533–1540 (1996)
 45. Zimmerman, S.: Slicing space. *Coll. Math. J.* **32**, 126–128 (2001)