Submodular Cost Allocation Problem and Applications*

Chandra Chekuri and Alina Ene

Dept. of Computer Science, University of Illinois, Urbana, IL 61801, USA
{chekuri,ene1}@cs.illinois.edu

Abstract. We study the Minimum Submodular-Cost Allocation problem (MSCA). In this problem we are given a finite ground set V and k non-negative submodular set functions $f_1, \ldots, f_k$ on V. The objective is to partition V into k (possibly empty) sets $A_1, \cdots, A_k$ such that the sum $\sum_{i=1}^{k} f_i(A_i)$ is minimized. Several well-studied problems such as the non-metric facility location problem, multiway-cut in graphs and hypergraphs, and uniform metric labeling and its generalizations can be shown to be special cases of MSCA. In this paper we consider a convex-programming relaxation obtained via the Lovász-extension for submodular functions. This allows us to understand several previous relaxations and rounding procedures in a unified fashion and also develop new formulations and approximation algorithms for related problems. In particular, we give a $(1.5 - 1/k)$-approximation for the hypergraph multiway partition problem. We also give a $\min\{2(1-1/k), H_\Delta\}$-approximation for the hypergraph multiway cut problem when Δ is the maximum hyperedge size. Both problems generalize the multiway cut problem in graphs and the hypergraph cut problem is approximation equivalent to the node-weighted multiway cut problem in graphs.

1 Introduction

We consider the following allocation problem with submodular costs.

Minimum Submodular-Cost Allocation (MSCA). Let V be a finite ground set and let $f_1, \cdots, f_k$ be k non-negative submodular set functions on V. That is, for $1 \leq i \leq k$, $f_i : 2^V \to \mathbb{R}_+$ and $f_i(A) + f_i(B) \geq f_i(A \cup B) + f_i(A \cap B)$ for all $A, B \subseteq V$. In the MSCA problem the goal is to partition the ground set V into k (possibly empty) sets $A_1, \cdots, A_k$ such that the sum $\sum_{i=1}^{k} f_i(A_i)$ is minimized.

We observe that the problem is interesting only if the f_i's are different for otherwise allocating all of V to f_1 is trivially an optimal solution. We assume that the functions f_i are given as a value oracle although in specific applications they may be available as explicit poly-time computable functions of some auxiliary

* This is an extended abstract without proofs. A longer version of the paper will be made available on the arXiv. The authors are supported in part by United States NSF grants CCF-0728782 and CCF-1016684.

input. The special case of this problem in which all of the functions are monotone ($f(A) \leq f(B)$ if $A \subseteq B$) has been previously considered by Svitkina and Tardos [21]. In this paper, we consider the problem with both monotone and non-monotone functions. We show that several well-studied problems such as non-metric facility location, multiway cut problems in graphs and hypergraphs, uniform metric labeling and its generalization to hub location among others can be cast as special cases of MSCA. In particular, we investigate the integrality gap of a simple and natural convex-programming relaxation for MSCA that is obtained via the use of the Lovász extension of a submodular function.

Lovász extension and a convex program for MSCA: Let V be a finite ground set of cardinality n. Each real-valued set function on V corresponds to a function $f : \{0,1\}^n \to \mathbb{R}$ on the vertices of the n-dimensional hypercube. The Lovász extension of f to the continuous domain $[0,1]^n$ denoted by $\hat{f}$ is defined as[1]

$$\hat{f}(\mathbf{x}) = \mathbb{E}_{\theta \in [0,1]}\left[f(\mathbf{x}^\theta)\right] = \int_0^1 f(\mathbf{x}^\theta) d\theta$$

where $\mathbf{x}^\theta \in \{0,1\}^n$ for a given vector $\mathbf{x} \in [0,1]^n$ is defined as: $x_i^\theta = 1$ if $x_i \geq \theta$ and 0 otherwise.

Lovász showed that $\hat{f}$ is convex if and only if f is a submodular set function [16]. Moreover, it is easy to see that, given $\mathbf{x}$, the value $\hat{f}(\mathbf{x})$ can be computed in polynomial-time by using a value oracle for f. Via this extension, we obtain a straightforward relaxation for MSCA with a convex objective function and linear constraints. Let $v_1, \cdots, v_n$ denote the elements of V. The relaxation has variables $x(v,i)$ for $v \in V$ and $1 \leq i \leq k$ with the interpretation that $x(v,i)$ is 1 if v is assigned to A_i and 0 otherwise. Let $\mathbf{x}_i = (x(v_1,i), \cdots, x(v_n,i))$. The relaxation is given below.

LE-Rel

$$\begin{aligned} \min \quad & \sum_{i=1}^k \hat{f}_i(\mathbf{x}_i) & \\ & \sum_{i=1}^k x(v,i) = 1 & \forall v \\ & x(v,i) \geq 0 & \forall v, i \end{aligned}$$

Throughout, we use OPT and OPT_{FRAC} to denote the value of an optimal integral and an optimal fractional solution to LE-Rel (respectively).

We remark that LE-Rel can be solved in time that is polynomial in n and $\log(\max_{i,S\subseteq V} f_i(S))$ via the ellipsoid method. Moreover, for some problems of interest the above convex program can be rewritten into an equivalent linear program. We now describe several problems that can be cast as special cases of MSCA, and also how some previously considered linear-programming relaxations can be seen as being equivalent to the convex program above.

---

[1] The definition is not the standard one but is equivalent to it, see [23]. This definition is convenient to us in describing and understanding rounding procedures.

1.1 Problems related to MSCA

Monotone MSCA (Monotone-MSCA) and Facility Location: In facility location, we have a set of facilities $\mathcal{F}$ and a set of clients or demands $\mathcal{D}$. There is a non-negative cost c_{ij} to connect facility i to client j (we do not necessarily assume that these costs form a metric). Opening facility $i \in \mathcal{F}$ costs f_i. The goal is to open a subset of the facilities and assign each client to an open facility so as to minimize the sum of the facility opening cost and the connection costs. Svitkina and Tardos [21] considered the setting where the cost of opening a facility i is a monotone submodular function g_i of the clients assigned to it, and gave an $(1 + \ln |\mathcal{D}|)$-approximation, and matching hardness via a reduction from set cover. We note that this problem is equivalent to MSCA when all the f_i are monotone submodular functions, which we refer to as Monotone-MSCA. In [21] a greedy algorithm via submodular function minimization is used to derive the approximation. Here we prove that the integrality gap of LE-Rel is $(1 + \ln |\mathcal{D}|)$, and describe how certain rounding algorithms achieve this bound. These algorithms are useful when considering functions that are not necessarily monotone.

Submodular Multiway Partition (Sub-MP): We define an abstract problem and then specialize to known problems. Let $f : 2^V \to \mathbb{R}_+$ be a submodular set function over V and let $S = \{s_1, s_2, \ldots, s_k\}$ be k terminals in V. The submodular multiway partition problem is to find a partition of V into $A_1, \ldots, A_k$ such that $s_i \in A_i$ and $\sum_{i=1}^{k} f(A_i)$ is minimized. This has been previously considered by Zhao, Nagamochi and Ibaraki [26]. This can be seen as a special case of MSCA as follows. Define the ground set to be $V' = V \setminus S$ and, for $1 \le i \le k$, $f_i : 2^{V'} \to \mathbb{R}_+$ is the function defined as $f_i(S) = f(S \cup \{s_i\})$. If in addition f is symmetric ($f(A) = f(V - A)$ for all A) we call this the symmetric Sub-MP problem (Sym-Sub-MP). Note that although the problem is based on a single function f, k different submodular functions (induced by the terminals) are needed to reduce it to MSCA. We now discuss some important special cases of this problem.

Multiway Cut in Graphs (Graph-MC): The input is an edge-weighted undirected graph $G = (V, E)$ and k terminal vertices $S = \{s_1, \ldots, s_k\}$; the goal is to remove a minimum-weight set of edges to disconnect the terminals. This can be seen as a special case of the symmetric submodular multiway partition problem by simply choosing f to be the cut-capacity function of G scaled down by a factor of 2. That is, $f(A) = \frac{1}{2} \sum_{e \in \delta(A)} w(e)$ where $w(e)$ is the weight of edge e. We observe that LE-Rel for this problem is equivalent to the well-known geometric LP relaxation of Calinescu, Karloff and Rabani [2], which led to significant improvements ($1.5 - 1/k$ in [2] and 1.3438 in [13]) over the $2(1 - 1/k)$-approximation obtained via the isolating-cut heuristic [4].

Multiway Cut and Partition in Hyper-Graphs: Given an edge-weighted hypergraph $\mathcal{G} = (V, \mathcal{E})$ and terminal set $S \subset V$, the Hypergraph Multiway Cut problem (Hypergraph-MC) (see [17, 25, 7]) asks for the minimum weight subset of hyperedges whose removal disconnects the terminals. This can be seen as a special case of Sub-MP [17]; this reduction requires some care and the underlying submodular function is *asymmetric*. A related problem is the Hypergraph Multiway Partition problem (Hypergraph-MP) introduced by Lawler [15] where the cost for hyperedge e is proportional to the number of non-trivial pieces it is partitioned into. This can be seen as a special case of the Sym-Sub-MP with f being the hypergraph cut capacity function. We note that Graph-MC is a special case of both Hypergraph-MC and Hypergraph-MP.

Node-weighted Multiway Cut in Graphs (Node-wt-MC): In this problem [8] the graph has weights on nodes instead of edges and the goal is to find a minimum weight subset of nodes whose removal disconnects a given set of terminals. It is not difficult to show that Hypergraph-MC and Node-wt-MC are approximation equivalent [17].

Zhao *et al.* [26] consider generalizations of the above problems where some set of terminals $S \subseteq V$ and k are specified and the goal is to partition V into k sets such that each set contains at least one terminal and the total cost of the partition is minimized. We do not discuss these further since they are not directly related to MSCA, although one can reduce them to MSCA if k is a fixed constant.

Uniform Metric Labeling and Submodular Cost Labeling (Sub-Label): The metric labeling problem was introduced by Kleinberg and Tardos [14] as a general classification problem. We are given an undirected edge-weighted graph $G = (V, E)$ and k labels and the goal is to assign a label to each vertex to minimize the sum of the labeling cost and the edge-cut cost. The labeling cost is given by functions $c_i : V \to \mathbb{R}_+$ and the edge-cut cost is given by a metric d on the label space; $d(ij)$ is the distance between labels i and j. Assigning label i to v incurs a cost $c_i(v)$ and if an edge uv of weight $w(uv)$ has u labeled with i and v labeled with j then the edge-cut cost incurred is $w(uv) \cdot d(ij)$. The uniform metric labeling problem is obtained when $d(ij) = 1$ for all $i \neq j$. We consider the following generalization that we call the Submodular Cost Labeling (Sub-Label) problem which is a special case of MSCA. The k labels correspond to the k functions $f_1, \ldots, f_k$. We define f_i as the sum of two functions, a monotone function g_i that models the label assignment cost, and a non-monotone function h that models the cut-cost. The goal then is to partition V into $A_1, \ldots, A_k$ to minimize $\sum_{i=1}^{k}(g_i(A_i) + h(A_i))$. Note that uniform metric labeling is the special case when g_i are modular and h is the graph cut function, which is symmetric. We are motivated to consider this generalization by problems that have been considered previously, such as metric labeling on hypergraphs, hub location problem [9], and the extension of metric labeling to handle label opening costs [5].

1.2 Overview of Results and Techniques

In this paper we examine the complexity of MSCA primarily through the "integrality gap" of the convex relaxation LE-Rel which can be optimized in polynomial time. All the problems we consider are NP-hard and our focus is on polynomial time approximation algorithms.

A significant portion of our contribution is to highlight the naturalness of MSCA and the Lovász-extension based relaxation LE-Rel by showing connections to previously studied problems, linear programming relaxations, and rounding strategies. Viewing these problems in the more abstract setting of submodularity gives insights into prior algorithms. In the process, we obtain new and interesting results. Although one would like to obtain a single unifying algorithm that achieves a good approximation for MSCA, it turns out that LE-Rel has a large integrality gap and we believe that MSCA is hard to approximate to a polynomial factor. However, it is fruitful to examine special cases of MSCA that admit good approximations via LE-Rel. We describe several applications below by summarizing our results; all of them are based on LE-Rel.

- The integrality gap of LE-Rel for Monotone-MSCA is $\Theta(\log n)$.
- There is a $(1.5 - 1/k)$-approximation for Hypergraph-MP.
- There is a $\min\{2(1-1/k), H_\Delta\}$-approximation for HMC, where Δ is the maximum hyperedge size and H_i is the i-th harmonic number. For $\Delta = 2$ this gives a 1.5-approximation and for $\Delta = 3$ this gives a 1.833-approximation.
- LE-Rel for HMC gives a new mathematical programming relaxation for Node-wt-MC and a new $2(1-1/k)$-approximation. Moreover, if all non-terminal nodes have degree at most 3 we obtain a 1.833-approximation improving upon the $2(1-1/k)$ known via the distance-based relaxation [8].
- The integrality gap of LE-Rel for Sym-Sub-MP is at most $2 - 2/k$; this gives an alternative approximation to previous combinatorial algorithms [18, 26]. We raise the question as to whether the integrality gap is at most 1.5.
- There is an $O(\log n)$ for Sub-Label when the cut function is symmetric. We derive results for other special cases of Sub-Label.

Rounding the convex relaxation: Recall that the objective function in LE-Rel is $\sum_{i=1}^{k} \hat{f}_i(\mathbf{x}_i)$, where $\hat{f}_i(\mathbf{x}_i) = \mathbb{E}_{\theta\in[0,1]}[f(\mathbf{x}_i^\theta)]$. How do we round while preserving the objective function? If we focus on a specific i, the objective function suggests that we pick θ randomly from $[0,1]$ and assign the elements in $\mathbf{x}_i^\theta$ to i; we call this θ-rounding. However, there are two issues to contend with. First, if we independently round for each i then the same element may be assigned multiple times. Second, we need to ensure that all elements are assigned, which is not guaranteed by the θ-rounding. We remark that there is an integrality gap example for hypergraph metric labeling that shows that there is no effective rounding strategy that works in general.

Our approach is to understand the rounding process by considering various special cases of interest. In particular, we consider monotone functions, symmetric functions, the hypergraph cut function (which is asymmetric), and combinations of such functions. Monotonocity helps in that if elements are assigned to a label i, they can be removed without increasing the fractional cost. Although one can use different strategies to obtain an $O(\log n)$-approximation and integrality gap, a useful strategy here is the rounding of Kleinberg and Tardos [14] that they introduced for metric labeling. This has the additional property of ensuring that an element u is assigned to i with probability exactly $x(u,i)$. We then consider the rounding process for Sub-MP, in particular the symmetric case Sym-Sub-MP. Here, we crucially take advantage of the fact that there is a single underlying function f, and moreover the fact that it is symmetric. We consider the CKR-Rounding strategy from [2] and show its effectiveness for hypergraphs by abstracting away some of the properties specific to graphs that were previously exploited in the analysis. In the process, we also observe that a variant is equally effective for graphs but is more insightful for Sym-Sub-MP.

Finally, Sub-Label combines a monotone function and a non-monotone function. Here, we resort to KT-Rounding since it is a reasonable strategy to approximately preserve the cost of the monotone component. For the uniform metric labeling problem, [14] showed that KT-Rounding approximately (to within a factor of 2) preserves the fractional connection cost in the case of graphs. We show bounds for hypergraph cut functions in an analogous fashion. Our insights enable us to develop a variant of the rounding that gives an $O(\log n)$-approximation for Sub-Label when the cut function is an arbitrary symmetric submodular function. Due to space constraints, we omit the results for metric labeling.

Other Related Work: Due to space constraints, in this extended abstract we restrict our attention to closely related work. There has been much recent interest in optimizing with submodular set functions. In particular, maximization problems have been examined via combinatorial techniques as well as the multilinear relaxation [1]. The submodular welfare problem [22] is similar in spirit to MSCA except that one is interested in maximizing the value of an allocation rather than minimizing the cost. Minimization problems with submodular costs have also received substantial attention [19, 11, 12, 10] with several negative results for basic problems as well as positive approximation results for problems such as the submodular cost vertex cover problem [12, 10]. Lovász-extension based convex programs have been effectively used for these problems. Various submodular cut and partition problems and their special cases such as the hypergraph cut and partition have been studied recently [26, 25, 17, 7]; however, these papers have typically focussed on greedy and divide-and-conquer based approaches while we use LE-Rel.

Recent Results for Sym-Sub-MP and Sub-MP: Very recently, building on the work in this paper and a non-trivial new technical theorem, we showed [3] that the integrality gap of SubMP-Rel is at most $1.5 - 1/k$ for Sym-Sub-MP and at most 2 for Sub-MP.

2 Monotone MSCA

KT-Rounding

```
let x be a solution to LE-Rel
S ← ∅        ⟨⟨set of all assigned vertices⟩⟩
⟨⟨set of vertices that are eventually assigned to i⟩⟩
A_i ← ∅ for all i (1 ≤ i ≤ k)
while S ≠ V
      pick i ∈ {1, 2, ··· , k} uniformly at random
      pick θ ∈ [0, 1] uniformly at random
      A_i ← A_i ∪ ({v | x(v, i) ≥ θ} − S)
      S ← S ∪ A_i
return (A_1, ··· , A_k)
```

In this section we consider Monotone-MSCA where $f_1, \ldots, f_k$ are monotone submodular functions. We will assume for simplicity that $f_i(\emptyset) = 0$ for all i. Svitkina and Tardos [21] considered this problem in the context of facility location and gave a $(1+\ln n)$-approximation and matching hardness via an approximation preserving reduction from set cover. Let $\alpha = \min_{S \subseteq V, 1 \leq i \leq k} f_i(S)/|S|$. The main observation in [21] is that $\alpha \leq \mathrm{OPT}/n$, and moreover a pair (S, i) such that $f_i(S)/|S| = \alpha$ can be computed in polynomial-time via submodular function minimization. One can then iterate using a greedy scheme, by using the monotonicity of the functions, to obtain a $(1 + \ln n)$-approximation. Using a similar argument, we can prove the following theorem.

Theorem 1. *The integrality gap of* LE-Rel *for* Monotone-MSCA *is at most* $(1 + \ln n)$. *In particular,* $\alpha \leq \mathrm{OPT}_{\text{FRAC}}/n$.

We consider other rounding algorithms that also achieve an $O(\log n)$-approximation. We focus on KT-Rounding derived from the work of Kleinberg and Tardos on metric labeling [14].

Theorem 2. KT-Rounding *achieves a randomized $O(\ln n)$-approximation for* Monotone MSCA.

3 Submodular Multiway Partition

We consider MSCA when the f_i can be non-monotone. We can show that the integrality gap of LE-Rel even for a special case of labeling on hypergraphs can be $\Omega(n)$, and we suspect that the problem is hard to approximate to a polynomial factor in n. We therefore focus on Submodular Multiway Partition (Sub-MP) and Submodular Cost Labeling (Sub-Label); these are broad special cases which capture several problems that have been considered previously.

The reduction of Sub-MP to MSCA requires one to work with the non-terminals V' as the ground set. It is however more convenient to work with the terminals and non-terminals. In particular, we work with the relaxation below. Recall that $\mathbf{x}_i = (x(v_1, i), \cdots, x(v_n, i))$.

SubMP-Rel

$$\begin{array}{llll} \min & \displaystyle\sum_{i=1}^{k} \hat{f}(\mathbf{x}_i) & & \\ & \displaystyle\sum_{i=1}^{k} x(v,i) & = 1 & \forall v \\ & x(s_i, i) & = 1 & \forall i \\ & x(v,i) & \geq 0 & \forall v, i \end{array}$$

As before, a starting point for rounding the relaxation is the basic θ-rounding that preserves the objective function. Suppose we do θ-rounding for each i to obtain sets $A(1, \theta), \ldots, A(k, \theta)$ where each $A(i, \theta) \subseteq V$. Here we could use independent random θ values for each i or the same θ. Note that the constraints ensure that $s_i \in A(j, \theta)$ iff $i = j$. However, the sets $A(1, \theta), \ldots, A(k, \theta)$ may intersect and also may not cover the entire set V, in which case we have to allocate the remaining elements in some fashion. First we show how to take advantage of the case when f is symmetric and then discuss how to obtain results for hypergraph problems that are special cases of Sub-MP.

A $2(1-1/k)$-approximation for Sym-Sub-MP: A $2(1-1/k)$-approximation for Sym-Sub-MP is known via greedy combinatorial algorithms [18, 26]. However, no mathematical programming formulation for the problem has been previously considered. Here we show that, on instances of Sym-Sub-MP, the integrality gap of LE-Rel is $2(1-1/k)$ by using an uncrossing property of symmetric functions.

The following lemma is standard and it has been used in previous work [20].

Lemma 1. *Let f be a symmetric submodular set function over V and let $A_1, \ldots, A_k$ be subsets of V. Then there exist sets $A'_1, \ldots, A'_k$ such that (i) $A'_i \subseteq A_i$ for $1 \leq i \leq k$, (ii) $A'_1, \ldots, A'_k$ are mutually disjoint (iii) $\cup_i A'_i = \cup_i A_i$ and (iv) $\sum_i f(A'_i) \leq \sum_i f(A_i)$. Moreover, given the A_i's a collection of sets A'_i satisfying the above properties can be found in polynomial time via a value oracle for f.*

Theorem 3. *The integrality gap of* LE-Rel *for* Sym-Sub-MP *is* $\leq 2(1-1/k)$.

In an earlier version of this paper, we raised the following question.

Question. Is the integrality gap of LE-Rel for Sym-Sub-MP at most 1.5?

As we already noted, we have shown in subsequent work [3].

Rounding for Hypergraph-MC and Hypergraph-MP: Calinescu *et al.* [2] gave a new geometric relaxation for Graph-MC, and a rounding procedure that gave a $(1.5 - 1/k)$-approximation; the integrality gap was subsequently improved to a bound of $1.3438 - \varepsilon_k$ in [13], while the best known lower bound is $8/(7+1/k-1)$ [6]. Calinescu *et al.* [2] derived their relaxation as a way to improve the integrality gap of $2(1-1/k)$ for a natural distance based linear programming relaxation; in fact, it often goes unnoticed that [2] shows the equivalence of their geometric relaxation to that of another relaxation obtained by adding valid strengthening constraints to the distance based relaxation. Interestingly, when we specialize MSCA to Graph-MC, LE-Rel becomes the geometric relaxation of [2]! The rounding procedure in [2] can be naturally extended to rounding LE-Rel for Sub-MP and we describe it below.

CKR-Rounding

let x be a solution to SubMP-Rel
pick a random permutation π of $\{1, 2, \cdots, k\}$
pick $\theta \in [0, 1)$ uniformly at random
$S \leftarrow \emptyset$ ⟨⟨*set of all assigned vertices*⟩⟩
for $i = 1$ to $k - 1$
  $A_{\pi(i)} \leftarrow (\{v \mid x(v, \pi(i)) \geq \theta\} - S)$
  $S \leftarrow S \cup A_{\pi(i)}$
$A_{\pi(k)} \leftarrow V - S$
return $(A_1, \cdots, A_k)$

CKR-Rounding uses the same θ for all i and a random permutation, both of which are crucially used in the 1.5-approximation analysis for Graph-MC. In this paper we investigate CKR-Rounding and other roundings for Hypergraph-MC and Hypergraph-MP.

Although Hypergraph-MC and Hypergraph-MP appear similar, their objective functions are different. The objective of Hypergraph-MC is to remove a minimum weight subset of hyperedges such that the terminals are separated, whereas the objective of Hypergraph-MP is to minimize $\sum_e w(e)p(e)$, where $p(e)$ is the number of non-trivial parts that e is partitioned into (a part is non-trivial if some vertex of e is in that part but not all of e). For graphs we have that either $p(e) = 0$ or $p(e) = 2$, and therefore the two problems Hypergraph-MC and Hypergraph-MP are equivalent; this is the reason why one can view Graph-MC as a partition problem as well. However, when the hyperedges can have size larger than 2, the objective function values are not related to each other (it is easy to see that the Hypergraph-MP objective is always larger).

Hypergraph-MP and Hypergraph-MC have been studied for their theoretical interest and their applications. It is easy to see from its definition that Hypergraph-MP is a special case of Sym-Sub-MP. It has been observed by a simple yet nice reduction [17] that Hypergraph-MC is a special case of Sub-MP. In addition, it has been observed that Hypergraph-MC is approximation-equivalent to the *node*-weighted multiway cut problem in graphs (Node-wt-MC) [8].

We show that CKR-Rounding gives a $(1.5-1/k)$-approximation to Hypergraph-MP and a tight H_Δ-approximation for Hypergraph-MC with maximum hyperedge size Δ. Note that when $\Delta = 2$, $H_\Delta = 1.5$ and when $\Delta = 3$, $H_\Delta \simeq 1.833$. For $\Delta > 3$, CKR-Rounding gives a worse than 2 bound while we give an alternate rounding which gives a $2(1-1/k)$-approximation. Our analysis of CKR-Rounding differs from that in [2] since we cannot use the "edge alignment" properties of the fractional solution that hold for graphs and were exploited in [2]; our analysis is inspired by the proof given by Williamson and Shmoys [24].

It is natural to wonder whether CKR-Rounding is crucial to obtaining a bound that is better than 2 for these problems, and in particular whether it gives a 1.5-approximation for Sym-Sub-MP. We show that a $1.5-1/k$-approximation for Hypergraph-MP (and hence Graph-MC also) can be obtained via a different algorithm as well; in particular, the crucial ingredient in CKR-Rounding for Graph-MC when viewed as a special case of Hypergraph-MP is the correlation provided by the use of the same θ for all i; one can replace the random permutation by the uncrossing scheme in Lemma 1. We describe this algorithm in the next section. However, for Hypergraph-MC, the random permutation is important in proving the H_Δ-bound.

3.1 A 1.5-approximation for Hypergraph Multiway Partition

We start by understanding the objective function of SubMP-Rel in the context of Hypergraph-MP. Let $\mathbf{x}$ be a feasible fractional solution, and let $\mathbf{x}_i = (x(v_1, i), \cdots, x(v_n, i))$ be the allocation to i. Recall that f here is the hypergraph cut function. What is $\sum_{i=1}^{n} \hat{f}(\mathbf{x}_i)$? For each terminal i and each hyperedge e, let $I(e,i) = [\min_{v \in e} x(v,i), \max_{v \in e} x(v,i)]$. Let $d(e,i)$ denote the length of $I(e,i)$, and let $d(e) = \sum_{i=1}^{k} d(e,i)$. Note that $d(e) \in [0, |e|]$.

Lemma 2. $\sum_{i=1}^{k} \hat{f}(\mathbf{x}_i) = \sum_e w(e)d(e)$.

A crucial technical lemma that we need is the following which states that the contribution of any i to $d(e)$ is at most $d(e)/2$.

Lemma 3. *For any* i, $d(e,i) \leq d(e)/2$.

```
SymSubMP-Rounding
  let x be a feasible solution to SubMP-Rel
  pick θ ∈ [0,1] uniformly at random
  A(i,θ) ← {v | x(v,i) ≥ θ} for each i (1 ≤ i ≤ k)
  ⟨⟨uncross A(1,θ), ··· , A(k,θ)⟩⟩
  A'_i ← A(i,θ) for each i (1 ≤ i ≤ k)
  while there exist i ≠ j such that A'_i ∩ A'_j ≠ ∅
      if (f(A'_i) + f(A'_j − A'_i) ≤ f(A'_i) + f(A'_j))
          A'_j ← A'_j − A'_i
      else
          A'_i ← A'_i − A'_j
  return (A'_1, ··· , A'_{k−1}, V − (A'_1 ∪ ··· A'_{k−1}))
```

The algorithm SymSubMP-Rounding that we analyze is described in the adjacent box. We can prove that CKR-Rounding gives the same bound; however, SymSubMP-Rounding and its analysis are perhaps more intuitive in the context of symmetric functions. The algorithm does θ-rounding to obtain sets $A(1,\theta), \ldots, A(k,\theta)$ and then uncrosses these sets to make them disjoint without increasing the expected cost (see Lemma 1).

Theorem 4. SymSubMP-Rounding *achieves an* $(1.5 - 1/k)$*-approximation for* Hypergraph-MP.

Lemma 4. *Let* i^* *be the index such that the interval* $I(e,i^*)$ *has the rightmost ending point among the intervals* $I(e,i)$*. More precisely,* $I(e,i^*)$ *is an interval such that* $\max_{v\in e} x(v,i^*) = \max_i \max_{v\in e} x(v,i)$*; if there are several such intervals, we choose one arbitrarily. Let* Z_e *be an indicator random variable equal to* 1 *iff* $e \in \delta(V - (A(1,\theta) \cup \cdots \cup A(k,\theta)))$*. Then* $\mathbb{E}[Z_e] \leq d(e,i^*)$.

Theorem 4 follows from Lemma 1 and Lemma 4.

3.2 Algorithms for Hypergraph Multiway Cut

Now we consider Hypergraph-MC. For each hyperedge e, pick an arbitrary representative node $r(e) \in e$. Define the function $f : 2^V \to \mathbb{R}_+$ as follows: for $A \subseteq V$, let $f(A) = \sum_{e: r(e)\in A, e \not\subseteq A} w(e)$ be the weight of hyperedges whose representatives are in A and they cross A. It is easy to verify that f is asymmetric and submodular. Sub-MP with this function f captures Hypergraph-MC [17].

Let $\mathbf{x}$ be a feasible fractional allocation and $\mathbf{x}_i$ be the allocation for i. For each hyperedge e and each terminal i, let $I(e,i) = [\min_{v\in e} x(v,i), \max_{v\in e} x(v,i)]$. Let $d(e,i) = x(r(e),i) - \min_{v\in e} x(v,i)$ and $d(e) = \sum_{i=1}^{k} d(e,i)$.

Lemma 5. $\sum_{i=1}^{k} \hat{f}(\mathbf{x}_i) = \sum_e w(e)d(e)$.

Half-Rounding

```
let x be a solution to SubMP-Rel
pick θ ∈ (1/2, 1] uniformly at random
for i = 1 to k − 1
      A(i, θ) ← {v | x(v, i) ≥ θ}
U(θ) ← V − (A(1, θ) ∪ · · · ∪ A(k − 1, θ))
return (A(1, θ), · · · , A(k − 1, θ), U(θ))
```

For Hypergraph-MC we show that Half-Rounding achieves a $2(1-1/k)$-approximation and that CKR-Rounding achieves an H_{Δ}-approximation where Δ is the maximum hyperedge size.

Theorem 5. *Let F be the set of all hyperedges crossing the partition returned by* Half-Rounding. *For each hyperedge e, $\Pr[e \in F] \leq 2d(e)$.*

Theorem 6. *Let F be the set of all hyperedges crossing the partition returned by* CKR-Rounding. *For each hyperedge e, $\Pr[e \in F] \leq H_{|e|} \cdot d(e)$. Moreover, this analysis is tight for the rounding.*

Acknowledgments: We thank Lisa Fleischer for suggesting that we contact Zoya Svitkina about MSCA and thank Zoya for pointing out her work in [21] on monotone MSCA. CC thanks Jan Vondrak for pointing out the interpretation of the Lovász extension from his paper [23] which was very helpful in thinking about rounding procedures. AE thanks Sungjin Im and Ben Moseley for discussions.

References

1. Gruia Calinescu, Chandra Chekuri, Martin Pál, and Jan Vondrák. Maximizing a submodular set function subject to a matroid constraint (extended abstract). In *Integer Programming and Combinatorial Optimization (IPCO)*, pages 182–196, 2007.
2. Gruia Calinescu, Howard J. Karloff, and Yuval Rabani. An improved approximation algorithm for multiway cut. *Journal of Computer and System Sciences*, 60(3):564–574, 2000. Preliminary version in STOC 1998.
3. Chandra Chekuri and Alina Ene. Approximation algorithms for submodular multiway partition. Manuscript, April 2011.
4. Elias Dahlhaus, David S. Johnson, Christos H. Papadimitriou, Paul D. Seymour, and Mihalis Yannakakis. The complexity of multiterminal cuts. *SIAM Journal on Computing*, 23(4):864–894, 1994. Preliminary version in STOC 1992.
5. Andrew Delong, Anton Osokin, Hossam N. Isack, and Yuri Boykov. Fast approximate energy minimization with label costs. In *IEEE Computer Vision and Pattern Recognition (CVPR)*, pages 2173–2180, 2010.
6. Ari Freund and Howard J. Karloff. A lower bound of 8/(7+(1/k)-1) on the integrality ratio of the calinescu-karloff-rabani relaxation for multiway cut. *Information Processing Letters*, 75(1-2):43–50, 2000.
7. Takuro Fukunaga. Computing Minimum Multiway Cuts in Hypergraphs from Hypertree Packings. *Integer Programming and Combinatorial Optimization (IPCO)*, pages 15–28, 2010.
8. Naveen Garg, Vijay V. Vazirani, and Mihalis Yannakakis. Multiway cuts in node weighted graphs. *Journal of Algorithms*, 50(1):49–61, 2004. Preliminary version in ICALP 1994.

9. Dongdong Ge, Yinyu Ye, and Jiawei Zhang. The Fixed-Hub Single Allocation Problem: A Geometric Rounding Approach. Preprint available at `http://www.stanford.edu/~yyye/revisedHub.pdf`, 2007.
10. Gagan Goel, Chinmay Karande, Pushkar Tripathi, and Lei Wang. Approximability of combinatorial problems with multi-agent submodular cost functions. In *IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 755–764, 2009.
11. Michel X. Goemans, Nicholas J. A. Harvey, Satoru Iwata, and Vahab S. Mirrokni. Approximating submodular functions everywhere. In *ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 535–544, 2009.
12. Satoru Iwata and Kiyohito Nagano. Submodular function minimization under covering constraints. In *IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 671–680, 2009.
13. David R. Karger, Philip N. Klein, Clifford Stein, Mikkel Thorup, and Neal E. Young. Rounding algorithms for a geometric embedding of minimum multiway cut. *Mathematics of Operations Research*, 29(3):436–461, 2004. Preliminary version in STOC 1999.
14. Jon M. Kleinberg and Éva Tardos. Approximation algorithms for classification problems with pairwise relationships: Metric labeling and Markov random fields. *Journal of the ACM (JACM)*, 49(5):616–639, 2002. Preliminary version in FOCS 1999.
15. Eugene L. Lawler. Cutsets and partitions of hypergraphs. *Networks*, 3(3):275–285, 1973.
16. László Lovász. Submodular functions and convexity. *Mathematical programming: the state of the art*, pages 235–257, 1983.
17. Kazumasa Okumoto, Takuro Fukunaga, and Hiroshi Nagamochi. Divide-and-conquer algorithms for partitioning hypergraphs and submodular systems. *Algorithmica*, pages 1–20, 2010. Preliminary version in ISAAC 2009.
18. Maurice Queyranne. Minimizing symmetric submodular functions. *Mathematical Programming*, 82(1):3–12, 1998. Preliminary version in SODA 1995.
19. Zoya Svitkina and Lisa Fleischer. Submodular approximation: Sampling-based algorithms and lower bounds. In *IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 697–706, 2008.
20. Zoya Svitkina and Éva Tardos. Min-max multiway cut. *Approximation, Randomization, and Combinatorial Optimization (APPROX-RANDOM)*, pages 207–218, 2004.
21. Zoya Svitkina and Éva Tardos. Facility location with hierarchical facility costs. *ACM Transactions on Algorithms (TALG)*, 6(2):1–22, 2010. Preliminary version in SODA 2006.
22. Jan Vondrák. Optimal approximation for the submodular welfare problem in the value oracle model. In *ACM Symposium on Theory of Computing (STOC)*, pages 67–74, 2008.
23. Jan Vondrák. Symmetry and Approximability of Submodular Maximization Problems. In *IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 651–670, 2010.
24. David P. Williamson and David B. Shmoys. The design of approximation algorithms. Preprint available at `http://www.designofapproxalgs.com`, 2010.
25. Mingyu Xiao. Finding minimum 3-way cuts in hypergraphs. *Information Processing Letters*, 110(14-15):554–558, 2010. Preliminary version in TAMC 2008.
26. Liang Zhao, Hiroshi Nagamochi, and Toshihide Ibaraki. Greedy splitting algorithms for approximating multiway partition problems. *Mathematical Programming*, 102(1):167–183, 2005.