A Local Search-Based Approach for Set Covering

Anupam Gupta^{*} Euiv

Euiwoong Lee[†]

Jason Li[‡]

Abstract

In the Set Cover problem, we are given a set system with each set having a weight, and we want to find a collection of sets that cover the universe, whilst having low total weight. There are several approaches known (based on greedy approaches, relax-and-round, and dual-fitting) that achieve a $H_k \approx \ln k + O(1)$ approximation for this problem, where the size of each set is bounded by k. Moreover, getting a $\ln k - O(\ln \ln k)$ approximation is hard.

Where does the truth lie? Can we close the gap between the upper and lower bounds? An improvement would be particularly interesting for small values of k, which are often used in reductions between Set Cover and other combinatorial optimization problems.

We consider a non-oblivious local-search approach: to the best of our knowledge this gives the first H_k -approximation for Set Cover using an approach based on localsearch. Our proof fits in one page, and gives a integrality gap result as well. Refining our approach by considering larger moves and an optimized potential function gives an $(H_k - \Omega(\log^2 k)/k)$ -approximation, improving on the previous bound of $(H_k - \Omega(1/k^8))$ (R. Hassin and A. Levin, SICOMP '05) based on a modified greedy algorithm.

^{*}Carnegie Mellon University, Pittsburgh PA 15217. Supported in part by NSF awards CCF-1955785, CCF-2006953, and CCF-2224718.

[†]University of Michigan, Ann Arbor MI 48105. Partially supported by Google.

[‡]Simons Institute for the Theory of Computing and UC Berkeley, Berkeley CA 94720.

1 Introduction

(Weighted) Set Cover is one of the most important problems in the approximation algorithms literature. Given a set system (U, S) where each set $S \in S$ has weight w(S) > 0, the Set Cover problem asks to find a subcollection $\mathcal{F} \subseteq S$ that covers the universe (i.e., $\bigcup_{S \in \mathcal{F}} S = U$) while minimizing the total weight $\sum_{S \in \mathcal{F}} w(S)$. This problem is NP-hard as long as the sets have size at least three (the edge-cover problem can be solved in polynomial time). As the flagship problem in two standard textbooks in approximation algorithms [Vaz01, WS11], and as an abstract setting capturing numerous covering problems, it has always been an important testbed for new algorithmic techniques.

Let k-Set Cover be the special case of Set Cover where every set has at most k elements. The simple greedy algorithm that iteratively selects the set maximizing the current density (i.e., the ratio of the number of uncovered elements in the set to its weight) guarantees an H_k -approximation, where $H_k = 1 + \frac{1}{2} + \dots + \frac{1}{k} = \ln k + O(1)$ is the kth harmonic number [Joh74, Lov75, Chv79]. It can be analyzed by the dual-fitting method, upper bounding the integrality gap of the standard LP relaxation by H_k as well. Another proof of the integrality gap comes via the relax-and-round approach ([You22], see also §B). These algorithms are almost optimal due to the $(1 - o(1)) \ln n$ -hardness of Feige [Fei98] and its refinement to the $\ln k - O(\ln \ln k)$ -hardness for k-Set Cover [Tre01], which even holds against $n^{f(k)}$ -time algorithms for any computable function f.

How about local search, one of the most intuitive and popular algorithm design techniques? It maintains a solution (a set cover in this case), and in each iteration, it tries to find a *local move* that swaps at most p sets between the current solution and the remaining sets. If there exists a local move that results in a better set cover, execute the local move; otherwise, output the current set system. While it has been successfully applied for many problems including bounded-degree network design [FR92, FR94], Facility Location [KPR00, CG05] and k-Median [AGK+04, GT08], the local search cannot yield any finite approximation ratio for Set Cover, at least when the local width is one; simply consider the example where the universe has k elements, the optimal solution contains k singleton sets of weight ε , but the current solution consists of the set containing all the elements but has large weight 1. As $\varepsilon \to 0$ for fixed k, the gap between the two solutions becomes unbounded.

Is there a way to "redeem" local search? One reason that the above example is bad for local search is that the *potential* that the standard local search is trying to optimize, which is the same as the total weight of the current solution, is too rigid; while adding a singleton set from the optimal solution can be seen as a *progress*, since the large set is still needed to cover all elements, adding the singleton set only worsens the total weight and is not executed. To fix this issue, *non-oblivious local search (NOLS)* tries to find a local move improving a carefully designed potential different from the objective function of the problem. Originally defined by Khanna, Motwani, Sudan, and Vazirani [KMSV98], it has been recently shown to work well for problems including Submodular Optimization [FW14], Tree Augmentation and Steiner Tree [TZ22] (which directly inspired this paper), Steiner forest [GGK⁺18] and *k*-Median [CAGH⁺22]. Our first result is the following "redemption" of local search for Set Cover showing that NOLS with a natural potential can exactly match the H_k -approximation guarantee, up to an arbitrarily small constant error $\varepsilon > 0$ that ensures that the local search terminates in polynomial time.

Theorem 1.1. For any $\varepsilon > 0$, there exists a width-1 non-oblivious local search algorithm that

can be implemented in time $poly(n, 1/\varepsilon)$ and yields an $H_k + \varepsilon$ -approximation for k-Set Cover.

Our proof also shows that any local optimum has weight at most H_k times an optimal solution to the LP relaxation for set cover, thereby giving yet another proof of its integrality gap.

We then explore the power of NOLS beyond the H_k -approximation. While Trevisan's $(\ln k - O(\ln \ln k))$ -hardness shows that we cannot dramatically improve it, there are still unanswered questions, especially for small values of k, which are important for hardness of other well-known problems including Steiner Tree [BP89, Thi01]. For unweighted k-Set Cover, there is a long series of works [GHY93, Hal95, Hal96, DF97, Lev09, ACK09, FY11] giving an $(H_k - \beta_k)$ -approximation where $\beta_k \geq 0.5$ for every $k \geq 3$ and approaches to 0.6402. (Some of these works even use a combination of oblivious local search to solve a packing problem, and greedy to extend the solution to a matching.)

But the status for weighted k-Set Cover—which is the problem we focus on—is understood far poorly. The best approximation ratio remains $H_k - \Omega(1/k^8)$ [HL05], obtained by a variant of the greedy algorithm. We make progress on this direction, and prove the following improved approximation guarantees for k-Set Cover.

Theorem 1.2. For any $\varepsilon > 0$, there exist width-2 and width-k non-oblivious local search algorithms for k-Set Cover that yield $(H_k - \Omega(1/k) + \varepsilon)$ and $(H_k - \Omega((\log k)^2/k) + \varepsilon)$ -approximations respectively.

Note that a width-2 local search can be implemented in time $poly(n, 1/\varepsilon)$ and a width-k one can be naïvely implemented in time $n^{O(k)} poly(1/\varepsilon)$. We present clean locality gap results for width 1, 2, and k swaps in Section 2, 3, 4 respectively and show how to implement them in polynomial time in Section 5. In Section 6, we provide matching lower bounds showing that these two results are tight for a large class of natural potentials.

2 Set Cover

Consider a weighted set system (U, S) with weights $w : S \to \mathbb{R}^+$ where each $S \in S$ has cardinality at most k. Define the downwards closure S^{\downarrow} of the set system as containing all sets $\{T \mid \exists S \in S, T \subseteq S\}$, where each subset has the same cost as the original set. Letting $S \leftarrow S^{\downarrow}$ does not change the optimal value, so we assume that S is downwards-closed. With this assumption, we can further assume the optimal solution \mathcal{F}^* forms a partition of the universe U, and our algorithms will maintain the solution \mathcal{F} that also forms a partition of U. (Letting $S \leftarrow S^{\downarrow}$ might significantly increase the number of sets. In Section 5, we show how to efficiently implement it.)

For any collection \mathcal{F} of sets that partition U, define the Rosenthal potential [Ros73]:

$$\Phi(\mathcal{F}) := \sum_{S \in \mathcal{F}} w(S) H_{|S|}.$$
(1)

For each element e, let $\bar{w}_{\mathcal{F}}(e) := \frac{w(S)}{|S|}$ for the set $S \in \mathcal{F}$ that covers e. (We omit the subscript \mathcal{F} if it is clear from context.) Then $\sum_{e} \bar{w}(e) = w(\mathcal{F})$.

2.1 An H_k -competitive Local Search Algorithm

Consider the following single (set) local moves:

Add in a single set $S \in S^{\downarrow}$, and then for each $T \in \mathcal{F}$ in the current solution, replace T by $T \setminus S$ to get back a new set cover solution that is a partition.

If this move decreases the Rosenthal potential—i.e., if this is an *improving* local move—we move to this resulting solution.

Theorem 2.1 (Single-Set Moves). Suppose \mathcal{F} is a local optimum, i.e., there are no improving local moves. Then $w(\mathcal{F}) \leq H_k \cdot w(\mathcal{F}^*)$.

Proof. To show the locality gap, we consider a specific set of local moves (called *test moves*). Since there are no improving local moves, each of these test moves do not reduce the potential, thereby giving us relationships between the costs of some solutions related to the local and optimal solution. Combining these then proves the theorem.

Indeed, consider using any of the sets in the optimal solution $S \in \mathcal{F}^*$ as a local move from \mathcal{F} . The resulting potential function change is

$$w(S)H_{|S|} - \sum_{T \in \mathcal{F}} w(T) \left[H_{|T|} - H_{|T \setminus S|} \right] \ge 0.$$

Since \mathcal{F} is a partition of U, the second term on the LHS is

$$\sum_{T \in \mathcal{F}} \sum_{i=0}^{|T \cap S| - 1} \frac{w(T)}{|T| - i} \ge \sum_{T \in \mathcal{F}} \frac{w(T)}{|T|} |T \cap S| = \sum_{T \in \mathcal{F}} \sum_{e \in S \cap T} \bar{w}(e) = \sum_{e \in S} \bar{w}(e).$$
(2)

Therefore we have for each $S \in \mathcal{F}^*$ that

$$w(S)H_k - \sum_{e \in S} \bar{w}(e) \ge 0.$$
(3)

Summing over all sets in \mathcal{F}^* , which we also imagine is a partition, we get

$$H_k \sum_{S \in \mathcal{F}^*} w(S) - \sum_e \bar{w}(e) \ge 0 \qquad \Longrightarrow \qquad w(\mathcal{F}) \le H_k w(\mathcal{F}^*). \qquad \square$$

2.2 An Integrality Gap Result

A small change bounds the cost against any solution to the standard linear programing relaxation:

$$\min \Big\{ \sum_{S} w(S) x_S \mid \sum_{S:e \in S} x_S \ge 1, x \ge 0 \Big\}.$$

Indeed, suppose x^* is any feasible solution, and we consider local moves with each of the sets sets in the support of x^* . Multiplying (3) with x_S^* and summing gives

$$\sum_{S} H_{|S|} w(S) x_{S}^{*} - \sum_{e} \bar{w}(e) \sum_{S:e \in S} x_{S}^{*} \ge 0.$$
(4)

But $\sum_{S:e\in S} x_S^* \ge 1$ by feasibility of the LP, so we infer that

$$\sum_{e} \bar{w}(e) = w(\mathcal{F}) \le \sum_{S \in \mathcal{F}^*} w(S) x_S^* H_{|S|} \le H_k \cdot (w^{\mathsf{T}} x^*).$$

3 An Improvement Using Double Moves

The above analysis suggests one avenue for improvement: if the move adding set $S \in \mathcal{F}^*$ removes more than one element from some set $T \in \mathcal{F}$, then the inequality (2) bounds the decrease in potential by $\frac{|T \cap S|}{|T|}$, whereas the actual decrease is $H_{|T|} - H_{|T \setminus S|}$, which is possibly greater. Concretely, if |T| = k and $|T \cap S| = 2$, then we claim an improvement of 2/k, whereas the actual improvement is 1/k + 1/k-1. In this section we show how this idea can be used to get an improvement.

The algorithm is now a natural "width-two" generalization of the above local search:

Add in two sets $S, S' \in S^{\downarrow}$ to \mathcal{F} , and replace each existing set $T \in \mathcal{F}$ by $T \setminus (S \cup S')$. If the resulting partition has a smaller potential value, move to it.

We allow S = S', which captures the case of adding a single set. Hence local optima with these moves have cost at most $H_k w(\mathcal{F}^*)$ by the previous section; we want to show a better bound. Let us first do this for a special case, and then show how to remove this assumption (in Lemma 3.2).

Theorem 3.1 (Double Moves). Consider a solution \mathcal{F} that is a local optimum for the above width-two local search with the Rosenthal potential. Let $\mathcal{F}_1 \subseteq \mathcal{F}$ be the subcollection of sets in \mathcal{F} having unit size, and suppose $w(\mathcal{F}_1) \leq 0.99 w(\mathcal{F})$. Then

$$w(\mathcal{F}) \le H_k(1 - \Theta(1/k^2)) \cdot w(\mathcal{F}^*).$$

Proof. The proof again goes via analyzing a collection of test moves; these try to add in at most two sets at a time. To get the *test moves*, consider a bipartite graph whose nodes are the sets in \mathcal{F}^* and those in \mathcal{F} , and there is an edge between $S \in \mathcal{F}^*$ and $T \in \mathcal{F}$ iff $S \cap T \neq \emptyset$. For a vertex S, let \mathcal{N}_S be the set of its neighbors. There are two kinds of test moves:

- 1. For a set $T \in \mathcal{F}$, let $S_0, S_1, \ldots, S_{\ell-1} \in \mathcal{F}^*$ be its neighbors in an arbitrary order. If $\ell = 1$, then try to add in S_0 twice. Else, for each index $0 \leq i < \ell$, try to add in S_i and $S_{(i+1) \mod \ell}$ together.
- 2. A set $S \in \mathcal{F}^*$ is added exactly $2|\mathcal{N}_S|$ times above, twice for each $T \in \mathcal{N}_S$. Add in S another $2k 2|\mathcal{N}_S|$ number of times.

The local optimality ensures that none of the moves above decreases the potential. Let us consider the total potential change caused by the above moves. Firstly, each set $S \in \mathcal{F}^*$ is added exactly 2k times, so the total potential increase by adding it is exactly $2kw(S)H_{|S|} \leq 2kH_k \cdot w(S)$. Moreover, let us consider the potential decrease due to the removal of elements from each set in \mathcal{F} . Indeed, for a set $T \in \mathcal{F}$, let ℓ denote its neighborhood size in the bipartite graph.

• If $\ell = 1$, then T is a subset of S_0 , its only neighbor. Thus the potential from T is decreased by $w(T)H_{|T|}$ twice, for a total of $2w(T)H_{|T|}$. If $\ell > 1$, then for each pair S, S' added together, the potential from T is decreased by $w(T)(H_{|T|} - H_{|T|-|(S \cup S') \cap T|})$. Since $|(S \cup S') \cap T| = |S \cap T| + |S' \cap T| \ge 2$, the average per-element decrease,

$$\frac{w(T)(H_{|T|} - H_{|T| - |(S \cup S') \cap T|})}{|(S \cup S') \cap T|},$$

is at least what it would be if $|(S \cup S') \cap T| = 2$, i.e., $w(T)(H_{|T|} - H_{|T|-2})/2$. So the overall decrease is at least

$$(|(S \cup S') \cap T|) w(T) (H_{|T|} - H_{|T|-2})/2.$$

Since the sum of $(|(S \cup S') \cap T|)$ over all added pairs S, S' is exactly 2|T|, the overall potential decrease is at least $|T|w(T)(H_{|T|} - H_{|T|-2})$.

If |T| = 1, then $\ell = 1$ and the potential decrease is $2w(T)H_{|T|} = 2w(T)$. If $|T| \ge 2$, then regardless of whether $\ell = 1$ or $\ell > 1$, the potential decrease is at least $|T|w(T)(H_{|T|} - H_{|T|-2})$.

• In total, each $S \in \mathcal{N}_T$ is added exactly 2k times, so each element of T is removed a total of 2k times. Two of these 2k removals are accounted above, so the remaining potential decrease over all elements is at least $w(T) \cdot (2k-2) \cdot |T| \cdot (H_{|T|} - H_{|T|-1}) = (2k-2)w(T)$.

The total potential decrease, which is at most 0, is at least

$$\sum_{T \in \mathcal{F}_1} 2w(T) + \sum_{T \in \mathcal{F} \setminus \mathcal{F}_1} |T| w(T) (H_{|T|} - H_{|T|-2}) + \sum_{T \in \mathcal{F}} (2k-2)w(T) - H_k \sum_{S \in \mathcal{F}^*} 2kw(S), \quad (5)$$

where \mathcal{F}_1 is the collection of sets in \mathcal{F} of unit size. Observe that

$$|T|(H_{|T|} - H_{|T|-2}) = |T|\left(\frac{1}{|T|} + \frac{1}{|T|-1}\right) = 2 + \frac{1}{|T|-1} \ge 2 + \frac{1}{k-1}.$$

By assumption, $w(\mathcal{F} \setminus \mathcal{F}_1) \geq 0.01 w(\mathcal{F})$, so the potential decrease is at least

$$\sum_{T \in \mathcal{F}} \left(2 + \frac{1}{100(k-1)} \right) w(T) + \sum_{T \in \mathcal{F}} (2k-2)w(T) - H_k \sum_{S \in \mathcal{F}^*} 2kw(S)$$
$$= \sum_{T \in \mathcal{F}} \left(2k + \frac{1}{100(k-1)} \right) w(T) - H_k \sum_{S \in \mathcal{F}^*} 2kw(S).$$

Since this decrease is at most 0, we conclude that

$$w(\mathcal{F}) \le \frac{2k}{2k + \frac{1}{(100(k-1))}} H_k \cdot w(\mathcal{F}^*) \le (1 - \Theta(\frac{1}{k^2})) H_k \cdot w(\mathcal{F}^*).$$

To get a better-than- H_k approximation for all instances, we can go two ways: the first approach is to post-process the locally optimal solution \mathcal{F} using the following lemma (which we prove in §A) to handle the case not handled by Theorem 3.1:

Lemma 3.2 (Post-processing). Given any solution \mathcal{F} with $w(\mathcal{F}) \leq H_k w(\mathcal{F}^*)$, let \mathcal{F}_1 be the subcollection of sets in \mathcal{F} of unit size. If $w(\mathcal{F}_1) \geq 0.99 w(\mathcal{F})$, there is an efficient algorithm that returns a new solution \mathcal{F}' with $w(\mathcal{F}') \leq 0.99 H_k \cdot w(\mathcal{F})$.

Hence, returning the better of the solutions \mathcal{F} produced by the local-search procedure, and \mathcal{F}' from using Lemma 3.2 applied to \mathcal{F} , gives a solution of cost at most

$$H_k \cdot (1 - \Theta(1/k^2)) \cdot w(\mathcal{F}^*).$$

The second—better and more principled approach—is to modify the potential function, which we do in the next section.

3.1 An Improved Analysis using a Custom Potential Function

Let us consider a somewhat generic potential function: define $f_1 := 1$, and let $f_i \ge 0$ be values to be fixed later, satisfying $f_i \ge f_{i+1}$ for all *i*. Let $F_i := \sum_{j=1}^i f_j$, and define the following potential

$$\Psi(\mathcal{F}) = \sum_{S \in \mathcal{F}} w(S) F_{|S|}.$$

We get back the Rosenthal potential by setting $f_i = 1/i$, but now we can optimize over settings of f_i to give better results. We again consider the two-set local search algorithm, trying to reduce the value of the new potential $\Psi(\mathcal{F})$. The test moves remain unchanged.

Moreover, the calculations remain essentially unchanged beyond replacing H_i by F_i ; the argument about the average per-element decrease being largest for $|(S \cup S') \cap T| = 2$ follows from the f_i values being non-increasing. Consequently, the total potential decrease, which is at most 0, is at least

$$\sum_{T \in \mathcal{F}_1} 2w(T) + \sum_{T \in \mathcal{F} \setminus \mathcal{F}_1} |T| w(T) (F_{|T|} - F_{|T|-2}) + \sum_{T \in \mathcal{F}} (2k-2) |T| w(T) f_{|T|} - F_k \sum_{S \in \mathcal{F}^*} 2k \, w(S).$$
(6)

This equation can be compared to (5), where we had used the fact that the Rosenthal potential satisfies $if_i = i(1/i) = 1$ and simplified the third summation above: $\sum_{T \in \mathcal{F}} (2k - 2)|T|w(T)f_{|T|} = \sum_{T \in \mathcal{F}} (2k - 2)w(T)$. Let us abstract (6) further: define

$$\alpha_t := \frac{1}{w(\mathcal{F})} \cdot \sum_{T \in \mathcal{F} : |T| = t} w(T),$$

and note that $\alpha = (\alpha_1, \ldots, \alpha_k)$ gives a probability distribution over the set sizes, and hence belongs to the probability simplex Δ_k . Dividing (6) through by 2k, simplifying slightly, and using that this decrease is at most zero gives

$$w(\mathcal{F}) \cdot \left(\alpha_1 + \sum_{t \ge 2} \underbrace{\frac{t(F_t - F_{t-2}) + (2k-2)tf_t}{2k}}_{=:\phi_t} \alpha_t\right) \le F_k w(\mathcal{F}^*).$$

Let $\phi_1 := 1$ and ϕ_t be the coefficient of α_t as shown above. Getting the best approximation becomes an optimization problem: we want to set f_i values to minimize $\max_{\alpha \in \Delta_k} \left\{ \frac{F_k}{\sum_t \phi_t \alpha_t} \right\}$, or equivalently, to minimize $F_k \cdot \max_t \left\{ \frac{1}{\phi_t} \right\}$. Recall that we require $f_1 = 1$; we set

$$f_t := \frac{1}{t} - \frac{1}{4kt(t-1)} \text{ for } t > 1.$$
(7)

Then for $t \geq 3$,

$$t(F_t - F_{t-2}) = t\left(\frac{1}{t} - \frac{1}{4kt(t-1)} + \frac{1}{t-1} - \frac{1}{4k(t-1)(t-2)}\right)$$
$$\ge 2 + \frac{1}{t-1} - \frac{1}{2k(t-1)} \ge 2 + \frac{1}{2(t-1)},$$

and the bound $t(F_t - F_{t-2}) \ge 2 + \frac{1}{2(t-1)}$ can be separately verified for t = 2. Observe that

 $\phi_1 = 1$ by definition, and for t > 1, we have

$$\phi_t = \frac{t(F_t - F_{t-2}) + (2k-2)tf_t}{2k} \ge \frac{1}{2k} \left(2 + \frac{1}{2(t-1)} + (2k-2)\left(1 - \frac{1}{4k(t-1)}\right)\right) \ge 1.$$

Hence, the approximation guarantee is at most $F_k \max_t(1/\phi_t) = F_k$. Finally, we bound $F_k \leq H_k - \frac{1}{8k}$ since for i = 2 alone, f_2 beats the corresponding term $\frac{1}{2}$ from H_k by $\frac{1}{8k}$.

Theorem 3.3 (Two-Sets Moves). Any local optimum for the two-sets local search using the potential Ψ using the f_i values from (8) satisfies $w(\mathcal{F}) \leq (H_k - 1/(8k)) \cdot w(\mathcal{F}^*)$.

4 Further Improvements: Moves with Width k

We now consider the "width-k" generalization: add in sets $S_1, S_2, \ldots, S_k \in S^{\downarrow}$ to \mathcal{F} , and replace each existing set $T \in \mathcal{F}$ by $T \setminus (S_1 \cup \cdots \cup S_k)$. If the resulting partition has a smaller potential value, move to it. (Once again, we allow repeats in the sets, or equivalently, we allow moves of fewer than k sets.)

Intuitively, if we choose sets S_1, S_2, \ldots, S_k that cover a set $T \in \mathcal{F}$ of size k, then T disappears from \mathcal{F} and the improvement to the Rosenthal potential is $1 + \frac{1}{2} + \cdots + \frac{1}{k}$, or an average of $\frac{1}{k} \cdot (1 + \frac{1}{2} + \cdots + \frac{1}{k})$, which is even better than the average $\frac{1}{2} \cdot (\frac{1}{k-1} + \frac{1}{k})$ in the width-2 case. We will actually use a custom potential function as before, but the Rosenthal potential provides a good baseline intuition.

We now define the test moves. For $T \in \mathcal{F}$, let $\mathcal{N}_T := \{S \in \mathcal{F}^* : |S \cap T| \neq 0\}$ and similarly, for $S \in \mathcal{F}^*$, let $\mathcal{N}_S := \{T \in \mathcal{F} : |S \cap T| \neq 0\}$.

- 1. For each $T \in \mathcal{F}$, add the sets in \mathcal{N}_T together.
- 2. For each $S \in \mathcal{F}^*$, add S. This move is multiplied by $(k |\mathcal{N}_S|)$ times so that each S participates in exactly k moves.

The local optimality ensures that none of the moves above decreases the potential. Let us consider the total potential change caused by the above moves. First, each $S \in \mathcal{F}^*$ is added exactly k times, so the total potential increase by adding it is exactly $kw(S)F_{|S|}$, where we define the custom potential function f_i later.

For $T \in \mathcal{F}$, we consider the two types of moves separately.

- For the move when \mathcal{N}_T is added together, T is removed from \mathcal{F} , so the potential from T is decreased by $w(T)F_{|T|}$.
- Other than this move, each $S \in \mathcal{N}_T$ is added exactly k-1 times, so each element of T is removed k-1 times more. Therefore, the total potential decrease for such moves is at least $w(T) \cdot (k-1) \cdot |T| \cdot (F_{|T|} F_{|T|-1})$.

Therefore, the total potential decrease, which is at most 0, is at least

$$\sum_{T \in \mathcal{F}} w(T)F_{|T|} + \sum_{T \in \mathcal{F}} (k-1)|T|w(T)f_{|T|} - F_k \sum_{S \in \mathcal{F}^*} k w(S).$$

We now follow the recipe from §3.1: dividing by k and defining the probability distribution $\alpha_t := \frac{1}{w(\mathcal{F})} \cdot \sum_{T \in \mathcal{F}: |T| = t} w(T)$ gives

$$w(\mathcal{F}) \cdot \left(\sum_{t \ge 1} \underbrace{\frac{F_t + (k-1)tf_t}{k}}_{=:\phi_t} \alpha_t\right) \le F_k w(\mathcal{F}^*).$$

Let ϕ_t be the coefficient of α_t as shown above. Getting the best approximation is again an optimization problem: we want to set f_i values to minimize $\max_{\alpha \in \Delta_k} \left\{ \frac{F_k}{\sum_t \phi_t \alpha_t} \right\}$, or equivalently, to minimize $F_k \cdot \max_t \left\{ \frac{1}{\phi_t} \right\}$. We set

$$f_t := \frac{1}{t} - \frac{\log t}{8kt} \text{ for } t \ge 1.$$
(8)

To verify that $\max_t \{1/\phi_t\} \leq 1$, we first bound F_t by

$$F_t = \sum_{i=1}^t f_i = H_t - \sum_{i=1}^t \frac{\log i}{8ki} \ge H_t - \int_{x=1}^{t+1} \frac{\log x}{8kx} dx = H_t - \frac{\log^2(t+1)}{16k} \ge H_t - \frac{\log(t+1)}{16k}$$

using $\log(t+1) \le t \le k$ for the last inequality. Therefore,

$$k\phi_t = F_t + (k-1)tf_t = H_t - \frac{\log(t+1)}{16} + (k-1)t\left(\frac{1}{t} - \frac{\log t}{8kt}\right)$$
$$\ge (k-1) + H_t - \frac{\log(t+1)}{16} - \frac{\log t}{8}$$
$$\ge k \text{ for } t \ge 2.$$

We can show $k\phi_1 \ge k$ separately for t = 1 using $F_1 = f_1 = 1$. Therefore, the approximation ratio is at most $F_k = H_k - \Theta((\log k)^2/k)$.

5 A Polynomial-Time Implementation

There are two issues with the running time of the above algorithms: (a) since we consider adding sets from the exponentially-large subset-closed family of sets (i.e., we assumed $S = S^{\downarrow}$), finding such a feasible local move may not naively be polynomial-time implementable. Moreover, (b) reaching a local optimum may not be feasible in polynomial time. The second issue can be handled using the standard technique of stopping when none of the local moves decrease the potential by more than a $\delta w(\mathcal{F})$ (see, e.g., [WS11, §9.1]). By setting $\delta = \varepsilon/|U|$ and changing the RHS of (3) from 0 to $\delta w(\mathcal{F})$ and using $|\mathcal{F}^*| \leq |U|$, we can ensure that $w(\mathcal{F}) \leq \frac{H_k}{1-\varepsilon}w(F^*)$ when there is no such improving move. Assuming the initial solution is poly(n)-approximate, one can ensure that the running time is poly(n, ε).

For issue (a), let S be the original collection of sets, not necessarily downwards closed. Suppose that we have the current solution \mathcal{F} and $S_1, \ldots, S_p \in S$, and want to find appropriate pairwise disjoint subsets $S'_1 \subseteq S_1, \ldots, S'_p \subseteq S_p$ such that the new solution that adds S'_1, \ldots, S'_p to \mathcal{F} (and subtracts their union from every $S \in \mathcal{F}$) has a low potential. We do not know of an efficient way to compute the optimal choice of S'_1, \ldots, S'_p . (One possible solution is, letting $T_1, \ldots, T_q \in \mathcal{F}$ be the sets intersecting $\bigcup_{i \in [p]} S_i$ (so $q \leq pk$), to (1) guess $|S'_i|$ for every $i \in [p]$, and $|T_j \cap (\bigcup_i S'_i)|$ for every $j \in [q]$ that exactly determine the potential change and (2) set up a network flow testing whether such S'_1, \ldots, S'_p exist, but it takes time $k^{O(pk)}$ poly(n).) A more efficient implementation without necessarily finding the optimal S'_1, \ldots, S'_p is this: in all our previous proofs, when we considered adding $S \in \mathcal{F}^*$ to the solution, the analysis always used $F_k w(S)$ as (an upper bound on) the increase of the potential by adding S instead of the exact increase $F_{|S|}w(S)$. This means that we can indeed run more conservative local search; given S_1, \ldots, S_p , let $A = \sum_{i \in [p]} w(S_i) F_k$ be the total increase from adding them, compute the total decrease B caused by removing $\bigcup_{i \in [p]} S_i$ from the current sets in \mathcal{F} , and only execute the local move when A is smaller than B (by $\delta w(\mathcal{F})$). All our analyses prove that we achieve the claimed approximation guarantees even when such a more conservative local move is not possible. Of course, if S_1, \ldots, S_p overlap, we can arbitrarily drop elements from them to ensure that the new solution is a partition as well; this further drops the potential.

6 Tight Lower Bounds

In this section, we show that the approximation ratios of $H_k - \Theta(1/k)$ and $H_k - \Theta((\log k)^2/k)$ achieved by the width-2 and width-k local search respectively are optimal. In fact, they are optimal under any potential of the form $\Phi(\mathcal{F}) = \sum_{S \in \mathcal{F}} w(S)F_{|S|}$ where $F_\ell = f_1 + \cdots + f_\ell$ for some $f_1 \geq \cdots \geq f_k$. (We believe that the monotonicity condition is unnecessary, but currently do not have a formal proof.)

6.1 Width-2 Lower Bounds

We construct a collection of lower bound instances \mathcal{I}_{ℓ} , one for each $\ell \in [k]$. The instance \mathcal{I}_{ℓ} is the following:

- 1. The universe is U.
- 2. The optimum \mathcal{F}^* partitions U into sets of size exactly k where each set has cost 1.
- 3. The local optimum \mathcal{F} partitions U into sets of size exactly ℓ where each set has cost α_{ℓ} , to be determined below.
- 4. The set system is $(U, \mathcal{F} \cup \mathcal{F}^*)$.
- 5. Let G_{ℓ} be a bipartite graph with \mathcal{F} and \mathcal{F}^* as two sides, where each element $u \in U$ corresponds an edge connecting the two sets that contain u. We can ensure that the girth of G_{ℓ} is at least a constant arbitrarily larger than k [FLS⁺95].

Let us determine the value of α_{ℓ} so that \mathcal{F} becomes a local optimum. When $\ell = 1$, $\alpha_1 = F_k/k$ suffices, which yields the approximation ratio F_k .

For $\ell \geq 2$, there are essentially two kinds of moves: there are width-2 local moves that add sets $S_1, S_2 \in \mathcal{F}^*$ such that $|(S_1 \cup S_2) \cap T| \leq 1$ for all $T \in \mathcal{F}$, or to add two distinct sets $S_1, S_2 \in \mathcal{F}^*$ that intersect a common $T \in \mathcal{F}$. (The girth condition ensures that there can be at most one such T, and $|S_1 \cap T| = |S_2 \cap T| = 1$.

- (i) The potential increase due to adding S_1, S_2 is $2F_k$.
- (ii) Either each set $T \in \mathcal{F}$ intersecting $S_1 \cup S_2$ loses one element, or some T loses two elements and all the other sets intersecting S_1 or S_2 lose exactly one element. So the potential decrease due to removing elements from sets in \mathcal{F} is

$$(f_{\ell} + f_{\ell-1}) + (2k-2)f_{\ell}.$$

(Here we use the fact that $f_{\ell-1} \ge f_{\ell}$.)

(iii) Therefore, \mathcal{F} is a local optimum as long as

$$\alpha_{\ell} = \frac{2F_k}{(2k-1)f_{\ell} + f_{\ell-1}},$$

and the approximation ratio in this case is

$$\frac{k\alpha_k}{\ell} = \frac{k}{\ell} \cdot \frac{2F_k}{(2k-1)f_\ell + f_{\ell-1}} = \frac{F_k}{\ell(f_\ell + f_{\ell-1} - f_\ell/2k)}.$$

Fixing $f_1 = 1$ and optimizing f_2, \ldots, f_k to minimize the worst-case approximation ratio over the k instances $\mathcal{I}_1, \ldots, \mathcal{I}_k$ shows that the best possible approximation ratio is determined by setting

$$f_{\ell} = \frac{1}{\ell} + \frac{1}{2k-1} \left(\frac{1}{\ell} - f_{\ell-1} \right)$$

for $\ell = 2, ..., k$. It yields $f_{\ell} = 1/\ell - \Theta(1/k\ell^2)$ as in the upper bound proof in Section 3, showing that no potential can guarantee strictly better than $H_k - \Theta(1/k)$.

6.2 Width-*p* Lower Bounds

The lower bound for the case of width-p follows the same framework. Fix $\ell \in [k]$ and consider the instance \mathcal{I}_{ℓ} defined in Section 6.1 (while ensuring that the girth $\gg p$), and determine the value of α_{ℓ} so that \mathcal{F} becomes a local optimum. When $\ell = 1$, $\alpha_1 = F_k/k$ suffices, which yields the approximation ratio F_k .

For $\ell \geq 2$, let us consider what the best local width-p moves would be. For S_1, \ldots, S_p from the optimal solution \mathcal{F}^* , consider the bipartite graph where the left vertices are S_1, \ldots, S_p , the right vertices are the sets from the current solution \mathcal{F} intersecting S_1, \ldots, S_p , and there is an edge if two sets intersect. Since the girth of the instance is much larger than W, this bipartite graph is a tree with exactly Wk edges and Wk + 1 vertices, so the number of right vertices is P = W(k-1) + 1. If we let $d_1 \geq \cdots \geq d_P$ be the degrees of the right vertices, the potential decrease from the current set is

$$\sum_{i=1}^{P} (F_{\ell} - F_{\ell-d_i}).$$

Since both P and $\sum_{i=1}^{P} d_i = Wk$ are fixed, the monotonicity of $f_1 \ge \cdots \ge f_k$ implies that the above is when the degree is *maximally skewed*; defining $t \in \mathbb{N}$ and $1 \le r < \ell - 1$ such that $p - 1 = t(\ell - 1) + r$, we have t right vertices have degree ℓ , one right vertex has degree r + 1, and the remaining p(k - 1) - t right vertices have degree 1. As a sanity check, note that $t \cdot \ell + (r + 1) + p(k - 1) - t = (t(\ell - 1) + r) + 1 + p(k - 1) = pk$.

With this move,

- (i) The potential increase due to adding $S_1, S_2, \ldots S_p$ is pF_k .
- (ii) The potential decrease due to removing elements from sets in \mathcal{F} is

$$tF_{\ell} + (F_{\ell} - F_{\ell-r-1}) + (p(k-1) - t)f_{\ell}.$$

(iii) Therefore, \mathcal{F} becomes a local optimum if

$$\alpha_{\ell} = \frac{pF_k}{tF_{\ell} + (F_{\ell} - F_{\ell-r-1}) + (p(k-1) - t)f_{\ell}},$$

and the approximation ratio in this case is $\alpha_{\ell} \cdot (k/\ell)$.

Again fixing $f_1 = 1$ and optimizing f_2, \ldots, f_k to minimize the approximation ratio for $\mathcal{I}_1, \ldots, \mathcal{I}_k$, yields $f_{\ell} = 1/\ell - \Theta(\log \ell/k\ell)$ just like we used in §4 for the upper bound for k-moves. Intuitively, setting $s = (p-1)/(\ell-1)$ so that $t(\ell-1) + r = s(\ell-1)$, the approximation ratio $\alpha \cdot (k/\ell)$ becomes

$$\begin{split} & \frac{k}{\ell} \cdot \frac{pF_k}{tF_\ell + (F_\ell - F_{\ell-r-1}) + (p(k-1) - t)f_\ell} \\ \geq & \frac{k}{\ell} \cdot \frac{pF_k}{sF_\ell + (pk - s\ell)f_\ell} \\ = & \frac{F_k}{\ell\left(\frac{sF_\ell}{pk} + (1 - s^{\ell/pk})f_\ell\right)} \\ = & \frac{F_k}{\Theta(\frac{F_\ell}{k}) + (1 - \Theta(1/k))\ell f_\ell}, \end{split}$$

so that with $F_{\ell} = \Theta(\log \ell)$, the denominator becomes at least 1 when $\ell f_{\ell} = (1 - \Theta(\log \ell/k))$. Therefore, no potential can guarantee strictly better than $H_k - \Theta((\log k)^2/k)$.

References

- [ACK09] Stavros Athanassopoulos, Ioannis Caragiannis, and Christos Kaklamanis. Analysis of approximation algorithms for k-set cover using factor-revealing linear programs. *Theory* of computing systems, 45(3):555–576, 2009.
- [AGK⁺04] Vijay Arya, Naveen Garg, Rohit Khandekar, Adam Meyerson, Kamesh Munagala, and Vinayaka Pandit. Local search heuristics for k-median and facility location problems. SIAM Journal on computing, 33(3):544–562, 2004.
- [BP89] Marshall Bern and Paul Plassmann. The Steiner problem with edge lengths 1 and 2. Information Processing Letters, 32(4):171–176, 1989.
- [CAGH⁺22] Vincent Cohen-Addad, Anupam Gupta, Lunjia Hu, Hoon Oh, and David Saulpic. An improved local search algorithm for k-median. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1556–1612. SIAM, 2022.
- [CG05] Moses Charikar and Sudipto Guha. Improved combinatorial algorithms for facility location problems. *SIAM Journal on Computing*, 34(4):803–824, 2005.
- [Chv79] Vasek Chvatal. A greedy heuristic for the set-covering problem. *Mathematics of operations* research, 4(3):233–235, 1979.
- [DF97] Rong-chii Duh and Martin Fürer. Approximation of k-set cover by semi-local optimization. In *Proceedings of the twenty-ninth annual ACM symposium on Theory of computing*, pages 256–264, 1997.
- [Fei98] Uriel Feige. A threshold of $\ln n$ for approximating set cover. Journal of the ACM (JACM), 45(4):634-652, 1998.

- [FLS+95] Z Furedi, Felix Lazebnik, A Seress, Vasiliy A Ustimenko, and Andrew J Woldar. Graphs of prescribed girth and bi-degree. Journal of Combinatorial Theory, Series B, 64(2):228– 239, 1995.
- [FR92] Martin Fürer and Balaji Raghavachari. Approximating the minimum degree spanning tree to within one from the optimal degree. In *Proceedings of the third annual ACM-SIAM symposium on Discrete algorithms*, pages 317–324, 1992.
- [FR94] Martin Furer and Balaji Raghavachari. Approximating the minimum-degree steiner tree to within one of optimal. *Journal of Algorithms*, 17(3):409–423, 1994.
- [FW14] Yuval Filmus and Justin Ward. Monotone submodular maximization over a matroid via non-oblivious local search. *SIAM Journal on Computing*, 43(2):514–542, 2014.
- [FY11] Martin Fürer and Huiwen Yu. Packing-based approximation algorithm for the k-set cover problem. In International Symposium on Algorithms and Computation, pages 484–493. Springer, 2011.
- [GGK⁺18] Martin Groß, Anupam Gupta, Amit Kumar, Jannik Matuschke, Daniel R. Schmidt, Melanie Schmidt, and José Verschae. A local-search algorithm for steiner forest. In *ITCS*, pages 31:1–31:17, Jan 2018.
- [GHY93] Olivier Goldschmidt, Dorit S Hochbaum, and Gang Yu. A modified greedy heuristic for the set covering problem with improved worst case bound. *Information processing letters*, 48(6):305–310, 1993.
- [GT08] Anupam Gupta and Kanat Tangwongsan. Simpler analyses of local search algorithms for facility location. arXiv preprint arXiv:0809.2554, 2008.
- [Hal95] Magnús M Halldórsson. Approximating discrete collections via local improvements. In SODA, volume 95, pages 160–169, 1995.
- [Hal96] Magnús M. Halldórsson. Approximating k-set cover and complementary graph coloring. In William H. Cunningham, S. Thomas McCormick, and Maurice Queyranne, editors, Integer Programming and Combinatorial Optimization, 5th International IPCO Conference, Vancouver, British Columbia, Canada, June 3-5, 1996, Proceedings, volume 1084 of Lecture Notes in Computer Science, pages 118–131. Springer, 1996.
- [HL05] Refael Hassin and Asaf Levin. A better-than-greedy approximation algorithm for the minimum set cover problem. *SIAM Journal on Computing*, 35(1):189–200, 2005.
- [Hoc82] Dorit S Hochbaum. Approximation algorithms for the set covering and vertex cover problems. *SIAM Journal on computing*, 11(3):555–556, 1982.
- [Joh74] David S Johnson. Approximation algorithms for combinatorial problems. Journal of computer and system sciences, 9(3):256–278, 1974.
- [KMSV98] Sanjeev Khanna, Rajeev Motwani, Madhu Sudan, and Umesh Vazirani. On syntactic versus computational views of approximability. SIAM Journal on Computing, 28(1):164– 191, 1998.
- [KPR00] Madhukar R Korupolu, C Greg Plaxton, and Rajmohan Rajaraman. Analysis of a local search heuristic for facility location problems. *Journal of algorithms*, 37(1):146–188, 2000.
- [Lev09] Asaf Levin. Approximating the unweighted k-set cover problem: greedy meets local search. SIAM Journal on Discrete Mathematics, 23(1):251–264, 2009.
- [Lov75] László Lovász. On the ratio of optimal integral and fractional covers. Discrete mathematics, 13(4):383–390, 1975.
- [Ros73] Robert W Rosenthal. A class of games possessing pure-strategy nash equilibria. International Journal of Game Theory, 2(1):65–67, 1973.

- [Svi04] Maxim Sviridenko. A note on maximizing a submodular set function subject to a knapsack constraint. Operations Research Letters, 32(1):41–43, 2004.
 [Thi01] Martin Thimm. On the approximability of the Steiner tree problem. In International Symposium on Mathematical Foundations of Computer Science, pages 678–689. Springer, 2001.
 [Tre01] Luca Trevisan. Non-approximability results for optimization problems on bounded degree instances. In Proceedings of the thirty-third annual ACM symposium on Theory of Science.
- gree instances. In Proceedings of the thirty-third annual ACM symposium on Theory of computing, pages 453–461, 2001.
- [TZ22] Vera Traub and Rico Zenklusen. Local search for weighted tree augmentation and Steiner tree. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms* (SODA), pages 3253–3272. SIAM, 2022.
- [Vaz01] Vijay V Vazirani. Approximation algorithms, volume 1. Springer, 2001.
- [WS11] David P Williamson and David B Shmoys. *The design of approximation algorithms*. Cambridge university press, 2011.
- [You22] Neal Young. Greedy set cover III: weighted H(d)-approximation via localizing. unpublished notes, 2022. https://algnotes.info/on/obliv/greedy/set-cover-local/.

A The Post-processing Algorithm

Lemma 3.2 (Post-processing). Given any solution \mathcal{F} with $w(\mathcal{F}) \leq H_k w(\mathcal{F}^*)$, let \mathcal{F}_1 be the subcollection of sets in \mathcal{F} of unit size. If $w(\mathcal{F}_1) \geq 0.99 w(\mathcal{F})$, there is an efficient algorithm that returns a new solution \mathcal{F}' with $w(\mathcal{F}') \leq 0.99 H_k \cdot w(\mathcal{F})$.

Proof. We assume that $k \geq 3$, else the resulting edge-cover problem can be solved exactly in polynomial time. Suppose $w(\mathcal{F}) \leq 0.99 H_k \cdot w(\mathcal{F}^*)$, then we are already done, so assume otherwise. The assumption of the lemma means $w(\mathcal{F}_1) \geq 0.98 H_k \cdot w(\mathcal{F}^*)$.

Note that adding $S \in \mathcal{F}^*$ immediately allows us to remove the singleton sets which are contained in S. For a collection of sets $\mathcal{C} \subseteq \mathcal{S}$, let $\mathcal{N}_{\mathcal{C},1} := \{T \in \mathcal{F}_1 \mid T \subseteq \bigcup_{S \in \mathcal{C}} S\}$ be the collection of singletons from \mathcal{F}_1 that can be removed from the solution by adding \mathcal{C} .

We set up an instance of Knapsack Cover: we seek a collection $\mathcal{C} \subseteq \mathcal{S}$ with $w(\mathcal{C}) \leq w(\mathcal{F}^*)$ that maximizes the saving $w(\mathcal{N}_{\mathcal{C},1})$. Since \mathcal{F}^* is a feasible solution with saving at least $w(\mathcal{F}_1) \geq 0.98 H_k \cdot w(\mathcal{F}^*)$, we can use an (1-1/e)-approximation algorithm [Svi04] for knapsack cover to find a feasible solution \mathcal{C} having weight at most $w(\mathcal{F}^*)$ and savings at least $(1-1/e) \cdot 0.98 H_k \cdot w(\mathcal{F}^*)$. This means $\mathcal{F} \cup \mathcal{C} \setminus \mathcal{N}_{\mathcal{C},1}$ is a set cover with cost at most

$$\left(1 + H_k(1 - 0.98(1 - 1/e))\right) w(\mathcal{F}^*) \le 0.99 H_k w(\mathcal{F}^*),$$

for all integers $k \geq 3$.

B The H_k Bound via Relax-and-Round

The traditional analysis of relax-and-round achieves a bound of $O(\ln k)$ [Hoc82], but one can tighten the bound to H_k . Consider the following algorithm:

Solve the LP relaxation to get solution x^* . Repeatedly pick sets S_1, S_2, \ldots from S, each time picking set S with probability $\frac{x_S^*}{\sum_T x_T^*}$. Finally, let \mathcal{F} be the sets S_i that cover elements not covered by previous sets S_1, \ldots, S_{i-1} .

The following claim is proved in lecture notes of Young [You22]:

Theorem B.1. The algorithm above incurs expected cost at most $H_k \cdot \sum_S w(S) x_S^*$.