

A Unified Approach to Approximating Partial Covering Problems

Jochen Könnemann^{1*}, Ojas Parekh², and Danny Segev³

¹ Department of Combinatorics and Optimization, University of Waterloo, Canada.
jochen@math.uwaterloo.ca

² Department of Mathematics and Computer Science, Emory University, USA.
ojas@mathcs.emory.edu

³ School of Mathematical Sciences, Tel-Aviv University, Israel.
segevd@post.tau.ac.il

Abstract. An instance of the *generalized partial cover* problem consists of a ground set U and a family of subsets $\mathcal{S} \subseteq 2^U$. Each element $e \in U$ is associated with a profit $p(e)$, whereas each subset $S \in \mathcal{S}$ has a cost $c(S)$. The objective is to find a minimum cost subcollection $\mathcal{S}' \subseteq \mathcal{S}$ such that the combined profit of the elements covered by \mathcal{S}' is at least P , a specified profit bound. In the *prize-collecting* version of this problem, there is no strict requirement to cover any element; however, if the subsets we pick leave an element $e \in U$ uncovered, we incur a penalty of $\pi(e)$. The goal is to identify a subcollection $\mathcal{S}' \subseteq \mathcal{S}$ that minimizes the cost of \mathcal{S}' plus the penalties of uncovered elements.

Although problem-specific connections between the partial cover and the prize-collecting variants of a given covering problem have been explored and exploited, a more general connection remained open. The main contribution of this paper is to establish a formal relationship between these two variants. As a result, we present a unified framework for approximating problems that can be formulated or interpreted as special cases of generalized partial cover. We demonstrate the applicability of our method on a diverse collection of covering problems, for some of which we obtain the first non-trivial approximability results.

1 Introduction

For over three decades the *set cover* problem and its ever-growing list of generalizations, variants, and special cases have attracted the attention of researchers in the fields of discrete optimization, complexity theory, and combinatorics. Essentially, these problems are concerned with identifying a minimum cost collection of sets that covers a given set of elements, possibly with additional side constraints. While such settings may appear to be very simple at first glance, they still capture computational tasks of great theoretical and practical importance, as the reader may verify by consulting directly related surveys [2, 13, 18, 28] and the references therein.

In the present paper we focus our attention on the *generalized partial cover* problem, whose input consists of a ground set of elements U and a family \mathcal{S} of subsets of U . In addition, each element $e \in U$ is associated with a profit $p(e)$, whereas each subset $S \in \mathcal{S}$

* Research supported by NSERC grant no. 288340-2004.

has a cost $c(S)$. The objective is to find a minimum cost subcollection $S' \subseteq S$ such that the combined profit of the elements covered by S' is at least P , a specified profit bound. When all elements are endowed with unit profits, we obtain the well-known *partial cover* problem, in which the goal is to cover a given number of elements by picking subsets of minimum total cost.

Numerous computational problems can be formulated or interpreted as special cases of generalized partial cover, although this fact may be well-hidden. For most of these problems, novel techniques in the design of approximation algorithms have emerged over the years, and it is clearly beyond the scope of this writing to present an exhaustive overview. However, from the abundance of greedy schemes, local-search heuristics, randomized methods, and LP-based algorithms a simple observation is revealed: There is currently no unified approach to approximating partial covering problems.

The suggested method: Preliminaries. The main contribution of this paper is to establish a formal relationship between the partial cover and the prize-collecting versions of a given covering problem. In the *prize-collecting set cover* problem there is no strict requirement to cover any element; however, if the subsets we pick leave an element $e \in U$ uncovered, we incur a penalty of $\pi(e)$. The objective is to find a subcollection $S' \subseteq S$ that minimizes the cost of S' plus the penalties of the uncovered elements. A polynomial-time algorithm for this problem is said to be *Lagrangian multiplier preserving with factor r* (henceforth, r -LMP) if for every instance I it constructs a solution that satisfies $C + r\Pi \leq r \cdot \text{OPT}(I)$, where C is the total cost of the subsets picked, and Π is the sum of penalties over all uncovered elements. We further denote by \mathcal{I}_r the family of weighted set systems (U, \mathcal{S}, c) that possess the following property: There is an r -LMP algorithm for all prize-collecting instances (U, \mathcal{S}, c, π) , $\pi : U \rightarrow \mathbb{Q}_+$. In other words, for every penalty function π the corresponding instance admits an r -LMP approximation.

The main result. At the heart of our method is an algorithm for the generalized partial cover problem that computes an approximate solution by making use of an r -LMP prize-collecting algorithm in a black-box fashion. Specifically, in Section 2 we prove the following theorem.

Theorem 1. *Let I be a generalized partial cover instance defined on an underlying weighted set system (U, \mathcal{S}, c) , and suppose that $(U, \mathcal{S}, c) \in \mathcal{I}_r$ for some $r \geq 1$. Then, for any $\epsilon > 0$, we can find a feasible solution to I whose cost is at most $(\frac{4}{3} + \epsilon)r$ times the optimum, within time polynomial in $|U|$, $|\mathcal{S}|^{1/\epsilon}$ and the input length of I .*

Here is a rough outline of how the proof of Theorem 1 will proceed. We begin by formulating the generalized partial cover problem as an integer program. Next, we dualize the complicating constraint that places a lower bound of P on the total profit. More precisely, we lift this constraint to the objective function multiplied by an auxiliary variable λ , and obtain its corresponding Lagrangian relaxation. For any fixed $\lambda \geq 0$, the new program describes, up to a constant term, a prize-collecting set cover instance with non-uniform penalties. We now conduct a binary search, using the r -LMP prize-collecting algorithm as a subroutine, to find sufficiently close $\lambda_1 \geq \lambda_2$ that satisfy: For λ_1 , the algorithm constructs a solution $S_1 \subseteq S$ such that the total profit of the elements covered by S_1 is at least P ; For λ_2 , it constructs a solution $S_2 \subseteq S$ with a total profit of at most P .

Although we can exploit the r -LMP property to show that the cost of \mathcal{S}_2 is within factor r of optimum, this solution is not necessarily feasible. The situation is quite the opposite with respect to \mathcal{S}_1 , which is a feasible solution whose cost may be arbitrarily large. Having observed these facts, we create an additional feasible solution \mathcal{S}_3 by augmenting \mathcal{S}_2 with a carefully chosen subset of \mathcal{S}_1 . The cost of this subset is bounded by extending the arguments used by Levin and Segev [24] and independently by Golovin, Nagarajan and Singh [17] for approximating the k -multicut problem. Finally, we establish Theorem 1 by proving that the cost of the cheaper of \mathcal{S}_1 and \mathcal{S}_3 is at most $(\frac{4}{3} + \epsilon)r$ times the cost of an optimal solution.

Designing LMP algorithms. At this point in time, the reader should bear in mind that the performance guarantee of our algorithm, as stated in Theorem 1, depends on the existence of an LMP prize-collecting algorithm for a given covering problem. Indeed, this dependence appears to be the primary factor limiting the employment of Lagrangian relaxations in most problems of interest. The latter drawback was pointed out by Chudak, Roughgarden and Williamson [6], who asked whether it is possible to devise more general variants of the Lagrangian relaxation framework that apply to a broader class of problems. We answer this question in the affirmative, by developing prize-collecting algorithms with the LMP property for some of the most fundamental integer covering problems. These results, along with a detailed description of previous work, are formally presented in Section 3.

2 The Generalized Partial Cover Algorithm

The main result of this section is a constructive proof of Theorem 1. Recall that a generalized partial cover instance I is defined with respect to an underlying weighted set system, consisting of a ground set U and a family of subsets $\mathcal{S} \subseteq 2^U$, where each $S \in \mathcal{S}$ has a cost $c(S)$. The additional ingredients of I are profits $p(e)$, specified for each element $e \in U$, and a requirement parameter P . Now suppose that $(U, \mathcal{S}, c) \in \mathcal{I}_r$ for some $r \geq 1$, meaning that there is an r -LMP algorithm \mathcal{A} for all prize-collecting instances (U, \mathcal{S}, c, π) , $\pi : U \rightarrow \mathbb{Q}_+$.

2.1 Preliminaries

The method we suggest and its analysis will be based on a natural integer programming formulation of the generalized partial cover problem. In the following, let $\mathcal{S}_e \subseteq \mathcal{S}$ be the collection of sets that contain $e \in U$, and let $P_U = \sum_{e \in U} p(e)$.

$$\text{minimize } \sum_{S \in \mathcal{S}} c(S)x_S \tag{GC}$$

$$\text{subject to } \sum_{S \in \mathcal{S}_e} x_S + z_e \geq 1 \quad \forall e \in U \tag{2.1}$$

$$\sum_{e \in U} p(e)z_e \leq P_U - P \tag{2.2}$$

$$x_S, z_e \in \{0, 1\} \quad \forall S \in \mathcal{S}, e \in U \tag{2.3}$$

In this formulation, the variable x_S indicates whether we pick the set S , whereas z_e indicates whether the element e is uncovered. Constraint (2.1) guarantees that we either pick at least one set that contains e , or specify that this element is uncovered by setting $z_e = 1$. Constraint (2.2) forces any feasible solution to cover elements with a total profit of at least P .

Essential to the subsequent analysis will be the fact that the LP-relaxation of (GC), obtained by replacing constraint (2.3) with $x_S \geq 0$ and $z_e \geq 0$, has an integrality gap of $O(r)$. Unfortunately, this prerequisite is not satisfied even in the case of unit profits, as the next example illustrates. Consider an instance in which the ground set U consists of n elements, and the family \mathcal{S} contains a single set $S = U$ with cost n . When we are required to cover at least one element, the integral optimum is clearly n . However, by setting $x_S = \frac{1}{n}$ and $z_e = 1 - \frac{1}{n}$ for every $e \in U$, we define a feasible fractional solution whose cost is 1. This example, as well as additional constructions of similar nature, demonstrate that an unbounded integrality gap may arise whenever a small number of sets in the optimal solution contribute a large fraction of its cost.

Therefore, an inevitable part of our algorithm is a preprocessing step in which, given a fixed accuracy parameter $\epsilon > 0$, we “guess” the $\lfloor \frac{1}{\epsilon} \rfloor$ most expensive sets in the optimal solution, whose cost we denote by OPT . More precisely, we enumerate all $O(|\mathcal{S}|^{1/\epsilon})$ subsets $\mathcal{S}' \subseteq \mathcal{S}$ of cardinality at most $\lfloor \frac{1}{\epsilon} \rfloor$, test each such subset as the correct guess, and return the best solution we find. For a given subset \mathcal{S}' , we include it as part of the solution to be constructed, eliminate the sets in \mathcal{S}' from \mathcal{S} , remove all covered elements from U and from the remaining sets, and update the profit requirement. Any set whose cost is greater than $\min_{S \in \mathcal{S}'} c(S)$ is also eliminated. Consequently, the cost of each remaining set is at most $\epsilon \cdot \text{OPT}$.

In the remainder of this section we will bypass the preprocessing step, and assume that the maximum cost of a set in \mathcal{S} is at most $\epsilon \cdot \text{OPT}$. For ease of presentation, we also assume that $c(S) > 0$ for every $S \in \mathcal{S}$ and that $p(e) > 0$ for every $e \in U$, since zero-cost sets can be picked in advance and zero-profit elements can be discarded.

2.2 Obtaining \mathcal{S}_1 and \mathcal{S}_2

We now dualize the profit constraint (2.2), and lift it to the objective function multiplied by $\lambda \geq 0$. The resulting Lagrangian relaxation is:

$$\begin{aligned} \text{LR}(\lambda) = \text{minimize} \quad & \sum_{S \in \mathcal{S}} c(S)x_S + \lambda \left(\sum_{e \in U} p(e)z_e - (P_U - P) \right) \\ \text{subject to} \quad & \sum_{S \in \mathcal{S}_e} x_S + z_e \geq 1 \quad \forall e \in U \\ & x_S, z_e \in \{0, 1\} \quad \forall S \in \mathcal{S}, e \in U \end{aligned}$$

We remark that, excluding the constant term of $-\lambda(P_U - P)$ in the objective function, $\text{LR}(\lambda)$ is an integer programming formulation of the prize-collecting set cover problem, in which each element $e \in U$ is associated with a penalty $\lambda p(e)$. We refer to this instance as I_λ , and use $\text{OPT}(I_\lambda)$ to denote its optimum value. It is not difficult to verify that $\text{LR}(\lambda) = \text{OPT}(I_\lambda) - \lambda(P_U - P)$ is at most OPT for any $\lambda \geq 0$, by observing that an

optimal solution to (GC) is also a feasible solution to $\text{LR}(\lambda)$, whose cost is at most OPT .

Since the underlying weighted set system of I_λ is identical to that of I , we may apply the prize-collecting algorithm \mathcal{A} to approximate I_λ . Let x^λ indicate which sets in \mathcal{S} were picked by the algorithm, and let z^λ indicate which elements were left uncovered. In terms of (x^λ, z^λ) , the r -LMP property of \mathcal{A} is equivalent to

$$\sum_{S \in \mathcal{S}} c(S)x_S^\lambda + r \sum_{e \in U} \lambda p(e)z_e^\lambda \leq r \cdot \text{OPT}(I_\lambda) , \quad (2.4)$$

an inequality that, in particular, leads to the following observation.

Lemma 2. *When $\lambda > \frac{1}{\min_{e \in U} p(e)} \sum_{S \in \mathcal{S}} c(S)$, the solution (x^λ, z^λ) covers all elements. On the other hand, (x^0, z^0) does not cover any element.*

Proof. Let $\lambda > \frac{1}{\min_{e \in U} p(e)} \sum_{S \in \mathcal{S}} c(S)$, and suppose that there is an element $\bar{e} \in U$ for which $z_{\bar{e}}^\lambda = 1$, that is, \bar{e} is not covered by any set the algorithm \mathcal{A} picks when we approximate I_λ . Then (x^λ, z^λ) no longer satisfies inequality (2.4), as

$$\sum_{S \in \mathcal{S}} c(S)x_S^\lambda + r \sum_{e \in U} \lambda p(e)z_e^\lambda \geq r\lambda p(\bar{e}) > r \frac{p(\bar{e})}{\min_{e \in U} p(e)} \sum_{S \in \mathcal{S}} c(S) \geq r \sum_{S \in \mathcal{S}} c(S) \geq r \cdot \text{OPT}(I_\lambda) ,$$

where the last inequality holds since \mathcal{S} is a feasible solution to I_λ .

Now let $\lambda = 0$, and note that each element of the instance I_0 has a zero penalty. Therefore, by deciding not to pick any set and instead pay all penalties we obtain a feasible solution with zero cost, implying that $\text{OPT}(I_0) = 0$. Since (x^0, z^0) satisfies inequality (2.4), it follows that this solution cannot pick any set, as all sets in \mathcal{S} have strictly positive costs by assumption. \square

This observation allows us to conduct a binary search over $[0, \frac{2}{\min_e p(e)} \sum_{S \in \mathcal{S}} c(S)]$, consisting of a polynomially-bounded number of calls to the prize-collecting algorithm \mathcal{A} , as a result of which we find $\lambda_1 \geq \lambda_2$ that satisfy:

1. $\lambda_1 - \lambda_2 \leq \frac{\epsilon c_{\min}}{P_U}$, where $c_{\min} = \min_{S \in \mathcal{S}} c(S) > 0$.
2. The elements covered by $(x^{\lambda_1}, z^{\lambda_1})$ have a total profit of $P_1 \geq P$, and at the same time those covered by $(x^{\lambda_2}, z^{\lambda_2})$ have a total profit of $P_2 \leq P$.

For ease of notation, we designate by \mathcal{S}_1 and \mathcal{S}_2 the subsets of \mathcal{S} that were picked by the solutions $(x^{\lambda_1}, z^{\lambda_1})$ and $(x^{\lambda_2}, z^{\lambda_2})$, respectively. Without loss of generality, $P_1 > P$, or otherwise \mathcal{S}_1 is already a feasible solution whose cost is at most $r \cdot \text{LR}(\lambda_1) \leq r \cdot \text{OPT}$. Similarly, we assume that $P_2 < P$. The analysis of our algorithm crucially depends on the next lemma, which is a consequence of the r -LMP property.

Lemma 3. *Let $\alpha = \frac{P-P_2}{P_1-P_2} \in (0, 1)$. Then, $\alpha c(\mathcal{S}_1) + (1-\alpha)c(\mathcal{S}_2) \leq r(1+\epsilon)\text{OPT}$.*

Proof. By combining inequality (2.4) with the fact that $\text{LR}(\lambda) = \text{OPT}(I_\lambda) - \lambda(P_U - P) \leq \text{OPT}$ for every $\lambda \geq 0$, we have

$$\begin{aligned} c(\mathcal{S}_1) &= \sum_{S \in \mathcal{S}} c(S)x_S^{\lambda_1} \leq r \left(\text{OPT}(I_{\lambda_1}) - \lambda_1 \sum_{e \in U} p(e)z_e^{\lambda_1} \right) = r(\text{OPT}(I_{\lambda_1}) - \lambda_1(P_U - P_1)) \\ &= r(\text{LR}(\lambda_1) + \lambda_1(P_1 - P)) \leq r(\text{OPT} + \lambda_1(P_1 - P)) . \end{aligned} \quad (2.5)$$

A similar argument shows that $c(\mathcal{S}_2) \leq r(\text{OPT} + \lambda_2(P_2 - P))$. Therefore,

$$\begin{aligned}
\alpha c(\mathcal{S}_1) + (1 - \alpha)c(\mathcal{S}_2) &\leq \alpha r(\text{OPT} + \lambda_1(P_1 - P)) + (1 - \alpha)r(\text{OPT} + \lambda_2(P_2 - P)) \\
&\leq r \cdot \text{OPT} + \alpha r \left(\lambda_2 + \frac{\epsilon c_{\min}}{P_U} \right) (P_1 - P) + (1 - \alpha)r\lambda_2(P_2 - P) \\
&= r \cdot \text{OPT} + r\lambda_2(\alpha(P_1 - P) + (1 - \alpha)(P_2 - P)) + r\alpha\epsilon c_{\min} \cdot \frac{P_1 - P}{P_U} \\
&\leq r \cdot \text{OPT} + r\epsilon c_{\min} \\
&\leq r(1 + \epsilon)\text{OPT} .
\end{aligned}$$

The second inequality follows from observing that $P_1 > P$ and $\lambda_1 \leq \lambda_2 + \frac{\epsilon c_{\min}}{P_U}$. The third inequality holds since $\alpha(P_1 - P) + (1 - \alpha)(P_2 - P) = 0$, $\alpha < 1$ and $P_1 - P \leq P_U$. \square

2.3 Composing an additional solution

Up until now, the only feasible solution we have at our possession is \mathcal{S}_1 , as this subset of \mathcal{S} covers elements with an overall profit of $P_1 > P$. Inequality (2.5) places an upper bound of $r \cdot \text{OPT} + r\lambda_1(P_1 - P)$ on the cost of \mathcal{S}_1 . However, the latter term may be arbitrarily large in comparison to OPT , implying that \mathcal{S}_1 cannot approximate the instance I by itself. The situation is quite the opposite with respect to \mathcal{S}_2 : Although this solution covers elements with an insufficient profit of $P_2 < P$, a similar bound of $r \cdot \text{OPT} + r\lambda_2(P_2 - P)$ on its cost actually yields the inequality $c(\mathcal{S}_2) \leq r \cdot \text{OPT}$, since in this case $r\lambda_2(P_2 - P) \leq 0$.

At this point, we are concerned with creating an additional feasible solution \mathcal{S}_3 , by augmenting \mathcal{S}_2 with a carefully chosen subset $\mathcal{S}' \subseteq \mathcal{S}_1$. To attain feasibility, we must ensure that of the elements that were left uncovered by \mathcal{S}_2 , a subcollection with a total profit of at least $P - P_2$ is covered by \mathcal{S}' . We construct this augmenting subset as follows. Let $U' \subseteq U$ be the collection of elements that are covered by \mathcal{S}_1 but not by \mathcal{S}_2 . We assign each element $e \in U'$ to an arbitrary set in $\mathcal{S}_1 \setminus \mathcal{S}_2$ that contains it, and denote by $\varphi(\mathcal{S})$ the total profit of the elements assigned to \mathcal{S} . Without loss of generality, we assume that $\mathcal{S}_1 \setminus \mathcal{S}_2 = \{\mathcal{S}_1, \dots, \mathcal{S}_k\}$, where these sets are indexed by non-decreasing order of the ratio $\frac{c(\mathcal{S}_i)}{\varphi(\mathcal{S}_i)}$. Finally, let $\mathcal{S}' = \{\mathcal{S}_1, \dots, \mathcal{S}_q\}$, where q is the minimal index for which $\sum_{i=1}^q \varphi(\mathcal{S}_i) \geq P - P_2$. Note that such an index exists, since $\sum_{i=1}^k \varphi(\mathcal{S}_i) \geq P_1 - P_2$ and $P_1 > P$. The next lemma bounds the cost of $\mathcal{S}_3 = \mathcal{S}_2 \cup \mathcal{S}'$.

Lemma 4. $c(\mathcal{S}_3) \leq c(\mathcal{S}_2) + \alpha c(\mathcal{S}_1 \setminus \mathcal{S}_2) + \epsilon \cdot \text{OPT}$.

Proof. By assumption, the cost of each set in \mathcal{S} is at most $\epsilon \cdot \text{OPT}$. Therefore, it is sufficient to prove that $c(\mathcal{S}' \setminus \{\mathcal{S}_q\}) = \sum_{i=1}^{q-1} c(\mathcal{S}_i) \leq \alpha c(\mathcal{S}_1 \setminus \mathcal{S}_2)$. To this end, consider a random variable K that takes the values $1, \dots, k$, such that $\mathbb{P}(K = i) = \frac{\varphi(\mathcal{S}_i)}{\sum_{l=1}^k \varphi(\mathcal{S}_l)}$, and let $R = \frac{c(\mathcal{S}_K)}{\varphi(\mathcal{S}_K)}$. Since the sets in $\mathcal{S}_1 \setminus \mathcal{S}_2$ are indexed by non-decreasing order of $\frac{c(\mathcal{S}_i)}{\varphi(\mathcal{S}_i)}$, we have $\mathbb{E}(R|K \leq q - 1) \leq \mathbb{E}(R)$. As $\alpha = \frac{P - P_2}{P_1 - P_2}$, this inequality implies $\sum_{i=1}^{q-1} c(\mathcal{S}_i) \leq \alpha c(\mathcal{S}_1 \setminus \mathcal{S}_2)$, since

$$\mathbb{E}(R) = \sum_{i=1}^k \frac{c(\mathcal{S}_i)}{\varphi(\mathcal{S}_i)} \cdot \frac{\varphi(\mathcal{S}_i)}{\sum_{l=1}^k \varphi(\mathcal{S}_l)} = \frac{1}{\sum_{l=1}^k \varphi(\mathcal{S}_l)} \sum_{i=1}^k c(\mathcal{S}_i) \leq \frac{1}{P_1 - P_2} c(\mathcal{S}_1 \setminus \mathcal{S}_2)$$

and

$$\mathbb{E}(R|K \leq q-1) = \sum_{i=1}^{q-1} \frac{c(\mathcal{S}_i)}{\varphi(\mathcal{S}_i)} \cdot \frac{\varphi(\mathcal{S}_i)}{\sum_{l=1}^{q-1} \varphi(\mathcal{S}_l)} = \frac{1}{\sum_{l=1}^{q-1} \varphi(\mathcal{S}_l)} \sum_{i=1}^{q-1} c(\mathcal{S}_i) \geq \frac{1}{P - P_2} \sum_{i=1}^{q-1} c(\mathcal{S}_i) .$$

The last inequality holds since $\sum_{l=1}^{q-1} \varphi(\mathcal{S}_l) < P - P_2$, by the minimality of q . \square

2.4 Deriving the approximation factor

We now conclude the proof of Theorem 1, by demonstrating that the cost of the cheaper of \mathcal{S}_1 and \mathcal{S}_3 is within factor $(\frac{4}{3} + O(\sqrt{\epsilon}))r$ of optimum. An appropriate choice of ϵ restores the original form of the theorem.

Lemma 5. $\min\{c(\mathcal{S}_1), c(\mathcal{S}_3)\} \leq (\frac{4}{3} + O(\sqrt{\epsilon}))r \cdot \text{OPT}$.

Proof. To simplify the analysis, we begin by introducing a new parameter, $\beta = \frac{c(\mathcal{S}_2)}{\text{OPT}} \in [0, r]$, and bound the cost of \mathcal{S}_1 and \mathcal{S}_3 in terms of OPT , α and β . We first observe that

$$c(\mathcal{S}_1) = \frac{\alpha c(\mathcal{S}_1)}{\alpha} \leq \frac{r(1+\epsilon)\text{OPT} - (1-\alpha)c(\mathcal{S}_2)}{\alpha} = \frac{r(1+\epsilon) - (1-\alpha)\beta}{\alpha} \text{OPT} ,$$

where the first inequality follows from Lemma 3, and the last equation is obtained by substituting $c(\mathcal{S}_2) = \beta \cdot \text{OPT}$. In addition, Lemma 4 implies that

$$\begin{aligned} c(\mathcal{S}_3) &\leq c(\mathcal{S}_2) + \alpha c(\mathcal{S}_1 \setminus \mathcal{S}_2) + \epsilon \cdot \text{OPT} \leq (1-\alpha)c(\mathcal{S}_2) + \alpha c(\mathcal{S}_1) + \alpha c(\mathcal{S}_2) + \epsilon \cdot \text{OPT} \\ &\leq r(1+\epsilon)\text{OPT} + \alpha c(\mathcal{S}_2) + \epsilon \cdot \text{OPT} = (r(1+\epsilon) + \alpha\beta + \epsilon)\text{OPT} , \end{aligned}$$

where the third inequality and the last equation follow from Lemma 3 and the definition of β , respectively. Finally, we bound the resulting approximation factor by considering the worst possible choice for the parameters α and β , to conclude that

$$\begin{aligned} \min\{c(\mathcal{S}_1), c(\mathcal{S}_3)\} &\leq \min \left\{ \frac{r(1+\epsilon) - (1-\alpha)\beta}{\alpha}, r(1+\epsilon) + \alpha\beta + \epsilon \right\} \text{OPT} \\ &\leq \max_{\substack{\alpha \in (0,1) \\ \beta \in [0,r]}} \min \left\{ \frac{r(1+\epsilon) - (1-\alpha)\beta}{\alpha}, r(1+\epsilon) + \alpha\beta \right\} \text{OPT} + \epsilon \cdot \text{OPT} \\ &= \left(\frac{4}{3} + O(\sqrt{\epsilon}) \right) r \cdot \text{OPT} . \end{aligned}$$

\square

3 Applications

In what follows, we demonstrate the applicability of our method on a diverse collection of covering problems, which is by no means exhaustive. Rather, the problems we have chosen to study are only meant to illustrate that the LMP property is applicable in a variety of settings. For the vast majority of these problems, we propose the first algorithm that approximates their generalized partial cover version. For others, our algorithms offer approximation guarantees that compete with the currently best known results. Due to space limitations, we defer the description of problem-specific prize-collecting sub-routines to the full version of this paper [23].

3.1 Set cover, in terms of Δ

Kearns [21, Thm. 5.15] seems to have been the first to study the partial cover problem, showing that the greedy set cover algorithm [20, 25] can be adapted to provide an approximation factor of $2H(|U|) + 3$. A slightly different algorithm was suggested by Slavík [32], who obtained a factor of $H(\min\{\Delta, k\})$, where Δ is the maximum size of a set in \mathcal{S} and k is the coverage requirement. We remark that the partial cover problem contains set cover as a special case, implying that it cannot be approximated within a factor of $(1 - \epsilon) \ln |U|$ for any $\epsilon > 0$, unless $\text{NP} \subset \text{TIME}(n^{O(\log \log n)})$ [10].

To the best of our knowledge, the greedy heuristic has not been studied in the context of generalized partial cover, and in fact no algorithm is currently known for this problem. In the full version of this paper [23], we prove that every weighted set system (U, \mathcal{S}, c) is in $\mathcal{I}_{H(\Delta)}$, where $\Delta = \max_{S \in \mathcal{S}} |S|$. The next theorem follows.

Theorem 6. *The generalized partial set cover problem can be approximated within a factor of $(\frac{4}{3} + \epsilon)H(\Delta)$, for any fixed $\epsilon > 0$.*

3.2 Set cover, in terms of f

Let f_e be the number of sets in \mathcal{S} that contain the element $e \in U$; f_e is also known as the *frequency* of e . A recent line of work, that was initiated by Bshouty and Burroughs [4] and Hochbaum [19] in the context of *partial vertex cover*, is approximating partial cover in terms of f , the maximum frequency of any element. Based on the local-ratio method, Bar-Yehuda [3] devised an algorithm for generalized partial cover whose approximation guarantee is f , a result that was independently obtained by Fujito [11] using a primal-dual algorithm. Gandhi, Khuller and Srinivasan [12] achieved a similar ratio for partial cover.

In the full version of this paper [23], we present a combinatorial f -LMP algorithm for the prize-collecting set cover problem, showing that every weighted set system (U, \mathcal{S}, c) is in \mathcal{I}_f , where $f = \max_{e \in U} f_e$. Combined with Theorem 1, this result allows us to approximate the generalized partial set cover problem within a factor of $(\frac{4}{3} + \epsilon)f$, which is slightly worse than the currently best.

3.3 Laminar cover

Let $G = (V, E)$ be an undirected graph, in which each edge $e \in E$ has a non-negative cost $c(e)$, and let $\mathcal{F} = \{V_1, \dots, V_k\} \subseteq 2^V$ be a *laminar family* of vertex sets, meaning that $V_i \cap V_j \in \{\emptyset, V_i, V_j\}$ for every $i \neq j$. We say that an edge e *covers* V_i if it has exactly one endpoint in V_i . The objective is to find a minimum cost set of edges that collectively cover all sets in \mathcal{F} . Note that every instance of this problem induces a weighted set system $(\mathcal{F}, \mathcal{S}, c)$, where for each edge $e \in E$ there is an analogous subset $S_e \in \mathcal{S}$, consisting of all vertex sets $V_i \in \mathcal{F}$ covered by e . Laminar cover can be approximated by applying various techniques, most of which actually deal with the more general *tree augmentation* problem, and produce solutions whose cost is within factor 2 of optimum. We refer the reader to a short survey of these results [9, Sec. 1]. For the unweighted case, Nagamochi [27] proposed a $(1.875 + \epsilon)$ -approximation for any fixed $\epsilon > 0$, a ratio that was later improved to $\frac{3}{2}$ by Even, Feldman, Kortsarz and Nutov [9].

In the *generalized partial laminar cover* problem, each $V_i \in \mathcal{F}$ is associated with a profit $p(V_i)$. The goal is to identify a minimum cost set of edges $E' \subseteq E$ such that the overall profit of the sets in \mathcal{F} covered by E' is at least P , a specified profit bound. We are not aware of any approximability result for this problem, even for the seemingly simple case of unit profits. In the full version of this paper [23], we prove that $(\mathcal{F}, \mathcal{S}, c) \in \mathcal{I}_2$ for every weighted set system induced by a laminar cover instance, to obtain the following theorem.

Theorem 7. *The generalized partial laminar cover problem can be approximated within a factor of $\frac{8}{3} + \epsilon$, for any fixed $\epsilon > 0$.*

3.4 Totally unimodular cover and k -interval cover

The *element-set incidence matrix* $\mathcal{M}_U^{\mathcal{S}}$ of a set system (U, \mathcal{S}) has a row for every element $e \in U$ and a column for every set $S \in \mathcal{S}$; its entry in row e and column S is 1 when $e \in S$ and 0 otherwise. Totally unimodular cover (TUC) is a special case of the set cover problem in which $\mathcal{M}_U^{\mathcal{S}}$ is totally unimodular, that is, every square submatrix of this matrix has determinant 0, 1 or -1 . We remark that although TUC is known to have integral LP solutions (see, for example, [7, Sec. 6.5]), this property does not extend to its partial covering version, which has not been explicitly studied yet. A particularly interesting problem captured by the latter variant is *partial bipartite vertex cover*: While the approximability of the unit-profit case is still open, arbitrary profits render the problem NP-hard, since it generalizes *minimum knapsack* even when the given graph is a star. We omit the straightforward reduction.

As illustrated in the full version of this paper [23], the prize-collecting set cover problem can be formulated as an integer program whose constraint matrix is $[\mathcal{M}_U^{\mathcal{S}}, I]$. Simple linear algebra arguments show that whenever $\mathcal{M}_U^{\mathcal{S}}$ is totally unimodular then so is $[\mathcal{M}_U^{\mathcal{S}}, I]$, implying that we obtain a 1-LMP algorithm by solving prize-collecting TUC to optimality as a linear program. The next theorem follows.

Theorem 8. *The generalized partial TUC problem can be approximated within a factor of $\frac{4}{3} + \epsilon$, for any fixed $\epsilon > 0$.*

We say that $\mathcal{M}_U^{\mathcal{S}}$ is a *k -interval matrix* if it contains at most k blocks of consecutive 1's in each row. The *k -interval cover problem* (k -IC) is a special case of set cover in which $\mathcal{M}_U^{\mathcal{S}}$ is a k -interval matrix. In the full version of this paper [23], we present a k -LMP rounding algorithm for the prize-collecting k -IC problem, that makes use of our 1-LMP algorithm for the corresponding variant of totally unimodular cover. We derive the following result as a corollary of Theorem 1.

Theorem 9. *The generalized partial k -IC problem can be approximated within a factor of $(\frac{4}{3} + \epsilon)k$, for any fixed $\epsilon > 0$.*

This provides, for instance, the first algorithm that approximates *partial rectangle stabbing* in \mathbb{R}^d , noting that the resulting factor of $(\frac{4}{3} + \epsilon)d$ nearly matches the d -approximation of Gaur, Ibaraki and Krishnamurti [16] for the full coverage version of this problem. In addition, we obtain an alternative, albeit non-combinatorial, $(\frac{4}{3} + \epsilon)f$ -approximation for partial set cover with maximum element frequency f .

3.5 Edge cover

Given an undirected graph $G = (V, E)$ with non-negative edge costs, *edge cover* is the problem of finding a minimum cost set of edges that contains at least one edge incident to each vertex. Clearly, this problem is equivalent to the special case of set cover in which each subset consists of exactly two elements. We note that edge cover is actually a matching problem in disguise, implying its polynomial time solvability [8, 26]. Plesník [30] proved that unit-profit partial edge cover, which is also known as the *k-edge cover* problem, can be solved to optimality by reducing it to standard edge cover. However, when arbitrary profits are allowed, this problem becomes NP-hard, as it generalizes minimum knapsack. Since Parekh [29, Sec. 2.3] suggested a polynomial-time algorithm for prize-collecting edge cover, we obtain the following theorem.

Theorem 10. *Generalized partial edge cover can be approximated within a factor of $\frac{4}{3} + \epsilon$, for any fixed $\epsilon > 0$.*

3.6 Multicut

On trees. The input to this problem consists of an edge-weighted tree $T = (V, E)$ and a collection of k distinct pairs of vertices, $\{s_1, t_1\}, \dots, \{s_k, t_k\}$. The objective is to find a minimum cost set of edges whose removal from T disconnects each of the given pairs. It is important to note that, once again, we are facing a special case of set cover: The elements to cover are the input pairs, and an edge $e \in E$ covers $\{s_i, t_i\}$ if it resides on the unique path in T connecting s_i and t_i . Garg, Vazirani and Yannakakis [15] presented a primal-dual 2-approximation for this problem, which was also shown to be at least as hard to approximate as vertex cover.

The corresponding partial cover problem, in which we are required to disconnect a specified number of pairs, has recently been studied by Levin and Segev [24] and independently by Golovin et al. [17], who achieved an approximation guarantee of $\frac{8}{3} + \epsilon$, for any fixed $\epsilon > 0$. Since the former authors provide a 2-LMP algorithm for the prize-collecting multicut problem, we immediately obtain the following theorem, extending the factor of $\frac{8}{3} + \epsilon$ to the case of arbitrary profits.

Theorem 11. *When the underlying graph is a tree, the generalized partial multicut problem can be approximated within a factor of $\frac{8}{3} + \epsilon$, for any fixed $\epsilon > 0$.*

General graphs. When the input graph is no longer restricted to be a tree, the multicut problem becomes significantly harder to approximate. While Garg et al. [14] devised an $O(\log k)$ -approximation using the region growing method, a hardness result of $\Omega(\log \log n)$ was given by Chawla, Krauthgamer, Kumar, Rabani and Sivakumar [5], assuming a stronger version of the Unique Games Conjecture [22]. Based on Räcke's hierarchical decomposition method [31], Alon, Awerbuch, Azar, Buchbinder and Naor [1] have shown how to simulate multicuts in general graphs by multicuts in the corresponding decomposition tree. As observed by Golovin et al. [17], this method extends to approximate the partial multicut problem within factor $O(\alpha \log^2 n \log \log n)$, given an α -approximation for the more restricted tree case. Their arguments can be easily combined with Theorem 11 to derive the next result for arbitrary profits.

Theorem 12. *On arbitrary graphs, the generalized partial multicut problem can be approximated within a factor of $O(\log^2 n \log \log n)$.*

References

1. N. Alon, B. Awerbuch, Y. Azar, N. Buchbinder, and J. Naor. A general approach to online network optimization problems. In *Proceedings of the 15th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 577–586, 2004.
2. E. Balas and M. Padberg. Set partitioning: A survey. *SIAM Review*, 18(4):710–760, 1976.
3. R. Bar-Yehuda. Using homogeneous weights for approximating the partial cover problem. *Journal of Algorithms*, 39(2):137–144, 2001.
4. N. H. Bshouty and L. Burroughs. Massaging a linear programming solution to give a 2-approximation for a generalization of the vertex cover problem. In *Proceedings of the 15th Annual Symposium on Theoretical Aspects of Computer Science*, pages 298–308, 1998.
5. S. Chawla, R. Krauthgamer, R. Kumar, Y. Rabani, and D. Sivakumar. On the hardness of approximating multicut and sparsest-cut. In *Proceedings of the 20th Annual IEEE Conference on Computational Complexity*, pages 144–153, 2005.
6. F. A. Chudak, T. Roughgarden, and D. P. Williamson. Approximate k -MSTs and k -Steiner trees via the primal-dual method and Lagrangean relaxation. *Mathematical Programming*, 100(2):411–421, 2004.
7. W. J. Cook, W. H. Cunningham, W. R. Pulleyblank, and A. Schrijver. *Combinatorial Optimization*. John Wiley and Sons, New York, 1997.
8. J. Edmonds and E. L. Johnson. Matching: A well-solved class of integer linear programs. In *Combinatorial Structures and their Applications*, pages 89–92. Gordon and Breach, New York, 1970.
9. G. Even, J. Feldman, G. Kortsarz, and Z. Nutov. A $3/2$ -approximation algorithm for augmenting the edge-connectivity of a graph from 1 to 2 using a subset of a given edge set. In *Proceedings of the 4th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems*, pages 90–101, 2001.
10. U. Feige. A threshold of $\ln n$ for approximating set cover. *Journal of the ACM*, 45(4):634–652, 1998.
11. T. Fujito. On approximation of the submodular set cover problem. *Operations Research Letters*, 25(4):169–174, 1999.
12. R. Gandhi, S. Khuller, and A. Srinivasan. Approximation algorithms for partial covering problems. *Journal of Algorithms*, 53(1):55–84, 2004.
13. R. S. Garfinkel and G. L. Nemhauser. Optimal set covering: A survey. In A. M. Geoffrion, editor, *Perspectives on Optimization*, pages 164–183, 1972.
14. N. Garg, V. V. Vazirani, and M. Yannakakis. Approximate max-flow min-(multi)cut theorems and their applications. *SIAM Journal on Computing*, 25(2):235–251, 1996.
15. N. Garg, V. V. Vazirani, and M. Yannakakis. Primal-dual approximation algorithms for integral flow and multicut in trees. *Algorithmica*, 18(1):3–20, 1997.
16. D. R. Gaur, T. Ibaraki, and R. Krishnamurti. Constant ratio approximation algorithms for the rectangle stabbing problem and the rectilinear partitioning problem. *Journal of Algorithms*, 43(1):138–152, 2002.
17. D. Golovin, V. Nagarajan, and M. Singh. Approximating the k -multicut problem. In *Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 621–630, 2006.
18. D. S. Hochbaum. Approximating covering and packing problems: Set cover, vertex cover, independent set, and related problems. In D. S. Hochbaum, editor, *Approximation Algorithms for NP-Hard Problems*, chapter 3, pages 94–143. PWS Publishing Company, 1997.

19. D. S. Hochbaum. The t -vertex cover problem: Extending the half integrality framework with budget constraints. In *Proceedings of the 1st International Workshop on Approximation Algorithms for Combinatorial Optimization*, pages 111–122, 1998.
20. D. S. Johnson. Approximation algorithms for combinatorial problems. *Journal of Computer and System Sciences*, 9(3):256–278, 1974.
21. M. J. Kearns. *The Computational Complexity of Machine Learning*. MIT Press, 1990.
22. S. Khot. On the power of unique 2-prover 1-round games. In *Proceedings of the 34th Annual ACM Symposium on Theory of Computing*, pages 767–775, 2002.
23. J. Könemann, O. Parekh, and D. Segev. A unified approach to approximating partial covering problems, 2006. Available at <http://www.math.tau.ac.il/~segev/>.
24. A. Levin and D. Segev. Partial multicuts in trees. In *Proceedings of the 3rd International Workshop on Approximation and Online Algorithms*, pages 320–333, 2005.
25. L. Lovász. On the ratio of optimal integral and fractional covers. *Discrete Mathematics*, 13:383–390, 1975.
26. K. G. Murty and C. Perin. A 1-matching blossom type algorithm for edge covering problems. *Networks*, 12:379–391, 1982.
27. H. Nagamochi. An approximation for finding a smallest 2-edge-connected subgraph containing a specified spanning tree. *Discrete Applied Mathematics*, 126(1):83–113, 2003.
28. M. W. Padberg. Covering, packing and knapsack problems. *Annals of Discrete Mathematics*, 4:265–287, 1979.
29. O. Parekh. *Polyhedral Techniques for Graphic Covering Problems*. PhD thesis, Department of Mathematical Sciences, Carnegie Mellon University, 2002.
30. J. Plesník. Constrained weighted matchings and edge coverings in graphs. *Discrete Applied Mathematics*, 92(2–3):229–241, 1999.
31. H. Räcke. Minimizing congestion in general networks. In *Proceedings of the 43rd Annual IEEE Symposium on Foundations of Computer Science*, pages 43–52, 2002.
32. P. Slavík. Improved performance of the greedy algorithm for partial cover. *Information Processing Letters*, 64(5):251–254, 1997.