

LAGRANGIAN RELAXATION AND PARTIAL COVER

JULIÁN MESTRE ¹

¹ Max-Planck-Institute für Informatik, Saarbrücken, Germany.
E-mail address: `jmestre@mpi-inf.mpg.de`

ABSTRACT. Lagrangian relaxation has been used extensively in the design of approximation algorithms. This paper studies its strengths and limitations when applied to Partial Cover.

We show that for Partial Cover in general no algorithm that uses Lagrangian relaxation and a Lagrangian Multiplier Preserving (LMP) α -approximation as a black box can yield an approximation factor better than $\frac{4}{3}\alpha$. This matches the upper bound given by Könemann *et al.* (*ESA 2006*, pages 468–479).

Faced with this limitation we study a specific, yet broad class of covering problems: Partial Totally Balanced Cover. By carefully analyzing the inner workings of the LMP algorithm we are able to give an almost tight characterization of the integrality gap of the standard linear relaxation of the problem. As a consequence we obtain improved approximations for the Partial version of Multicut and Path Hitting on Trees, Rectangle Stabbing, and Set Cover with ρ -Blocks.

1. Introduction

Lagrangian relaxation has been used extensively in the design of approximation algorithms for a variety of problems such as TSP [20, 21], k -MST [13, 2, 8, 12], partial vertex cover [23], k -median [26, 6, 1], MST with degree constraints [32] and budgeted MST [36].

In this paper we study the strengths and limitations of Lagrangian relaxation applied to the Partial Cover problem. Let \mathcal{S} be collection of subsets of a universal set U with cost $c : \mathcal{S} \rightarrow R_+$ and profit $p : U \rightarrow R_+$, and let P be a target coverage parameter. A set $\mathcal{C} \subseteq \mathcal{S}$ is a *partial cover* if the overall profit of elements covered by \mathcal{C} is at least P . The objective is to find a minimum cost partial cover.

The high level idea behind Lagrangian relaxation is as follows. In an IP formulation for Partial Cover, the constraint enforcing that at least P profit is covered is *relaxed*: The constraint is multiplied by a parameter λ and lifted to the objective function. This relaxed IP corresponds, up to a constant factor, to the prize-collecting version of the underlying covering problem in which there is no requirement on how much profit to cover but a penalty of $\lambda p(i)$ must be paid if we leave element $i \in U$ uncovered. An approximation algorithm

2000 ACM Subject Classification: G.2.1.

Key words and phrases: Lagrangian Relaxation, Partial Cover, Primal-Dual Algorithms.

Supported by NSF Award CCF 0430650, a University of Maryland Dean's Dissertation Fellowship, and partly by an Alexander von Humboldt Fellowship.

for the prize-collecting version having the Lagrangian Multiplier Preserving (LMP) property¹ is used to obtain values λ_1 and λ_2 that are close together for which the algorithm produces solutions \mathcal{C}_1 and \mathcal{C}_2 respectively. These solutions are such that \mathcal{C}_1 is inexpensive but unfeasible (covering less than P profit), and \mathcal{C}_2 is feasible (covering at least P profit) but potentially very expensive. Finally, these two solutions are combined to obtain a cover that is both inexpensive and feasible.

Broadly speaking there are two ways to combine \mathcal{C}_1 and \mathcal{C}_2 . One option is to treat the approximation algorithm for the prize-collecting version as a black box, only making use of the LMP property in the analysis. Another option is to focus on a particular LMP algorithm and exploit additional structure that it may offer. Not surprisingly, the latter approach has yielded better approximation guarantees. For example, for k -median compare the 6-approximation of Jain and Vazirani [26] to the 4-approximation of Charikar and Guha [6]; for k -MST compare the 5-factor to the 3-factor approximation due to Garg [13].

The results in this paper support the common belief regarding the inherent weakness of the black-box approach. First, we show a lower bound on the approximation factor achievable for Partial Cover in general using Lagrangian relaxation and the black-box approach that matches the recent upper bound of Könemann et al. [31]. To overcome this obstacle, we concentrate on Kolen’s algorithm for Prize-Collecting Totally Balanced Cover [30]. By carefully analyzing the algorithm’s inner workings we identify structural similarities between \mathcal{C}_1 and \mathcal{C}_2 , which we later exploit when combining the two solutions. As a result we derive an almost tight characterization of the integrality gap of the standard linear relaxation for Partial Totally Balanced Cover. This in turn implies improved approximation algorithms for a number of related problems.

1.1. Related Work

Much work has been done on covering problems because of both their simple and elegant formulation, and their pervasiveness in different application areas. In its most general form the problem, also known as Set Cover, cannot be approximated within $(1 - \epsilon) \ln |U|$ unless $NP \subseteq DTIME(|U|^{\log \log |U|})$ [10]. Due to this hardness, easier, special cases have been studied.

A general class of covering problems that can be solved efficiently are those whose element-set incidence matrix is balanced. A $0, 1$ matrix is *balanced* if it does not contain a square submatrix of odd order with row and column sums equal to 2. These matrices were introduced by Berge [5] who showed that if A is balanced then the polyhedron $\{x \geq 0 : Ax \geq 1\}$ is integral. A $0, 1$ matrix is *totally balanced* if it does not contain a square submatrix with row and column sums equal to 2 and no identical columns. Kolen [30] gave a simple primal-dual algorithm that solves optimally the covering problem defined by a totally balanced matrix. A $0, \pm 1$ matrix is *totally unimodular* if every square submatrix has determinant 0 or ± 1 . Although totally balanced and totally unimodular matrices are subclasses of balanced matrices, the two classes are neither disjoint nor one is included in the other.

Beyond this point, even minor generalizations can make the covering problem hard. For example, consider the *vertex cover* problem: Given a graph $G = (V, E)$ we are to choose a minimum size subset of vertices such that every edge is incident on at least one of the chosen vertices. If G is bipartite, the element-set incidence matrix for the problem is totally unimodular; however, if G is a general graph the problem becomes NP-hard [29]. Numerous

¹The definition of the LMP property is outlined in Section 2.

approximation algorithms have been developed for vertex cover [24]. The best known approximation factor for general graphs is $2 - o(1)$ [4, 18, 28]; yet, after 25 years of study, the best constant factor approximation for vertex cover remains 2 [9, 3, 22]. This lack of progress has led researchers to seek generalizations of vertex cover that can still be approximated within twice of optimum. One such generalization is the *multicut* problem on trees: Given a tree T and a collection of pairs of vertices, a cover is formed by a set of edges whose removal separates all pairs. The problem was first studied by Garg et al. [14] who gave an elegant primal-dual 2-approximation.

A notable shortcoming of the standard set cover formulation is that certain hard-to-cover elements, also known as *outliers* [7], can render the optimal solution very expensive. Motivated by the presence of outliers, the unit-profit partial version calls for a collection of sets covering not all but a specified number k of elements. Partial Multicut, a.k.a. k -Multicut, was recently studied independently by Levin and Segev [33] and by Golovin et al. [17], who gave a $\frac{8}{3} + \epsilon$ approximation algorithm. This scheme was generalized by Könemann et al. [31] who showed how to design a $\frac{4}{3}\alpha + \epsilon$ approximation for any covering problem using Lagrangian relaxation and an α -LMP approximation as a black box. (Their algorithm runs in time polynomial on $|U|, |S|^{\frac{1}{\epsilon}}$ and the running time of the α -LMP approximation.)

1.2. Our Results and Outline of the Paper

Section 3 shows that for Partial Cover in general no algorithm that uses Lagrangian relaxation and an α -LMP approximation as a black box can yield an approximation factor better than $\frac{4}{3}\alpha$. In Section 4 we give an almost tight characterization of the integrality gap of the standard LP for Partial Totally Balanced Cover, settling a question posed by Golovin et al. [17]. Our approach is based on Lagrangian relaxation and Kolen's algorithm. We prove that $\text{IP} \leq \left(1 + \frac{1}{3^{k-1}}\right) \text{LP} + k c_{\max}$ for any $k \geq 1$, where IP and LP are the costs of the optimal integral and fractional solutions respectively and c_{\max} is the cost of the most expensive set in the instance. The trade-off between additive and multiplicative error is not an artifact of our analysis or a shortcoming of our approach. On the contrary, this is precisely how the integrality gap behaves. More specifically, we show a family of instances where $\text{IP} > \left(1 + \frac{1}{3^{k-1}}\right) \text{LP} + \frac{k}{2} c_{\max}$. In other words, there is an unbounded additive gap in terms of c_{\max} but as it grows the multiplicative gap narrows exponentially fast.

Finally, in Section 5 we show how the above result can be applied, borrowing ideas from [15, 19, 17], to get a $\rho + \epsilon$ approximation or a quasi-polynomial time ρ -approximation for covering problems that can be expressed with a suitable combination of ρ totally-balanced matrices. This translates into improved approximations for a number of problems: a $2 + \epsilon$ approximation for the Partial Multicut on Trees [33, 17], a $4 + \epsilon$ approximation for Partial Path Hitting on Trees [35], a 2-approximation for Partial Rectangle Stabbing [15], and a ρ approximation for Partial Set-Cover with ρ -blocks [19]. In addition, the ϵ can be removed from the first two approximation guarantees if we allow quasi-polynomial time. It is worth noting that prior to this work, the best approximation ratio for all these problems could be achieved with the framework of Könemann et al. [31]. In each case our results improve the approximation ratio by a $\frac{4}{3}$ multiplicative factor.

2. Lagrangian relaxation

Let $\mathcal{S} = \{1, \dots, m\}$ be a collection of subsets of a universal set $U = \{1, \dots, n\}$. Each set has a cost specified by $c \in R_+^m$, and each element has a profit specified by $p \in R_+^n$. Given a target coverage P , the objective of the Partial Cover problem is to find a minimum cost solution $\mathcal{C} \subseteq \mathcal{S}$ such that $p(\mathcal{C}) \geq P$, where the notation $p(\mathcal{C})$ denotes the overall profit of elements covered by \mathcal{C} . The problem is captured by the IP below. Matrix $A = \{a_{ij}\} \in \{0, 1\}^{n \times m}$ is an element-set incidence matrix, that is, $a_{ij} = 1$ if and only if element $i \in U$ belongs to set $j \in \mathcal{S}$; variable x_j indicates whether set j is chosen in the solution \mathcal{C} ; variable r_i indicates whether element i is left uncovered.

Lagrangian relaxation is used to get rid of the constraint bounding the profit of uncovered elements to be at most $p(U) - P$. The constraint is multiplied by the parameter λ , called Lagrange Multiplier, and is lifted to the objective function. The resulting IP corresponds, up to the constant $\lambda(p(U) - P)$ factor in the objective function, to the prize-collecting version of the covering problem, where the penalty for leaving element i uncovered is λp_i .

$$\begin{array}{ccc}
 \min c \cdot x & & \min c \cdot x + \lambda p \cdot r - \lambda(p(U) - P) \\
 Ax + Ir \geq 1 & & Ax + Ir \geq 1 \\
 p \cdot r \leq p(U) - P & \xrightarrow{\text{Lagrangian Relaxation}} & r_i, x_j \in \{0, 1\} \\
 r_i, x_j \in \{0, 1\} & &
 \end{array}$$

Let OPT be the cost of an optimal partial cover and $\text{OPT-PC}(\lambda)$ be the cost of an optimal prize-collecting cover for a given λ . Let \mathcal{A} be an α -approximation for the prize-collecting variant of the problem. Algorithm \mathcal{A} is said to have the Lagrangian Multiplier Preserving (LMP) property if it produces a solution \mathcal{C} such that

$$c(\mathcal{C}) + \alpha \lambda (p(U) - p(\mathcal{C})) \leq \alpha \text{OPT-PC}(\lambda). \quad (2.1)$$

Note that $\text{OPT-PC}(\lambda) \leq \text{OPT} + \lambda(p(U) - P)$. Thus,

$$c(\mathcal{C}) \leq \alpha \left(\text{OPT} + \lambda(p(\mathcal{C}) - P) \right). \quad (2.2)$$

Therefore, if we could find a value of λ such that \mathcal{C} covers exactly P profit then \mathcal{C} is α -approximate. However, if $p(\mathcal{C}) < P$, the solution is not feasible, and if $p(\mathcal{C}) > P$, equation (2.2) does not offer any guarantee on the cost of \mathcal{C} . Unfortunately, there are cases where no value of λ produces a solution covering exactly P profit. Thus, the idea is to use binary search to find two values λ_1 and λ_2 that are close together and are such that $\mathcal{A}(\lambda_1)$ covers less, and $\mathcal{A}(\lambda_2)$ covers more than P profit. The two solutions are then combined in some fashion to produce a feasible cover.


3. Limitations of the black-box approach

A common way to combine the two solutions returned by the α -LMP is to treat the algorithm as a black box, solely relying on the LMP property (2.1) in the analysis. More formally, an algorithm for Partial Cover that uses Lagrangian relaxation and an α -LMP approximation \mathcal{A} as a black box is as follows. First, we are allowed to run \mathcal{A} with as many different values of λ as desired; then, the solutions thus found are combined to produce a

feasible partial cover. No computational restriction is placed on the second step, except that only sets returned by \mathcal{A} may be used.

Theorem 3.1. *In general, the Partial Cover problem cannot be approximated better than $\frac{4}{3}\alpha$ using Lagrangian relaxation and an α -LMP algorithm \mathcal{A} as a black box.*

Let A_1, \dots, A_q and B_1, \dots, B_q be sets as depicted on the right. For each i and j the intersection $A_i \cap B_j$ consists of a cluster of q elements. There are q^2 clusters. Set A_i is made up of q clusters; set B_i is made up of q clusters and two additional elements (the leftmost and rightmost elements in the picture.) Thus $|A_i| = q^2$ and $|B_i| = q^2 + 2$. In addition, there are sets O_1, \dots, O_q , which are not shown in the picture. Set O_i contains one element from each cluster and the leftmost element of B_i . Thus $|O_i| = q^2 + 1$. The cost of O_i is $\frac{1}{q}$, the cost of A_i is $\frac{2\alpha}{3q}$, and the cost of B_i is $\frac{4\alpha}{3q}$. Every element has unit profit and the target coverage is $P = q^3 + q$. It is not hard to see that O_1, \dots, O_q is an optimal partial cover with a cost of 1. Furthermore, for any value of λ the optimal prize-collecting cover uses sets of one kind.



Lemma 3.2. *For all values of λ , the optimal prize-collecting cover in our instance is either the empty cover or A_1, \dots, A_q or B_1, \dots, B_q or O_1, \dots, O_q .*

Proof. Suppose that some subsets of the B -sets and the O -sets have already been chosen. Every A -set has the same marginal benefit (the penalty of elements not-yet-covered by the set minus its cost) the independent of the other A -sets. Thus, in an optimal solution either all the A -sets are chosen or none is.


Now suppose that the O -sets and the A -sets have already been chosen. In this case there are two marginal benefits for the B -sets, depending on whether an O -set already covers the leftmost element of the B -set or not. Thus, an optimal strategy for the B -sets is either to choose none, all, or the complement of the O -sets, i.e., B_i is chosen if and only if O_i is not chosen. A solution where B -sets and O -sets complement each other is always worse than either choosing all the O -sets and no B -sets, or vice versa. Thus, in an optimal solution either all the B -sets are chosen or none is. It follows that the same holds for the O -sets; i.e., they are either all in or all out.

Notice that if the B -sets are chosen then there is no reason to choose any of the remaining sets. If the B -sets are not chosen and the O -sets are chosen then there is no reason to choose the A -sets. The only possibility left is to choose only the A -sets, or the empty cover. ■


The α -LMP approximation algorithm we use has the unfortunate property that it never returns sets from the optimal solution.

Lemma 3.3. *There exists an α -LMP approximation \mathcal{A} that for the above instance and any value of λ outputs either \emptyset or A_1, \dots, A_q or B_1, \dots, B_q .*

Proof. Let us denote each of the four alternatives in Lemma 3.2 by E , A , B and O . The prize-collecting cost of these covers is $(q^3 + 2q)\lambda$, $\frac{2}{3}\alpha + 2q\lambda$, $\frac{4}{3}\alpha$ and $1 + q\lambda$ respectively. For a fixed value of λ , the minimum of these four quantities corresponds to the cost of the optimal prize-collecting cover. The line below shows which solution is optimal as a function of λ .



If α is big enough then the A -interval disappears. Namely, if $\frac{2\alpha}{3q^3} > \frac{3-2\alpha}{3q}$ then the line looks as follows



We are now ready to describe the α -LMP algorithm. Given a value of λ , we need to decide whether to output the empty cover, the A -sets or the B -sets. If λ falls in the interval corresponding to one of these three solutions then output that solution; since the cover output is optimal, the LMP property (2.1) follows trivially.

If λ falls in the O -interval and $\lambda \leq \frac{1}{3q}$ then output the A -sets; the LMP property holds since

$$c(A) + \alpha \lambda \overline{p(A)} = \frac{2}{3}\alpha + \alpha \lambda 2q \leq \alpha(1 + \lambda q) = \alpha \left(c(O) + \lambda(p(U) - p(O)) \right).$$

If λ falls in the O -interval and $\lambda > \frac{1}{3q}$ then output the B -sets; the LMP property holds since

$$c(B) + \alpha \lambda \overline{p(B)} = \frac{4}{3}\alpha < \alpha(1 + \lambda q) = \alpha \left(c(O) + \lambda(p(U) - p(O)) \right).$$

■

Hence, if we use \mathcal{A} as a black box we must build a partial cover with the sets A_1, \dots, A_q and B_1, \dots, B_q . Note that in order to cover $q^2 + q$ elements either all A -sets, or all B -sets must be used. In the first case $\frac{q}{2}$ additional B -sets are needed to attain feasibility, and the solution has cost $\frac{4}{3}\alpha$; in the second case the solution is feasible but again has cost $\frac{4}{3}\alpha$. Theorem 3.1 follows.

One assumption usually made in the literature [2, 11, 31] is that $c_{\max} = \max_j c_j \leq \epsilon \text{OPT}$, for some constant $\epsilon > 0$, or more generally an additive error in terms of c_{\max} is allowed. This does not help in our construction as c_{\max} can be made arbitrarily small by increasing q .

Admittedly, our lower bound example belongs to a specific class of covering problem (every element belongs to at most three sets) and although the example can be embedded into a partial totally unimodular covering problem (see Appendix A), it is not clear how to embed the example into other classes. Nevertheless, the $\frac{4}{3}\alpha$ upper bound of Koneman et al. [31] makes no assumption about the underlying problem, only using the LMP property (2.1) in the analysis. It was entirely conceivable that the $\frac{4}{3}\alpha$ factor could be improved using a different merging strategy—Theorem 3.1 precludes this possibility.

4. Partial Totally Balanced Cover

In order to overcome the lower bound of Theorem 3.1, one must concentrate on a specific class of covering problems or make additional assumptions about the α -LMP algorithm. In this section we focus on covering problems whose IP matrix A is totally balanced. More specifically, we study the integrality gap of the standard linear relaxation for Partial Totally Balanced Cover (P-TBC) shown below.

Theorem 4.1. *Let IP and LP be the cost of the optimal integral and fractional solutions of an instance of P-TBC. Then $\text{IP} \leq \left(1 + \frac{1}{3^{k-1}}\right) \text{LP} + k c_{\max}$ for any $k \in \mathbb{Z}_+$. Furthermore, for any large enough $k \in \mathbb{Z}_+$ there exists an instance where $\text{IP} > \left(1 + \frac{1}{3^{k-1}}\right) \text{LP} + \frac{k}{2} c_{\max}$.*

$$\begin{array}{ccc}
 \min c \cdot x & & \max 1 \cdot y - (p(U) - P) \lambda \\
 Ax + Ir \geq 1 & & A^T y \leq c \\
 p \cdot r \leq p(U) - P & \xrightarrow{\text{LP Duality}} & y \leq \lambda p \\
 r_i, x_e \geq 0 & & y_i, \lambda \geq 0
 \end{array}$$

The rest of this section is devoted to proving Theorem 4.1 Our approach is based on Lagrangian relaxation and Kolen's algorithm for Prize-Collecting Totally Balanced Cover (PC-TBC). The latter exploits the fact that a totally balanced matrix can be put into greedy standard form by permuting the order of its rows and columns; in fact, the converse is also true [25]. A matrix is in standard greedy form if it does not contain as an induced submatrix

$$\begin{bmatrix} 1 & 1 \\ 1 & 0 \end{bmatrix} \tag{4.1}$$

There are polynomial time algorithms that can transform a totally balanced matrix into greedy standard form [37] by shuffling the rows and columns of A . Since this transformation does not affect the underlying covering problem, we assume that A is given in standard greedy form.

4.1. Kolen's algorithm for Prize-Collecting Totally Balanced Cover

For the sake of completeness we describe Kolen's primal-dual algorithm for PC-TBC. The algorithm finds a dual solution y and a primal solution C , which is then pruned in a reverse-delete step to obtain the final solution \hat{C} . The linear and dual relaxations for PC-TBC appear below.

$$\begin{array}{ccc}
 \min c \cdot x + \lambda p \cdot r & & \max 1 \cdot y \\
 Ax + Ir \geq 1 & & A^T y \leq c \\
 r_i, x_e \geq 0 & \xrightarrow{\text{LP Duality}} & y \leq \lambda p \\
 & & y_i \geq 0
 \end{array}$$

The residual cost of the set j w.r.t. y is defined as $c'_j = c_j - \sum_{i:a_{ij}=1} y_i$. The algorithm starts from the trivial dual solution $y = 0$, and processes the elements in increasing column

KOLEN(A, c, p, λ)	
1 // Dual update	7 // Reverse delete
2 $y \leftarrow 0, C \leftarrow \emptyset, \widehat{C} \leftarrow \emptyset$	8 while $C \neq \emptyset$ do
3 for $i \leftarrow 1$ to n do	9 $j \leftarrow$ largest set index in C
4 $\delta \leftarrow \min\{c'_j \mid a_{ij} = 1\}$	10 $\widehat{C} \leftarrow \widehat{C} + j$
5 $y_i \leftarrow \min\{\lambda p_i, \delta\}$	11 $C \leftarrow C \setminus \{j' \mid j \text{ dominates } j' \text{ or } j = j'\}$
6 $C \leftarrow \{j \mid c'_j = 0\}$	12 return (\widehat{C}, y)

order of A^T . Let i the index of the current element. Its corresponding dual variable, y_i , is increased until either the residual cost of some set j containing i equals 0 (we say set j becomes tight), or y_i equals λp_i (Lines 3-5).

Let $C = \{j \mid c'_j = 0\}$ be the set of tight sets after the dual update is completed. As it stands the cover C may be too expensive to be accounted for using the lower bound provided by $1 \cdot y$ because a single element may belong to multiple sets in C . The key insight is that some of the sets in C are redundant and can be pruned.

Definition 4.2. Given sets j_1, j_2 we say that j_1 *dominates* j_2 in y if $j_1 > j_2$ and there exists an item i such that $y_i > 0$ and i belongs to j_1 and j_2 , that is, $a_{ij_1} = a_{ij_2} = 1$.

The reverse-delete step iteratively identifies the largest index j in C , adds j to \widehat{C} , and removes j and all the sets it dominates. This is repeated until no set is left in C (Lines 8–11).

Notice that all sets $j \in C$ are tight, thus we can pay for set j by charging the dual variables of items that belong to j . Because of the reverse-delete step if $y_i > 0$ then i belongs to at most one set in \widehat{C} ; thus in paying for \widehat{C} we charge covered items at most once. Using the fact A is in standard greedy form, it can be shown [30] that if i was left uncovered then we can afford its penalty, i.e., $y_i = \lambda p_i$. The solution \widehat{C} is optimal for PC-TBC since

$$\sum_{j \in \widehat{C}} c_j + \sum_{\substack{i \in U \text{ s.t.} \\ \nexists j \in \widehat{C} : a_{ij} = 1}} \lambda p_i = \sum_{\substack{i \in U \text{ s.t.} \\ \exists j \in \widehat{C} : a_{ij} = 1}} y_i + \sum_{\substack{i \in U \text{ s.t.} \\ \nexists j \in \widehat{C} : a_{ij} = 1}} y_i = \sum_{i \in U} y_i. \quad (4.2)$$

If we could find a value of λ such that KOLEN(A, c, p, λ) returns a solution (\widehat{C}, y) covering *exactly* P profit, we are done since from (4.2) it follows that

$$\sum_{j \in \widehat{C}} c_j = \sum_{i \in U} y_i - \lambda(p(U) - P). \quad (4.3)$$

Notice that (y, λ) is a feasible for the dual relaxation of P-TBC and its cost is precisely the right hand side of (4.3). Therefore for this instance IP=DL=LP and Theorem 4.1 follows.

Unfortunately, there are cases where no such value of λ exists. Nonetheless, we can always find a *threshold value* λ such that for any infinitesimally small $\delta > 0$, $\lambda^- = \lambda - \delta$ and $\lambda^+ = \lambda + \delta$ produce solutions covering less and more than P profit respectively. A threshold value can be found using Megiddo's parametric search [34] by making $O(n \log m)$ calls to the procedure KOLEN. For completeness the technique is sketched in Appendix B.

Let y (y^-) be the dual solution and C (C^-) the set of tight sets when KOLEN is run on λ (λ^-). Without loss of generality assume \widehat{C} covers more than P profit. (The case where \widehat{C} covers less than P profit is symmetrical: we work with y^+ and C^+ instead of y^- and C^- .)

Our plan to prove Theorem 4.1 is to devise an algorithm to merge \widehat{C} and \widehat{C}^- in order to obtain a cheap solution covering at least P profit.

4.2. Merging two solutions

Before describing the algorithm we need to establish some important properties regarding these two solutions and their corresponding dual solutions.

For any i , the value of y_i^- is a linear function of δ for all i . This follows from the fact that δ is infinitesimally small. Furthermore, the constant term in this linear function is y_i .

Lemma 4.3. *For each $i \in U$ there exists $a \in \mathbb{Z}$, independent of δ , such that $y_i^- = y_i + a\delta$.*

Proof. By induction on the number of iteration of the dual update step of KOLEN, using the fact that the same property holds for the residual cost of the sets. ■

A useful corollary of Lemma 4.3 is that $C^- \subseteq C$, since if the residual cost of a set is non-zero in y it must necessarily be non-zero in y^- . The other way around may not hold.

At the heart of our approach is the notion of a merger graph $G = (V, E)$. The vertex set of G is made up of sets from the two solutions, i.e., $V = \widehat{C} \oplus \widehat{C}^-$. The edges of G are directed and given by

$$E = \left\{ (j_1, j_2) \mid \begin{array}{l} j_1 \in \widehat{C}^- \setminus \widehat{C}, j_2 \in \widehat{C} \setminus \widehat{C}^- \text{ s.t. } j_1 \text{ dominates } j_2 \text{ in } y^-, \text{ or} \\ j_1 \in \widehat{C} \setminus \widehat{C}^-, j_2 \in \widehat{C}^- \setminus \widehat{C} \text{ s.t. } j_1 \text{ dominates } j_2 \text{ in } y \end{array} \right\} \quad (4.4)$$

This graph has a lot of structure that can be exploited when merging the solutions.

Lemma 4.4. *The merger graph $G = (V, E)$ of \widehat{C}^- and \widehat{C} is a forest of out-branchings.*

Proof. First note that G is acyclic, since if $(j_1, j_2) \in E$ then necessarily $j_1 > j_2$. Thus, it is enough to show that the in-degree of every $j \in V$ is at most one. Suppose otherwise, that is, there exist $j_1, j_2 \in V$ such that $(j_1, j), (j_2, j) \in E$. Assume that $j_1 < j_2$ and $j \in \widehat{C}$ (the remaining cases are symmetrical).

By definition (4.4), we know that $j_1(j_2) \in \widehat{C}^-$ and that there exists $i_1(i_2)$ that belongs to j and $j_1(j_2)$ such that $y_{i_1}^- > 0$ ($y_{i_2}^- > 0$). Since A^T is in standard greedy form we can infer that i_2 belongs to j_1 if $i_1 < i_2$, or i_1 belongs to j_2 if $i_1 > i_2$: The diagram on the right shows how, using the fact that A^T does not contain (4.1) as an induced submatrix, we can infer that the boxed entries must be 1. In either case we get that j_2 dominates j_1 in y^- , which contradicts the fact that both belong to \widehat{C}^- . ■

$$\begin{array}{cc|cc} & i_1 & i_2 & i_2 & i_1 \\ j & 1 & 1 & 1 & 1 \\ j_1 & 1 & \boxed{1} & & 1 \\ j_2 & & 1 & 1 & \boxed{1} \end{array}$$

The procedure MERGE starts from the unfeasible solution $D = \widehat{C}^-$ and guided by the merger graph G , it modifies D step by step until feasibility is attained. The operation used to update D is to take the symmetric difference of D and a subtree of G rooted at a vertex $r \in V$, which we denote by T_r . For each root r of an out-branchings of G we set $D \leftarrow D \oplus T_r$, until $p(D \oplus T_r) > P$. At this point we return the solution produced by INCREASE(r, D).

Notice that after setting $D \leftarrow D \oplus T_r$ in Line 5, the solution D “looks like” \widehat{C} within T_r . Indeed, if all roots are processed then $D = \widehat{C}$. Therefore, at some point we are bound to have $p(D \oplus T_r) > P$ and to make the call INCREASE(r, D) in Line 6. Before describing

```

MERGE( $\widehat{C}^-$ ,  $\widehat{C}$ )
1 let  $G$  be the merger graph for  $\widehat{C}^-$  and  $\widehat{C}$ 
2  $D \leftarrow \widehat{C}^-$ 
3 for each root  $r$  in  $G$  do
4   if  $p(D \oplus T_r) \leq P$ 
5     then  $D \leftarrow D \oplus T_r$ 
6     else return INCREASE( $r, D$ )

```

INCREASE we need to define a few terms. Let the *absolute benefit* of set j , which we denote by b_j , be the profit of elements uniquely covered by set j , that is,

$$b_j = p\left(\{i \in U \mid \forall j' \in \widehat{C} \cup \widehat{C}^- : a_{ij'} = 1 \text{ iff } j' = j\}\right). \quad (4.5)$$

Let $D \subseteq \widehat{C} \cup \widehat{C}^-$. Note that if $j \in D$, the removal of j decreases the profit covered by D by at least b_j ; on the other hand, if $j \notin D$, its addition increases the profit covered by at least b_j . This notion of benefit can be extended to subtrees,

$$\Delta(T_j, D) = \sum_{j' \in T_j \setminus D} b_{j'} - \sum_{j' \in T_j \cap D} b_{j'}. \quad (4.6)$$

We call this quantity the *relative benefit* of T_j with respect to D . It shows how the profit of uniquely covered elements changes when we take $D \oplus T_j$. Note that $\Delta(T_j, D)$ can be positive or negative.

Everything is in place to explain INCREASE(j, D). The algorithm assumes the input solution is unfeasible but can be made feasible by adding some sets in T_j ; more precisely, we assume $p(D) \leq P$ and $P < p(D) + \Delta(T_j, D)$. If adding j to D makes the solution feasible then return $D + j$ (Lines 2-3). If there exists a child c of j that can be used to propagate the call down the tree then do that (Lines 4-5). Otherwise, *split* the subtree T_j : Add j to D and process the children of c , setting $D \leftarrow D \oplus T_c$ until D becomes feasible (Lines 6-9). At this point $p(D) > P$ and $p(D \oplus T_c) \leq P$. If $P - p(D \oplus T_c) < p(D) - P$ then call INCREASE($c, D \oplus T_c$) else call DECREASE(c, D) and let D' be the cover returned by the recursive call (Lines 10-12). Finally, return the cover with minimum cost between D and D' .

The twin procedure DECREASE(j, D) is essentially symmetrical: Initially the input is feasible but can be made unfeasible by removing some sets in T_j ; more precisely $p(D) \geq P$ and $P < p(D) + \Delta(T_c, D)$.

At a very high level, the intuition behind the INCREASE/DECREASE scheme is as follows. In each call one of three things must occur:

- (i) A feasible cover with a small coverage excess is found (Lines 2-3), or
- (ii) The call is propagated down the tree at no cost (Lines 4-5), or
- (iii) A subtree T_j is split (Lines 6-9). In this case, the cost c_j cannot be accounted for, but the offset in coverage $|P - p(D)|$ is reduced at least by a factor of 3.

If the INCREASE/DECREASE algorithms split many subtrees (incurring a high extra cost) then the offset in coverage must have been very high at the beginning, which means the cost of the dual solution is high and so the splitting cost can be charged to it. In order to flesh out

<pre> INCREASE(j, D) 1 // assume $p(D) \leq P < p(D) + \Delta(T_j, D)$ 2 if $p(D + j) \geq P$ 3 then return $D + j$ 4 if \exists child c of $j : p(D) + \Delta(T_c, D) > P$ 5 then return INCREASE(c, D) 6 $D \leftarrow D + j$ 7 while $p(D) \leq P$ do 8 $c \leftarrow$ child of j maximizing $\Delta(T_c, D)$ 9 $D \leftarrow D \oplus T_c$ 10 if $P - p(D \oplus T_c) < p(D) - P$ 11 then $D' \leftarrow$ INCREASE($c, D \oplus T_c$) 12 else $D' \leftarrow$ DECREASE(c, D) 13 return min cost $\{D, D'\}$ </pre>	<pre> DECREASE(j, D) 1 // assume $p(D) \geq P > p(D) + \Delta(T_j, D)$ 2 if $p((D \oplus T_j) + j) \geq P$ 3 then return $(D \oplus T_j) + j$ 4 if \exists child c of $j : p(D) + \Delta(T_c, D) < P$ 5 then return DECREASE(c, D) 6 $D \leftarrow D + j$ 7 while $p(D) \geq P$ do 8 $c \leftarrow$ child of j minimizing $\Delta(T_c, D)$ 9 $D \leftarrow D \oplus T_c$ 10 if $p(D \oplus T_c) - P < P - p(D)$ 11 then $D' \leftarrow$ INCREASE(c, D) 12 else $D' \leftarrow$ DECREASE($c, D \oplus T_c$) 13 return min cost $\{D \oplus T_c, D'\}$ </pre>
---	--

these ideas into a formal proof we need to establish some crucial properties of the merger graph and the algorithms.

Lemma 4.5. *If $y_i < \lambda p_i$ then there exist $j' \in \widehat{C}$ and $j'' \in \widehat{C}^-$ such that either $j' = j''$ or $(j', j'') \in E$ or $(j'', j') \in E$.*

Proof. Since $y_i < \lambda p_i$, by Lemma 4.3 we get that $y_i^- < \lambda p_i$ as well. Thus, there exists a set $j \in C^-$ such that $a_{ij} = 1$ that becomes tight right after processing i in the dual update of KOLEN. Due to the reverse-delete step either $j \in \widehat{C}^-$ or there exists $j_1 \in \widehat{C}^-$ that dominates j in y^- through some element $i_1 \in U$, i.e., $y_{i_1} > 0$ and $a_{i_1 j} = a_{i_1 j_1} = 1$. In the latter case, since the set j is tight after i is processed, it follows that $i_1 \leq i$. Because A is in standard greedy form we infer that $a_{i_1 j_1} = 1$. By Lemma 4.3, we have $C^- \subseteq C$, thus $j \in C$. A similar reasoning as above shows that either $j \in \widehat{C}$ or there exists $j_2 \in \widehat{C}$ that dominates j in y through some element $i_2 \in U$ such that $a_{i_2 j} = 1$ and $i_2 \leq i$.

	i_1	i_2	i	i_2	i_1	i
j	1	1	1	1	1	1
j_1	1	$\boxed{1}$	$\boxed{1}$		1	$\boxed{1}$
j_2		1	$\boxed{1}$	1	$\boxed{1}$	$\boxed{1}$

If there exists a set in $\widehat{C}^- \cap \widehat{C}$ covering i the lemma follows, so suppose otherwise. If $j \in \widehat{C}^-$ ($j \in \widehat{C}$) then there exists $j_2 \in \widehat{C}$ ($j_1 \in \widehat{C}^-$) that dominates j in y^- (y), and again the lemma holds. Finally, consider the case $j \notin \widehat{C}^-$ and $j \notin \widehat{C}$. Assume $j_1 < j_2$, the other case is symmetrical. Because A is in standard greedy form we get that $a_{i_2 j_1} = 1$ and $(j_2, j_1) \in E$ if $i_1 < i_2$, or $a_{i_1 j_2} = 1$ and $(j_1, j_2) \in E$ if $i_2 < i_1$. In either case the lemma follows. The diagram on the right shows a summary of the entries that were inferred using the fact that A does not contain (4.1) as an induced submatrix. ■

Lemma 4.6. *Let (j, D) be the input of INCREASE/DECREASE. Then at the beginning of each call we have $j' \in D$ or $j'' \in D$ for all $(j', j'') \in E$. Furthermore, if $j' \in D$ and $j'' \in D$ then j' or j'' must have been split in a previous call.*

Proof. We prove the lemma by induction on the number of calls to INCREASE/DECREASE. Clearly, the lemma is true on the first call when $|\{j', j''\} \cap D| = 1$ for all $(j', j'') \in E$ and no vertex has been split yet.

Suppose the lemma holds at the beginning of a call to INCREASE/DECREASE, we argue that it also holds at the beginning of the next call. If the next call happens in Line 5 then the lemma holds trivially since D did change. Otherwise, j is split; namely, j is added to D (Line 6) and D is updated by taking $D \oplus T_c$ (Line 9) for a number of children c of j . Before splitting j , for every child c of j and every $(j', j'') \in T_c$ we have $|\{j', j''\} \cap D| = 1$, and the same holds after j is split. Hence, the lemma follows. ■

Lemma 4.7. *Let (j, D) be the input of INCREASE/DECREASE. For INCREASE we always have $p(D) \leq P < p(D) + \Delta(T_j, D)$, and for DECREASE we have $p(D) \geq P > p(D) + \Delta(T_j, D)$.*

Proof. By induction on the number of calls to INCREASE/DECREASE. At the base case the call is made by MERGE and thus we have $p(D) \leq P < p(D \oplus T_j)$. We claim that

$$p(D \oplus T_j) - p(D) = \Delta(T_j, D), \quad (4.7)$$

from which the lemma follows. We argue that every $i \in U$ contributes the same amount to each side of (4.7). From now on we assume that i is not covered by $D \setminus T_j$, otherwise its contribution to (4.7) is zero.

Suppose i is uniquely covered by some set $j \in T_j$. Then its contribution to both sides of (4.7) is either p_i or $-p_i$ depending on whether $i \in D \oplus T_j$ or $i \in D$.

Now consider the case when i is covered by multiple sets in T_j . Recall that $|\{j', j''\} \cap D| = 1$ for all $(j', j'') \in T_j$ because $D \cap T_j = \widehat{C}^- \cap T_j$. If $y_i < \lambda p_i$ then by Lemma 4.5 is covered by both D and $D \oplus T_j$. If $y_i = \lambda p_i$ then i covered by at most one set in \widehat{C}^- and at most one set in \widehat{C} . It follows that i must be covered exactly by one set in D and another set in $D \oplus T_j$. Hence, the contribution of i to both sides of (4.7) is zero.

For the inductive step suppose the lemma holds at the beginning of this call. Clearly, if the next call is made in Line 3 the lemma holds. Suppose that the call is made in Lines 11-12. Note that after adding j to D (Line 7) we have for every child c of j

$$p(D \oplus T_c) - p(D) = \Delta(T_c, D).$$

Therefore, by inductive hypothesis, we are bound to exit the while loop and the lemma holds in the next call. ■

Recall that y is also a feasible solution for the dual relaxation of P-TBC and its cost is given by $DL = \sum_{i=1}^n y_i - (p(U) - P)\lambda$. The following lemma proves the upper bound of Theorem 4.1.

Lemma 4.8. *Suppose MERGE outputs D . Then $c(D) \leq (1 + \frac{1}{3^{k-1}})DL + k c_{\max}$ for all $k \in \mathbb{Z}_+$.*

Proof. Let us digress for a moment for the sake of exposition. Suppose that in Line 6 of MERGE, instead of calling INCREASE, we return $D' = D \oplus T_r$. Notice every arc in the merger graph has exactly one endpoint in D' . By Lemma 4.5, any element i not covered by D' must have $y_i = \lambda p_i$. Furthermore, if $y_i > 0$ then there exists at most one set in D' that covers i ; if two such sets exist, one must dominate the other in y and y^- , which is not possible. Hence,

$$c(D) = \sum_{j \in D'} \sum_{i: a_{ij}=1} y_i = \sum_{\substack{i \text{ s.t.} \\ \exists j \in D': a_{ij}=1}} y_i = \sum_{i \in U} y_i - (p(U) - p(D'))\lambda \leq DL + (p(D') - P)\lambda \quad (4.8)$$

In the fortunate case that $(p(D') - P)\lambda \leq k c_{\max}$, the lemma would follow. Of course, this need not happen and this is why we make the call to INCREASE instead of returning D' .

Let j_q be the root of the q^{th} subtree split by INCREASE/DECREASE. Also let D_q the solution right before splitting T_{j_q} , and D'_q and D''_q be the unfeasible/feasible pair of solutions

after the splitting, which are used as parameters in the recursive calls (Lines 11-12). Suppose Lines 7-9 processed only one child of j_q , this can only happen in INCREASE, in which case $p(D''_q) > P$ but $p(D''_q) - b_{j_q} < P$. The same argument used to derive (4.8) gives us

$$c(D''_q \setminus \{j_{\leq q}\}) \leq \sum_{i \in U} y_i - (p(U) - p(D''_q) + b_{j_q}) \lambda \leq \text{DL} \quad (4.9)$$

The cost of the missing sets is $c(\{j_{\leq q}\}) \leq q c_{\max}$, thus if $q \leq k$ the lemma follows. A similar bound can be derived if the recursive call ends in Line 3 before splitting the k^{th} subtree. Finally, the last case to consider is when Lines 7-9 process two or more children j_q for all $q \leq k$. In this case

$$|p(D_q) - P| \geq 3 \min \{|p(D'_q) - P|, |p(D''_q) - P|\} = 3 |p(D_{q+1}) - P|, \quad (4.10)$$

which implies $|p(D_1) - P| \geq 3^{k-1} |p(D_k) - P| \geq 3^{k-1} |p(D''_k) - P|$. Also, $\lambda(P - p(D_1)) \leq \text{DL}$ since all elements i not covered by D_1 must be such that $y_i = \lambda p_i$. Hence, as before

$$c(D''_k \setminus \{j_{\leq k}\}) \leq \text{DL} + \lambda (p(D''_k) - P) \leq \text{DL} + \lambda \frac{P - p(D_1)}{3^{k-1}} \leq \left(1 + \frac{1}{3^{k-1}}\right) \text{DL} \quad (4.11)$$

Adding the cost of $\{j_{\leq k}\}$ we get the lemma. \blacksquare

4.3. Integrality gap example

To finish the proof of Theorem 4.1 we now show a family of instances of P-TBC exhibiting an integrality gap of $\text{IP} > \left(1 + \frac{1}{3^{k-1}}\right) \text{LP} + \frac{k}{2} c_{\max}$ for large enough k

Let T be a rooted tree and $\{(s_i, t_i)\}_{i=1}^n$ be a collection of pairs of nodes of T , each defining a unique path in T . Let A be the incidence matrix of paths to edges of T . The covering problem defined by A is the well-know Multicut problem where the objective is to find a minimum cost set of edges whose removal separates all pairs. If for every (s_i, t_i) pair s_i is the ancestor of t_i or vice-versa then A is totally balanced.

Our tree T is made up of a complete binary tree with height $2q$ plus a 2-path coming out of the apex of the binary tree going up into the real root of T ; thus the tree has $2^{2q+1} + 1$ nodes. The reader is referred Figure 1 for a picture of the instance. The cost of every edge is 3. There are two kinds of paths: *internal* and *fringe* paths. For every node in the binary tree there is an internal path of length two coming out of the node going up; there are $2^{2q+1} - 1$ such paths each having a profit of 4^q . For each leaf and the root there is a fringe path of length 1 incident on it; there are $2^{2q} + 1$ such paths each having a profit of 2. The target coverage is given by $\bar{P} = 2(4^{q-1} + \dots + 4^0 + 1) = \frac{2^{2q+1} + 4}{3}$, where the shorthand notation \bar{X} stands for $p(U) - X$.

Consider the dual solution y where every internal path gets a dual value of 1 and every fringe path gets a dual value of 2. The solution is feasible for $\lambda = 1$ and has cost

$$\text{DL} = \sum_{i \in U} y_i - \bar{P} \lambda = \frac{10 \cdot 4^q - 1}{3} \quad (4.12)$$

To show that y is optimal, we construct a primal (fractional) solution x with the same cost, which is a convex combination of two integral solutions \tilde{x}_1 and \tilde{x}_2 . Let \tilde{x}_1 consist of edges in every other level of T starting at the leaf level and let \tilde{x}_2 be the complement of \tilde{x}_1 .

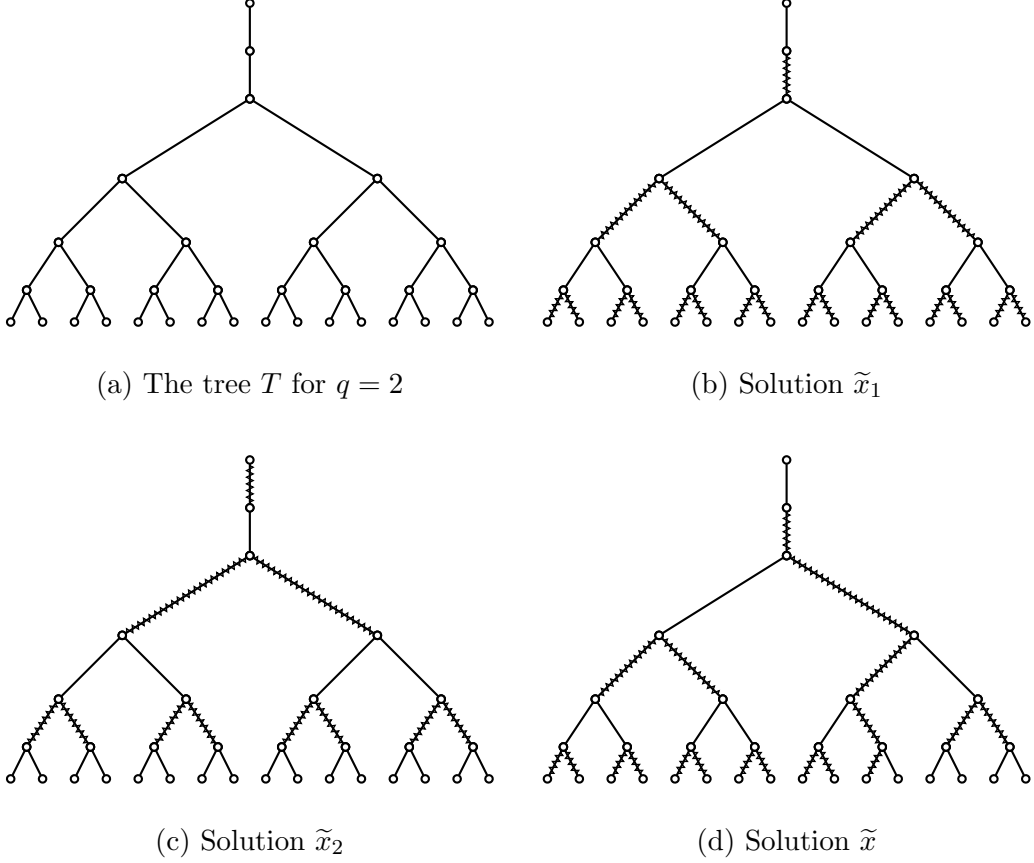


Figure 1: Integrality gap example for $q = 2$. In (b-d) the wiggly edges belong to the corresponding solution

(See Figure 1. Note that $\overline{p(\tilde{x}_1)} = 2$ and $\overline{p(\tilde{x}_2)} = 2^{2q+1}$. Consider the convex combination $\alpha\overline{p(\tilde{x}_1)} + \beta\overline{p(\tilde{x}_2)} = \overline{P}$ and let $x = \alpha\tilde{x}_1 + \beta\tilde{x}_2$. Its cost is given by

$$\begin{aligned}
 c(x) &= \alpha c(\tilde{x}_1) + \beta c(\tilde{x}_2) \\
 &= \alpha \left(\sum_i y_i - \lambda \overline{p(\tilde{x}_1)} \right) + \beta \left(\sum_i y_i - \lambda \overline{p(\tilde{x}_2)} \right) - \lambda \overline{P} \\
 &= \text{DL}.
 \end{aligned}$$

Clearly, x is a feasible fractional solution. Therefore, it is optimal.

Let \tilde{x} be the solution defined as follows. For edges incident on a leaf and leave out the $\frac{\overline{P}}{2} - 1$ rightmost ones and choose remaining ones. For other edges, choose the edge only if one of the edges immediately below is not chosen. (See Figure 1). If we try to pay for \tilde{x} using the dual cost we will charge twice $2q - 1$ internal paths whose both edges are chosen in \tilde{x} . In other words,

$$c(\tilde{x}) = \text{DL} + 2q - 1 \tag{4.13}$$

Due to their high profit, internal paths cannot be left uncovered. Using this fact we can infer that \tilde{x} is indeed an optimal integral solution covering P profit. Choosing $k = q \log_3 4$ we get the lower bound of Theorem 4.1. That is, for large enough q ,

$$c(\tilde{x}) > \left(1 + \frac{1}{3^{k-1}}\right) c(x) + c_{\max} \frac{k}{2}. \quad (4.14)$$

For smaller values of k the slightly weaker bound with $c_{\max} \frac{k-5}{2}$ additive error holds. It is worth noting that the example can be adapted to yield the same bound for instances with unit profits.

5. Applications

In this section we show how Theorem 4.1 implies better approximation algorithms for a number of covering problems that can be expressed with a suitable combination of ρ totally-balanced matrices.

Definition 5.1. Matrix B is said to be row-induced by a collection of matrices $A_1, \dots, A_k \in R^{n \times m}$ if for all i , the i th row of B equals the i th row of A_j for some $1 \leq j \leq k$.

Definition 5.2. Matrix $A \in \{0, 1\}^{n \times m}$ is said to be ρ -separable if there exist matrices $A_1, \dots, A_\rho \in \{0, 1\}^{n \times m}$ such that $A = \sum_q A_q$ and every matrix row-induced by A_1, \dots, A_ρ is totally balanced.

Our algorithms make us of following lemma to absorb the additive error in our bounds.

Lemma 5.3. Let \mathcal{A} be an algorithm for a given partial covering problem (U, \mathcal{S}, P) that produces a solution with cost at most $\alpha \text{OPT} + k c_{\max}$, where OPT is the cost of the optimal solution. Then there exists an α -approximation that makes $|U|^{\frac{k}{\alpha-1}}$ calls to \mathcal{A} .

Proof. The idea is to run \mathcal{A} on a modified instance (U', \mathcal{S}', P') . Let X be the $\frac{k}{\alpha-1}$ most expensive sets in the optimal cover for (U, \mathcal{S}, P) . Let $\mathcal{S}' = \mathcal{S} \setminus \{j \mid c_j > \min_{j' \in X} c_{j'}\}$, $U' = U \setminus \{i \mid \text{covered by } X\}$, and $P' = P - p(X)$. The optimal solution in the new instance has cost $\text{OPT}' = \text{OPT} - c(X)$.

Adding X to the solution returned by $\mathcal{A}(U', \mathcal{S}', P')$ gives us a feasible solution, for the original instance, with cost at most

$$\alpha \text{OPT}' + k c'_{\max} + c(X) \leq \alpha \text{OPT}' + k \frac{\alpha-1}{k} c(X) + c(X) = \alpha \text{OPT}.$$

Unfortunately we do not know which sets comprise X . Therefore \mathcal{A} is run on every choice of X and the best cover found is returned. The number of calls to \mathcal{A} needed is $|U|^{\frac{k}{\alpha-1}}$. \blacksquare

We are ready to describe our approximation algorithms for covering problems that can be described with a ρ -separable matrix. We assume the decomposition is given to us. For an arbitrary matrix finding such a decomposition, or even testing for its existence, may be hard. However, for our application problems it is easy to find the ρ matrices using the problem definition.

Theorem 5.4. Let A be ρ -separable into matrices A_1, \dots, A_ρ where $\rho > 1$. For any constant $\epsilon > 0$ there is a $(\rho + \epsilon)$ -approximation and a quasi-polynomial time ρ -approximation for the partial covering problem defined by A .

Proof. Our algorithm is based on the approach of [17, 19]. First, we find an optimal fractional solution (x, r) for the partial covering problem defined by A . Let a_i^q be the i th row of A_q . Notice that for each i there must exist a q_i such that $a_i^{q_i} \cdot x \geq \frac{1-r_i}{\rho}$. Second, we construct a matrix B by choosing $a_i^{q_i}$ as the i th row of B . Note that B is totally balanced. Finally, we find a threshold value λ^* for B and invoke MERGE to find a cover D .

Any feasible solution for B is also feasible for A , thus D is a feasible cover for A . Note that $(\rho x, r)$ is a feasible fractional solution for B . Let OPT be the cost optimal solution for A . By Lemma 4.8 and letting $k \geq \log_3 \frac{\rho}{\epsilon} + 1$ we get,

$$c(D) \leq \left(1 + \frac{1}{3^{k-1}}\right) c(\rho x) + k c_{\max} = \left(1 + \frac{1}{3^{k-1}}\right) \rho c(x) + k c_{\max} \leq (\rho + \epsilon) \text{OPT} + k c_{\max} \quad (5.1)$$

This solution is $\rho + \epsilon$ approximate with an additive error of $k c_{\max}$ that can be absorbed using Lemma 5.3

For the quasi-polynomial time ρ -approximation, setting $k \geq \log \rho |U| + 1$ we get

$$c(D) \leq \left(1 + \frac{1}{\rho^{|U|}}\right) c(\rho x) + k c_{\max} \leq \rho c(x) + (k + 1) c_{\max} \leq \rho \text{OPT} + (k + 1) c_{\max}. \quad (5.2)$$

And the theorem follows. \blacksquare

This implies improved approximation algorithm for the partial version of Multicut and Path Hitting on Trees, and Rectangle Stabbing. To show this, we use the following fact about totally balanced matrices. Let T be a rooted tree. An s - t path in T is said to be *descending* if s is an ancestor of t . Let \mathcal{P} and \mathcal{Q} be collections of descending paths in T , and let A be the \mathcal{P} - \mathcal{Q} incidence matrix $A = \{a_{i,j}\}$, where $a_{i,j} = 1$ if and only if the i th path in \mathcal{P} intersects the j th path in \mathcal{Q} . It is known that A is totally balanced: To put the matrix into Greedy Standard form arrange the columns and rows of A so that the paths in \mathcal{P} and \mathcal{Q} appear in non-increasing distance from the root.

Corollary 5.5. *There is a $2 + \epsilon$ approximation and a quasi-polynomial time 2-approximation for Partial Multicut on Trees.*

Proof. The input of Partial Multicut is a tree T and a collection of paths \mathcal{P} in T , the problem is defined by the \mathcal{P} - $E[T]$ incidence matrix A . Even though a path in \mathcal{P} may not be descending, we can always split such a path into two separate descending paths. Therefore, A is 2-separable. \blacksquare

Corollary 5.6. *There is a $4 + \epsilon$ approximation and a quasi-polynomial time 4-approximation for Path Hitting on Trees.*

Proof. The input of Partial Path Hitting is a tree T and two collections of paths \mathcal{P} and \mathcal{Q} in T , the covering problem is defined by the \mathcal{P} - \mathcal{Q} incidence matrix A . Parekh and Segev [35] noted that if we split each path in \mathcal{Q} into two descending paths the cost of the optimal solution increases by at most a factor of 2. The matrix of this modified problem is 2-separable. \blacksquare

Our last application problem is Rectangle Stabbing. This is a special case of Set Cover with ρ -Blocks, a broad class of covering problems introduced by Hassin and Segev [19], where the incidence matrix defining the problem is such that every row has ρ blocks of contiguous 1's.

Theorem 5.7. *Let A be a matrix defining an instance of Set Cover with ρ -Blocks. Then there exists a polynomial time ρ -approximation algorithm for the partial covering problem defined by A .*

Proof. We proceed as in Theorem 5.4 to reduce A to a matrix B . This new matrix is not only totally balanced, but each row consists of a single block of consecutive 1's. For such matrices the merger graph used in Section 4.2 is in fact a path. In this case, at most one subtree is split in the execution of INCREASE. Therefore we get the stronger guarantee that $IP \leq LP + c_{\max}$. Plugging in this into the proof of Theorem 5.4 gives the desired result. ■

Corollary 5.8. *There is a 2-approximation for Partial Rectangle Stabbing and a d -approximation for Partial d -dimensional Rectangle Stabbing.*

6. Concluding remarks and open problems

The results in this paper suggest that Lagrangian relaxation is a powerful technique for designing approximation algorithms for partial covering problems, even though the black-box approach may not be able to fully realize its potential.

It would be interesting to extend this study on the strengths and limitation of Lagrangian relaxation to other problems. The obvious candidate is the k -Median problem. Jain and Vazirani [26] designed a 2α -approximation for k -Median using as a black box an α -LMP approximation for Facility Location. Later, Jain et al. [27] gave a 2-LMP approximation for Facility Location. Is the algorithm in [26] optimal in the sense of Theorem 3.1? Can the algorithm in [27] be turned into a 2-approximation for k -Median by exploiting structural similarities when combining the two solutions?

The class of totally unimodular matrices is undoubtedly the most important subclass of balanced matrices. An open problem is to establish good approximations for Partial Totally Unimodular Cover (P-TUC). The matrix used in Section 4.3 is also totally unimodular, so the lower bound on the integrality gap applies for P-TUC as well. Does the upper bound of Theorem 4.1 also hold P-TUC?

Acknowledgments: I am indebted to Danny Segev for sharing an early draft of [31] and for pointing out Kolen's work. Also thanks to Mohit Singh and Arie Tamir for helpful discussions and to Elena Zotenko for suggesting deriving the result of Section 3.

References

- [1] A. Archer, R. Rajagopalan, and D. B. Shmoys. Lagrangian relaxation for the k -median problem. In *Proceedings of the 11th Annual European Symposium on Algorithms (ESA)*, pages 31–42, 2003.
- [2] S. Arora and G. Karakostas. A $2 + \epsilon$ approximation algorithm for the k -MST problem. In *Proceedings of the 11th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 754–759, 2000.
- [3] R. Bar-Yehuda and S. Even. A linear time approximation algorithm for approximating the weighted vertex cover. *Journal of Algorithms*, 2:198–203, 1981.
- [4] R. Bar-Yehuda and S. Even. A local-ratio theorem for approximating the weighted vertex cover problem. *Annals of Discrete Mathematics*, 25:27–46, 1985.
- [5] C. Berge. Balanced matrices. *Mathematical Programming*, 2:19–31, 1972.
- [6] M. Charikar and S. Guha. Improved combinatorial algorithms for the facility location and k -median problems. In *Proceedings of the 40th Annual IEEE Symposium on Foundations of Computer Science (FOCS'99)*, pages 378–388, 1999.

- [7] M. Charikar, S. Khuller, D. M. Mount, and G. Narasimhan. Algorithms for facility location problems with outliers. In *Proceedings of the 12th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'01)*, pages 642–651, 2001.
- [8] F. A. Chudak, T. Roughgarden, and D. P. Williamson. Approximate k-MSTs and k-Steiner trees via the primal-dual method and Lagrangean relaxation. In *Proceedings of the 9th Integer Programming and Combinatorial Optimization Conference (IPCO)*, pages 60–70, 2001.
- [9] K. L. Clarkson. A modification of the greedy algorithm for vertex cover. *Information Processing Letters*, 16(1):23–25, 1983.
- [10] U. Feige. A threshold of $\ln n$ for approximating set cover. *J. of the ACM*, 45(4):634–652, 1998.
- [11] R. Gandhi, S. Khuller, and A. Srinivasan. Approximation algorithms for partial covering problems. *J. Algorithms*, 53(1):55–84, 2004.
- [12] N. Garg. Saving an epsilon: a 2-approximation for the k-MST problem in graphs. In *Proceedings of the 37th annual ACM symposium on Theory of computing (STOC)*, pages 396–402, 2005.
- [13] N. Garg. A 3-approximation for the minimum tree spanning k vertices. In *Proceedings of the 37th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 302–309, 1996.
- [14] 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.
- [15] 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.
- [16] A. Ghouilla-Houri. Caractérisations des matrices totalment unimodulaires. *C. R. Acad. Sc. Paris*, 254:1192–1193, 1962.
- [17] D. Golovin, V. Nagarajan, and M. Sing. Approximating the k-multicut problem. In *7th Annual ACM-SIAM Symposium on Discrete Algorithms*, 2006.
- [18] E. Halperin. Improved approximation algorithms for the vertex cover problem in graphs and hypergraphs. *SIAM Journal on Computing*, 31(5):1608–1623, 2002.
- [19] R. Hassin and D. Segev. Rounding to an integral program. In *Proceedings of the 4th International Workshop on Efficient and Experimental Algorithms (WEA'05)*, pages 44–54, 2005.
- [20] M. Held and R. Karp. The traveling salesman problem and minimum spanning trees. *Operations Research*, 18:1138–1162, 1970.
- [21] M. Held and R. Karp. The traveling salesman problem and minimum spanning trees: Part II. *Mathematical Programming*, 1:6–25, 1971.
- [22] D. S. Hochbaum. Approximation algorithms for the set covering and vertex cover problems. *SIAM Journal on Computing*, 11(3):555–556, 1982.
- [23] 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 Problems (APPROX)*, pages 111–122, 1998.
- [24] D. S. Hochbaum, editor. *Approximation Algorithms for NP-hard Problems*. PWS Publishing Company, 1997.
- [25] A. J. Hoffman, A. Kolen, and M. Sakarovitch. Totally-balanced and greedy matrices. *SIAM Journal on Algebraic and Discrete Methods*, 6:721–730, 1985.
- [26] K. Jain and V. V. Vazirani. Approximation algorithms for metric facility location and k-median problems using the primal-dual schema and Lagrangian relaxation. *Journal of the ACM*, 48(2): 274–296, 2001.
- [27] K. Jain, M. Mahdian, E. Markakis, A. Saberi, and V. V. Vazirani. Greedy facility location algorithms analyzed using dual fitting with factor-revealing LP. *Journal of the ACM*, 50(6): 795–824, 2003.
- [28] G. Karakostas. A better approximation ratio for the vertex cover problem. In *Proceedings of the 15th International Colloquium on Automata, Languages, and Programming (ICALP'05)*, pages 1043–1050, 2005.

- [29] R. M. Karp. Reducibility among combinatorial problems. In *Complexity of Computer Computations*, pages 85–103. Plenum Press, 1972.
- [30] A. Kolen. *Location problems on trees and in the rectilinear plane*. PhD thesis, Mathematisch Centrum, Amsterdam, 1982.
- [31] J. Könemann, O. Parekh, and D. Segev. A unified approach to approximating partial covering problems. In *Proceedings of the 14th Annual European Symposium on Algorithms (ESA)*, pages 468–479, 2006.
- [32] J. Könemann and R. Ravi. A matter of degree: Improved approximation algorithms for degree-bounded minimum spanning trees. *SIAM Journal of Computing*, 31:1783–1793, 2002.
- [33] A. Levin and D. Segev. Partial multicuts in trees. In *3rd International Workshop on Approximation and Online Algorithms (WAOA)*, 2005.
- [34] N. Megiddo. Combinatorial optimization with rational objective functions. In *Proceedings of the tenth annual ACM symposium on Theory of computing (STOC)*, pages 1–12, 1978.
- [35] O. Parekh and D. Segev. Path hitting in acyclic graphs. In *Proceedings of the 14th Annual European Symposium on Algorithms (ESA)*, pages 564–575, 2006.
- [36] R. Ravi and M. X. Goemans. The constrained minimum spanning tree problem (extended abstract). In *5th Scandinavian Workshop on Algorithm Theory (SWAT)*, pages 66–75, 1996.
- [37] J. P. Spinard. Doubly lexical ordering of dense 0-1 matrices. *Information Processing Letters*, 45: 229–235, 1993.

Appendix A. Partial Totally Unimodular Cover

Theorem A.1. *Partial totally unimodular cover cannot be approximated better than $\frac{4}{3}$ using Lagrangian relaxation and a 1-LMP algorithm \mathcal{A} as a black box.*

Proof. The instance is similar to that used in Theorem 3.1: The A -sets and the B -sets are the same; for each i we define O_i as B_i minus the rightmost element. The cost of each A , B and O set is $\frac{2}{3}$, $\frac{4}{3}$ and 1 respectively. The target coverage parameter is again $P = q(q^2 + 1)$.

It is straightforward to check that Lemmas 3.2 and 3.3 still holds for our new instance and $\alpha = 1$. Then the same argument used in the proof of Theorem 3.1 gives us a lower bound of $\frac{4}{3}$.

It only remains to show that the resulting element-set incidence matrix A is totally unimodular. A matrix A is totally unimodular if and only if every submatrix A' of A has an equitable coloring [16]. An equitable coloring of a 0,1 matrix A' is a partition of its columns into red and blue columns such that in every row of A' the number of blue 1's and red 1's differs by at most one. Let us construct an equitable coloring for A' : all the A -sets are colored blue; for each i , if B_i and O_i are present in A' then color one red and the other blue, and if only one is present then color it red. Clearly the coloring is equitable; thus, A is totally unimodular. ■

Appendix B. Finding a threshold value

The idea is to use parametric search treating λ as an unknown which lies in a certain range (λ_l, λ_r) . Initially $\lambda_l = 0$ and $\lambda_r = \max_{i,j} \frac{c_j}{p_i}$. Residual capacities and dual variables are kept as a linear function of λ . We maintain the invariant that λ_l^+ separates less than P profit and λ_r^- covers more than P profit. Suppose that in the interval (λ_l, λ_r) the algorithm agrees on the first i elements. By this we mean that if we run the algorithm with any value $\lambda \in (\lambda_l, \lambda_r)$ the value of the dual variables of these i elements (as a function of λ) is always

the same. In each iteration we either find a threshold value or we narrow the interval such that the algorithm agrees on one more element, while maintaining the invariant. This cannot go on forever because the algorithm will eventually behave the same throughout the interval and the invariant would be violated. If at some point along the way we find a value of λ covering exactly P profit we stop as the solution is optimal. For simplicity, from now on we assume that this never happens.




Figure 2: Narrowing the interval for λ .

Suppose that KOLEN agrees on the first $i - 1$ elements in the interval (λ_l, λ_r) . Note that the residual costs and i 's penalty are linear functions of λ . As a result, which set has the minimum residual cost, and thus which one becomes tight, if any, varies with λ . Our goal is to narrow the interval such that the set that becomes tight is always the same, or $y_i = p_i \lambda$ within the new interval. If we draw the lines corresponding to the residual costs of set that i belongs to and $p_i \lambda$, the segments on the lower envelope correspond to the next tight event, either a set or element i . See Figure 2. Let $\lambda_1, \dots, \lambda_s$ correspond to the intersection points of the lower envelope, and let $\lambda_0 = \lambda_l$ and $\lambda_{s+1} = \lambda_r$. For every $0 \leq a \leq s$, within the interval $(\lambda_a, \lambda_{a+1})$ the algorithm agrees on i . Note that either one of the λ_a is a threshold value, or there exists an a such that λ_a^+ covers less than P profit and λ_{a+1}^- covers more than P profit. Given the latter we update $\lambda_l = \lambda_a$ and $\lambda_r = \lambda_{a+1}$ and repeat.

Theorem B.1. *A threshold value can be found by making $O(|U| \log |\mathcal{S}|)$ calls to KOLEN.*

Proof. The above discussion outlines the algorithm for finding a threshold value. Regarding the time complexity, when searching for the next tight event, instead of trying every λ_a value, we can use binary search. Thus, only $\log |\mathcal{S}| + 1$ calls to KOLEN are needed to find the right value of a to narrow the interval. ■