Tight Approximation Results for General Covering Integer Programs

Stavros G. Kolliopoulos[1]
McMaster University
Department of Computing and Software
Hamilton, ON, Canada
stavros@mcmaster.ca

Neal E. Young
Akamai Technologies
Cambridge, MA
neal@acm.org

Abstract

In this paper we study approximation algorithms for solving a general covering integer program. *An n-vector x of nonnegative integers is sought, which minimizes $c^T \cdot x$, subject to $Ax \geq b, x \leq d$. The entries of A, b, c are nonnegative. Let m be the number of rows of A. Covering problems have been heavily studied in combinatorial optimization. We focus on the effect of the* multiplicity constraints, $x \leq d$, *on approximability. Two longstanding open questions remain for this general formulation with upper bounds on the variables.*

(i) *The integrality gap of the standard LP relaxation is arbitrarily large. Existing approximation algorithms that achieve the well-known $O(\log m)$-approximation with respect to the LP value do so at the expense of violating the upper bounds on the variables by the same $O(\log m)$ multiplicative factor. What is the smallest possible violation of the upper bounds that still achieves cost within $O(\log m)$ of the standard LP optimum?*

(ii) *The best known approximation ratio for the problem has been $O(\log(\max_j \sum_i A_{ij}))$ since 1982. This bound can be as bad as polynomial in the input size. Is an $O(\log m)$-approximation, like the one known for the special case of Set Cover, possible?*

We settle these two open questions. To answer the first question we give an algorithm based on the relatively simple new idea of randomly rounding variables to smaller-than-integer units. To settle the second question we give a reduction from approximating the problem while respecting multiplicity constraints to approximating the problem with a bounded violation of the multiplicity constraints.

[1]Research partially supported by NSERC Grant 227809-00 and a CFI New Opportunities Award

1. Introduction

In this paper we examine approximation algorithms for the general formulation of a *covering integer program (CIP)*. We also study the integrality gaps of two related linear relaxations. Here is a formal definition of CIPs.

Definition 1.1 *Given $A \in R_+^{m \times n}$, $b \in R_+^m$, $c \in R_+^n$ and $d \in R_+^n$, a CIP $\mathcal{P} = (A, b, c, d)$ seeks to minimize $c^T \cdot x$ subject to $Ax \geq b$, $x \in Z_+^n$, and $x \leq d$. If $A \in \{0,1\}^{m \times n}$, each entry of b is assumed integral and the CIP is called* $(0,1)$.

The *dilation* α of a CIP is the maximum number of constraints any variable appears in.

Definition 1.2 *Given a CIP $\mathcal{P} = (A, b, c, d)$, the* standard linear relaxation *of $\mathcal{P}$, seeks to minimize $c^T \cdot x$ subject to $A'x \geq b$, $x \leq d$ and x nonnegative, where $A'_{ij} = \min\{A_{ij}, b_i\}$. The* width W *of the standard linear relaxation is defined as* $\min_{i,j \mid A'_{ij} \neq 0} b_i / A'_{ij}$.

Covering integer programs form a large subclass of integer programs (IPs) encompassing such classical NP-hard problems as Minimum Knapsack and Set Cover. Set Cover is a $(0,1)$-CIP with $b^T \in \{1\}^m$. The n columns of the matrix A correspond to the sets and the m rows to the elements to be covered. In Set Cover the upper bound of 1 on all the variables is implicit in the minimization of the objective, in that it never helps to set a variable above 1. However in a general-form CIP constraints of the type $x \leq d$, as given in Definition 1.1, have to be explicitly included. We call these inequalities, which disallow arbitrarily large variable values, *multiplicity constraints.* They express a natural resource limitation: a restricted number of copies is available of each covering object, thus imposing an upper bound on the multiplicity of the latter in the final cover. Consider for instance the natural generalization of Set Cover where the i-th element needs to be covered $b_i > 1$ times. Setting d_j to a vector of ones means that each set can be chosen only once.

Our paper focuses on the effect of multiplicity constraints on approximability.

Background. Let CIP_∞ denote the problem of solving a covering integer program without multiplicity constraints or equivalently with trivially large upper bounds on the variables. A CIP_∞ instance is hence defined by a triple $\mathcal{P} = (A, b, c)$. There is a long line of research investigating approximation algorithms for CIP_∞. Most of this work uses the value of the linear relaxation LP_∞ as a lower bound on the optimum, even when a fractional solution is not explicitly computed as is the case with the greedy algorithm for Set Cover [10, 3]. In other work the error is analyzed directly with respect to some estimate on the integral optimum [8, 4, 6]. Among the most recent work some relevant references are [15, 12, 22, 20, 19]. The reader is referred to the survey in [7] for a discussion of the extensive literature on covering problems. It is well known that the integrality gap of LP_∞ is $\Theta(1 + \log(m)/W)$. Moreover Raz and Safra [16] showed that it is NP-hard to obtain an $o(\ln m)$ approximation algorithm. Srinivasan's work [20, 19] yields the currently best known approximation results for CIP_∞ with existential improvements on the $\log m$ factor.

When the formulation contains multiplicity constraints there are two natural avenues of investigation: (1) find an approximation with respect to the optimum y_* of the standard LP relaxation and (2) find an approximation with respect to some other estimate of the integral optimum. We describe what is known from both perspectives.

Let ρ, l be scalars greater than or equal to 1. Define as a (ρ, l)-*approximate solution w.r.t. the standard LP optimum,* an integral vector x that meets the covering constraints $Ax \geq b$ and has the following two properties: (i) $c^T x \leq \rho y_*$ and (ii) for all j, $x_j \leq \lceil l d_j \rceil$. In words, a (ρ, l)-approximate solution achieves a bicriteria approximation with respect to the cost and the violation of the upper bounds on the variables.

For $(0, 1)$-CIPs Rajagopalan and Vazirani [15] give an efficient algorithm to find an $(O(\log \alpha), 1)$-approximate solution. An algorithm by Srinivasan and Teo [21] yields an $(O(1 + \log(m)/W), 1 + \varepsilon)$-approximate solution for CIPs where each $c_j = 1$. Kolliopoulos studied in [9] *column-restricted CIPs (CCIPs)* where all non-zero entries of the j-th column of A have the same value ρ_j. The algorithm in [9] obtains an $(O(\log m), 12)$-approximate solution.

Simple as they appear, multiplicity constraints make covering problems much harder. The recent paper of Carr et al. [1] gives a simple instance of a Minimum Knapsack problem (trivially a CCIP), for which the integrality gap of the standard linear relaxation can be made arbitrarily large if multiplicity constraints are to be respected[1]. The LP below has an integrality gap of at least $M > 0$:

$$\begin{array}{lrl} \text{minimize} & x_2 & \\ & (M-1)x_1 + Mx_2 & \geq M \\ & 0 \leq x_1, x_2 \leq 1 & \end{array}$$

However if one sets the right hand side of the multiplicity constraints to 2, an integral solution of zero cost becomes possible. This example demonstrates that for any finite ρ a $(\rho, 1)$-approximate solution w.r.t. the standard LP optimum is impossible for general CIPs. A second negative result from [9] shows that the integrality gap for a $(0, 1)$-CIP is $\Omega(\log m)$ even when allowing arbitrarily large values for the variables. Hence for any l, a (ρ, l)-approximate solution with $\rho = o(\log m)$ is also impossible. The two negative results imply that finding an $(O(\log m), l)$-approximate solution with l as close to 1 as possible has a natural significance: by allowing a small increase on the number of copies available from each covering object, one is able to achieve a cost guarantee which would have otherwise been impossible to attain.

Some of the existing algorithms for the CIP_∞ problem without multiplicity constraints, such as standard randomized rounding [14] or Srinivasan's algorithms [20, 19], are easily seen to produce $(O(\log m), O(\log m))$-approximate solutions when applied to CIPs. This was the best known tradeoff between cost and violation of the multiplicity constraints prior to our work. The first basic question our paper addresses is the following:

Question 1: What is the smallest possible violation of the multiplicity constraints that still allows an integer solution of cost within $O(\log m)$ of the optimum of the standard LP relaxation?

Next we discuss what is known about approximations of the integer optimum using methods other than the standard LP relaxation. In what follows, a *ρ-approximation algorithm* means one that produces an integer solution meeting the multiplicity constraints and having cost at most ρ times the integer optimum of the CIP. Dobson in 1982 [4] gave an $H(\max_{1 \leq j \leq n} \sum_{1 \leq i \leq m} A_{ij})$-approximation algorithm, where $H(t)$ is the harmonic series with t terms. Recently, Carr, Fleischer, Leung and Phillips [1] gave a p-approximation algorithm, where p denotes the maximum number of variables in any constraint. Their algorithm is based on a new linear relaxation LP-KC that is stronger than the standard relaxation.

One would like to achieve for general CIPs the same approximation as for Set Cover. The second basic question we address is the following:

Question 2: Is there an $O(\log m)$-approximation algorithm for general CIPs?

[1]In [15] an $O(\log m)$ integrality gap was erroneously claimed for general CIPs.

Our results. In this paper we settle the two open questions above.

For Question 1, we improve the previously known logarithmic violation of the multiplicity constraints to constant. We give an algorithm that produces an $(O(1 + \log(m)/W), 1+\varepsilon)$-approximate solution w.r.t. the standard LP optimum for any $\varepsilon > 0$. (The constant in the order notation depends on $1/\varepsilon$.) For the case where $\max_j d_j$ is bounded by a constant, we obtain solutions that violate the multiplicity constraints by at most an *additive* 1. The $\log m$ in the performance guarantee can be strengthened to $\log \alpha$ with a more complicated analysis based on Srinivasan's results [19]. Recall α is the maximum number of constraints any variable appears in.

Observe that when finding an integral solution for a CIP we can replace every d_j by $\lfloor d_j \rfloor$. We chose to allow fractional d_j's in Definitions 1.1 and 1.2 to obtain a larger solution space for the standard LP relaxation and therefore prove a stronger result for the cost integrality gap.

The new idea underlying our algorithm is simple — we randomly round variables to *smaller-than-integer units* and then deterministically round the resulting values to integers.

For Question 2, we give an $O(1 + \log(m)/W)$-approximation algorithm. The algorithm works by reducing the problem of approximating a CIP (in which multiplicity constraints must be respected) to finding a (ρ, l)-approximate solution to an auxiliary covering problem, for appropriate ρ and l. The reduction uses the LP-KC relaxation of Carr et al. [1]. As in Question 1, the $\log m$ term can be improved to $\log \alpha$.

Preliminaries. We denote $\sum_j A_{ij} x_j$ by $(Ax)_i$ and $\max_j d_j$ by Δ. The symbol $\mathsf{vc}(t)$ denotes a vector whose coordinates are all equal to t. The dimension of such a vector will be clear from the context. We use $\lceil t \rceil_{1/k}$ to denote the smallest integer multiple of $1/k$ that is greater than or equal to t. We overload notation by denoting by $\lceil v \rceil_{1/k}$, v a vector, the vector obtained by applying the operation componentwise. Similarly for the floor operation. The well-known Chernoff bound [2] is central to our result. The form we will use was given by Raghavan [13].

Theorem 1.1 *[13] Let $X = \sum_{i=1}^{N} X_i$ be the sum of N independent random variables in $[0, U]$ with $E(X) \geq \mu$. Let $\varepsilon > 0$. Then $Pr[X \leq (1-\varepsilon)\mu] < exp(-\mu\varepsilon^2/(2U))$.*

2. The Rounding Argument

Let $\mathcal{P} = (A, b, c, d)$ be a given CIP and $\bar{x}$ a fractional optimal solution. Our goal is to show that an integer solution $\hat{x}$ with nice properties exists with nonzero probability. "Nice" refers to a solution of cost within $O(\log m)$ of $c^T \bar{x}$, and small violation of the multiplicity constraints. Based on Definition 1.2 we assume throughout Sections 2, 3 that $A_{ij} \leq b_i$ for all i, j.

We compute $\hat{x}$ in two steps. First randomly round $\bar{x}$ to a vector whose coordinates are integer multiples of some integer parameter $k \geq 1$. Next, round up each coordinate of the resulting vector to the closest integer. The details of the rounding experiment follow. It takes two parameters: k and $\varepsilon > 0$.

ALGORITHM FINELY_ROUND $(\mathcal{P}, \bar{x}, k, \varepsilon)$
STEP 0: Set $x' := \bar{x}/(1-\varepsilon)$.
STEP 1: Randomly round each coordinate of x' to its upper or lower multiple of $1/k$: set $x'' := \lfloor x' \rfloor_{1/k} + r/k$, where the j-th entry of random vector r takes the value 1 with probability $k(x'_j - \lfloor x'_j \rfloor_{1/k})$ and the value 0 otherwise. Note that $E(x'') = \bar{x}/(1-\varepsilon)$.
STEP 2: Return $\hat{x} := \lceil x'' \rceil$.

The ensuing lemma analyzes Step 1 of the rounding experiment.

Lemma 2.1 *For any $\delta, \varepsilon > 0$ the probability that x'' fails to satisfy*

$$c^T x'' \leq \delta c^T \bar{x}/(1-\varepsilon), \quad Ax'' \geq b, \quad x'' \leq \lceil \bar{x}/(1-\varepsilon) \rceil$$

is less than $1/\delta + m\ exp(-\varepsilon^2 kW/2)$.

Proof. Think of the fixed part of each x''_j, that is, the $\lfloor x'_j \rfloor_{1/k}$ part, as a sum of independent random variables, each in $[0, 1/k]$, happening to take the value $1/k$ with probability 1. Then for all i the quantity $(Ax'')_i/b_i$ is a sum of independent random variables with value in $[0, 1/(kW)]$. Observe that by Step 1, $E[x''_j] = \bar{x}_j/(1-\varepsilon)$. Accordingly $E[(Ax'')_i/b_i] \geq 1/(1-\varepsilon)$. By the Chernoff bound of Theorem 1.1

$$Pr[(Ax'')_i/b_i \leq 1] < exp(-\varepsilon^2 kW/2).$$

By the Markov inequality,

$$Pr[c^T x'' \geq \delta c^T \bar{x}/(1-\varepsilon)] \leq 1/\delta.$$

By the union bound, the probability that x'' fails to satisfy $c^T x'' \leq \delta c^T \bar{x}/(1-\varepsilon)$ and $Ax'' \geq b$, is at most $1/\delta + m\ exp(-\varepsilon^2 kW/2)$. The observation that by Step 1, $x''_j \leq \lceil \bar{x}_j/(1-\varepsilon) \rceil$, $\forall j$, completes the proof. ∎

We now analyze Step 2.

Lemma 2.2 *For any $\delta, \varepsilon > 0$ the probability that $\hat{x}$ fails to satisfy*

$$c^T \hat{x} \leq \frac{k\delta}{1-\varepsilon} c^T \bar{x}, \quad A\hat{x} \geq b, \quad \hat{x} \leq \lceil \bar{x}/(1-\varepsilon) \rceil$$

is less than $1/\delta + m\ exp(-\varepsilon^2 kW/2)$.

Proof. The factor of k comes into the cost because each x''_j is a multiple of $1/k$, so $\lceil x''_i \rceil \le kx''_j$. The remaining inequalities follow directly from Lemma 2.1. ∎

Theorem 2.1 *Given a CIP $\mathcal{P} = (A, b, c, d)$, and an $\varepsilon' > 0$, one can obtain in polynomial time an $\left(O\left(\max\{1, 1/\varepsilon'^2\}[1 + \log(m)/W]\right), 1+\varepsilon'\right)$-approximate solution.*

Proof. Set $\varepsilon = \min\{1/2, \varepsilon'/(1+\varepsilon')\}$. We choose δ and k so that the performance ratio $k\delta/(1-\varepsilon)$ is $O(1 + \log(m)/W)$ and the bound, $1/\delta + m\, exp(-\varepsilon^2 kW/2)$, on the probability of failure is at most 1. Choose $\delta = 2$, $k = \lceil 2\log(2m)/(W\varepsilon^2)\rceil$. By Lemma 2.2 the integer vector $\hat{x}$ returned by FINELY_ROUND satisfies with nonzero probability the conditions:

$$\begin{aligned} c^T\hat{x} &\le \min\{2, 1+\varepsilon'\}2kc^T\bar{x} && (1)\\ A\hat{x} &\ge b && (2)\\ \hat{x} &\le \lceil (1+\varepsilon')\bar{x}\rceil && (3)\end{aligned}$$

Derandomizing the construction to compute the solution deterministically can be done in a standard fashion using the method of conditional probabilities [5, 18, 13]. ∎

Recall that Δ denotes $\max_j d_j$. In Theorem 2.1, by choosing $\varepsilon' < 1/(2d_j)$, for all j, we can obtain an additive 1 guarantee. But then the performance guarantee on the cost becomes proportional to Δ^2.

Theorem 2.2 *Given a CIP $\mathcal{P} = (A, b, c, d)$, denote by y_* the optimum of the standard linear relaxation. One can obtain in deterministic polynomial time an integral vector $\hat{x}$ of cost at most $O(\Delta^2[1 + \log(m)/W]y_*)$, such that $A\hat{x} \ge b$ and $\hat{x} \le d + \mathsf{vc}(1)$.*

We can easily deduce from Theorem 2.1 a bicriteria-type result for the CIP_∞ problem. Solve the standard linear relaxation for CIP_∞ and then round the fractional solution $\bar{x}$ introducing multiplicity bound $d = \bar{x}$.

Corollary 2.1 *Given a CIP_∞ instance $\mathcal{P} = (A, b, c)$, denote by $\bar{x}$ the optimal solution to the standard linear relaxation and by y_* the value $c^T\bar{x}$. For any $\varepsilon' > 0$, we can obtain in polynomial time an integral solution $\hat{x}$ of cost $O(\max\{1, 1/\varepsilon'^2\}[1 + \log(m)/W]y_*)$ such that for all j, $\hat{x}_j \le \lceil (1+\varepsilon')\bar{x}_j\rceil$.*

Cost guarantees depending on the dilation

Recall that the dilation α of a CIP is the maximum number of constraints any variable appears in. We sketch briefly how to improve the performance guarantee so that it depends on α instead of m. The basic idea is to combine our idea of rounding to finer than integer units with Srinivasan's derandomization in [19] of the standard randomized rounding for CIP_∞ :

Theorem 2.3 *[19] Given a CIP_∞ instance $\mathcal{P} = (A, b, c)$, of width W, and a fractional solution $\bar{x}$ one can obtain in deterministic polynomial time an integer feasible solution of cost $L \doteq 1 + O(\max\{\ln(\alpha+1)/W, \sqrt{\ln(\alpha+1)/W}\})$ times the fractional cost.*

By inspection of Srinivasan's analysis [19] the obtained integer solution $\hat{x}$ satisfies $\hat{x} \le \lceil L\bar{x}\rceil$. By scaling by k the following corollary is immediate:

Corollary 2.2 *Given a CIP instance $\mathcal{P} = (A, b, c, d)$, of width W, a fractional solution $\bar{x}$, and integer $k > 0$, let $L_k \doteq 1 + O(\max\{\ln(\alpha+1)/(kW), \sqrt{\ln(\alpha+1)/(kW)}\})$. One can obtain in deterministic polynomial time a feasible solution x'' where each x''_j is an integer multiple of $1/k$, and*

$$c^T x'' \le L_k c^T\bar{x},\ Ax'' \ge b,\ x'' \le \lceil L_k\bar{x}\rceil.$$

Comparing this to Lemma 2.1 we see that if we modify the FINELY_ROUND procedure to use Corollary 2.2 in place of Steps 0 and 1, we can choose k to obtain the following theorem.

Theorem 2.4 *Given a CIP $\mathcal{P} = (A, b, c, d)$, and an $\varepsilon' > 0$, one can obtain in polynomial time an $\left(O\left(\max\{1, 1/\varepsilon'^2\}[1 + \log(\alpha)/W]\right), 1+\varepsilon'\right)$-approximate solution.*

3. Generalized multiplicity constraints

The rounding argument from Section 2 can deal with a more general form of multiplicity constraints. Consider an integer program $\mathcal{P}_G = (A, B, b, c, d)$ of the form

$$\begin{aligned} \min\ & c^T x \\ Ax &\ge b \\ Bx &\le d \\ x &\in Z_+^n \end{aligned}$$

where all the coefficients are nonnegative and B is an $r \times n$ matrix, $r \ge 1$. Matrix A has again dimension $m \times n$. When B is the identity $n \times n$ matrix we obtain a CIP as a special case. Generalized multiplicity constraints capture natural additional constraints. Consider the situation when B is a $(0, 1)$ matrix. Row l of the matrix B can be seen as corresponding to a subset S_l of the covering objects. We seek a minimum-cost cover but no more than d_l objects can be chosen from each set S_l. CIPs with (generalized) multiplicity constraints fall in the class of Mixed Packing Covering IPs.

Denote by T_l the sum of coefficients at the l-th row of B. The algorithm is very similar to the FINELY_ROUND so we only outline the differences. After Step 1 of the rounding scheme, the equivalent of Lemma 2.1 goes through with

$(Bx'')_l \le d_l/(1-\varepsilon) + T_l/k$. After Step 2, the generalized multiplicity constraints yield $(B\hat{x})_l \le \lceil d_l/(1-\varepsilon)+T_l/k \rceil$. To bound the probability of failure below 1 we choose again $k = \lceil 2\log(2m)/(W\varepsilon^2) \rceil$.

Theorem 3.1 *Let $\mathcal{P}_G = (A, B, b, c, d)$ be a mixed packing covering integer program as defined above. For any $\varepsilon' > 0$, one can obtain in deterministic polynomial time an integral solution $\hat{x}$ of cost $O(\left(\max\{1, 1/\varepsilon'^2\}[1 + \log(m)/W]\right))$ times the optimum of the standard LP relaxation, which satisfies $A\hat{x} \ge b$ and $\forall 1 \le l \le r$,*

$$(B\hat{x})_l \le \lceil \ (1+\varepsilon')d_l + O\left(\min\{\varepsilon'^2, 1\}T_l W/(\log m)\right) \ \rceil.$$

Picking different values for ε' gives tradeoffs between the violation of the multiplicity constraints and the cost approximation. Observe that when $\max_l T_l = o(\log m/W)$ a multiplicative 2-violation of the generalized multiplicity constraints becomes possible. If a larger approximation for the cost can be tolerated, higher values of $\max_l T_l$ can still lead to small multiplicity constraint violation.

4. An approximation with respect to the integral optimum

In this section we show how to obtain an $O(\log m)$ cost-approximation with respect to the integral optimum of a CIP, while respecting the multiplicity constraints. Careful examination of Theorems 2.1, 2.2 reveals that if a value $\bar{x}_j$ in the fractional solution is significantly smaller than the upper bound d_j, (e.g., $\bar{x}_j \le d_j/2$) the rounding scheme can be tuned to respect the corresponding multiplicity constraint. It suffices to focus on the residual covering problem $\mathcal{P}'$ defined on those small variables — the large ones can be rounded up to d_j at low cost. The rounding scheme requires the variable coefficients to be at most the residual covering requirement b'_i on each constraint. Enforcing this on $\mathcal{P}'$ leads to a potential decrease of the original variable coefficients A_{ij}. Then the original fractional values are not feasible any more and one has to solve the LP-relaxation of $\mathcal{P}'$ from scratch. The repetition of this process over t iterations would yield a cost guarantee proportional to t.

Let N denote the set of all variables. To address the difficulty described above we use as a starting point for the rounding scheme the solution $\bar{x}'$ to a new linear relaxation introduced recently by Carr et al. [1]. The latter formulation guarantees the following strong decomposition property for any fractional solution $\bar{x}$: if all variables in a certain set $F \subset N$ are set to their upper bounds, one can formulate an IP $\mathcal{P}_{N-F}$ for the residual covering problem over the variables in $N - F$ such that (i) $\mathcal{P}_{N-F}$ meets the CIP Definition 1.1 and (ii) the fractional optimum of $\mathcal{P}_{N-F}$ is no more than $\sum_{j \in N-F} c_j \bar{x}'_j$.

Here is the new LP formulation. Given a CIP $\mathcal{P} = (A, b, c, d)$ we generate a set S_i of valid inequalities, called Knapsack Cover (KC) inequalities in [1], associated with the i-th constraint $(Ax)_i \ge b_i$. define $b_i(F) \doteq \max\{0, b_i - \sum_{j\in F} A_{ij} d_j\}$. In words, $b_i(F)$ denotes the residual covering requirement of the i-th constraint once all variables in F have been set to their upper bounds. Define also $A^F_{ij} \doteq \min\{A_{ij}, b_i(F)\}$. The set S_i consists of the following KC inequalities:

$$\sum_{j \in N \setminus F} A^F_{ij} x_j \ge b_i(F), \quad \forall F \subset N.$$

Definition 4.1 *Given a CIP $\mathcal{P} = (A, b, c, d)$ the* LP-KC *relaxation of $\mathcal{P}$, seeks to minimize $c^T \cdot x$ subject to $Ax \ge b$, $x \le d$, the set of constraints $\cup_{i=1,\ldots,m} S_i$, and x nonnegative.*

We are not aware of an algorithm that solves LP-KC exactly in polynomial time. Carr et al. define the following type of solutions, which are adequate for our purpose. For $\lambda > 1$, call a vector x a *λ-relaxed solution* to LP-KC if it has cost at most the optimum of LP-KC and satisfies (i) $Ax \ge b$ (ii) $x \le d$ and (iii) the KC inequalities defined for the set $F_\lambda = \{j | x_j \ge d_j/\lambda\}$. The following theorem is an immediate corollary of the results by Carr et al. [1] and the properties of the ellipsoid method [11].

Theorem 4.1 *For any $\lambda > 1$, a λ-relaxed solution to an* LP-KC *formulation with rational coefficients can be found in polynomial time.*

We reduce the problem of approximating $\mathcal{P}$ while *respecting* the multiplicity constraints to the problem of finding an approximate solution with *bounded violation* of the multiplicity constraints. The following algorithm implements the reduction. It is parameterized by a subroutine $\mathcal{A}$, which we assume finds a (ρ, l)-approximate solution to a CIP w.r.t. the standard LP optimum.

ALGORITHM KC_ROUND $(\mathcal{P}, \mathcal{A}, l)$
STEP 1: Set $d := \lfloor d \rfloor$. Find an l-relaxed solution $\bar{x}'$ to the LP-KC relaxation for $\mathcal{P}$.
STEP 2: Let $H = \{j \mid \bar{x}'_j \ge d_j/l\}$. Set $\hat{x}_j := d_j$ for all $j \in H$.
STEP 3: Define a CIP $\mathcal{P}_{N-H} = (A', b', c, d')$ over the set of variables $N - H$. For notational convenience define A' as a full-dimensional $m \times n$ matrix with $A'_{ij} := 0$ if $j \in H$ and $A'_{ij} := A^H_{ij}$ otherwise. For $1 \le i \le m$ define $b'_i := b_i(H)$. Define $\bar{x}''_j := 0$ if $j \in H$ and $\bar{x}'_j$ otherwise. For $1 \le j \le n$ define $d_j := \bar{x}''_j$.
STEP 4: Round $\bar{x}'_j$, $j \in N - H$, to $\hat{x}_j$ by invoking the (ρ, l)-approximate algorithm $\mathcal{A}$ on $\mathcal{P}'$.

In the new IP $\mathcal{P}'$ defined in Step 3, the coefficient of a variable x_j from $N - H$ in the i-th constraint has been potentially reduced from A_{ij} so as not to exceed the residual covering requirement b'_i. This normalization of the coefficients is essential for the success of the randomized rounding experiment. The advantage of the LP-KC formulation is that even after the coefficients are reduced, the projection of the fractional solution $\bar{x}'$ on $N - H$ is feasible for the residual problem without requiring an increase of the variables. In fact all the constraints in $\mathcal{P}'$ which result from normalization are KC inequalities for the set F_l. The following lemma is a direct consequence of the definition of $\mathcal{P}'$.

Lemma 4.1 *The vector $\bar{x}''$ is a feasible fractional solution to the standard linear relaxation of $\mathcal{P}'$. The objective value of $\bar{x}''$ is $\sum_{j \in N-H} c_j \bar{x}'_j$.*

We are now ready to state the reduction.

Theorem 4.2 *If one can obtain in polynomial time a (ρ, l)-approximate solution to a CIP problem with rational coefficients then one can also obtain in polynomial time an integral solution that respects the multiplicity constraints and has cost $\max\{\rho, l\}$ times the integral optimum.*

Remark 4.1 *The assumption on rationality is necessary only to ensure polynomial-time, based on Theorem 4.1. The reduction itself works without this assumption.*

Proof. The assumption of the theorem is equivalent to assuming the existence of the subroutine $\mathcal{A}$ invoked in KC_ROUND. Taking the floor of every d_j in Step 1 does not change the space of integral solutions of $\mathcal{P}$. The variables in H trivially satisfy their upper bounds when the algorithm terminates. For $j \in N - H$, since $\mathcal{A}$ is (ρ, l)-approximate, $\hat{x}_j \leq \lceil l d'_j \rceil \leq d_j$. Hence all the multiplicity constraints are satisfied. The contribution of the variables from H to the objective is at most $l \sum_{j \in H} c_j \bar{x}'_j$. The rounding scheme applied to the variables from $N - H$ yields an integral vector of cost at most $\rho \sum_{j \in N-H} c_j \bar{x}'_j$. Given that $\sum_{j \in N} c_j \bar{x}'_j$ is at most the integral optimum the theorem follows. Observe that the approximation guarantees ρ and l for $\mathcal{P}'$ may depend on the CIP parameters (e.g., dimension, dilation); the respective parameters of $\mathcal{P}$ are always higher. ∎

As mentioned, standard randomized rounding gives an algorithm which is $(O(\log m), O(\log m))$-approximate.

Corollary 4.1 *Given a CIP $\mathcal{P} = (A, b, c, d)$, one can compute in polynomial time an integral solution that respects the multiplicity constraints and has cost $O(\log m)$ times the optimum.*

Using the dilation bound of Srinivasan [19] (cf. Theorem 2.3) the cost guarantee can be improved.

Corollary 4.2 *Given a CIP $\mathcal{P} = (A, b, c, d)$, one can compute in polynomial time an integral solution that respects the multiplicity constraints and has cost $O(\log \alpha)$ times the optimum.*

Corollary 4.3 *The integrality gap of the* LP-KC *relaxation is $O(\log \alpha)$.*

Observe that KC_ROUND can be easily changed to use an $(O(\log m), 2)$-approximate subroutine such as FINELY_ROUND. In this case, the contribution of the H-variables to the cost will be at most doubled instead of being scaled up by $\log m$.

5. Open questions

Some of the open questions resulting from this work are as follows. First, fine tune the constants in the asymptotic guarantees. Second, give an additive 1 violation of the multiplicity constraints and logaritmic cost guarantee w.r.t. the standard LP optimum for general d values. Finally, develop a simple greedy algorithm for general CIPs that achieves the $O(\log m)$ approximation.

Acknowledgements. Thanks to Vijay Vazirani for a clarification on [15]. Thanks to Lisa Fleischer and Aravind Srinivasan for comments on a draft of this paper.

References

[1] R. D. Carr, L. K. Fleischer, V. J. Leung, and C. A. Phillips. Strengthening integrality gaps for capacitated network design and covering problems. In *Proceedings of the 11th ACM-SIAM Symposium on Discrete Algorithms*, 2000.

[2] H. Chernoff. A measure of the asymptotic efficiency for tests of a hypothesis based on sum of observations. *Ann. Math. Stat.*, 23:493–509, 1952.

[3] V. Chvátal. A greedy heuristic for the set-covering problem. *Mathematics of Operations Research*, 4(3):233–235, Aug. 1979.

[4] G. Dobson. Worst-case analysis of greedy heuristics for integer programming with non-negative data. *Mathematics of Operations Research*, 7(4):515–531, 1982.

[5] P. Erdős and J. L. Selfridge. On a combinatorial game. *Journal of Combinatorial Theory A*, 14:298–301, 1973.

[6] M. L. Fisher and L. A. Wolsey. On the greedy heuristic for continuous covering and packing problems. *SIAM Journal of Algebraic and Discrete Methods*, 3:584–591, 1982.

[7] 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*, pages 94–143. PWS, 1997.

[8] D. S. Johnson. Approximation algorithms for combinatorial problems. *Journal of Computer and System Sciences*, 9:256–278, 1974.

[9] S. G. Kolliopoulos. Approximating covering integer programs with multiset constraints. Submitted for journal publication, September 2000.

[10] L. Lovász. On the ratio of optimal and fractional covers. *Discrete Mathematics*, 13:383–390, 1975.

[11] L. Lovász. *An algorithmic theory of numbers, graphs and complexity*. SIAM, Philadelphia, PA, 1986.

[12] S. Plotkin, D. B. Shmoys, and É.. Tardos. Fast approximation algorithms for fractional packing and covering problems. *Mathematics of Operations Research*, 20:257–301, 1995.

[13] P. Raghavan. Probabilistic construction of deterministic algorithms: approximating packing integer programs. *Journal of Computer and System Sciences*, 37:130–143, 1988.

[14] P. Raghavan and C. D. Thompson. Randomized rounding: a technique for provably good algorithms and algorithmic proofs. *Combinatorica*, 7:365–374, 1987.

[15] S. Rajagopalan and V. V. Vazirani. Primal-dual RNC approximation algorithms for (multi)-set (multi)-cover and covering integer programs. In *Proceedings of the 34th Annual Symposium on Foundations of Computer Science*, pages 322–331, 1993.

[16] R. Raz and S. Safra. A sub-constant error-probability low-degree test, and a sub-constant error-probability PCP characterization of NP. In *Proceedings of the 29th Annual ACM Symposium on Theory of Computing*, pages 475–484, 1997.

[17] D. B. Shmoys, E. Tardos, and K. I. Aardal. Approximation algorithms for facility location problems. In *Proceedings of the 29th Annual ACM Symposium on Theory of Computing*, pages 265–274, 1997.

[18] J. Spencer. *Ten Lectures on the Probabilistic Method*. SIAM, Philadelphia, 1987.

[19] A. Srinivasan. An extension of the Lovász Local Lemma and its applications to integer programming. In *Proceedings of the 7th ACM-SIAM Symposium on Discrete Algorithms*, pages 6–15, 1996.

[20] A. Srinivasan. Improved approximations guarantees for packing and covering integer programs. *SIAM J. Comput.*, 29:648–670, 1999. Preliminary version in Proc. STOC 95.

[21] A. Srinivasan and C.-P. Teo. A constant-factor approximation algorithm for packet routing and balancing local vs. global criteria. *SIAM J. Comput.*, 30:2051–2068, 2001. Preliminary version in Proc. STOC 97.

[22] N. E. Young. Randomized rounding without solving the linear program. In *Proceedings of the 6th ACM-SIAM Symposium on Discrete Algorithms*, pages 170–178, 1995.