




ℓ_1 -Sparsity approximation bounds for packing integer programs

Chandra Chekuri¹ · Kent Quanrud² · Manuel R. Torres³ 

Received: 31 May 2019 / Accepted: 23 January 2020

© Springer-Verlag GmbH Germany, part of Springer Nature and Mathematical Optimization Society 2020

Abstract

We consider approximation algorithms for packing integer programs (PIPs) of the form $\max\{\langle c, x \rangle : Ax \leq b, x \in \{0, 1\}^n\}$ where A , b and c are nonnegative. We let $W = \min_{i,j} b_i / A_{i,j}$ denote the width of A which is at least 1. Previous work by Bansal et al. (Theory Comput 8(24):533–565, 2012) obtained an $\Omega(\frac{1}{\Delta_0^{1/|W|}})$ -approximation ratio where Δ_0 is the maximum number of nonzeros in any column of A (in other words the ℓ_0 -column sparsity of A). They raised the question of obtaining approximation ratios based on the ℓ_1 -column sparsity of A (denoted by Δ_1) which can be much smaller than Δ_0 . Motivated by recent work on covering integer programs (Chekuri and Quanrud, in: Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, pp 1596–1615. SIAM, 2019; Chen et al., in: Proceedings of the Twenty-seventh Annual ACM-SIAM Symposium on Discrete Algorithms, pp 1984–2003. SIAM, 2016) we show that simple algorithms based on randomized rounding followed by alteration, similar to those of Bansal et al. (Theory Comput 8(24):533–565, 2012) (but with a twist), yield approximation ratios for PIPs based on Δ_1 . First, following an integrality gap example from (Theory Comput 8(24):533–565, 2012), we observe that the case of $W = 1$ is as hard as maximum independent set even when $\Delta_1 \leq 2$. In sharp contrast to this negative result, as soon as width is strictly larger than one, we obtain positive results via the natural LP relaxation. For PIPs with width $W = 1 + \epsilon$ where $\epsilon \in (0, 1]$, we obtain an $\Omega(\epsilon^2 / \Delta_1)$ -approximation. In the large width regime, when $W \geq 2$, we obtain an $\Omega((\frac{1}{1+\Delta_1/W})^{1/(W-1)})$ -approximation. We also obtain a $(1 - \epsilon)$ -approximation when $W = \Omega(\frac{\log(\Delta_1/\epsilon)}{\epsilon^2})$. Viewing the rounding

C. Chekuri and K. Quanrud supported in part by NSF Grant CCF-1526799. M. Torres supported in part by fellowships from NSF and the Sloan Foundation. A preliminary version of this work appeared in the proceedings of the 20th conference on Integer Programming and Combinatorial Optimization (IPCO 2019).

✉ Manuel R. Torres
manuelt2@illinois.edu

Extended author information available on the last page of the article

algorithms as contention resolution schemes, we obtain approximation algorithms in the more general setting when the objective is a non-negative submodular function.

Keywords Approximation algorithms · Sparse packing integer programs · Randomized rounding · Submodular optimization

Mathematics Subject Classification Primary 68W25 · Secondary 90C59

1 Introduction

Packing integer programs (abbr. PIPs) are an expressive class of integer programs of the form:

$$\text{maximize } \langle c, x \rangle \text{ over } x \in \{0, 1\}^n \text{ s.t. } Ax \leq b,$$

where $A \in \mathbb{R}_{\geq 0}^{m \times n}$, $b \in \mathbb{R}_{\geq 0}^m$ and $c \in \mathbb{R}_{\geq 0}^n$ all have nonnegative entries¹. Many important problems in discrete and combinatorial optimization can be cast as special cases of PIPs. These include the maximum independent set in graphs and hypergraphs, set packing, matchings and b -matchings, knapsack (when $m = 1$), and the multi-dimensional knapsack. The maximum independent set problem (MIS), a special case of PIPs, is NP -hard and unless $P = NP$ there is no $n^{1-\epsilon}$ -approximation where n is the number of nodes in the graph [14,23]. For this reason it is meaningful to consider special cases and other parameters that control the difficulty of PIPs. Motivated by the fact that MIS admits a simple $\frac{1}{\Delta(G)}$ -approximation where $\Delta(G)$ is the maximum degree of G , previous work considered approximating PIPs based on the maximum number of nonzeros in any column of A (denoted by Δ_0); note that when MIS is written as a PIP, Δ_0 coincides with $\Delta(G)$. As another example, when maximum weight matching is written as a PIP, $\Delta_0 = 2$. Bansal et al. [1] obtained a simple and clever algorithm that achieved an $\Omega(1/\Delta_0)$ -approximation for PIPs via the natural LP relaxation; this improved previous work of Pritchard [17,18] who was the first to obtain an approximation for PIPs only as a function of Δ_0 . Moreover, the rounding algorithm in [1] can be viewed as a contention resolution scheme which allows one to get similar approximation ratios even when the objective is submodular [1,8]. It is well-understood that PIPs become easier when the entries in A are small compared to the packing constraints b . To make this quantitative we consider the well-studied notion called the *width* defined as $W := \min_{i,j:A_{i,j}>0} b_i/A_{i,j}$. Bansal et al. obtain an $\Omega((\frac{1}{\Delta_0})^{1/|W|})$ -approximation which improves as W becomes larger. Although they do not state it explicitly, their approach also yields a $(1 - \epsilon)$ -approximation when $W = \Omega(\frac{1}{\epsilon^2} \log(\Delta_0/\epsilon))$.

Δ_0 is a natural measure for combinatorial applications such as MIS and matchings where the underlying matrix A has entries from $\{0, 1\}$. However, in some applications of PIPs such as knapsack and its multi-dimensional generalization which are

¹ We can allow the variables to have general integer upper bounds instead of restricting them to be boolean. As observed in [1], one can reduce this more general case to the $\{0, 1\}$ case without too much loss in the approximation.

more common in resource-allocation problems, the entries of A are arbitrary rational numbers (which can be assumed to be from the interval $[0, 1]$ after scaling). In such applications it is natural to consider another measure of column-sparsity based on the ℓ_1 norm. Specifically we consider Δ_1 , the maximum column sum of A . Unlike Δ_0 , Δ_1 is not scale invariant so one needs to be careful in understanding the parameter and its relationship to the width W . For this purpose we normalize the constraints $Ax \leq b$ as follows. Let $W = \min_{i,j:A_{i,j}>0} b_i/A_{i,j}$ denote the width as before (we can assume without loss of generality that $W \geq 1$ since we are interested in integer solutions). We can then scale each row A_i of A separately such that, after scaling, the i 'th constraint reads as $A_i x \leq W$. After scaling all rows in this fashion, entries of A are in the interval $[0, 1]$, and the maximum entry of A is equal to 1. Note that this scaling process does not alter the original width. We let Δ_1 denote the maximum column sum of A after this normalization and observe that $1 \leq \Delta_1 \leq \Delta_0$. In many settings of interest $\Delta_1 \ll \Delta_0$. We also observe that Δ_1 is a more robust measure than Δ_0 ; small perturbations of the entries of A can dramatically change Δ_0 while Δ_1 changes minimally.

Bansal et al. raised the question of obtaining an approximation ratio for PIPs as a function of only Δ_1 . They observed that this is not feasible via the natural LP relaxation by describing a simple example where the integrality gap of the LP is $\Omega(n)$ while Δ_1 is a constant. Their example essentially shows the existence of a simple approximation preserving reduction from MIS to PIPs such that the resulting instances have $\Delta_1 \leq 2$; thus no approximation ratio that depends only on Δ_1 is feasible for PIPs unless $P = NP$. These negative results seem to suggest that pursuing bounds based on Δ_1 is futile, at least in the worst case. However, the starting point of this paper is the observation that both the integrality gap example and the hardness result are based on instances where the width W of the instance is arbitrarily close to 1. We demonstrate that these examples are rather brittle and obtain several positive results when we consider $W \geq 1 + \epsilon$ for any fixed $\epsilon > 0$.

1.1 Our results

Our first result is on the hardness of approximation for PIPs that we already referred to. The hardness result suggests that one should consider instances with $W > 1$. Recall that after normalization we have $\Delta_1 \geq 1$ and $W \geq 1$ and the maximum entry of A is 1. We consider three regimes of W and obtain the following results, all via the natural LP relaxation, which also establish corresponding upper bounds on the integrality gap.

- (i) $1 < W \leq 2$. For $W = 1 + \epsilon$ where $\epsilon \in (0, 1]$ we obtain an $\Omega(\frac{\epsilon^2}{\Delta_1})$ -approximation.
- (ii) $W \geq 2$. We obtain an $\Omega((\frac{1}{1+\frac{\Delta_1}{W}})^{1/(W-1)})$ -approximation which can be simplified to $\Omega((\frac{1}{1+\Delta_1})^{1/(W-1)})$ since $W \geq 1$.
- (iii) A $(1 - \epsilon)$ -approximation when $W = \Omega(\frac{1}{\epsilon^2} \log(\Delta_1/\epsilon))$.

Our results establish approximation bounds based on Δ_1 that are essentially the same as those based on Δ_0 as long as the width is not too close to 1. When the matrix A is a $\{0, 1\}$ -matrix, $\Delta_1 = \Delta_0$, and previous integrality gap results based on Δ_0 [1] show that the bounds we obtain are essentially tight modulo constant factors. We describe randomized algorithms which can be derandomized via standard techniques.

All our algorithms are based on a simple randomized rounding plus alteration framework that has been successful for both packing and covering problems. Our scheme is similar to that of Bansal et al. at a high level but we make a simple but important change in the algorithm and its analysis. This is inspired by recent work on covering integer programs [5] where ℓ_1 -sparsity based approximation bounds from [9] were simplified.

The rounding algorithms can be viewed as contention resolution schemes, and via known techniques [1,8], we also obtain approximation algorithms for submodular objectives. We present the results for this generalization in Sect. 6.

1.2 Other related work

We note that PIPs are equivalent to the multi-dimensional knapsack problem. When $m = 1$ we have the classical knapsack problem, which admits a very efficient FPTAS (see [3]). There is a PTAS for any fixed m [12] but unless $P = NP$ an FPTAS does not exist for $m = 2$. Approximation algorithms for PIPs in their general form were considered initially by Raghavan and Thompson [19] and refined substantially by Srinivasan [20]. Srinivasan obtained approximation ratios of the form $\Omega(1/n^{1/(W+1)})$ when A had entries from $\{0, 1\}$, and a ratio of the form $\Omega(1/n^{1/W})$ when A had entries from $[0, 1]$. Pritchard [17] was the first to obtain a bound for PIPs based solely on the column sparsity parameter Δ_0 . He used iterated rounding and his initial bound was improved in [18] to $\Omega(1/\Delta_0^2)$. The current state of the art is due to Bansal et al. [1]. Previously we ignored constant factors when describing the ratio. In fact [1] obtains a ratio of $\frac{1}{e\Delta_0 + o(\Delta_0)}$ by strengthening the basic LP relaxation.

In terms of hardness of approximation, PIPs generalize MIS and hence one cannot obtain a ratio better than $n^{1-\epsilon}$ unless $P = NP$ [14,23]. Building on MIS, [4] shows that PIPs are hard to approximate within a $n^{\Omega(1/W)}$ factor for any constant width W . Hardness of MIS in bounded degree graphs [21] and hardness for k -set-packing [15] imply that PIPs are hard to approximate to within $\Omega(1/\Delta_0^{1-\epsilon})$ and to within $\Omega((\log \Delta_0)/\Delta_0)$ when Δ_0 is a sufficiently large constant. These hardness results are based on $\{0, 1\}$ matrices for which Δ_0 and Δ_1 coincide.

There is a large literature on deterministic and randomized rounding algorithms for packing and covering integer programs and connections to several topics and applications including discrepancy theory. ℓ_1 -sparsity guarantees for covering integer programs (CIPs) were first obtained by Chen, Harris and Srinivasan [9] partly inspired by [13].

Recent (and extensive) work on submodular function maximization has demonstrated that several approximation algorithms for modular objectives can be generalized to non-negative submodular objectives. Of particular relevance is the approach via the multilinear relaxation followed by rounding via contention resolution schemes. We refer the reader to [8] for the framework. This allows us to extend our results to submodular objectives.

2 Hardness of approximating PIPs as a function of Δ_1

Bansal et al. [1] showed that the integrality gap of the natural LP relaxation for PIPs is $\Omega(n)$ even when Δ_1 is a constant. One can use essentially the same construction to show the following theorem.

Theorem 1 *There is an approximation preserving reduction from MIS to instances of PIPs with $\Delta_1 \leq 2$.*

Proof Let $G = (V, E)$ be an undirected graph without self-loops and let $n = |V|$. Let $A \in [0, 1]^{n \times n}$ be indexed by V . For all $v \in V$, let $A_{v,v} = 1$. For all $uv \in E$, let $A_{u,v} = A_{v,u} = 1/n$. For all the remaining entries in A that have not yet been defined, set these entries to 0. Consider the following PIP:

$$\text{maximize } \langle x, \mathbf{1} \rangle \text{ over } x \in \{0, 1\}^n \text{ s.t. } Ax \leq \mathbf{1}. \quad (1)$$

Let S be the set of all feasible integral solutions of (1) and \mathcal{I} be the set of independent sets of G . Define $g : S \rightarrow \mathcal{I}$ where $g(x) = \{v : x_v = 1\}$. To show g is surjective, consider a set $I \in \mathcal{I}$. Let y be the characteristic vector of I . That is, y_v is 1 if $v \in I$ and 0 otherwise. Consider the row in A corresponding to an arbitrary vertex u where $y_u = 1$. For all $v \in V$ such that v is a neighbor to u , $y_v = 0$ as I is an independent set. Thus, as the nonzero entries in A of the row corresponding to u are, by construction, the neighbors of u , it follows that the constraint corresponding to u is satisfied in (1). As u is an arbitrary vertex, it follows that y is a feasible integral solution to (1) and as $I = \{v : y_v = 1\}$, $g(y) = I$.

Define $h : S \rightarrow \mathbb{N}_0$ such that $h(x) = |g(x)|$. It is clear that $\max_{x \in S} h(x)$ is equal to the optimal value of (1). Let I_{max} be a maximum independent set of G . As g is surjective, there exists $z \in S$ such that $g(z) = I_{max}$. Thus, $\max_{x \in S} h(x) \geq |I_{max}|$. As $\max_{x \in S} h(x)$ is equal to the optimum value of (1), it follows that a β -approximation for PIPs implies a β -approximation for maximum independent set.

Furthermore, we note that for this PIP, $\Delta_1 \leq 2$, thus concluding the proof. \square

Unless $P = NP$, MIS does not admit a $n^{1-\epsilon}$ -approximation for any fixed $\epsilon > 0$ [14,23]. Hence the preceding theorem implies that unless $P = NP$ one cannot obtain an approximation ratio for PIPs solely as a function of Δ_1 .

3 Round and alter framework

The algorithms in this paper have the same high-level structure. The algorithms first scale down the fractional solution x by some factor α , and then randomly round each coordinate independently. The rounded solution x' may not be feasible for the constraints. The algorithm alters x' to a feasible x'' by considering each constraint separately in an arbitrary order; if x' is not feasible for constraint i , some subset S of variables are chosen to be set to 0. Each constraint corresponds to a knapsack problem and the framework (which is adapted from [1]) views the problem as the intersection of several knapsack constraints. A formal template is given in Fig. 1. To make the

```

Round-and-Alter Framework: input  $A$ ,  $b$ , and  $\alpha$ 
let  $x$  be the optimum fractional solution of the natural LP relaxation
for  $j \in [n]$ , set  $x'_j$  to be 1 independently with probability  $\alpha x_j$  and 0 otherwise
 $x'' \leftarrow x'$ 
for  $i \in [m]$  do
  find  $S \subseteq [n]$  such that setting  $x'_j = 0$  for all  $j \in S$  would satisfy  $\langle e_i, Ax' \rangle \leq b_i$ 
  for all  $j \in S$ , set  $x''_j = 0$ 
end for
return  $x''$ 

```

Fig. 1 Randomized rounding with alteration framework

framework into a formal algorithm, one must define α and how to choose S in the for loop. These parts will depend on the regime of interest.

For an algorithm that follows the round-and-alter framework, the expected output of the algorithm is $\mathbb{E}[\langle c, x'' \rangle] = \sum_{j=1}^n c_j \cdot \Pr[x''_j = 1]$. Independent of how α is defined or how S is chosen, $\Pr[x''_j = 1] = \Pr[x'_j = 1 | x'_j = 1] \cdot \Pr[x'_j = 1]$ since $x''_j \leq x'_j$. Then we have

$$\mathbb{E}[\langle c, x'' \rangle] = \alpha \sum_{j=1}^n c_j x_j \cdot \Pr[x'_j = 1 | x'_j = 1].$$

Let E_{ij} be the event that x''_j is set to 0 when ensuring constraint i is satisfied in the for loop. As x''_j is only set to 0 if at least one constraint sets x'_j to 0, we have

$$\Pr[x''_j = 0 | x'_j = 1] = \Pr \left[\bigcup_{i \in [m]} E_{ij} | x'_j = 1 \right] \leq \sum_{i=1}^m \Pr[E_{ij} | x'_j = 1].$$

Combining these two observations, we have the following lemma, which applies to all of our subsequent algorithms.

Lemma 1 *Let \mathcal{A} be a randomized rounding algorithm that follows the round-and-alter framework given in Fig. 1. Let x' be the rounded solution obtained with scaling factor α . Let E_{ij} be the event that x''_j is set to 0 by constraint i . If for all $j \in [n]$ we have $\sum_{i=1}^m \Pr[E_{ij} | x'_j = 1] \leq \gamma$, then \mathcal{A} is an $\alpha(1 - \gamma)$ -approximation for PIPs.*

We will refer to the quantity $\Pr[E_{ij} | x'_j = 1]$ as the *rejection probability* of item j in constraint i . We will also say that constraint i *rejects* item j if x''_j is set to 0 in constraint i .

3.1 Concentration inequalities and other useful inequalities

In the subsequent sections, for particular regimes of interest, we rely on Chernoff bounds to upper bound rejection probabilities. The following standard Chernoff bound is used to obtain a more convenient bound in Lemma 3. The proof of Lemma 3 follows

directly from choosing δ such that $(1 + \delta)\mu = W - \beta$ and applying Lemma 2. We include the proof for convenience.

Lemma 2 ([16]) *Let X_1, \dots, X_n be independent random variables where X_i is defined on $\{0, \beta_i\}$, where $0 < \beta_i \leq \beta \leq 1$ for some β . Let $X = \sum_i X_i$ and denote $\mathbb{E}[X]$ as μ . Then for any $\delta > 0$,*

$$\Pr[X \geq (1 + \delta)\mu] \leq \left(\frac{e^\delta}{(1 + \delta)^{1+\delta}} \right)^{\mu/\beta}$$

Lemma 3 *Let $X_1, \dots, X_n \in [0, \beta]$ be independent random variables for some $0 < \beta \leq 1$. Suppose $\mu = \mathbb{E}[\sum_i X_i] \leq \alpha W$ for some $0 < \alpha < 1$ and $W \geq 1$ where $(1 - \alpha)W > \beta$. Then*

$$\Pr \left[\sum_i X_i > W - \beta \right] \leq \left(\frac{\alpha e^{1-\alpha} W}{W - \beta} \right)^{(W-\beta)/\beta}.$$

Proof Since the right-hand side is increasing in α , it suffices to assume $\mu = \alpha W$. Choose δ such that $(1 + \delta)\mu = W - \beta$. Then $\delta = (W - \beta - \mu)/\mu$. Because $\mu = \alpha W$ and since $(1 - \alpha)W > \beta$, we have $\delta = ((1 - \alpha)W - \beta)/\mu > 0$. We apply the standard Chernoff bound in Lemma 3 to obtain

$$\Pr \left[\sum_i X_i > W - \beta \right] = \Pr \left[\sum_i X_i > (1 + \delta)\mu \right] \leq \left(\frac{e^\delta}{(1 + \delta)^{1+\delta}} \right)^{\mu/\beta}.$$

Because $1 + \delta = (W - \beta)/\mu$ and $\delta = (W - \beta - \mu)/\mu$,

$$\left(\frac{e^\delta}{(1 + \delta)^{1+\delta}} \right)^{\mu/\beta} = \left(\frac{e^{W-\beta-\mu}}{((W - \beta)/\mu)^{W-\beta}} \right)^{1/\beta}.$$

Exponentiating the denominator,

$$\left(\frac{e^{W-\beta-\mu}}{((W - \beta)/\mu)^{W-\beta}} \right)^{1/\beta} = \exp \left(\frac{1}{\beta} \left(W - \beta - \mu + (W - \beta) \ln \left(\frac{\mu}{W - \beta} \right) \right) \right)$$

As $\mu = \alpha W$,

$$\begin{aligned} & \exp \left(\frac{1}{\beta} \left(W - \beta - \mu + (W - \beta) \ln \left(\frac{\mu}{W - \beta} \right) \right) \right) \\ &= \exp \left(\frac{1}{\beta} \left((1 - \alpha)W - \beta + (W - \beta) \ln \left(\frac{\alpha W}{W - \beta} \right) \right) \right) \end{aligned}$$

We can rewrite the exponent to show that

$$\exp\left(\frac{1}{\beta}\left((1-\alpha)W - \beta - (W-\beta)\ln\left(\frac{W-\beta}{\alpha W}\right)\right)\right) \leq \left(\frac{\alpha e^{1-\alpha}W}{W-\beta}\right)^{(W-\beta)/\beta}.$$

□

The following three lemmas are used in the proofs bounding the rejection probabilities for different regimes of width. The inequalities are easily verified via calculus. The proofs are included for the sake of completeness.

Lemma 4 *Let $x \in (0, 1]$. Then $(1/e^{1/e})^{1/x} \leq x$.*

Proof Taking logs of both sides of the stated inequality and rearranging, it suffices to show that $\ln(1/e^{1/e}) \leq x \ln x$ for $x > 0$. $x \ln x$ is convex and its minimum is $-1/e$ at $x = 1/e$. Since $\ln(1/e^{1/e}) = -1/e$, the inequality holds. □

Lemma 5 *Let $y \geq 2$ and $x \in (0, 1]$. Then $x/y \geq (1/e^{2/e})^{y/2x}$.*

Proof We start with a simple rewriting of the statement. After taking logs and rearranging, it is sufficient to show

$$(x/y) \ln(x/y) \geq (1/2) \ln(1/e^{2/e}) = -1/e.$$

Replacing x/y with z , we see that it suffices to prove $z \ln z \geq -1/e$ for $0 < z \leq 1/2$. We note that $x \ln x$ is convex and its minimum is $-1/e$ at $x = 1/e$. Thus, $z \ln z \geq -1/e$. This concludes the proof. □

Lemma 6 *Let $0 < \epsilon \leq 1$ and $x \in (0, 1]$. Then $\epsilon x/2 \geq (\epsilon/e^{2/e})^{1/x}$.*

Proof Let $d = e^{2/e}/2$ and observe that $d > 1$. We first do a change of variables, replacing $\epsilon/2$ with ϵ and x with x/ϵ . If we take a log of both sides, then our reformulated goal is to show that

$$x \ln x \geq \epsilon \ln(\epsilon/d)$$

for $0 < \epsilon \leq 1/2$ and $x \in (0, \epsilon]$. Letting $f(y) = y \ln y$ and $g(y) = y \ln(y/d)$, we want to show that $f(x) \geq g(\epsilon)$. We will proceed by cases.

First, suppose $0 < \epsilon \leq d/e$. It is easy to show that f is decreasing on $(0, 1/e]$ and increasing on $[1/e, \infty)$ and that g is decreasing on $(0, d/e]$ and increasing on $[d/e, \infty)$. As f is decreasing on $(0, 1/e]$, for $0 < \epsilon \leq 1/e$, we have $f(x) \geq f(\epsilon)$ as $x \leq \epsilon$. As $d > 1$, it follows that $f(\epsilon) \geq g(\epsilon)$. Therefore, $f(x) \geq g(\epsilon)$ for $0 < \epsilon \leq 1/e$. Furthermore, as g is decreasing on $[1/e, d/e]$ and f is increasing on $[1/e, d/e]$, we have $f(x) \geq g(\epsilon)$ for $0 < \epsilon \leq d/e$.

For the second case, suppose $d/e < \epsilon \leq 1/2$. Note that the minimum of f on the interval $(0, 1/2]$ is $f(1/e) = -1/e$. Thus, it would suffice to show that $g(\epsilon) \leq -1/e$. As we noted previously that g is increasing on $[d/e, 1/2]$, it would suffice to show that $g(1/2) \leq -1/e$. By definition of g , we see $g(1/2) = -1/e$. Therefore, $f(x) \geq g(\epsilon)$. This concludes the proof. □


```

round-and-alter-by-sorting( $A, b, \alpha_1$ ):
  let  $x$  be an optimum fractional solution of the natural LP relaxation
  for  $j \in [n]$ , set  $x'_j$  to be 1 independently with probability  $\alpha_1 x_j$  and 0 otherwise
   $x'' \leftarrow x'$ 
  for  $i \in [m]$  do
    sort and renumber such that  $A_{i,1} \leq \dots \leq A_{i,n}$ 
     $s \leftarrow \max\{\ell \in [n] : \sum_{j=1}^{\ell} A_{i,j} x'_j \leq b_i\}$ 
    for each  $j \in [n]$  such that  $j > s$ , set  $x''_j = 0$ 
  end for
  return  $x''$ 
    
```

Fig. 2 Round-and-alter in the large width regime. Each constraint sorts the coordinates in increasing size and greedily picks a feasible set and discards the rest

4 The large width regime: $W \geq 2$

In this section, we consider PIPs with width $W \geq 2$. Recall that we assume $A \in [0, 1]^{m \times n}$ and $b_i = W$ for all $i \in [m]$. Therefore we have $A_{i,j} \leq W/2$ for all i, j and from a knapsack point of view all items are “small”. We apply the round-and-alter framework in a simple fashion where in each constraint i the coordinates are sorted by the coefficients in that row and the algorithm chooses the largest prefix of coordinates that fit in the capacity W and the rest are discarded. We emphasize that this sorting step is crucial for the analysis and differs from the scheme in [1]. Figure 2 describes the formal algorithm.

The key property for the analysis. The analysis relies on obtaining a bound on the rejection probability of coordinate j by constraint i . Let X_j be the indicator variable for j being chosen in the first step. We show that $\Pr[E_{ij} \mid X_j = 1] \leq c A_{ij}$ for some c that depends on the scaling factor α . Thus coordinates with smaller coefficients are less likely to be rejected. The total rejection probability of j , $\sum_{i=1}^m \Pr[E_{ij} \mid X_j = 1]$, is proportional to the column sum of coordinate j which is at most Δ_1 .

4.1 An $\Omega(1/\Delta_1)$ -approximation algorithm

We show that round-and-alter-by-sorting yields an $\Omega(1/\Delta_1)$ -approximation if we set the scaling factor $\alpha_1 = \frac{1}{c_1 \Delta_1}$ where $c_1 = 4e^{1+1/e}$.

The rejection probability is captured by the following main lemma.

Lemma 7 *Let $\alpha_1 = \frac{1}{c_1 \Delta_1}$ for $c_1 = 4e^{1+1/e}$. Let $i \in [m]$ and $j \in [n]$. Then we have $\Pr[E_{ij} \mid X_j = 1] \leq \frac{A_{i,j}}{2\Delta_1}$ in the algorithm round-and-alter-by-sorting(A, b, α_1).*

Proof At iteration i of round-and-alter-by-sorting, after the set $\{A_{i,1}, \dots, A_{i,n}\}$ is sorted, the indices are renumbered so that $A_{i,1} \leq \dots \leq A_{i,n}$. Note that j may now be a different index j' , but for simplicity of notation we will refer to j' as j . Let $\xi_\ell = 1$ if $x'_\ell = 1$ and 0 otherwise. Let $Y_{ij} = \sum_{\ell=1}^{j-1} A_{i,\ell} \xi_\ell$.

If E_{ij} occurs, then $Y_{ij} > W - A_{i,j}$, since x''_j would not have been set to zero by constraint i otherwise. That is,

$$\Pr[E_{ij} \mid X_j = 1] \leq \Pr[Y_{ij} > W - A_{i,j} \mid X_j = 1].$$

The event $Y_{ij} > W - A_{i,j}$ does not depend on x'_j . Therefore,

$$\Pr[Y_{ij} > W - A_{i,j} | X_j = 1] \leq \Pr[Y_{ij} \geq W - A_{i,j}].$$

To upper bound $\mathbb{E}[Y_{ij}]$, we have

$$\mathbb{E}[Y_{ij}] = \sum_{\ell=1}^{j-1} A_{i,\ell} \cdot \Pr[X_\ell = 1] \leq \alpha_1 \sum_{\ell=1}^n A_{i,\ell} x_\ell \leq \alpha_1 W.$$

As $A_{i,j} \leq 1$, $W \geq 2$, and $\alpha_1 < 1/2$, we have $\frac{(1-\alpha_1)W}{A_{i,j}} > 1$. Using the fact that $A_{i,j}$ is at least as large as all entries $A_{i,j'}$ for $j' < j$, we satisfy the conditions to apply the Chernoff bound in Lemma 3. This implies

$$\Pr[Y_{ij} > W - A_{i,j}] \leq \left(\frac{\alpha_1 e^{1-\alpha_1} W}{W - A_{i,j}} \right)^{(W-A_{i,j})/A_{i,j}}.$$

Note that $\frac{W}{W-A_{i,j}} \leq 2$ as $W \geq 2$. Because $e^{1-\alpha_1} \leq e$ and by the choice of α_1 , we have

$$\left(\frac{\alpha_1 e^{1-\alpha_1} W}{W - A_{i,j}} \right)^{(W-A_{i,j})/A_{i,j}} \leq (2e\alpha_1)^{(W-A_{i,j})/A_{i,j}} = \left(\frac{1}{2e^{1/e} \Delta_1} \right)^{(W-A_{i,j})/A_{i,j}}.$$

Then we prove the final inequality in two parts. First, we see that $W \geq 2$ and $A_{i,j} \leq 1$ imply that $\frac{W-A_{i,j}}{A_{i,j}} \geq 1$. This implies

$$\left(\frac{1}{2\Delta_1} \right)^{(W-1)/A_{i,j}} \leq \frac{1}{2\Delta_1}.$$

Second, we see that

$$(1/e^{1/e})^{(W-A_{i,j})/A_{i,j}} \leq (1/e^{1/e})^{1/A_{i,j}} \leq A_{i,j}$$

for $A_{i,j} \leq 1$, where the first inequality holds because $W - A_{i,j} \geq 1$ and the second inequality holds by Lemma 4. This concludes the proof. \square

Theorem 2 When setting $\alpha_1 = \frac{1}{c_1 \Delta_1}$ where $c_1 = 4e^{1+1/e}$, round-and-alter-by-sorting (A, b, α_1) is a randomized $(\alpha_1/2)$ -approximation algorithm for PIPs with width $W \geq 2$.

Proof Fix $j \in [n]$. By Lemma 7 and the definition of Δ_1 , we have

$$\sum_{i=1}^m \Pr[E_{ij} | X_j = 1] \leq \sum_{i=1}^m \frac{A_{i,j}}{2\Delta_1} \leq \frac{1}{2}.$$

By Lemma 1, which shows that upper bounding the sum of the rejection probabilities by γ for every item leads to an $\alpha_1(1 - \gamma)$ -approximation, we get the desired result. \square

4.2 An $\Omega\left(\frac{1}{(1+\Delta_1/W)^{1/(W-1)}}\right)$ -approximation

We improve the bound from the previous section by setting $\alpha_1 = \frac{1}{c_2(1+\Delta_1/W)^{1/(W-1)}}$ where $c_2 = 4e^{1+2/e}$. Note that the scaling factor becomes larger as W increases.

Lemma 8 *Let $\alpha_1 = \frac{1}{c_2(1+\Delta_1/W)^{1/(W-1)}}$ for $c_2 = 4e^{1+2/e}$. Let $i \in [m]$ and $j \in [n]$. Then in the algorithm *round-and-alter-by-sorting*(A, b, α_1), we have $\Pr[E_{ij}|X_j = 1] \leq \frac{A_{i,j}}{2\Delta_1}$.*

Proof The proof proceeds similarly to the proof of Lemma 7. Since $\alpha_1 < 1/2$, everything up to and including the application of the Chernoff bound there applies. This gives that for each $i \in [m]$ and $j \in [n]$,

$$\Pr[E_{ij}|X_j = 1] \leq (2e\alpha_1)^{(W-A_{i,j})/A_{i,j}} = \left(\frac{1}{2e^{2/e}(1+\Delta_1/W)^{1/(W-1)}}\right)^{(W-A_{i,j})/A_{i,j}},$$

where the equality is by choice of α_1 . We prove the final inequality in two parts. First, note that $\frac{W-A_{i,j}}{A_{i,j}} \geq W-1$ since $A_{i,j} \leq 1$. Thus,

$$\left(\frac{1}{2(1+\Delta_1/W)^{1/(W-1)}}\right)^{(W-A_{i,j})/A_{i,j}} \leq \frac{1}{2^{W-1}(1+\Delta_1/W)} \leq \frac{W}{2\Delta_1}.$$

Second, we see that

$$\left(\frac{1}{e^{2/e}}\right)^{(W-A_{i,j})/A_{i,j}} \leq \left(\frac{1}{e^{2/e}}\right)^{W/(2A_{i,j})} \leq \frac{A_{i,j}}{W}$$

for $A_{i,j} \leq 1$, where the first inequality holds because $W \geq 2$ and the second inequality holds by Lemma 5. \square

If we replace Lemma 7 with Lemma 8 in the proof of Theorem 2, we obtain the following stronger guarantee.

Theorem 3 *When setting $\alpha_1 = \frac{1}{c_2(1+\Delta_1/W)^{1/(W-1)}}$ where $c_2 = 4e^{1+2/e}$, for PIPs with width $W \geq 2$, *round-and-alter-by-sorting*(A, b, α_1) is a randomized $(\alpha_1/2)$ -approximation.*

4.3 A $(1 - O(\epsilon))$ -approximation when $W \geq \Omega\left(\frac{1}{\epsilon^2} \ln\left(\frac{\Delta_1}{\epsilon}\right)\right)$

In this section, we give a randomized $(1 - O(\epsilon))$ -approximation for the case when $W \geq \Omega\left(\frac{1}{\epsilon^2} \ln\left(\frac{\Delta_1}{\epsilon}\right)\right)$. We use the algorithm *round-and-alter-by-sorting* in Fig. 2 with

the scaling factor $\alpha_1 = 1 - \epsilon$. The analysis follows the same structure as the analyses for the lemmas bounding the rejection probabilities from the previous sections.

Lemma 9 *Let $0 < \epsilon < \frac{1}{e}$, $\alpha_1 = 1 - \epsilon$, and $W = \frac{2}{\epsilon^2} \ln(\frac{\Delta_1}{\epsilon}) + 1$. Let $i \in [m]$ and $j \in [n]$. Then in round-and-alter-by-sorting(A, b, α_1), we have $\Pr[E_{ij} | X_j = 1] \leq e \cdot \frac{\epsilon A_{i,j}}{\Delta_1}$.*

Proof Renumber indices so that $A_{i,1} \leq \dots \leq A_{i,n}$ and if the index of j changes to j' , we still refer to j' as j . Let $Y_{ij} = \sum_{\ell=1}^{j-1} A_{i,\ell} \xi_\ell$ where $\xi_\ell = 1$ if $x'_\ell = 1$ and 0 otherwise. We first note that

$$\Pr[E_{ij} | X_j = 1] \leq \Pr[Y_{ij} > W - A_{i,j}].$$

By the choice of α_1 and the fact that $A_{i,j} \leq 1$ and $W = \frac{2}{\epsilon^2} \ln(\frac{\Delta_1}{\epsilon}) + 1$, we have $((1 - \alpha_1)W)/A_{i,j} \geq \epsilon W = \frac{2}{\epsilon} \ln(\frac{\Delta_1}{\epsilon}) + \epsilon$. A direct argument via calculus shows $\frac{2}{\epsilon} \ln(\frac{\Delta_1}{\epsilon}) + \epsilon > 1$ for $\epsilon \in (0, \frac{1}{e})$. Thus, $(1 - \alpha_1)W > A_{i,j}$.

By the LP constraints, $\mathbb{E}[Y_{ij}] \leq \alpha_1 W$. Then as $A_{i,j'} \leq A_{i,j}$ for all $j' < j$, we can apply the Chernoff bound in Lemma 3 to obtain

$$\Pr[Y_{ij} \geq W - A_{i,j}] \leq \left(\frac{\alpha_1 e^{1-\alpha_1} W}{W - A_{i,j}} \right)^{(W-A_{i,j})/A_{i,j}}.$$

We bound the right-hand side in two steps. First, as $A_{i,j} \leq 1$,

$$\left(\frac{W}{W - A_{i,j}} \right)^{(W-A_{i,j})/A_{i,j}} \leq \left(\frac{W}{W - 1} \right)^{W-1} \leq e,$$

where the last inequality follows from the fact that $(1 - 1/z)^{z-1} \geq 1/e$ for all $z \geq 1$.

Second, by the choice of α_1 ,

$$(\alpha_1 e^{1-\alpha_1})^{(W-A_{i,j})/A_{i,j}} = ((1 - \epsilon)e^\epsilon)^{(W-A_{i,j})/A_{i,j}}.$$

For $0 < \epsilon < \frac{1}{e}$, we have $1 - \epsilon \leq \exp(-\epsilon - \frac{\epsilon^2}{2})$. As $W = \frac{2}{\epsilon^2} \ln(\frac{\Delta_1}{\epsilon}) + 1$ and $A_{i,j} \leq 1$,

$$((1 - \epsilon)e^\epsilon)^{(W-A_{i,j})/A_{i,j}} \leq \left(e^{-\epsilon^2/2} \right)^{\frac{2}{\epsilon^2} \ln(\frac{\Delta_1}{\epsilon})} \leq \exp\left(-\frac{\ln(\frac{\Delta_1}{\epsilon})}{A_{i,j}} \right).$$

Observe that $\frac{1}{A_{i,j}} - \ln(\frac{e}{A_{i,j}}) \geq 0$. For $A_{i,j} \in [0, 1]$, a direct argument shows $\frac{\ln(t)}{A_{i,j}} - \ln(\frac{t}{A_{i,j}})$ is increasing in t for $t \geq e$. As $\Delta_1/\epsilon > e$, we have $\frac{\ln(\frac{\Delta_1}{\epsilon})}{A_{i,j}} \geq \ln(\frac{\Delta_1}{\epsilon A_{i,j}})$. Therefore,

$$\exp\left(-\frac{\ln(\frac{\Delta_1}{\epsilon})}{A_{i,j}} \right) \leq \exp\left(-\ln\left(\frac{\Delta_1}{\epsilon A_{i,j}} \right) \right) = \frac{\epsilon A_{i,j}}{\Delta_1}.$$

This concludes the proof. \square

Lemma 9 implies that we can upper bound the sum of the rejection probabilities for any item j by $e\epsilon$, leading to the following theorem.

Theorem 4 *Let $0 < \epsilon < \frac{1}{e}$ and $W = \frac{2}{\epsilon^2} \ln(\frac{\Delta_1}{\epsilon}) + 1$. When setting $\alpha_1 = 1 - \epsilon$ and $c = e + 1$, round-and-alter-by-sorting(A, b, α_1) is a randomized $(1 - c\epsilon)$ -approximation algorithm.*

Proof Fix $j \in [n]$. By Lemma 9 and the definition of Δ_1 ,

$$\sum_{i=1}^m \Pr[E_{ij} | X_j = 1] \leq \sum_{i=1}^m \frac{e\epsilon A_{i,j}}{\Delta_1} \leq e\epsilon.$$

By Lemma 1, which shows that an upper bound on the rejection probabilities of γ leads to an $\alpha_1(1 - \gamma)$ -approximation, we have an $\alpha_1(1 - e\epsilon)$ -approximation. Then note that $\alpha_1(1 - e\epsilon) = (1 - \epsilon)(1 - e\epsilon) \geq 1 - (e + 1)\epsilon$. This concludes the proof. \square

5 The small width regime: $W = 1 + \epsilon$

We now consider the regime when the width is small. Let $W = 1 + \epsilon$ for some $\epsilon \in (0, 1]$. We cannot apply the simple sorting based scheme that we used for the large width regime. We borrow the idea from [1] in splitting the coordinates into big and small in each constraint; now the definition is more refined and depends on ϵ . Moreover, the small coordinates and the big coordinates have their own reserved capacity in the constraint. This is crucial for the analysis. We provide more formal details below.

We set α_2 to be $\frac{\epsilon^2}{c_3 \Delta_1}$ where $c_3 = 8e^{1+2/e}$. The alteration step differentiates between “small” and “big” coordinates as follows. For each $i \in [m]$, let $S_i = \{j : A_{i,j} \leq \epsilon/2\}$ and $B_i = \{j : A_{i,j} > \epsilon/2\}$. We say that an index j is *small* for constraint i if $j \in S_i$. Otherwise we say it is *big* for constraint i when $j \in B_i$. For each constraint, the algorithm is allowed to pack a total of $1 + \epsilon$ into that constraint. The algorithm separately packs small indices and big indices. In an ϵ amount of space, small indices that were chosen in the rounding step are sorted in increasing order of size and greedily packed until the constraint is no longer satisfied. The big indices are packed by arbitrarily choosing one and packing it into the remaining space of 1. The rest of the indices are removed to ensure feasibility. Figure 3 gives pseudocode for the randomized algorithm round-alter-small-width which yields an $\Omega(\epsilon^2/\Delta_1)$ -approximation.

It remains to bound the rejection probabilities. Recall that for $j \in [n]$, we define X_j to be the indicator random variable $\mathbb{1}(x'_j = 1)$ and E_{ij} is the event that j was rejected by constraint i .

We first consider the case when index j is big for constraint i . Note that it is possible that there may not exist any big indices for a given constraint. The same holds true for small indices.

```

round-alter-small-width( $A, b, \epsilon, \alpha_2$ ):
  let  $x$  be an optimum fractional solution of the natural LP relaxation
  for  $j \in [n]$ , set  $x'_j$  to be 1 independently with probability  $\alpha_2 x_j$  and 0 otherwise
   $x'' \leftarrow x'$ 
  for  $i \in [m]$  do
    if  $|S_i| = 0$  then
       $s \leftarrow 0$ 
    else
      sort and renumber such that  $A_{i,1} \leq \dots \leq A_{i,n}$ 
       $s \leftarrow \max \left\{ \ell \in S_i : \sum_{j=1}^{\ell} A_{i,j} x'_j \leq \epsilon \right\}$ 
    end if
    if  $|B_i| = 0$ , then  $t = 0$ , otherwise let  $t$  be an arbitrary element of  $B_i$ 
    for each  $j \in [n]$  such that  $j > s$  and  $j \neq t$ , set  $x''_j = 0$ 
  end for
  return  $x''$ 

```

Fig. 3 By setting the scaling factor $\alpha_2 = \frac{\epsilon^2}{c_3 \Delta_1}$ for a sufficiently large constant c , round-alter-small-width is a randomized $\Omega(\epsilon^2/\Delta_1)$ -approximation for PIPs with width $W = 1 + \epsilon$ for some $\epsilon \in (0, 1]$ (see Theorem 5)

Lemma 10 Let $\epsilon \in (0, 1]$ and $\alpha_2 = \frac{\epsilon^2}{c_3 \Delta_1}$ where $c_3 = 8e^{1+2/e}$. Let $i \in [m]$ and $j \in B_i$. Then in round-alter-small-width(A, b, ϵ, α_2), we have $\Pr[E_{ij} | X_j = 1] \leq \frac{A_{i,j}}{2\Delta_1}$.

Proof Let \mathcal{E} be the event that there exists $j' \in B_i$ such that $j' \neq j$ and $X_{j'} = 1$. Observe that if E_{ij} occurs and $X_j = 1$, then it must be the case that at least one other element of B_i was chosen in the rounding step. Thus,

$$\Pr[E_{ij} | X_j = 1] \leq \Pr[\mathcal{E}] \leq \sum_{\substack{\ell \in B_i \\ \ell \neq j}} \Pr[X_\ell = 1] \leq \alpha_2 \sum_{\ell \in B_i} x_\ell,$$

where the second inequality follows by the union bound. Observe that for all $\ell \in B_i$, we have $A_{i,\ell} > \epsilon/2$. By the LP constraints, we have $1 + \epsilon \geq \sum_{\ell \in B_i} A_{i,\ell} x_\ell > \frac{\epsilon}{2} \cdot \sum_{\ell \in B_i} x_\ell$. Thus, $\sum_{\ell \in B_i} x_\ell \leq \frac{1+\epsilon}{\epsilon/2} = 2/\epsilon + 2$.

Using this upper bound for $\sum_{\ell \in B_i} x_\ell$, we have

$$\alpha_2 \sum_{\ell \in B_i} x_\ell \leq \frac{\epsilon^2}{c_3 \Delta_1} \left(\frac{2}{\epsilon} + 2 \right) \leq \frac{4\epsilon}{c_3 \Delta_1} \leq \frac{A_{i,j}}{2\Delta_1},$$

where the second inequality utilizes the fact that $\epsilon \leq 1$ and the third inequality holds because $c_3 \geq 16$ and $A_{i,j} > \epsilon/2$. □

Next we consider the case when index j is small for constraint i . The analysis here is similar to that in the preceding section with width at least 2.

Lemma 11 Let $\epsilon \in (0, 1]$ and $\alpha_2 = \frac{\epsilon^2}{c_3 \Delta_1}$ where $c_3 = 8e^{1+2/e}$. Let $i \in [m]$ and $j \in S_i$. Then in round-alter-small-width(A, b, ϵ, α_2), we have $\Pr[E_{ij} | X_j = 1] \leq \frac{A_{i,j}}{2\Delta_1}$.

Proof Renumbering as in the proof of Lemma 7 and defining $Y_{ij} = \sum_{\ell=1}^{j-1} A_{i,\ell} \cdot \mathbb{1}[x'_\ell = 1]$ in the same manner, we have

$$\Pr[E_{ij}|X_j = 1] \leq \Pr[Y_{ij} \geq \epsilon - A_{i,j}].$$

Let $A'_{i,\ell} = \frac{2}{\epsilon} \cdot A_{i,\ell}$ for $\ell \in [j]$. As $A_{i,\ell} \leq \epsilon/2$ for all $\ell \in [j]$, we have $A'_{i,\ell} \in [0, 1]$. Let $Y'_{ij} = \sum_{\ell=1}^{j-1} A'_{i,\ell} \xi_\ell$. Then

$$\Pr[Y_{ij} \geq \epsilon - A_{i,j}] = \Pr[Y'_{ij} \geq 2 - A'_{i,j}].$$

To upper bound $\mathbb{E}[Y'_{ij}]$, we use the LP constraints and the value of α_2 to see that $\mathbb{E}[Y'_{ij}] \leq \frac{2\epsilon(1+\epsilon)}{c_3\Delta_1}$. Let $\alpha'_2 = \frac{2\epsilon}{c_3\Delta_1}$ and $W = 2$. Then $\mathbb{E}[Y'_{ij}] \leq \alpha'_2 W$. With these parameter choices, we see that $(1 - \alpha'_2)W > A'_{i,j}$. Therefore, as $A'_{i,\ell} \leq A'_{i,j}$ for all $\ell < j$, we can apply the Chernoff bound in Lemma 3 to obtain

$$\Pr[Y'_{ij} \geq 2 - A'_{i,j}] \leq \left(\frac{\alpha'_2 e^{1-\alpha'_2} W}{W - A'_{i,j}} \right)^{(W-A'_{i,j})/A'_{i,j}}.$$

Observe that $e^{1-\alpha'_2} \leq e$ and $\frac{W}{W-A'_{i,j}} \leq 2$ since $W = 2$ and $A'_{i,j} \leq 1$. By our choice of α'_2 ,

$$\left(\frac{\alpha'_2 e^{1-\alpha'_2} W}{W - A'_{i,j}} \right)^{(W-A'_{i,j})/A'_{i,j}} \leq (2e\alpha'_2)^{(W-A'_{i,j})/A'_{i,j}} = \left(\frac{\epsilon}{2e^{2/e}\Delta_1} \right)^{(W-A'_{i,j})/A'_{i,j}}$$

Then note that $\frac{W-A'_{i,j}}{A'_{i,j}} \geq \frac{1}{A'_{i,j}} \geq 1$ since $W = 2$ and $A'_{i,j} \leq 1$. So

$$\left(\frac{\epsilon}{2e^{2/e}\Delta_1} \right)^{(W-A'_{i,j})/A'_{i,j}} = \left(\frac{1}{2\Delta_1} \right)^{(W-A'_{i,j})/A'_{i,j}} \cdot \left(\frac{\epsilon}{e^{2/e}} \right)^{(W-A'_{i,j})/A'_{i,j}} \leq \frac{1}{2\Delta_1} \cdot \frac{\epsilon A'_{i,j}}{2},$$

where the inequality follows by Lemma 6. We have shown $\Pr[E_{ij}|X_j = 1] \leq \frac{\epsilon A'_{i,j}}{4\Delta_1}$. Since $A'_{i,j} = A_{i,j} \cdot \frac{2}{\epsilon}$, the result follows. \square

As Lemma 11 shows that the rejection probability is small, we can prove the following approximation guarantee much like in Theorems 2 and 3.

Theorem 5 Let $\epsilon \in (0, 1]$. When setting $\alpha_2 = \frac{\epsilon^2}{c_3\Delta_1}$ for $c_3 = 8e^{1+2/e}$, for PIPs with width $W = 1 + \epsilon$, round-alter-small-width(A, b, ϵ, α_2) is a randomized $(\alpha_2/2)$ -approximation algorithm.

Proof Fix $j \in [n]$. Then by Lemmas 10 and 11 and the definition of Δ_1 , we have

$$\sum_{i=1}^m \Pr[E_{ij}|X_j = 1] \leq \sum_{i=1}^m \frac{A_{i,j}}{2\Delta_1} \leq \frac{1}{2}.$$

Recall that Lemma 1 gives an $\alpha_2(1 - \gamma)$ -approximation where γ is an upper bound on the sum of the rejection probabilities for any item. This concludes the proof. \square

An upper bound on the integrality gap We showed in Theorem 5 that the integrality gap is $\Omega(\epsilon^2/\Delta_1)$. The example that we used in the proof of Theorem 1 can be easily adapted to show that the gap is $O(\epsilon)$ when $\Delta_1 = O(1)$. It is an interesting open problem to resolve the integrality gap as a function of ϵ .

6 Approximating with a submodular objective

We show that the results from the previous sections can be generalized to the case where the objective function is a nonnegative submodular set function. Recall that a real-valued set function $f : 2^N \rightarrow \mathbb{R}$ over a finite ground set N is submodular iff $f(A) + f(B) \geq f(A \cup B) + f(A \cap B)$ for all $A, B \subseteq N$. Here we are interested in non-negative submodular functions $f : 2^N \rightarrow \mathbb{R}_+$. Also of interest are monotone submodular functions which satisfy the additional property that $f(A) \leq f(B)$ for all $A \subset B$.

The formal problem we are interested in is the following. Let $N = \{1, 2, \dots, n\}$ be the ground set and let $f : 2^N \rightarrow \mathbb{R}_+$ be a nonnegative submodular set function. We assume that one has oracle access to f ; given any set $S \subseteq N$, we can obtain $f(S)$ in constant time. Our goal is to approximate the following problem:

$$\text{maximize } f(S) \text{ over } S \subseteq N \text{ s.t. } \sum_{j \in S} A_{i,j} \leq b_i, \forall i \in [m]. \quad (2)$$

Equivalently, if we let $\mathbf{1}_S \in \{0, 1\}^n$ denote the characteristic vector of a set $S \subseteq N$ then we wish to approximate the problem: $\max f(S)$ s.t. $A\mathbf{1}_S \leq b$.

The rounding algorithms from the previous sections fall under the framework of contention resolution schemes (CR schemes) that allow one to extend the results to submodular objectives via the multilinear relaxation approach. We briefly outline this framework and follow [8].

Multilinear relaxation and CR scheme based rounding. Let N be a finite ground set and $f : 2^N \rightarrow \mathbb{R}_+$ be a non-negative submodular set function. Let $\mathcal{I} \subseteq 2^N$ be a downward closed family of sets² which abstractly models some constraints. We are then interested in the optimization problem $\max_{S \in \mathcal{I}} f(S)$. The multilinear relaxation approach for approximating this problem is to solve a continuous optimization problem $\max_{x \in P_{\mathcal{I}}} F(x)$ and then round it. Here $P_{\mathcal{I}} \supseteq \text{conv}\{\mathbf{1}_S : S \in \mathcal{I}\}$ is a convex set that serves as a relaxation for the constraint set, and F is a continuous extension of f to $[0, 1]^n$. Specifically, $F : [0, 1]^n \rightarrow \mathbb{R}_+$ is the multilinear extension of f and is defined as

$$F(x) = \sum_{S \subseteq N} f(S) \prod_{i \in S} x_i \prod_{j \notin S} (1 - x_j).$$

² We say a family of subsets $\mathcal{I} \subseteq 2^N$ is downward closed if for all $A \subseteq B \subseteq N$, if $B \in \mathcal{I}$, then $A \in \mathcal{I}$.

We say that $P_{\mathcal{I}}$ is solvable if one can efficiently optimize linear functions over it. For a scalar $\alpha \in (0, 1]$ we use $\alpha P_{\mathcal{I}}$ to denote the set $\{\alpha x \mid x \in P_{\mathcal{I}}\}$. Let OPT denote the value of the relaxation $\max_{x \in P_{\mathcal{I}}} F(x)$ which provides an upper bound on the optimum integral solution. Finding OPT is NP -Hard even for the simple cardinality constraint, however, randomized constant factor approximations are known whenever $P_{\mathcal{I}}$ is solvable and f is a nonnegative submodular function. In particular for any $\alpha \in (0, 1]$ one can obtain a point $x \in \alpha P_{\mathcal{I}}$ such that $F(x) \geq (1 - 1/e^\alpha)\text{OPT}$ when f is monotone [22] and such that $F(x) \geq \alpha e^{-\alpha}\text{OPT}$ when f is non-negative [11]³.

The second step in devising an algorithm is to round the fractional solution. We focus on a particular strategy based on CR schemes. It is motivated by the definition of $F(x)$ as the expected value $\mathbb{E}[f(R(x))]$, where $R(x)$ is a random set obtained by independently picking each $i \in N$ with probability x_i . Thus randomly rounding x preserves the objective $F(x)$ in expectation. However the resulting set $R(x)$ can violate the constraints. Thus we would like to alter $R(x)$ to a feasible set R' while not losing too much in the objective. We would also like to scale x down by an α factor for some $\alpha \in (0, 1]$ and work with $R(\alpha x)$ since this can be useful in the alteration step as we have seen for PIPs.

In [8] a formal and abstract definition of alteration schemes called CR schemes was provided, and it was shown that these schemes can be used in conjunction with the multilinear relaxation. In this setting an alteration scheme for $P_{\mathcal{I}}$ is viewed as a (potentially randomized) algorithm \mathcal{A} that takes as input $x \in P_{\mathcal{I}}$ and a set $B \subseteq N$ satisfying $B \subseteq \text{support}(x)$ ⁴ and outputs a feasible set $\mathcal{A}(B, x) \in \mathcal{I}$.

Definition 1 Let $\alpha, \beta \in (0, 1]$. We say that an algorithm \mathcal{A} is a (randomized) (α, β) -balanced CR scheme for $P_{\mathcal{I}}$ if \mathcal{A} returns a (random) set $\mathcal{A}(B, x)$ for all inputs $x \in \alpha P_{\mathcal{I}}$ and $B \subseteq N$ and satisfies the following properties:

1. for all $B \subseteq N$, with probability 1, $\mathcal{A}(B, x) \subseteq B \cap \text{support}(x)$ and $\mathcal{A}(B, x) \in \mathcal{I}$.
2. for all $i \in \text{support}(x)$, $\Pr[i \in \mathcal{A}(B, x) \mid i \in B] \geq \beta$.
3. for all $B_1 \subseteq B_2 \subseteq N$ and $i \in B_1$, $\Pr[i \in \mathcal{A}(B_1, x)] \geq \Pr[i \in \mathcal{A}(B_2, x)]$.

The first property guarantees that the output of \mathcal{A} is a feasible set. The second property gives a lower bound on the probability of an element being in the output conditioned on it being chosen in the first randomized rounding step. The last property requires that the alteration scheme is monotone from each element's perspective.

The following is a paraphrased version of the results in [8] that combines the algorithms for solving the multilinear relaxation followed by rounding with a CR scheme.

Theorem 6 ([8]) Consider the problem $\max_{S \in \mathcal{I}} f(S)$ for a non-negative submodular function $f : 2^N \rightarrow \mathbb{R}_+$ and its multilinear relaxation $\max_{x \in P_{\mathcal{I}}} F(x)$. Combining an $(\alpha e^{-\alpha})$ -approximation to $\max_{x \in \alpha P_{\mathcal{I}}} F(x)$ and an (α, β) -balanced CR scheme, one can obtain a randomized $(\alpha e^{-\alpha} \beta)$ -approximation to $\max_{S \in \mathcal{I}} f(S)$. The approximation ratio improves to $(1 - \frac{1}{e^\alpha})\beta$ when f is additionally monotone.

³ For non-negative functions there have been subsequent improvements in the approximation ratio [2,10] but the dependence on α is unclear and since the precise approximation ratios are not the focus in this paper, we confine ourselves to the simpler algorithm and bound from [11].

⁴ For $x \in [0, 1]^n$ we use $\text{support}(x)$ to denote the set $\{i \mid x_i > 0\}$.

Applying to packing constraints. We now apply the preceding theorem to the setting of packing constraints formalized earlier as (2). We have $\mathcal{I} = \{S \subseteq N \mid A\mathbf{1}_S \leq b\}$ and we use the natural LP relaxation $P_{\mathcal{I}} = \{x \in [0, 1]^n \mid Ax \leq b\}$. Clearly $P_{\mathcal{I}}$ is solvable via linear programming.

We now interpret the rounding algorithms in the preceding sections as CR schemes. In terms of the language used in the previous sections, α is the scaling factor used in obtaining the rounded solution and $1 - \beta$ is the upper bound on the rejection probability of every item. We note that the alteration schemes are deterministic.

Let $\text{round-and-alter-by-sorting-sub}(A, b, \alpha_1, x, B)$ be the same algorithm as $\text{round-and-alter-by-sorting}$ but now $\text{round-and-alter-by-sorting-sub}$ is given a fractional solution x to $\max_{x \in \alpha_1 P_{\mathcal{I}}} F(x)$ and B will be the rounded solution $R(\alpha_1 x)$. We make the same changes to $\text{round-alter-small-width}$ to obtain the algorithm $\text{round-alter-small-width-sub}(A, b, \epsilon, \alpha_2, x, B)$.

- Lemma 12** 1. For $W \geq 2$, $\text{round-and-alter-by-sorting-sub}(A, b, \alpha_1, x, B)$ is an $(\alpha_1, \frac{1}{2})$ -balanced CR scheme where $\alpha_1 = \frac{1}{c_2(1+\Delta_1/W)^{1/(W-1)}}$ and $c_2 = 4e^{1+2/e}$.
 2. For $W = 1 + \epsilon$, $\text{round-alter-small-width-sub}(A, b, \epsilon, \alpha_2, x, B)$ is an $(\alpha_2, \frac{1}{2})$ -balanced CR scheme where $\alpha_2 = \frac{\epsilon^2}{c_3\Delta_1}$ and $c_3 = 8e^{1+2/e}$.

Proof We only prove (1) as the proof for (2) is similar. $\text{round-alter-small-width-sub}$ takes an integral solution and guarantees feasibility by satisfying each constraint individually by setting variables to 0 if necessary, so (1) of Definition 1 is satisfied. Let E_{ij} be the event that item j is rejected by constraint i . By Lemma 8, the rejection probability of item j is at most $\sum_{i=1}^m \Pr[E_{ij} \mid x'_j = 1] \leq \sum_{i=1}^m \frac{A_{i,j}}{2\Delta_1} \leq \frac{1}{2}$. In the notation of Definition 1, this implies (2) of Definition 1 is satisfied with $\beta = 1/2$. (3) of Definition 1 is also satisfied as the probability that an item is rejected only decreases if less items are chosen in the rounding step.

Combining Lemma 12 and Theorem 6, we immediately get the following result.

Theorem 7 Let $\mathcal{I} = \{S \subseteq N \mid A\mathbf{1}_S \leq b\}$ and let $f : 2^N \rightarrow \mathbb{R}_+$ be a nonnegative submodular set function.

1. Assume $W \geq 2$. Let $\alpha_1 = \frac{1}{c_2(1+\Delta_1/W)^{1/(W-1)}}$ where $c_2 = 4e^{1+2/e}$. There exists an $((\alpha_1 e^{-\alpha_1}) \frac{1}{2})$ -approximation to $\max_{S \in \mathcal{I}} f(S)$. Assuming f is also monotone, there exists a $((1 - \frac{1}{e^{\alpha_1}}) \frac{1}{2})$ -approximation.
2. Let $\epsilon \in (0, 1]$ and assume $W = 1 + \epsilon$. Let $\alpha_2 = \frac{\epsilon^2}{c_3\Delta_1}$ where $c_3 = 8e^{1+2/e}$. There exists an $(\alpha_2 e^{-\alpha_2}) \frac{1}{2}$ -approximation to $\max_{S \in \mathcal{I}} f(S)$. Assuming f is also monotone, there exists a $((1 - \frac{1}{e^{\alpha_2}}) \frac{1}{2})$ -approximation.

We close this section with two remarks. First, CR schemes for different classes of constraints can be composed gracefully [8] and hence the ones here could be useful in conjunction with schemes for other constraints. Second, certain dependent rounding schemes for matroid and matroid intersection constraints satisfy concentration bounds, similar to Chernoff bounds, for non-negative sums; since our analysis for PIPs relied essentially only on Chernoff bounds, one can extend the analysis even under dependent

rounding. This allows for some applications when one combines packing constraints with matroid or matroid intersection type constraints. We refer the reader to [6,7] for more details.


References

1. Bansal, N., Korula, N., Nagarajan, V., Srinivasan, A.: Solving packing integer programs via randomized rounding with alterations. *Theory Comput.* **8**(24), 533–565 (2012). <https://doi.org/10.4086/toc.2012.v008a024>
2. Buchbinder, N., Feldman, M.: Constrained submodular maximization via a nonsymmetric technique. *Math. Oper. Res.* **44**(3), 988–1005 (2019)
3. Chan, T.M.: Approximation schemes for 0–1 knapsack. In: 1st Symposium on Simplicity in Algorithms (2018)
4. Chekuri, C., Khanna, S.: On multidimensional packing problems. *SIAM J. Comput.* **33**(4), 837–851 (2004)
5. Chekuri, C., Quanrud, K.: On approximating (sparse) covering integer programs. In: Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 1596–1615. SIAM (2019)
6. Chekuri, C., Vondrak, J., Zenklusen, R.: Dependent randomized rounding via exchange properties of combinatorial structures. In: 2010 IEEE 51st Annual Symposium on Foundations of Computer Science, pp. 575–584. IEEE (2010)
7. Chekuri, C., Vondrák, J., Zenklusen, R.: Multi-budgeted matchings and matroid intersection via dependent rounding. In: Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 1080–1097. SIAM (2011)
8. Chekuri, C., Vondrák, J., Zenklusen, R.: Submodular function maximization via the multilinear relaxation and contention resolution schemes. *SIAM J. Comput.* **43**(6), 1831–1879 (2014)
9. Chen, A., Harris, D.G., Srinivasan, A.: Partial resampling to approximate covering integer programs. In: Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 1984–2003. SIAM (2016)
10. Ene, A., Nguyen, H.L.: Constrained submodular maximization: beyond $1/e$. In: 2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS), pp. 248–257. IEEE (2016)
11. Feldman, M., Naor, J., Schwartz, R.: A unified continuous greedy algorithm for submodular maximization. In: 2011 IEEE 52nd Annual Symposium on Foundations of Computer Science, pp. 570–579. IEEE (2011)
12. Frieze, A., Clarke, M.: Approximation algorithms for the m -dimensional 0–1 knapsack problem: worst-case and probabilistic analyses. *Eur. J. Oper. Res.* **15**(1), 100–109 (1984)
13. Harvey, N.J.: A note on the discrepancy of matrices with bounded row and column sums. *Discrete Math.* **338**(4), 517–521 (2015)
14. Håstad, J.: Clique is hard to approximate within $n^{1-\epsilon}$. *Acta Math.* **182**(1), 105–142 (1999)
15. Hazan, E., Safra, S., Schwartz, O.: On the complexity of approximating k -set packing. *Comput. Complex.* **15**(1), 20–39 (2006)
16. Mitzenmacher, M., Upfal, E.: *Probability and Computing: Randomized Algorithms and Probabilistic Analysis*. Cambridge University Press, Cambridge (2005)
17. Pritchard, D.: Approximability of sparse integer programs. In: European Symposium on Algorithms, pp. 83–94. Springer (2009)
18. Pritchard, D., Chakrabarty, D.: Approximability of sparse integer programs. *Algorithmica* **61**(1), 75–93 (2011)
19. Raghavan, P., Thompson, C.D.: Randomized rounding: a technique for provably good algorithms and algorithmic proofs. *Combinatorica* **7**(4), 365–374 (1987)
20. Srinivasan, A.: Improved approximation guarantees for packing and covering integer programs. *SIAM J. Comput.* **29**(2), 648–670 (1999)
21. Trevisan, L.: Non-approximability results for optimization problems on bounded degree instances. In: Proceedings of the Thirty-Third Annual ACM Symposium on Theory of Computing, pp. 453–461. ACM (2001)

22. Vondrák, J.: Optimal approximation for the submodular welfare problem in the value oracle model. In: Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing, pp. 67–74. ACM (2008)
23. Zuckerman, D.: Linear degree extractors and the inapproximability of max clique and chromatic number. In: Proceedings of the Thirty-Eighth Annual ACM Symposium on Theory of Computing, pp. 681–690. ACM (2006)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Affiliations

Chandra Chekuri¹ · Kent Quanrud² · Manuel R. Torres³ 

Chandra Chekuri
chekuri@illinois.edu

Kent Quanrud
krq@purdue.edu

¹ Department of Computer Science, University of Illinois at Urbana-Champaign, 3228 Siebel Center, 201 N. Goodwin Ave, Urbana, IL, USA

² Department of Computer Science, Purdue University, 305 North University Street, Room #1211, West Lafayette, IN 47907, USA

³ Department of Computer Science, University of Illinois at Urbana-Champaign, 3219 Siebel Center, 201 N. Goodwin Ave, Urbana, IL, USA