

Approximation Algorithms for Integer Programming with Resource Augmentation

Hauke Brinkop ✉ 

Department of Computer Science, Kiel University, Germany

Hua Chen ✉ 

School of Mathematical Sciences, Zhejiang University of Technology, Hangzhou, China

Lin Chen ✉ 

College of Computer Science and Technology, Zhejiang University, Hangzhou, China

Klaus Jansen ✉ 

Department of Computer Science, Kiel University, Germany

Guochuan Zhang ✉ 

College of Computer Science and Technology, Zhejiang University, Hangzhou, China

Abstract

Solving a general integer program (IP) is NP-hard. The classic algorithm [Papadimitriou, J.ACM '81] for IPs has a running time $n^{\mathcal{O}(m)}(m \cdot \max\{\Delta, \|\mathbf{b}\|_\infty\})^{\mathcal{O}(m^2)}$, where m is the number of constraints, n is the number of variables, and Δ and $\|\mathbf{b}\|_\infty$ are, respectively, the largest absolute values among the entries in the constraint matrix and the right-hand side vector of the constraint. The running time is exponential in m , and becomes pseudo-polynomial if m is a constant. In recent years, there has been extensive research on FPT (fixed parameter tractable) algorithms for the so-called n -fold IPs, which may possess a large number of constraints, but the constraint matrix satisfies a specific block structure. It is remarkable that these FPT algorithms take as parameters Δ and the number of rows and columns of some small submatrices. If Δ is not treated as a parameter, then the running time becomes pseudo-polynomial even if all the other parameters are taken as constants.

This paper explores the trade-off between time and accuracy in solving an IP. We show that, for arbitrary small $\varepsilon > 0$, there exists an algorithm for IPs with m constraints that runs in $f(m, \varepsilon) \cdot \text{poly}(|I|)$ time, and returns a near-feasible solution that violates the constraints by at most $\varepsilon\Delta$. Furthermore, for n -fold IPs, we establish a similar result – our algorithm runs in time that depends on the number of rows and columns of small submatrices together with $1/\varepsilon$, and returns a solution that slightly violates the constraints. Meanwhile, both solutions guarantee that their objective values are no worse than the corresponding optimal objective values satisfying the constraints. As applications, our results can be used to obtain additive approximation schemes for multidimensional knapsack as well as scheduling.

2012 ACM Subject Classification Theory of computation → Design and analysis of algorithms

Keywords and phrases Approximation algorithms, Resource augmentation, Integer programs, n -fold IPs

Digital Object Identifier 10.4230/LIPIcs.STACS.2026.20

Related Version *Full Version:* <http://arxiv.org/abs/2512.24302>

Funding *Hauke Brinkop:* Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) (No. 335406402).

Hua Chen: Research supported in part by NSFC (No. 12401423).

Lin Chen: Research supported in part by NSFC (No. 62572428).

Klaus Jansen: Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) (No. 528381760).

Guochuan Zhang: Research supported in part by NSFC (No. 12271477).



© Hauke Brinkop, Hua Chen, Lin Chen, Klaus Jansen, and Guochuan Zhang; licensed under Creative Commons License CC-BY 4.0

43rd International Symposium on Theoretical Aspects of Computer Science (STACS 2026).

Editors: Meena Mahajan, Florin Manea, Annabelle McIver, and Nguyễn Kim Thăng

Article No. 20; pp. 20:1–20:20



Leibniz International Proceedings in Informatics

LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



1 Introduction

Integer programs (IPs) find widespread applications in many combinatorial optimization problems. Generally, an integer program can be formulated as follows:

$$\min \{ \mathbf{w}\mathbf{x} : H\mathbf{x} = \mathbf{b}, \ell \leq \mathbf{x} \leq \mathbf{u}, \mathbf{x} \in \mathbb{Z}^n \}, \quad (1)$$

where $H \in \mathbb{Z}^{m \times n}$ is the constraint matrix, $\mathbf{b} \in \mathbb{Z}^m$ is the vector on the right-hand side of the equality constraint, $\mathbf{w} \in \mathbb{Z}^n$ is the coefficient vector of the objective function, and $\ell, \mathbf{u} \in \mathbb{Z}^n$ are the lower and upper bounds of the variables, respectively. Karp [53] showed that solving IP (1) is NP-hard in 1972. Subsequent studies thus focus on finding efficient algorithms for IPs where the constraint matrix has some specific structure.

Lenstra [51] studied IPs with few variables and presented a $2^{\mathcal{O}(n^3)} \cdot \text{poly}(|I|)$ -time algorithm, where $\text{poly}(|I|)$ denotes a polynomial in the input length $|I|$. The running time has been improved in several subsequent works [20, 23, 52, 72]. Very recently, Reis and Rothvoss [72] proved that IP (1) can be solved in $(\log n)^{\mathcal{O}(n)} \cdot \text{poly}(|I|)$ -time.

Papadimitriou [70] considered IPs with few constraints and provided an algorithm running in $n^{\mathcal{O}(m)} (m \cdot \max\{\Delta, \|\mathbf{b}\|_\infty\})^{\mathcal{O}(m^2)}$ -time, where $\Delta := \|H\|_\infty$ represents the largest absolute value among all entries of H , and $\|\mathbf{b}\|_\infty$ represents the largest absolute value among all entries of \mathbf{b} . The running time was later improved by Eisenbrand and Weismantel [32], and Jansen and Rohwedder [49]. In particular, for the special case when $\ell = \mathbf{0}$ and $\mathbf{u} = \infty$, IP (1) can be solved in $\mathcal{O}(\sqrt{m}\Delta)^{2m} \cdot \log(\|\mathbf{b}\|_\infty) + \mathcal{O}(nm)$ time [49], and lower bounds on the running times have also been studied by [34, 49, 59].

The results mentioned above assume that the constraint matrix H has either few rows or few columns. There has also been extensive research on IPs where the number of rows and columns of H can be large but H possesses some block structure; see, e.g., [3, 18, 19, 24, 30, 37, 39, 43, 54, 60]. One most widely studied structure is the so-called n -fold, where H consists of small submatrices A^i, D^i of the following form

$$H := \begin{pmatrix} D^1 & D^2 & \dots & D^n \\ A^1 & 0 & & 0 \\ 0 & A^2 & & 0 \\ \vdots & & \ddots & \\ 0 & 0 & & A^n \end{pmatrix},$$

where $A^i \in \mathbb{Z}^{s_A \times t}$, and $D^i \in \mathbb{Z}^{s_D \times t}$ for all i . Many parameterized algorithms for n -fold IPs have been developed [1, 17, 19, 29, 30, 37, 44, 47]. In particular, Cslovjecsek et al. [17] presented an algorithm of running time $2^{\mathcal{O}(s_A^2 s_D)} (s_D s_A \Delta)^{\mathcal{O}(s_A^2 + s_A s_D)} \cdot (nt)^{1+o(1)}$, where $\Delta = \|H\|_\infty$, which is currently the best known FPT algorithm parameterized by Δ, s_A, s_D . Here, an FPT algorithm refers to an algorithm with running time $g(k) \cdot \text{poly}(|I|)$, where g is a computable function, $|I|$ is the input length, and k is a parameter independent of $|I|$. Note that their approach relies on proximity, which is (necessarily) dependent on Δ and thus not usable for our use-case; if there was a proximity bound of the form $f(s_A, s_D) \cdot \text{poly}(|I|)$, this would imply an n -fold algorithm (using the approach in [32]) that is FPT with parameters s_A, s_D . Finally, this would imply a polynomial time algorithm for subset sum [62].

We see that the running time is either exponential in the number of variables or dependent on Δ and is thus pseudo-polynomial. Such a situation is unavoidable if we are looking for an exact algorithm that solves IP (1) optimally, since all algorithms for bounded number of constraints must also depend on Δ as it could otherwise solve subset sum in polynomial time. However, a polynomial time algorithm may be possible if we are allowed to solve IP (1) approximately in the sense that the constraints can be violated slightly.

One most well-known example is Knapsack, where H contains one row, i.e., $m = 1$, $\ell = \mathbf{0}$, $\mathbf{u} = \mathbf{1}$, and all entries of \mathbf{w} and H are nonnegative¹. Knapsack admits an FPTAS whose running time is polynomial in the input length $|I|$ and $1/\varepsilon$ instead of Δ (see, e.g., [10, 15, 25, 40, 50, 66]). Alternatively, it admits an algorithm of the same running time that returns a solution whose objective value is at least OPT (where OPT denotes the optimal objective value satisfying the constraint), while violating the constraint by a factor of $1 + \varepsilon$.

Another example is due to Dadush [21], who presented an algorithm of running time $2^{\mathcal{O}(n)}(1/\varepsilon^2)^n$ which either asserts that the convex body K described by the constraints does not contain any integer points, or finds an integer point in the body stemming from K scaled by $1 + \varepsilon$ from its center of gravity.

Motivated by the above two examples, this paper aims to explore approximation algorithms for IPs. We are interested in an algorithm that returns a near-feasible solution that may violate the constraints by a factor of $1 + \varepsilon$, but runs in polynomial time in the input. Such relaxation of constraints is also referred to as *resource augmentation*, which has received much attention in various combinatorial optimization problems, including Knapsack [6, 42], Subset Sum [14, 67], Square Packing [33], etc.

1.1 Our contribution

Our first result is an approximation scheme for IPs with few constraints.

► **Theorem 1.** *Given is an integer program $\min\{\mathbf{w}\mathbf{x} : H\mathbf{x} = \mathbf{b}, \ell \leq \mathbf{x} \leq \mathbf{u}, \mathbf{x} \in \mathbb{Z}^n\}$ with optimal objective value OPT , where $H \in \mathbb{Q}^{m \times n}$. Then for arbitrary small $\varepsilon > 0$, there exists an algorithm of running time $f(m, \varepsilon) \cdot \text{poly}(|I|)$ which returns a near-feasible solution $\tilde{\mathbf{x}}$ such that $\tilde{\mathbf{x}} \in \{\mathbf{x} : \|H\mathbf{x} - \mathbf{b}\|_\infty \leq \varepsilon\Delta, \ell \leq \mathbf{x} \leq \mathbf{u}, \mathbf{x} \in \mathbb{Z}^n\}$, and $\mathbf{w}\tilde{\mathbf{x}} \leq \text{OPT}$, where $\Delta = \|H\|_\infty$.*

In particular, the precise running time in Theorem 1 is $2^{(\frac{m}{\varepsilon})^{\mathcal{O}(m)}} \cdot \text{poly}(|I|)$. We remark that OPT always refers to the optimal objective value **without** violating the constraints.

Theorem 1 extends Papadimitriou's algorithm in the direction of approximation. When $\varepsilon < 1/\Delta$, for example, $\varepsilon = 1/(2\Delta)$, Theorem 1 implies an exact FPT algorithm to solve IP (1) with running time $2^{(m\Delta)^{\mathcal{O}(m)}} \cdot \text{poly}(|I|)$. A special case of IP (1) is the classic multidimensional knapsack problem, where all the entries in $\mathbf{w}, H, \ell, \mathbf{u}$ are nonnegative. Thus Theorem 1 gives an EPTAS for the multidimensional knapsack problem with resource augmentation (i.e., by slightly relaxing the multidimensional knapsack constraints). Approximation schemes with resource augmentation (or weak approximation schemes) for the knapsack problem have been studied extensively in recent years [6, 16, 67]. Note that if the constraints cannot be relaxed, then there exists a PTAS of running time $O(n^{\lceil m/\varepsilon \rceil})$ for m -dimensional knapsack [8], and such a running time is essentially the best possible assuming ETH (Exponential Time Hypothesis) [26]. In particular, 2-dimensional knapsack does not admit an EPTAS assuming $\text{W}[1] \neq \text{FPT}$ [61]. Moreover, it does not admit a PTAS of running time $n^{o(1/\varepsilon)}$ assuming ETH [45].

We then extend our result to n -fold IPs and obtain the following.

► **Theorem 2.** *Given is an integer program $\min\{\mathbf{w}\mathbf{x} : \sum_{i=1}^n D^i \mathbf{x}^i = \mathbf{b}^0, A^i \mathbf{x}^i = \mathbf{b}^i, 1 \leq i \leq n, \mathbf{0} \leq \mathbf{x} \leq \mathbf{u}, \mathbf{x} \in \mathbb{Z}^{nt}\}$ with optimal objective value OPT , where $A^i \in \mathbb{Q}_{\geq 0}^{s_A \times t_A}$, $D^i \in \mathbb{Q}_{\geq 0}^{s_D \times t_D}$, and $t_A = t_D = t$. Then for arbitrary small $\varepsilon > 0$, there exists an algorithm of running time $f(s_A, s_D, t, \varepsilon) \cdot \text{poly}(|I|)$ which returns a near-feasible solution $\tilde{\mathbf{x}}$ such that $\tilde{\mathbf{x}} \in \{\mathbf{x} : (1 - \varepsilon)\mathbf{b}^0 \leq \sum_{i=1}^n D^i \mathbf{x}^i \leq (1 + \varepsilon)\mathbf{b}^0, (1 - \varepsilon)\mathbf{b}^i \leq A^i \mathbf{x}^i \leq (1 + \varepsilon)\mathbf{b}^i, 1 \leq i \leq n, \mathbf{0} \leq \mathbf{x} \leq \mathbf{u}, \mathbf{x} \in \mathbb{Z}^{nt}\}$, and $\mathbf{w}\tilde{\mathbf{x}} \leq \text{OPT}$.*

¹ Knapsack requires $H\mathbf{x} \leq \mathbf{b}$ instead of $H\mathbf{x} = \mathbf{b}$, but we can add dummy items of 0 profit and unit weight to enforce the equation.

In particular, the running time is $2^{2^{(s_D/\varepsilon)^{\mathcal{O}(s_D)}} \cdot (s_A t/\varepsilon)^{\mathcal{O}(t)}} \cdot \text{poly}(|I|)$.

Theorem 2 is a bit weak in the sense that it requires all the entries of the submatrices A^i, D^i as well as the variables to be nonnegative, and it gives an additive error of $\varepsilon \mathbf{b}^i$ instead of $\varepsilon \Delta$. In many combinatorial optimization problems where n -fold IPs are applicable, all the input parameters are indeed nonnegative and $(1 + \varepsilon) \mathbf{b}^i$ yields a standard multiplicative $(1 + \varepsilon)$ -factor. Nevertheless, we show that it is possible to get rid of such a weakness if the local constraint $A^i \mathbf{x}^i = \mathbf{b}^i$ has few solutions, as implied by the following Theorem 3.

► **Theorem 3.** *Given is an integer program $\min\{\mathbf{w}\mathbf{x} : \sum_{i=1}^n D^i \mathbf{x}^i = \mathbf{b}^0, \mathbf{x}^i \in \mathcal{P}^i, 1 \leq i \leq n, \mathbf{x} \in \mathbb{Z}^{nt}\}$ with optimal objective value OPT, where $D^i \in \mathbb{Q}^{s \times t}$ and \mathcal{P}^i is an arbitrary set of integer vectors. Let $\kappa = \max_{\mathbf{u} \in \cup_{i=1}^n \mathcal{P}^i} \|\mathbf{u}\|_\infty$, i.e., κ is the largest ℓ_∞ -norm among all integer vectors in $\cup_{i=1}^n \mathcal{P}^i$. Then for arbitrary small $\varepsilon > 0$, there exists an algorithm of running time $f(s, t, \kappa, \varepsilon) \cdot \text{poly}(|I|)$ which returns a near-feasible solution $\tilde{\mathbf{x}}$ such that $\tilde{\mathbf{x}} \in \{\mathbf{x} : \|\sum_{i=1}^n D^i \mathbf{x}^i - \mathbf{b}^0\|_\infty \leq \varepsilon \Delta, \mathbf{x}^i \in \mathcal{P}^i, 1 \leq i \leq n, \mathbf{x} \in \mathbb{Z}^{nt}\}$, and $\mathbf{w}\tilde{\mathbf{x}} \leq \text{OPT}$, where $\Delta = \max_{i \in [n]} \|D^i\|_\infty$.*

In particular, the running time is $2^{(\frac{s\kappa t}{\varepsilon})^{\mathcal{O}(s\kappa t)}} \cdot \text{poly}(|I|)$.

Theorem 2 and Theorem 3 can be viewed as FPT approximation schemes for n -fold IPs with different parameters. One application is to provide additive approximation schemes for the unrelated machine scheduling problem $Rm|C_{\max}$: Given are n jobs and m machines, with the processing time of job i on machine h as p_{ih} . The goal is to assign jobs to machines such that the makespan is at most some target value C_{\max} . Let $x_h^i \in \{0, 1\}$ denote whether job i is assigned to machine h , i.e., if job i is assigned to machine h , then $x_h^i = 1$; otherwise, $x_h^i = 0$. Then the natural IP formulation is the following feasibility test:

$$\left\{ \mathbf{x} : \sum_{i=1}^n p_{ih} x_h^i \leq C_{\max}, \forall 1 \leq h \leq m, \mathbf{x}^i \in \mathcal{P}^i, 1 \leq i \leq n, \mathbf{x} \in \mathbb{Z}^{nt} \right\},$$

where polytope $\mathcal{P}^i := \{\mathbf{x}^i : x_1^i + x_2^i + \dots + x_m^i = 1, x_h^i \in \{0, 1\}\}$ implies that exactly one x_h^i is 1. Observe that $D^i \in \mathbb{Z}^{m \times m}$ and $\kappa = 1$, so Theorem 3 implies a PTAS of running time $f(m, \varepsilon) \cdot n^{\mathcal{O}(1)}$ for finding a feasible schedule $\tilde{\mathbf{x}}$ such that the makespan is at most $C_{\max} + \varepsilon \cdot \max_{i \in [n], h \in [m]} p_{ih}$. In comparison, the current best-known PTAS for unrelated machine scheduling runs in $2^{\mathcal{O}(m \log(m/\delta))}$ -time [48], and returns a solution with makespan $(1 + \delta)C_{\max}$. Using this algorithm to derive an additive PTAS would require setting $\delta = \mathcal{O}(\varepsilon/n)$, yielding a non-polynomial running time.

Theorem 3 can also be used to obtain PTASes for other variants of scheduling problems. Consider the generalized assignment problem, where the input is the same as the above-mentioned unrelated machine scheduling problem, except that there is a cost c_{ih} for scheduling job i on machine h . The goal is a bi-criteria study on the makespan C_{\max} and the total scheduling cost T . Taking the total scheduling cost as a constraint, a classic result of Shmoys and Tardos [74] gave a 2-approximation algorithm. Angel et al. [2] proposed a PTAS for $Rm|c_{ih}|C_{\max}$ of running time $\mathcal{O}(n(n/\varepsilon)^m)$. Very recently Li et al. [64] showed a PTAS for identical machine scheduling $Pm|c_{ih}|C_{\max}$ of an improved running time $n(m/\varepsilon)^{\mathcal{O}(m)}$, which is an EPTAS when taking m as a parameter. Using Theorem 3, an EPTAS for unrelated machines $Rm|c_{ih}|C_{\max}$ of running time $2^{(\frac{m}{\varepsilon})^{\mathcal{O}(m)}} \cdot n^{\mathcal{O}(1)}$ directly follows. It is worth mentioning that researchers have studied the bi-objective for various other models (e.g., different machine environments, specific cost functions like $c_{ih} = a_i \cdot b_h$ or $c_{ih} = a_i + b_h$, etc.) [63, 65], and we believe Theorem 3 may also be applicable to some of them.

1.2 Technical overview

We first briefly discuss our technique for proving Theorem 1. Let $H = (\mathbf{h}_1, \mathbf{h}_2, \dots, \mathbf{h}_n)$, whereas $H\mathbf{x} = \sum_{j=1}^n \mathbf{h}_j x_j$. Observe that $\mathbf{h}_j \in [-\Delta, \Delta]^m$. If n is small, say $n = m^{\mathcal{O}(1)}$, then Lenstra’s algorithm [51] works directly and solves the IP in $n^{\mathcal{O}(n)} = 2^{m^{\mathcal{O}(1)}}$ time, concluding Theorem 1. So we may assume that n is sufficiently large, which means there are sufficiently many \mathbf{h}_j ’s within the box $[-\Delta, \Delta]^m$. Consequently, many \mathbf{h}_j ’s are very similar to each other in the sense that their difference lies in a small hypercube of side length $\mathcal{O}(\varepsilon\Delta)$. Now suppose \mathbf{h}_1 and \mathbf{h}_2 are similar. To approximately solve IP (1), we observe that it is not necessary to keep two integer variables x_1 and x_2 for \mathbf{h}_1 and \mathbf{h}_2 respectively. Instead, we may introduce a hyper-integer variable y that represents $x_1 + x_2$, and relax x_1 and x_2 to fractional variables. By doing so, we transform the original IP to a mixed integer program (MIP) with a constant number of integer variables, which is solvable in polynomial time. Then we round fractional variables to integer variables by maintaining the overall error incurred. The challenge here is that, to bound the overall error, we need to make sure that among all the fractional variables, only a few of them can take fractional values. Interestingly, while these fractional variables need to satisfy a large number of constraints, the specific structure of the constraint matrix allows us to conclude that very few variables may take a fractional value in a vertex solution. Towards showing this, we introduce some necessary notions and lemmas regarding vertex solutions in preliminaries.

The proof strategy for Theorem 2 can be viewed as a combination of that for Theorem 1 and Theorem 3, so we briefly discuss our technique for proving Theorem 3. For simplicity, let us assume that $\mathcal{P}^i = \mathcal{P} = \{\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_\tau\}$ for all i , where $\tau := \max_{i \in [n]} |\mathcal{P}^i| = \kappa^t$. Hence, we see that the vector $D^i \mathbf{x}^i$ is essentially one column out of the precomputed matrix $\mathcal{D}^i := \{D^i \mathbf{p}_1, D^i \mathbf{p}_2, \dots, D^i \mathbf{p}_\tau\}$. If D^i and D^j (and hence \mathcal{D}^i and \mathcal{D}^j) are close to each other, then we do not want to have separate integer variables \mathbf{x}^i and \mathbf{x}^j for them, instead, we want to introduce a hyper-integer variable \mathbf{y} representing $\mathbf{x}^i + \mathbf{x}^j$, and relax \mathbf{x}^i and \mathbf{x}^j to fractional variables. It turns out that we need to go a bit further: We need τ hyper-integer variables \mathbf{y}_1 to \mathbf{y}_τ since there are τ choices out of \mathcal{P}^i . Going a bit further does not affect much when we transform the original IP into an MIP with a constant number of integer variables; however, it causes the established MIP to contain way more constraints. Our next goal is still to argue that in a vertex solution of the LP (induced by the MIP by fixing the integer variables), only a few variables may take a fractional value. To achieve this, we observe that the constraint matrix of the LP can be decomposed into two parts (two submatrices), one having a small rank, and the other having a large rank but also having a nice diagonal block structure. We show that the property of the vertex solution to such a two-part LP can be studied by arguing on the two parts separately.

1.2.1 Comparison to some relevant prior works

Our algorithmic results on block-structured IPs (Theorem 2 and 3) are closely related to their FPT algorithms. In addition to n -fold IPs, researchers have also studied FPT algorithms for 2-stage stochastic IPs and 4-block n -fold IPs, whose constraint matrices can be written as:

$$H_{2\text{-stage}} := \begin{pmatrix} B^1 & A^1 & 0 & & 0 \\ B^2 & 0 & A^2 & & 0 \\ \vdots & & & \ddots & \\ B^n & 0 & 0 & & A^n \end{pmatrix}, \quad H_{4\text{-block}} := \begin{pmatrix} C & D^1 & D^2 & \dots & D^n \\ B^1 & A^1 & 0 & & 0 \\ B^2 & 0 & A^2 & & 0 \\ \vdots & & & \ddots & \\ B^n & 0 & 0 & & A^n \end{pmatrix}.$$

These FPT algorithms typically find an exact solution within a running time of $f(\Delta) \cdot \text{poly}(|I|)$ for some computable function f (see Section 1.3 for detailed related works). By taking ε to be sufficiently small (e.g., $\varepsilon < 1/\Delta$ in Theorem 1 and Theorem 3), our approximation algorithms can return an exact solution, albeit that they may have a larger running time compared with the existing FPT algorithms. The merit of our approximation algorithms is that they offer a more flexible control over the running time vs. accuracy.

Very recently, Eisenbrand and Rothvoss [31] studied 4-block n -fold IPs and obtained a result of a similar flavor. They showed that, in time $f(\max_i \|A^i\|_\infty) \cdot \text{poly}(|I|)$, one can find a solution $(\mathbf{x}^0, \mathbf{x}^1, \dots, \mathbf{x}^n)$ whose objective value is no worse than OPT, and which satisfies all constraints except the first one, i.e., $C\mathbf{x}^0 + \sum_{i=1}^n D^i \mathbf{x}^i \neq \mathbf{b}^0$. In particular, the error can be bounded as $\|C\mathbf{x}^0 + \sum_{i=1}^n D^i \mathbf{x}^i - \mathbf{b}^0\|_\infty \leq g(\max_i \|D^i\|_\infty)$ for some computable function g . As n -fold IPs are special cases of 4-block n -fold IPs, such a result also holds for n -fold IPs. Such a result is parallel to our work since the running time of the algorithm provided by Eisenbrand and Rothvoss depends on Δ , and its error can be way larger than Δ .

In terms of techniques, our algorithms need to round a fractional solution to an integral solution. During the rounding procedure, we will need to reduce the number of variables that take a fractional value, and this resembles the work of Cslovjcek et al. [17]. In particular, Cslovjcek et al. considered an LP of the form $A_1 \mathbf{x}^1 + \dots + A_n \mathbf{x}^n = \mathbf{b}$, where $\mathbf{x}^i \in Q_i$ and each $Q_i \subseteq \mathbb{R}_{\geq 0}^{t_i}$ is an integral polyhedron. They showed that in a vertex solution, all but m of the \mathbf{x}^i 's are vertices of the Q_i 's, respectively, and thus all but m of the \mathbf{x}^i 's are integral, where m is the number of rows of A_i 's. Their result applies directly to n -fold IPs where each A_i is a small submatrix of fixed size. Our rounding procedure also deals with an LP of the form $A_1 \mathbf{x}^1 + \dots + A_n \mathbf{x}^n = \mathbf{b}$; however, each A_i will have a huge size. The huge size is mainly due to a configuration-LP type of reformulation of the input IP. While the result of Cslovjcek et al. implies that in a vertex solution only a few \mathbf{x}^i 's are integral, one \mathbf{x}^i may still contain too many fractional variables. We will need to leverage the special structure of the A_i 's to reduce the number of fractional variables (See Section 3.2 for details). Overall, while we are following the general framework of arguing on the number of fractional variables in a vertex solution to a specific n -fold-like LP, we are not aware of a prior result that is applicable directly.

1.3 Related work

The approximation algorithms of IPs are rare in the literature. Dadush [21] studied *approximate integer programming* in 2014. Later, Dadush et al. [22] investigated how to reduce the exact integer programming problem to the approximate version. In addition, Shmoys and Swamy [73] showed that one could derive approximation algorithms for most of the stochastic IPs considered in [27, 35, 41, 71] by adopting a natural LP rounding approach. Swamy and Shmoys [75] gave the first approximation algorithms for a variety of k -stage generalizations of basic combinatorial optimization problems including the set cover, vertex cover, multicut on trees, facility location, and multicommodity flow problems.

We give additive approximation schemes for IPs, which yield additive approximation schemes for several scheduling problems. While most of the existing approximation schemes for scheduling problems give a multiplicative factor of $1 + \varepsilon$, additive approximation schemes have been studied very recently by Buchem et al. [7], who presented an algorithm for $Pm||C_{\max}$ with a running time of $m^2 n^{\mathcal{O}(1/\varepsilon)}$ that computes a solution with makespan at most $\text{OPT} + \varepsilon \cdot \max_{j \in [n]} p_j$. It is interesting to study additive approximation schemes for other classic scheduling problems.

In addition to scheduling problems, n -fold IPs have also been used to model a variety of other combinatorial optimization problems, based on which FPT algorithms are obtained; see, e.g., [24, 37, 37, 38, 44, 46, 55–58, 69]. We believe that the study of approximation schemes for n -fold type of IPs can also facilitate the design of approximation schemes for relevant problems.

FPT algorithms for IPs with respect to other structure parameters have also been studied extensively, such as 2-stage stochastic IPs [3, 18, 30, 39, 43, 54, 60] and 4-block n -fold IPs [11–13, 36, 68], as well as other structured IPs [5, 9, 28, 31]. Particularly, Cslovjecsek et al. [19] proved that the feasibility problems for 2-stage stochastic IPs can be solved in time $f(\gamma_1, \max_i \|A^i\|_\infty) \cdot |I|$ for some computable function f , where γ_1 is an upper bound of the dimensions of all blocks A^i, B^i . They also demonstrated that the optimization problems for n -fold IPs with $D^1 = D^2 = \dots = D^n$ can be solved in time $f(\gamma_2, \max_i \|A^i\|_\infty) \cdot \text{poly}(|I|)$, where γ_2 is an upper bound of the dimensions of all blocks A^i, D^i . Very recently, Eisenbrand and Rothvoss [31] solved an open problem raised by Cslovjecsek et al. [19] and showed that the optimization problems for 2-stage stochastic IPs can be solved in time $f(\gamma_1, \max_i \|A^i\|_\infty) \cdot |I|$.

It is worth mentioning that block-structured IPs with fractional variables have also been explored in the literature. In particular, Brand et al. [5] introduced the concept of *fractionality*, which denotes the minimum over the maxima of denominators in optimal solutions of an MILP instance. They showed that the fractionality of n -fold and 2-stage stochastic mixed-IPs is FPT-bounded.

1.4 Preliminaries

Notations. We write column vectors in boldface, e.g., \mathbf{x}, \mathbf{y} , and their entries in normal font, e.g., x_j, y_j . Let $[i] := \{1, 2, \dots, i\}$. For a vector or a matrix, let $\|\cdot\|_\infty$ denote the maximum absolute value of its elements. Let $\text{poly}(|I|)$ denote a polynomial in the input length $|I|$, f be a computable function, and OPT denote the optimal objective value.

The support of a vector \mathbf{x} is the set of indices $I = \text{supp}(\mathbf{x})$ such that $x_j = 0 \iff j \notin I$. Let V and W be vector spaces and $\phi: V \rightarrow W$. The kernel of ϕ , denoted $\ker(\phi)$, is the set of vectors mapping to zero, i.e., $\ker(\phi) = \{x \in V : \phi(x) = \mathbf{0}\}$. The nullity of ϕ , denoted $\text{nullity}(\phi)$, is the dimension of its kernel, that is, the dimension of the smallest dimensional vector space containing $\ker(\phi)$. The image of ϕ is defined as $\text{img}(\phi) = \{\phi(x) : x \in V\} \subseteq W$. The rank of ϕ , denoted $\text{rank}(\phi)$, is, analogously to the nullity, the dimension of the smallest dimensional vector space containing $\text{img}(\phi)$. If ϕ is a linear map, the definitions of rank and nullity simplify to

$$\text{nullity}(\phi) = \dim(\ker(\phi)), \quad \text{rank}(\phi) = \dim(\text{img}(\phi)).$$

The definitions are also used for matrices as a matrix $A \in \mathbb{R}^{m \times n}$ represents the linear map $\phi: \mathbb{R}^n \rightarrow \mathbb{R}^m$, $\mathbf{x} \mapsto A\mathbf{x}$. The kernel of a matrix $A \in \mathbb{R}^{m \times n}$ is $\ker(A) = \{\mathbf{x} \in \mathbb{R}^n : A\mathbf{x} = \mathbf{0}\}$, the nullity is its dimension, the image is $\{A\mathbf{x} \in \mathbb{R}^m : \mathbf{x} \in \mathbb{R}^n\}$ and the rank is its dimension. The latter coincides with the alternative rank definition of matrices as the largest number k such that any k columns of A are linearly independent.

We make use of the well-known rank-nullity theorem, stating that a linear map splits the dimension of the domain into the dimension of the kernel and the dimension of the image (cf. the splitting lemma).

► **Theorem 4 (Rank-Nullity Theorem).** *Let V, W be vector spaces, where $\dim(V) < \infty$, and $\phi: V \rightarrow W$ be a linear map. Then*

$$\dim(V) = \text{nullity}(\phi) + \text{rank}(\phi).$$

We are generally interested in the number of variables in an LP that attain non-integral values. Thus, analogously to the support of a vector, we define the non-integral support.

► **Definition 5** (Non-Integral Support).

$$\text{nisupp}(\mathbf{x}) := \{j \in \text{supp}(\mathbf{x}) : x_j \notin \mathbb{Z}\}.$$

We will need the following proposition, which generalizes the fact that for a system $A\mathbf{x} = \mathbf{b}$, $\mathbf{x} \geq 0$, the support of a vertex solution consists of linearly independent columns, to systems with (additional) lower and upper bounds on the variables:

► **Proposition 6** (Proposition 3.3.3, [4]). *Let \mathbf{x}^* be a vertex solution to an LP of the form $A\mathbf{x} = \mathbf{b}$, $\ell \leq \mathbf{x} \leq \mathbf{u}$, $\ell, \mathbf{u} \in \mathbb{Q}^n$. Then $(A_j)_{j \in \text{supp}(\mathbf{x}^*) : \ell_j < x_j < u_j}$ is non-singular.*

Given that a vertex solution can be computed in polynomial time, the following lemma is straightforward via Proposition 6.

► **Lemma 7.** *Consider an LP with the objective of minimizing or maximizing $\mathbf{w}\mathbf{x}$ subject to $\{\mathbf{x} : A\mathbf{x} = \mathbf{b}, \ell \leq \mathbf{x} \leq \mathbf{u}, \ell, \mathbf{u} \in \mathbb{Z}^n\}$. In polynomial time we can compute an optimal solution \mathbf{x}^* such that $(A_j)_{j \in \text{supp}(\mathbf{x}^*) : \ell_j < x_j < u_j}$ is non-singular.*

2 Approximation algorithms for IPs with a constant number of constraints – Proof of Theorem 1

This section is devoted to the proof of Theorem 1. We follow the general method we described in “Technical overview”. The whole proof is divided into three steps.

Step 1: Transform IP (1) to a mixed IP with $\mathcal{O}(1)$ integer variables

Recall that $H = (\mathbf{h}_1, \mathbf{h}_2, \dots, \mathbf{h}_n)$ where $\mathbf{h}_j \in [-\Delta, \Delta]^m$. We cut $[-\Delta, \Delta]^m$ into $(2/\delta)^m$ small boxes (called δ -boxes), $\prod_{i=1}^m [(\lambda_i - 1)\delta\Delta, \lambda_i\delta\Delta]$, where $\delta > 0$ is an arbitrary small constant and $\lambda_i \in \{-1/\delta + 1, -1/\delta + 2, \dots, 1/\delta\}$. We arbitrarily index all the δ -boxes, and let I_k be the set of indices j 's such that \mathbf{h}_j is in the k -th δ -box, where $k = 1, \dots, (2/\delta)^m$. Note that if some \mathbf{h}_j belongs to multiple δ -boxes, let it be in arbitrary one such δ -box. For the k -th δ -box, we pick one fixed vector \mathbf{v}_k and call it the *canonical vector*. For each $j \in I_k$, we define $\tilde{\mathbf{h}}_j = \mathbf{h}_j - \mathbf{v}_k$, so it is clear that $\|\tilde{\mathbf{h}}_j\|_\infty \leq \delta\Delta$.

We say that the vectors in the same δ -box are similar. By introducing a new variable y_k to the k -th δ -box, where $y_k := \sum_{j \in I_k} x_j$, we observe that $H\mathbf{x} = \sum_{j=1}^n \mathbf{h}_j x_j$ can be rewritten as $\sum_{k=1}^{(2/\delta)^m} \sum_{j \in I_k} (\mathbf{v}_k + \tilde{\mathbf{h}}_j) x_j$, thus obtaining the following MIP:

$$\begin{aligned} (\text{MIP}_1) \quad \min \quad & \sum_{j=1}^n w_j x_j \\ & \sum_{k=1}^{(2/\delta)^m} \mathbf{v}_k y_k + \sum_{k=1}^{(2/\delta)^m} \sum_{j \in I_k} \tilde{\mathbf{h}}_j x_j = \mathbf{b} \end{aligned} \tag{2a}$$

$$\sum_{j \in I_k} x_j = y_k, \quad \forall 1 \leq k \leq (2/\delta)^m \tag{2b}$$

$$l_j \leq x_j \leq u_j, \quad \forall j \in I_k, 1 \leq k \leq (2/\delta)^m$$

$$y_k \in \mathbb{Z}, x_j \in \mathbb{R}, \quad \forall 1 \leq j \leq n, 1 \leq k \leq (2/\delta)^m$$

Notice that (MIP₁) relaxes $x_j \in \mathbb{Z}$ to $x_j \in \mathbb{R}$, and only contains $(2/\delta)^m$ integer variables y_k . Hence, applying Kannan's algorithm [52], an optimal solution to (MIP₁) can be computed in $(2/\delta)^{\mathcal{O}(m(2/\delta)^m)} \cdot |I|$ time.

Let \mathbf{y}^* and \mathbf{x}^* be the optimal solution to (MIP₁). In the following step, we will round fractional variables such that most fractional variables take integral values, and all constraints are still satisfied.

Step 2: Obtain a feasible solution with $\mathcal{O}(m)$ variables taking a fractional value

Towards rounding x_j^* 's without changing the constraints in (MIP₁), it suffices to consider the following LP.

$$\begin{aligned}
 (\text{LP}_2) \quad \min \quad & \sum_{j=1}^n w_j x_j \\
 & \sum_{k=1}^{(2/\delta)^m} \sum_{j \in I_k} \tilde{\mathbf{h}}_j x_j = \sum_{k=1}^{(2/\delta)^m} \sum_{j \in I_k} \tilde{\mathbf{h}}_j x_j^* \tag{3a} \\
 & \sum_{j \in I_k} x_j = \sum_{j \in I_k} x_j^*, \quad \forall 1 \leq k \leq (2/\delta)^m \tag{3b} \\
 & l_j \leq x_j \leq u_j, x_j \in \mathbb{R}, \quad \forall j \in [n]
 \end{aligned}$$

Let $\bar{\mathbf{x}} = (\bar{x}_j)_{j=1}^n$ be an optimal solution to the above (LP₂) computed via Lemma 7. We will argue on the number of variables in $\bar{\mathbf{x}}$ taking a fractional value. Towards this, let $\eta := |\text{nisupp}(\bar{\mathbf{x}})|$.

It is straightforward to show that $\eta \leq \mathcal{O}(\delta^{-m})$. However, this is not sufficient to bound the rounding error, since in this case the error can reach $\mathcal{O}(\delta^{-m}) \cdot \delta \Delta = \mathcal{O}(\delta^{-m+1} \Delta)$, where $\delta > 0$ is an arbitrary small constant. The following claim gives a sharper bound.

▷ **Claim 8.** $\eta \leq 2m$.

Proof. For simplicity, let the constraints of (LP₂) be $A\mathbf{x} = \mathbf{b}$. Let $B = (A_j)_{j \in \text{nisupp}(\bar{\mathbf{x}})}$ be the submatrix induced by $\text{nisupp}(\bar{\mathbf{x}})$. Note that by Lemma 7, B is non-singular.

Observe that for every k we either have $|\text{nisupp}(\bar{\mathbf{x}}) \cap I_k| = 0$ or $|\text{nisupp}(\bar{\mathbf{x}}) \cap I_k| \geq 2$ due to Constraint (3b). If $|\text{nisupp}(\bar{\mathbf{x}}) \cap I_k| = 0$, Constraint (3b) becomes a zero-row in B . Let C be the matrix B after removing all zero-rows. Then C is non-singular (as we only removed zero rows from B). Moreover, since I_k 's are disjoint, there are at most $\eta/2$ constraints in Constraint (3b) which do not become a zero-row. Hence, there are at most $m + \frac{\eta}{2}$ constraints and η columns in C . If C has more columns than rows, then C is singular, which is a contradiction. So we have $m + \frac{\eta}{2} \geq \eta$, or equivalently $\eta \leq 2m$. ◁

Finally, we round the $\mathcal{O}(m)$ variables in $\bar{\mathbf{x}}$ that take fractional values.

Step 3: Round the $\mathcal{O}(m)$ fractional variables

Let $I_k^* \subseteq I_k$, $k = 1, \dots, (2/\delta)^m$, be the indices of variables that still take a fractional value after Step 2. Without loss of generality, we may assume that for all j 's in a fixed I_k^* , w_j 's are in non-decreasing order.

Let $\gamma_k := \sum_{j \in I_k^*} (\bar{x}_j - \lfloor \bar{x}_j \rfloor) = \sum_{j \in I_k^*} (\bar{x}_j - \lfloor \bar{x}_j \rfloor)$. Constraint (3b) guarantees that γ_k is an integer. Therefore, we can round up the first γ_k fractional variables in I_k^* (i.e., variables of the smallest indices) to $\lceil \bar{x}_j \rceil$, and meanwhile round down the other fractional variables in I_k^* to $\lfloor \bar{x}_j \rfloor$. This rounding method preserves the constraint $\sum_{j \in I_k} x_j = y_k$, and meanwhile does not increase the objective value. However, this rounding will introduce $\mathcal{O}(m\delta\Delta)$ error to Constraint (2a). Thus, Theorem 1 follows by taking $\delta = \mathcal{O}(\frac{\epsilon}{m})$.

Running time.

Step 1: $(2/\delta)^{\mathcal{O}(m(2/\delta)^m)} \cdot |I| = 2^{(\frac{m}{\varepsilon})^{\mathcal{O}(m)}} \cdot |I|$

Step 2: $n^4 \cdot |I|$

Step 3: $(2/\delta)^m \cdot n \log n$

All in all, the total time is $2^{(\frac{m}{\varepsilon})^{\mathcal{O}(m)}} \cdot \text{poly}(|I|)$, which has the form of $f(m, \varepsilon) \cdot \text{poly}(|I|)$.

3 Proof of Theorem 3

To prove Theorem 2, we first need to prove Theorem 3. We restate the theorem below.

► **Theorem 3.** *Given is an integer program $\min\{\mathbf{w}\mathbf{x} : \sum_{i=1}^n D^i \mathbf{x}^i = \mathbf{b}^0, \mathbf{x}^i \in \mathcal{P}^i, 1 \leq i \leq n, \mathbf{x} \in \mathbb{Z}^{nt}\}$ with optimal objective value OPT, where $D^i \in \mathbb{Q}^{s \times t}$ and \mathcal{P}^i is an arbitrary set of integer vectors. Let $\kappa = \max_{\mathbf{u} \in \cup_{i=1}^n \mathcal{P}^i} \|\mathbf{u}\|_\infty$, i.e., κ is the largest ℓ_∞ -norm among all integer vectors in $\cup_{i=1}^n \mathcal{P}^i$. Then for arbitrary small $\varepsilon > 0$, there exists an algorithm of running time $f(s, t, \kappa, \varepsilon) \cdot \text{poly}(|I|)$ which returns a near-feasible solution $\tilde{\mathbf{x}}$ such that $\tilde{\mathbf{x}} \in \{\mathbf{x} : \|\sum_{i=1}^n D^i \mathbf{x}^i - \mathbf{b}^0\|_\infty \leq \varepsilon \Delta, \mathbf{x}^i \in \mathcal{P}^i, 1 \leq i \leq n, \mathbf{x} \in \mathbb{Z}^{nt}\}$, and $\mathbf{w}\tilde{\mathbf{x}} \leq \text{OPT}$, where $\Delta = \max_{i \in [n]} \|D^i\|_\infty$.*

To prove Theorem 3, we follow the same general strategy as the proof of Theorem 1. That is, we first reformulate the input IP into an MIP via some appropriate δ -boxes, and then we solve the MIP, fix the values of integer variables and obtain an LP, and argue that the vertex solution of this LP only contains a small number of variables taking a fractional value. Finally we round those fractional variables. The main challenge lies in the argument for the vertex solution, which builds upon a specific block structure of the constraint matrix. To facilitate this argument, we need to further introduce some notions.

3.1 General mechanisms for LP that consists of two parts

In this section, we develop some preliminaries tailored to our setting.

We will encounter the case where the constraint matrix can be partitioned into two parts/subsystems: One with small rank and one where for a given solution either the non-integral support is small or we can easily find linearly independent kernel elements acting only on the non-integral support. Assume we have an LP of the following form:

$$\mathbf{A}\mathbf{x} = \mathbf{b}_A, \quad \mathbf{C}\mathbf{x} = \mathbf{b}_C, \quad \mathbf{x} \geq \mathbf{0},$$

where A is any matrix, possibly with large rank, but C has only a few (say, ι) rows. Consider some fixed solution \mathbf{x}^* and we want to reduce $\text{nisupp}(\mathbf{x}^*)$ as much as possible. If A has large rank, the combined system also has large rank, so it is not really helpful trying to obtain a solution with a small non-integral support through arguing a small rank. However, assume that we are able to find a set \mathcal{K} , $|\mathcal{K}| \geq \iota + 1$, of linearly independent non-trivial kernel elements of A , that are only non-zero on variables in $\text{nisupp}(\mathbf{x}^*)$. Then we can “combine” those kernel elements to a non-trivial kernel element of the combined system, using Theorem 4.

More formally, the following lemma allows us to show non-zero nullity of the combined system if enough kernel elements of the first system (more than the rank of the second one) can be found.

► **Lemma 9.** *Consider vector spaces V, W where $\dim(V) < \infty$ and two linear maps $\phi_1, \phi_2: V \rightarrow W$. If $\text{nullity}(\phi_1) > \text{rank}(\phi_2)$, then $\dim(\ker(\phi_1) \cap \ker(\phi_2)) > 0$.*

Proof. Let $\psi: \ker(\phi_1) \rightarrow W$, $x \mapsto \phi_2(x)$ be the restriction of ϕ_2 to $\ker(\phi_1)$. Then $\text{rank}(\psi) \leq \text{rank}(\phi_2)$ and $\ker(\phi_1) \cap \ker(\phi_2) = \ker(\psi)$, and thus

$$\dim(\ker(\phi_1) \cap \ker(\phi_2)) = \text{nullity}(\psi) \stackrel{\text{Theorem 4}}{=} \dim(\ker(\phi_1)) - \text{rank}(\psi) \stackrel{\text{rank}(\psi) \leq \text{rank}(\phi_2)}{\geq} \text{nullity}(\phi_1) - \text{rank}(\phi_2) \stackrel{\text{assumption}}{>} 0. \blacktriangleleft$$

The scenario mentioned above is then covered by applying this lemma with $V = \ker(A)$, $W = \text{img}(C)$, $\phi_1: \mathbf{x} \mapsto A\mathbf{x}$, $\phi_2: \mathbf{x} \mapsto C\mathbf{x}$: Given that $\text{nullity}(A) > \text{rank}(C)$, there is a non-trivial kernel element of the combined system. Note that $\text{nullity}(A) > \text{rank}(C)$ translates to “there are at least $\text{rank}(C) + 1$ linearly independent non-trivial kernel elements of A ”.

If the constraint matrix of an LP is totally unimodular and the right-hand side is integral, it is well-known that vertex solutions are integers. Based on this, we can easily obtain the following proposition.

► **Proposition 10.** *Let \mathbf{x} be a solution to an LP of the form $A\mathbf{x} = \mathbf{b}$, $\ell \leq \mathbf{x} \leq \mathbf{u}$, $\ell, \mathbf{u} \in \mathbb{Q}^m$, $\mathbf{b} \in \mathbb{Z}^m$, where the constraint matrix A is totally unimodular. Then if $\text{nisupp}(\mathbf{x}) \neq \emptyset$, we have $\text{nullity}((A_j)_{j \in \text{nisupp}(\mathbf{x})}) > 0$.*

Proof. Let B be the matrix A restricted to the columns from $\text{nisupp}(\mathbf{x})$. Since A is totally unimodular, so is B . Let \mathbf{y} be the vector \mathbf{x} restricted to entries in $\text{nisupp}(\mathbf{x})$. Then $B\mathbf{y}$ is integral, as $A\mathbf{x}$ is integral and only x_j 's that are integral have been deleted from \mathbf{x} to obtain \mathbf{y} . Let $\gamma := A\mathbf{x} - B\mathbf{y}$. Then γ is an integer vector as $A\mathbf{x}$ and $B\mathbf{y}$ also are. Equivalently, $B\mathbf{y} = A\mathbf{x} - \gamma$. Consider the LP $B\mathbf{z} = A\mathbf{x} - \gamma$ with \mathbf{z} being variables. It is feasible (due to \mathbf{y}). Further, there is an integer vector \mathbf{y}^* such that $B\mathbf{y}^* = A\mathbf{x} - \gamma$, since B is totally unimodular. By construction, \mathbf{y} has no integral components, so $\mathbf{y} - \mathbf{y}^*$ is a non-trivial kernel element of B , directly implying B , and hence also A , has non-zero nullity. ◀

Rephrasing the above: A solution to a totally unimodular system with integral right-hand side is either integral or the subsystem induced by its non-integral variables is underdetermined.

3.2 Proof of Theorem 3

Now we are ready to prove Theorem 3. Let $\tau := \max_{i \in [n]} |\mathcal{P}^i| = \kappa^t$. By adding dummy vectors, we may assume $|\mathcal{P}^i| = \tau$ for all i and let $\mathcal{P}^i = \{\mathbf{p}_1^i, \mathbf{p}_2^i, \dots, \mathbf{p}_\tau^i\}$. Consider the matrix $\mathcal{D}^i = (D^i \mathbf{p}_1^i, D^i \mathbf{p}_2^i, \dots, D^i \mathbf{p}_\tau^i)$ and notice that \mathcal{D}^i is fixed. Let $\Delta = \max_{i \in [n]} \|\mathcal{D}^i\|_\infty$. It is clear that $\|\mathcal{D}^i\|_\infty = \Delta \kappa^t$.

To prove Theorem 3, we shall follow the general strategy for proving Theorem 1. We first cut the box $[-\|\mathcal{D}^i\|_\infty, \|\mathcal{D}^i\|_\infty]^s$ into a constant number of small boxes, introduce one hyper-integer variable for each small box, and obtain an MIP, which can be solved in polynomial time. Finally, we round fractional variables to integer variables.

Towards the proof, we first rewrite the input IP in a form that is easier for our analysis.

Step 0: Rewrite the input IP

We cut the box $[-\|\mathcal{D}^i\|_\infty, \|\mathcal{D}^i\|_\infty]^s$ into $(\frac{2}{\delta})^s$ small δ -boxes of the form $\prod_{i=1}^s [(\lambda_i - 1)\delta\Delta\kappa^t, \lambda_i\delta\Delta\kappa^t]$, where $\lambda_i \in \{-1/\delta + 1, -1/\delta + 2, \dots, 1/\delta\}$, and δ is a sufficiently small constant. In particular, it will become clear later that taking $\delta = \mathcal{O}(\frac{\epsilon}{st\kappa^{t+1}})$ suffices.

We pick a fixed vector within each δ -box and let it be the canonical vector of this δ -box. Each $D^i \mathbf{p}_j^i$ lies in some δ -box and corresponds to the canonical vector of this δ -box. Hence, $\mathcal{D}^i = \mathcal{V}^i + \tilde{\mathcal{D}}^i$ where \mathcal{V}^i is the matrix consisting of the canonical vectors corresponding to the column vectors of \mathcal{D}^i , and $\tilde{\mathcal{D}}^i$ is the “residue matrix” satisfying that $\|\tilde{\mathcal{D}}^i\|_\infty \leq \delta\Delta\kappa^t$.

Since each \mathcal{V}^i has τ columns, it is easy to see that there are $\rho := ((\frac{2}{\delta})^s)^\tau = (\frac{2}{\delta})^{(s\kappa^t)}$ different types of canonical matrices \mathcal{V}^i 's. For $k = 1, \dots, \rho$, let I_k denote the set of indices i 's (of \mathcal{D}^i 's) whose canonical matrix is of type k . The constraint $\sum_{i=1}^n D^i \mathbf{x}^i = \mathbf{b}^0$ essentially asks for selecting some column from each \mathcal{D}^i such that they add up to \mathbf{b}^0 , so we rewrite it with the help of auxiliary variables \mathbf{z} which denote such a choice: Let $z_{ij} = 1$ represent that we select the j -th column of \mathcal{D}^i and $z_{ij} = 0$ otherwise. We have the following equivalent IP:

$$(IP_3) \quad \min \quad \sum_{i=1}^n \left(\sum_{j=1}^{\tau} \mathbf{w}^i \mathbf{p}_j^i z_{ij} \right) \quad (4a)$$

$$\sum_{k=1}^{\rho} \sum_{i \in I_k} \sum_{j=1}^{\tau} (D^i \mathbf{p}_j^i) z_{ij} = \mathbf{b}^0 \quad (4b)$$

$$\sum_{j=1}^{\tau} z_{ij} = 1, \quad \forall 1 \leq i \leq n \quad (4c)$$

$$z_{ij} \in \{0, 1\}, \quad \forall 1 \leq i \leq n, 1 \leq j \leq \tau \quad (4d)$$

Step 1: Transform (IP₃) to a mixed IP with $\mathcal{O}(1)$ integer variables

Recall that there are ρ distinct types of canonical matrices \mathcal{V}^i 's. We let $(\mathbf{v}_1^k, \mathbf{v}_2^k, \dots, \mathbf{v}_\tau^k)$ be the k -th type. Suppose that the canonical matrix of \mathcal{D}^i is of type k , and then $\tilde{\mathcal{D}}^i = (D^i \mathbf{p}_1^i - \mathbf{v}_1^k, D^i \mathbf{p}_2^i - \mathbf{v}_2^k, \dots, D^i \mathbf{p}_\tau^i - \mathbf{v}_\tau^k)$.

Similar to the proof of Theorem 1, we introduce $y_{kj} = \sum_{i \in I_k} z_{ij}$ as the group variable that counts how many j -th columns we have selected among all \mathcal{D}^i 's whose canonical matrix is of type k . There are $\rho\tau = (\frac{2}{\delta})^{(s\kappa^t)} \kappa^t$ group variables. We establish (MIP₄) below, where y_{kj} 's are integer variables and z_{ij} 's are fractional variables.

$$(MIP_4) \quad \min \quad \sum_{i=1}^n \left(\sum_{j=1}^{\tau} \mathbf{w}^i \mathbf{p}_j^i z_{ij} \right) \quad (5a)$$

$$\sum_{k=1}^{\rho} \sum_{j=1}^{\tau} \mathbf{v}_j^k y_{kj} + \sum_{k=1}^{\rho} \sum_{i \in I_k} \sum_{j=1}^{\tau} (D^i \mathbf{p}_j^i - \mathbf{v}_j^k) z_{ij} = \mathbf{b}^0 \quad (5b)$$

$$y_{kj} = \sum_{i \in I_k} z_{ij}, \quad \forall 1 \leq j \leq \tau, 1 \leq k \leq \rho \quad (5c)$$

$$\sum_{j=1}^{\tau} z_{ij} = 1, \quad \forall i \in I_k, 1 \leq k \leq \rho \quad (5d)$$

$$z_{ij} \in [0, 1], y_{kj} \in \mathbb{Z}, \quad \forall 1 \leq j \leq \tau, i \in I_k, 1 \leq k \leq \rho \quad (5e)$$

Notice that in (MIP₄), the coefficients $\mathbf{w}^i \mathbf{p}_j^i$, \mathbf{v}_j^k , and $(D^i \mathbf{p}_j^i - \mathbf{v}_j^k)$ are fixed values instead of variables. Observe that (MIP₄) only contains $\rho\tau$ integer variables y_{kj} . Thus, applying Kannan's algorithm [52], an optimal solution to (MIP₄) can be computed in $(\rho\tau)^{\mathcal{O}(\rho\tau)} \cdot |I|$ time, which is polynomial.

Let $y_{kj} = y_{kj}^* \in \mathbb{Z}$ and $z_{ij} = z_{ij}^* \in [0, 1]$ be the optimal solution to (MIP₄).

Step 2: Obtain a feasible solution with $\mathcal{O}(s\tau)$ variables taking a fractional value

Fix $\mathbf{y} = \mathbf{y}^*$, and then (MIP₄) becomes an LP. We see that this LP is a combined system: We denote Constraint (5b) as $C\mathbf{z} = \mathbf{b}_C$, and Constraints (5c) and (5d) as $A_\Sigma \mathbf{z} = \mathbf{b}_A$. We denote by LP _{Σ} the linear system obtained by removing $C\mathbf{z} = \mathbf{b}_C$. That is, the constraint matrix of LP _{Σ} is exactly A_Σ .

Observe that $A_\Sigma = \begin{pmatrix} A_1 & & \\ & \ddots & \\ & & A_\rho \end{pmatrix}$ is a block diagonal matrix, as for each $k \in [\rho]$, the

set of variables $I_k \times [\tau]$ appearing in the constraints is disjoint with those for any other $k' \in [\rho] \setminus \{k\}$. Moreover, each $A_k \in \{0, 1\}^{(|I_k|+\tau) \times (|I_k| \cdot \tau)}$ is the constraint matrix of an assignment problem and thus totally unimodular (and thus so is A_Σ). More precisely, for $k = 1, \dots, \rho$,

$$A_k = \left(\begin{array}{ccc|ccc|ccc} 1 & & & 1 & & & & & 1 \\ & \ddots & & & \ddots & & & & \\ & & & & & & \dots & & \ddots \\ \hline & & & 1 & & & & 1 & \\ \hline 1 & \dots & 1 & & & & & & 1 \\ \hline & & & 1 & \dots & 1 & & & \\ & & & & & & \ddots & & \\ & & & & & & & 1 & \dots & 1 \end{array} \right) \left. \begin{array}{l} \vphantom{A_k} \\ \vphantom{A_k} \\ \vphantom{A_k} \end{array} \right\} \begin{array}{l} |I_k| \text{ rows} \\ \tau \text{ rows} \end{array}$$

$\underbrace{\hspace{1.5cm}}_{|I_k| \text{ columns}} \quad \underbrace{\hspace{1.5cm}}_{|I_k| \text{ columns}} \quad \underbrace{\hspace{1.5cm}}_{|I_k| \text{ columns}}$

where the first $|I_k|$ rows are for Constraint (5d), and the last τ rows are for Constraint (5c).

► **Remark.** It is clear that Constraint (5b) together with $A_\Sigma \mathbf{z} = \mathbf{b}_A$ form an n -fold structure. However, the number of rows and columns of each A_k can be prohibitively large. We mentioned in Section 1.2.1 before that Cslovjcek et al. [17] established a general result which implies that in a vertex solution satisfying $C\mathbf{z} = \sum_k C_k \mathbf{z}^k = \mathbf{b}_C$ and $A_\Sigma \mathbf{z} = \mathbf{b}_A$, at most s vectors out of $\mathbf{z}^1, \mathbf{z}^2, \dots, \mathbf{z}^\rho$ may take fractional values, but each \mathbf{z}^k contains $|I_k| \cdot \tau$ fractional variables, which means the total number of fractional variables is still too large. Our goal is to upper bound the number of fractional variables to $\mathcal{O}(s\tau)$, which requires a more fine-grained analysis that leverages the special structure of A_k 's.

Interestingly, for A_k it is easy to find a comparably small kernel element if solutions have enough non-integral components (for integral right-hand sides):

► **Lemma 11.** *For $k \in [\rho]$, consider a vertex solution \mathbf{z} to $A_k \mathbf{z} = \mathbf{b}$ for some integer vector \mathbf{b} . Let B be the matrix A_k restricted to the columns in $\text{nisupp}(\mathbf{z})$. Then $\text{rank}(B) \leq 2\tau$.*

Proof. If $|\text{nisupp}(\mathbf{z})| \leq 2\tau$, the statement is trivial. Thus, assume $|\text{nisupp}(\mathbf{z})| \geq 2\tau + 1$. For the sake of contradiction, assume $\text{rank}(B) > 2\tau$. Let $I \subseteq \text{nisupp}(\mathbf{z})$ be a set of $|I| = 2\tau + 1$ column indices. Let B' be the matrix B restricted to the columns in I , where afterwards all zero-rows have been removed. Clearly, $2\tau + 1 = \text{rank}(B') \leq \text{rank}(B)$ as we only removed zero-rows. Observe that every row in $\text{rank}(B')$ that corresponds to Constraint (5d) has to contain at least two non-zero variables due to the integral right-hand side. Moreover, all the constraints in (5d) are variable-disjunct. So the number of rows in B' is at most

$$\tau + \frac{|I|}{2} = \tau + \frac{2\tau + 1}{2} < 2\tau + 1.$$

However, we have $|I| = 2\tau + 1$ columns, so B' has non-zero nullity, which, by Theorem 4, is a contradiction to $\text{rank}(B') = 2\tau + 1$. ◀

► **Lemma 12.** *Let $\hat{\mathbf{z}}$ be a solution to LP_Σ , and B_Σ the matrix A_Σ restricted to the columns $\text{nisupp}(\hat{\mathbf{z}})$. Then we have*

$$\text{nullity}(B_\Sigma) \geq \frac{|\text{nisupp}(\hat{\mathbf{z}})|}{2\tau + 1}.$$

Proof. Observe that since A_Σ is a block diagonal matrix, to prove this lemma, it suffices to show for every $k \in [\rho]$ that B_k , the restriction of A_k to $\text{nisupp}(\hat{\mathbf{z}}) \cap I_k$, we have (where $\eta := |\text{nisupp}(\hat{\mathbf{z}}) \cap I_k|$):

$$\text{nullity}(B_k) \geq \frac{\eta}{2\tau + 1}.$$

If $\eta = 0$, the statement is trivial. If $0 < \eta \leq 2\tau + 1$, we only have to show that the nullity is non-zero, which follows directly from Proposition 10. For $2\tau + 1 < \eta$, by Lemma 11 and rank-nullity theorem, the nullity is at least $\eta - 2\tau$. Observe that for any \bar{t} we have

$$\eta - \bar{t} \geq \frac{\eta}{2\tau + 1} \iff \eta - \frac{\eta}{2\tau + 1} \geq \bar{t} \iff \frac{(2\tau + 1)\eta - \eta}{2\tau + 1} \geq \bar{t} \iff \frac{2\tau\eta}{2\tau + 1} \geq \bar{t}.$$

Furthermore, as $\eta > 2\tau + 1$, then $\frac{2\tau\eta}{2\tau + 1} \geq \frac{2\tau(2\tau + 2)}{2\tau + 1} > 2\tau$, i.e., $\eta - \bar{t} \geq \frac{\eta}{2\tau + 1}$ holds in particular for $\bar{t} := 2\tau$, which eventually leads to

$$\frac{\eta}{2\tau + 1} \leq \eta - 2\tau \leq \text{nullity}(B_k). \quad \blacktriangleleft$$

Now, let us consider vertex solution γ to the combined linear system, computed by Lemma 7. Let B_Σ be the restriction of A_Σ to the columns in $\text{nisupp}(\gamma)$. We claim that $s \geq \frac{|\text{nisupp}(\gamma)|}{2\tau + 1}$. Assume otherwise. By Lemma 12, $\frac{|\text{nisupp}(\gamma)|}{2\tau + 1} \leq \text{nullity}(B_\Sigma)$. By Lemma 9, having $\text{nullity}(B_\Sigma) > s$ suffices to prove the existence of a non-trivial kernel element on $\text{nisupp}(\gamma)$ of the combined system, which is the one that includes Constraint (5b). This is a contradiction to Lemma 7 which postulates that the constraint matrix restricted to $\text{nisupp}(\gamma)$ is non-singular.

Therefore, it follows that $|\text{nisupp}(\gamma)| \leq s(2\tau + 1) = \mathcal{O}(s\tau)$.

Step 3: Round the $\mathcal{O}(s\tau)$ fractional variables

We can round the fractional variables for each k , respectively, since the constraints for different k 's cannot affect each other. For a fixed k , by Constraint (5c), we have $y_{kj} = \sum_{i \in I_k} z_{ij}, \forall 1 \leq j \leq \tau$, and by Constraint (5d), it holds that $\sum_{j=1}^{\tau} z_{ij} = 1, \forall i \in I_k$. We assume $I_k = \{i_1, i_2, \dots, i_\ell\}$, and then Constraints (5c) and (5d) are shown as Figure 1.

$$\begin{pmatrix} z_{i_1,1} \\ z_{i_1,2} \\ \vdots \\ z_{i_1,\tau} \end{pmatrix} \quad \begin{pmatrix} z_{i_2,1} \\ z_{i_2,2} \\ \vdots \\ z_{i_2,\tau} \end{pmatrix} \quad \begin{pmatrix} z_{i_3,1} \\ z_{i_3,2} \\ \vdots \\ z_{i_3,\tau} \end{pmatrix} \quad \begin{pmatrix} z_{i_4,1} \\ z_{i_4,2} \\ \vdots \\ z_{i_4,\tau} \end{pmatrix} \quad \dots \quad \begin{pmatrix} z_{i_\ell,1} \\ z_{i_\ell,2} \\ \vdots \\ z_{i_\ell,\tau} \end{pmatrix}$$

- The sum of each column is equal to 1.
- The sum of the j -th row is equal to y_{kj} .

■ **Figure 1** Constraints (5c) and (5d) for a fixed k .

Check all fractional variables z_{ij} 's, and pick up the fractional variable with the smallest coefficient in the objective function (5a). Without loss of generality, we let z_0 be this fractional variable. See Figure 1 as an illustration. Then we round up the value of z_0 to 1. By the constraint $\sum_{j=1}^{\tau} z_{ij} = 1, \forall i \in I_k$, we round down the values of these variables in the same column as z_0 to zeros. This preserves the constraint $\sum_{j=1}^{\tau} z_{ij} = 1$.

Next, for the fractional variables in the same row as z_0 , we apply a greedy rounding similar as that in Step 3 of Theorem 1. We assume j_0 is the row number that z_0 has. Since $y_{kj} = \sum_{i \in I_k} z_{ij}, \forall 1 \leq j \leq \tau$, without loss of generality we assume that z_{ij_0} 's where $i \in I_k$ in

the summation are ordered by the non-decreasing coefficients in the objective function (5a). We compute $\sum_{i \in I_k} (z_{ij_0} - \lfloor z_{ij_0} \rfloor)$. Let $\gamma_{kj_0} := \sum_{i \in I_k} (z_{ij_0} - \lfloor z_{ij_0} \rfloor)$. Then we pick out the first γ_{kj_0} variables, and round up them to 1 and round down the remaining fractional variables z_{ij_0} 's to zeros. This rounding method will preserve the constraint $y_{kj_0} = \sum_{i \in I_k} z_{ij_0}$.

After the previous two processes, we delete the column and row where z_0 is located. Iteratively find the fractional variable with the smallest coefficient in the objective function from the current fractional variables, and repeat the above operations. Clearly, the constraints $\sum_{j=1}^{\tau} z_{ij} = 1, \forall i \in I_k$ and $y_{kj} = \sum_{i \in I_k} z_{ij}, \forall 1 \leq j \leq \tau$ are always satisfied.

We round up or down all fractional variables by preserving Constraints (5c) and (5d) and meantime not increasing the objective, while this introduces $\mathcal{O}(s\tau\delta\Delta\kappa t)$ error to Constraint (5b). Thus, Theorem 3 follows by taking $\delta = \mathcal{O}(\frac{\varepsilon}{s\tau\kappa t}) = \mathcal{O}(\frac{\varepsilon}{st\kappa^{t+1}})$.

Running time.

Step 1: $(\rho\tau)^{\mathcal{O}(\rho\tau)} \cdot |I| = ((\frac{2}{\delta})^{(s\kappa^t)} \kappa^t)^{\mathcal{O}((\frac{2}{\delta})^{(s\kappa^t)} \kappa^t)} \cdot |I| = 2^{(\frac{s\kappa^t}{\varepsilon})^{\mathcal{O}(s\kappa^t)}} \cdot |I|$

Step 2: $(\rho n)^4 \cdot |I| = (\frac{2}{\delta})^{\mathcal{O}(s\kappa^t)} n^4 \cdot |I|$

Step 3: $\rho\tau^2 n^2 \log n = ((\frac{2}{\delta})^s)^{\tau} \tau^2 n^2 \log n = (\frac{2}{\delta})^{\mathcal{O}(s\kappa^t)} n^2 \log n$

All in all, the total time is $2^{(\frac{s\kappa^t}{\varepsilon})^{\mathcal{O}(s\kappa^t)}} \cdot \text{poly}(|I|)$, which has the form of $f(s, t, \kappa, \varepsilon) \cdot \text{poly}(|I|)$.

4 Approximation algorithms for n -fold IPs – Proof of Theorem 2

This section is devoted to the proof of Theorem 2. We restate the theorem below.

► **Theorem 2.** *Given is an integer program $\min\{\mathbf{w}\mathbf{x} : \sum_{i=1}^n D^i \mathbf{x}^i = \mathbf{b}^0, A^i \mathbf{x}^i = \mathbf{b}^i, 1 \leq i \leq n, \mathbf{0} \leq \mathbf{x} \leq \mathbf{u}, \mathbf{x} \in \mathbb{Z}^{nt}\}$ with optimal objective value OPT, where $A^i \in \mathbb{Q}_{\geq 0}^{s_A \times t_A}, D^i \in \mathbb{Q}_{\geq 0}^{s_D \times t_D}$, and $t_A = t_D = t$. Then for arbitrary small $\varepsilon > 0$, there exists an algorithm of running time $f(s_A, s_D, t, \varepsilon) \cdot \text{poly}(|I|)$ which returns a near-feasible solution $\tilde{\mathbf{x}}$ such that $\tilde{\mathbf{x}} \in \{\mathbf{x} : (1-\varepsilon)\mathbf{b}^0 \leq \sum_{i=1}^n D^i \mathbf{x}^i \leq (1+\varepsilon)\mathbf{b}^0, (1-\varepsilon)\mathbf{b}^i \leq A^i \mathbf{x}^i \leq (1+\varepsilon)\mathbf{b}^i, 1 \leq i \leq n, \mathbf{0} \leq \mathbf{x} \leq \mathbf{u}, \mathbf{x} \in \mathbb{Z}^{nt}\}$, and $\mathbf{w}\tilde{\mathbf{x}} \leq \text{OPT}$.*

The above theorem is closely related to Theorem 3. In particular, since all A^i 's are positive matrices, if all entries in A^i is large compared to \mathbf{b}^i , then we know each coordinate of \mathbf{x}^i must be small. Consequently, \mathbf{x}^i may only take $\mathcal{O}(1)$ possible values, in other words, $\mathbf{x}^i \in \mathcal{P}^i$ where $|\mathcal{P}^i|$ is fixed, and Theorem 3 is applicable. What if some entries of A^i are small, such that some coordinate of \mathbf{x}^i , say, x_j^i , can take a large value? In this case, we observe that when the variable x_j^i increases or decreases by a small amount, the equation $A^i \mathbf{x}^i = \mathbf{b}^i$ may only be violated by some small amount. This motivates us to split \mathbf{x}^i into two parts, the major part and the minor part, where the major part can only take $\mathcal{O}(1)$ possible values, and the minor part can be handled like a fractional variable.

More precisely, let A_j^i be the j -th column of matrix A^i . By the constraint $A^i \mathbf{x}^i = \mathbf{b}^i$, we have

$$A_1^i x_1^i + A_2^i x_2^i + \dots + A_t^i x_t^i = \mathbf{b}^i.$$

We classify all A_j^i 's into two types:

- If there is a component in vector A_j^i with a value at least ψ , then the vector A_j^i is defined as **big**.
- Otherwise, if the values on all components of vector A_j^i are less than ψ , then the vector A_j^i is defined as **small**.

Here, ψ is a parameter to be decided later. In fact, we will take $\psi := \frac{\varepsilon}{2t}$. We divide the proof of Theorem 2 into two cases.

Case 1. If all A_j^i 's are big vectors in $A^i \mathbf{x}^i = \mathbf{b}^i$ for all i 's, then we have $\sum_{j=1}^t x_j^i = \mathcal{O}(s_A/\psi)$, which implies that $\|\mathbf{x}^i\|_\infty \leq \|\mathbf{x}^i\|_1 = \mathcal{O}(s_A/\psi)$. By Theorem 3, there exists an algorithm of running time $f(s_A, s_D, t, \varepsilon) \cdot \text{poly}(|I|)$ which returns a near-feasible solution $\tilde{\mathbf{x}}$ such that $\tilde{\mathbf{x}} \in \{\mathbf{x} : (1 - \varepsilon)\mathbf{b}^0 \leq \sum_{i=1}^n D^i \mathbf{x}^i \leq (1 + \varepsilon)\mathbf{b}^0, A^i \mathbf{x}^i = \mathbf{b}^i, \mathbf{0} \leq \mathbf{x} \leq \mathbf{u}, \mathbf{x} \in \mathbb{Z}^{nt}\}$, and $\mathbf{w}\tilde{\mathbf{x}} \leq \text{OPT}$. Theorem 2 follows.

Case 2. Otherwise, for every A_j^i we can multiply it with some λ_j^i to make it big. If some A_j^i is initially big, let its λ_j^i be 1. Then the constraint $A^i \mathbf{x}^i = \mathbf{b}^i$ can be written as

$$A_j^i x_j^i = \lambda_j^i A_j^i \left\lfloor \frac{x_j^i}{\lambda_j^i} \right\rfloor + A_j^i \left(x_j^i - \lambda_j^i \left\lfloor \frac{x_j^i}{\lambda_j^i} \right\rfloor \right),$$

where λ_j^i is the smallest integer such that $\lambda_j^i A_j^i$ is big. Denote $(x_j^i)' := \lfloor \frac{x_j^i}{\lambda_j^i} \rfloor$, and $(x_j^i)'' := x_j^i - \lambda_j^i \lfloor \frac{x_j^i}{\lambda_j^i} \rfloor$. Then we have

$$A_j^i x_j^i = \lambda_j^i A_j^i (x_j^i)' + A_j^i (x_j^i)''.$$

Note that the coefficients of $(x_j^i)'$ are big, while that of $(x_j^i)''$ are not. We can bound $\sum_{j=1}^t A_j^i (x_j^i)''$ by a negligible value due to a simple analysis. It implies that all possible values of $(x_j^i)'$ can be enumerated.

Similarly, the other constraint $\sum_{i=1}^n D^i \mathbf{x}^i = \mathbf{b}^0$ can be written as

$$D_j^i x_j^i = \lambda_j^i D_j^i (x_j^i)' + D_j^i (x_j^i)''.$$

Then we formulate a new IP taking $(x_j^i)'$ and $(x_j^i)''$ as variables. In the new IP, the constraint $A^i \mathbf{x}^i = \mathbf{b}^i$ can be safely replaced by $(\mathbf{x}^i)' \in \mathcal{P}^i$. We shall deal with the two summands $\sum_{i=1}^n \sum_{j=1}^t \lambda_j^i D_j^i (x_j^i)'$ and $\sum_{i=1}^n \sum_{j=1}^t D_j^i (x_j^i)''$ separately. The way we deal with the former is the same as Step 1 of Theorem 3, and the way we deal with the latter is the same as Step 1 of Theorem 1. Afterwards, we apply the rounding procedures specified in Theorem 3 and Theorem 1.

The detailed proof of Theorem 2 is provided in the full version.

5 Conclusion

We study approximation algorithms for IPs with a constant number of constraints and n -fold IPs. For IPs with a constant number of constraints, our algorithm returns a solution that violates the constraints by an additive error of $\mathcal{O}(\varepsilon\Delta)$. For n -fold IPs, our algorithm returns a solution within a multiplicative error of $1 + \varepsilon$. It is not clear whether similar results can be obtained for other block-structured IPs such as 2-stage stochastic IPs or 4-block n -fold IPs. It is also an interesting open problem whether an additive approximation scheme can be obtained for n -fold IPs.

References

- 1 K. Altmanová, D. Knop, and M. Koutecký. Evaluating and tuning n -fold integer programming. *Journal of Experimental Algorithmics*, 24(1):1–22, 2019.
- 2 E. Angel, E. Bampis, and A. Kononov. A FPTAS for approximating the unrelated parallel machines scheduling problem with costs. In *Proceedings of the 9th Annual European Symposium on Algorithms (ESA)*, volume 2161 of *LIPIcs*, pages 194–205, 2001. doi: 10.1007/3-540-44676-1_16.

- 3 M. Aschenbrenner and R. Hemmecke. Finiteness theorems in stochastic integer programming. *Foundations of Computational Mathematics*, 7(2):183–227, 2007. doi:10.1007/S10208-005-0174-1.
- 4 D. Bertsekas, A. Nedic, and A. Ozdaglar. *Convex analysis and optimization*, volume 1. Athena Scientific, 2003.
- 5 C. Brand, M. Koutecký, and S. Ordyniak. Parameterized algorithms for MILPs with small treedepth. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence*, volume 35(14), pages 12249–12257, 2021.
- 6 K. Bringmann and A. Cassis. Faster knapsack algorithms via bounded monotone min-plus-convolution. In *Proceedings of the 49th International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 229 of *LIPICs*, pages 31:1–31:21, 2022. doi:10.4230/LIPICs.ICALP.2022.31.
- 7 M. Buchem, L. Rohwedder, T. Vredeveld, and A. Wiese. Additive approximation schemes for load balancing problems. In *Proceedings of the 48th International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 198 of *LIPICs*, pages 42:1–42:17, 2021. doi:10.4230/LIPICs.ICALP.2021.42.
- 8 A. Caprara, H. Kellerer, U. Pferschy, and D. Pisinger. Approximation algorithms for knapsack problems with cardinality constraints. *European Journal of Operational Research*, 123(2):333–345, 2000. doi:10.1016/S0377-2217(99)00261-1.
- 9 T. F. Chan, J. W. Cooper, M. Koutecký, D. Král, and K. Pekarkova. Matrices of optimal tree-depth and a row-invariant parameterized algorithm for integer programming. *SIAM Journal on Computing*, 51(3):664–700, 2022. doi:10.1137/20M1353502.
- 10 T. M. Chan. Approximation schemes for 0-1 knapsack. In *1st Symposium on Simplicity in Algorithms (SOSA)*, pages 5–1, 2018.
- 11 H. Chen, L. Chen, and G. Zhang. Block-structured integer programming: Can we parameterize without the largest coefficient? *Discrete Optimization*, 46:100743, 2022. doi:10.1016/J.DISOPT.2022.100743.
- 12 H. Chen, L. Chen, and G. Zhang. FPT algorithms for a special block-structured integer program with applications in scheduling. *Mathematical Programming*, 208:463–496, 2024. doi:10.1007/S10107-023-02046-Z.
- 13 L. Chen, M. Koutecký, L. Xu, and W. Shi. New bounds on augmenting steps of block-structured integer programs. In *Proceedings of the 28th Annual European Symposium on Algorithms (ESA)*, volume 173 of *LIPICs*, pages 33:1–33:19, 2020. doi:10.4230/LIPICs.ESA.2020.33.
- 14 L. Chen, J. Lian, Y. Mao, and G. Zhang. Approximating partition in near-linear time. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing (STOC)*, pages 307–318, 2024.
- 15 L. Chen, J. Lian, Y. Mao, and G. Zhang. A nearly quadratic-time fptas for knapsack. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing (STOC)*, pages 283–294, 2024.
- 16 L. Chen, J. Lian, Y. Mao, and G. Zhang. Weakly approximating knapsack in subquadratic time. In *Proceedings of the 52nd International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 334 of *LIPICs*, pages 51:1–51:18, 2025. doi:10.4230/LIPICs.ICALP.2025.51.
- 17 J. Cslovjecssek, F. Eisenbrand, C. Hunkenschröder, L. Rohwedder, and R. Weismantel. Block-structured integer and linear programming in strongly polynomial and near linear time. In *Proceedings of the 32nd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1666–1681. SIAM, 2021.
- 18 J. Cslovjecssek, F. Eisenbrand, M. Pilipczuk, M. Venzin, and R. Weismantel. Efficient sequential and parallel algorithms for multistage stochastic integer programming using proximity. In *Proceedings of the 29th Annual European Symposium on Algorithms (ESA)*, volume 204 of *LIPICs*, pages 33:1–33:14, 2021. doi:10.4230/LIPICs.ESA.2021.33.

- 19 J. Cslovjceksek, M. Koutecký, A. Lassota, M. Pilipczuk, and A. Polak. Parameterized algorithms for block-structured integer programs with large entries. In *Proceedings of the 35th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 740–751. SIAM, 2024.
- 20 D. Dadush. *Integer programming, lattice algorithms, and deterministic volume estimation*. Georgia Institute of Technology, 2012.
- 21 D. Dadush. A randomized sieving algorithm for approximate integer programming. *Algorithmica*, 70(2):208–244, 2014. doi:10.1007/S00453-013-9834-8.
- 22 D. Dadush, F. Eisenbrand, and T. Rothvoss. From approximate to exact integer programming. *Mathematical Programming*, pages 1–19, 2024.
- 23 D. Dadush, C. Peikert, and S. Vempala. Enumerative lattice algorithms in any norm via M-ellipsoid coverings. In *Proceedings of the 52nd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 580–589. IEEE, 2011.
- 24 J. A. De Loera, R. Hemmecke, S. Onn, and R. Weismantel. N-fold integer programming. *Discrete Optimization*, 5(2):231–241, 2008. doi:10.1016/J.DISOPT.2006.06.006.
- 25 M. Deng, C. Jin, and X. Mao. Approximating knapsack and partition via dense subset sums. In *Proceedings of the 34th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2961–2979. SIAM, 2023.
- 26 I. Doron-Arad, A. Kulik, and P. Manurangsi. Fine grained lower bounds for multidimensional knapsack. *arXiv preprint arXiv:2407.10146*, 2024.
- 27 S. Dye, L. Stougie, and A. Tomagard. The stochastic single resource service-provision problem. *Naval Research Logistics*, 50(8):869–887, 2003.
- 28 E. Eiben, R. Ganian, D. Knop, S. Ordyniak, M. Pilipczuk, and M. Wrochna. Integer programming and incidence treedepth. In *Proceedings of the 20th International Conference on Integer Programming and Combinatorial Optimization (IPCO)*, pages 194–204. Springer, 2019.
- 29 F. Eisenbrand, C. Hunkenschröder, and K. M. Klein. Faster algorithms for integer programs with block structure. In *Proceedings of the 45th International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 107 of *LIPICs*, pages 49:1–49:13, 2018. doi:10.4230/LIPICs.ICALP.2018.49.
- 30 F. Eisenbrand, C. Hunkenschröder, K. M. Klein, M. Koutecký, A. Levin, and S. Onn. An algorithmic theory of integer programming. *arXiv preprint arXiv:1904.01361*, 2019.
- 31 F. Eisenbrand and T. Rothvoss. A parameterized linear formulation of the integer hull. *arXiv preprint arXiv:2501.02347*, 2025.
- 32 F. Eisenbrand and R. Weismantel. Proximity results and faster algorithms for integer programming using the Steinitz Lemma. *ACM Transactions on Algorithms*, 16(1):1–14, 2019. doi:10.1145/3340322.
- 33 A. V. Fishkin, O. Gerber, K. Jansen, and R. Solis-Oba. On packing squares with resource augmentation: Maximizing the profit. In *Proceedings of the 2005 Australasian Symposium on Theory of Computing-Volume 41*, pages 61–67, 2005.
- 34 F. V. Fomin, F. Panolan, M. S. Ramanujan, and S. Saurabh. On the optimality of pseudo-polynomial algorithms for integer programming. *Mathematical Programming*, 198(1):561–593, 2023. doi:10.1007/S10107-022-01783-x.
- 35 A. Gupta, M. Pál, R. Ravi, and A. Sinha. Boosted sampling: Approximation algorithms for stochastic optimization. In *Proceedings of the 36th Annual ACM Symposium on Theory of Computing (STOC)*, pages 417–426, 2004.
- 36 R. Hemmecke, M. Köppe, and R. Weismantel. Graver basis and proximity techniques for block-structured separable convex integer minimization problems. *Mathematical Programming*, 145(1-2):1–18, 2014. doi:10.1007/S10107-013-0638-Z.
- 37 R. Hemmecke, S. Onn, and L. Romanchuk. N-fold integer programming in cubic time. *Mathematical Programming*, 137(1-2):325–341, 2013. doi:10.1007/S10107-011-0490-Y.
- 38 R. Hemmecke, S. Onn, and R. Weismantel. N-fold integer programming and nonlinear multi-transshipment. *Optimization Letters*, 5:13–25, 2011. doi:10.1007/S11590-010-0231-9.

- 39 R. Hemmecke and R. Schultz. Decomposition of test sets in stochastic integer programming. *Mathematical Programming*, 94(2-3):323–341, 2003. doi:10.1007/S10107-002-0322-1.
- 40 O. H. Ibarra and C. E. Kim. Fast approximation algorithms for the knapsack and sum of subset problems. *Journal of the ACM*, 22(4):463–468, 1975. doi:10.1145/321906.321909.
- 41 N. Immorlica, D. Karger, M. Minkoff, and V. S. Mirrokni. On the costs and benefits of procrastination: Approximation algorithms for stochastic combinatorial optimization problems. In *Proceedings of the 15th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 691–700, 2004.
- 42 K. Iwama and G. Zhang. Online knapsack with resource augmentation. *Information Processing Letters*, 110(22):1016–1020, 2010. doi:10.1016/J.IPL.2010.08.013.
- 43 K. Jansen, K. M. Klein, and A. Lassota. The double exponential runtime is tight for 2-stage stochastic ILPs. In *Proceedings of the 22nd International Conference on Integer Programming and Combinatorial Optimization (IPCO)*, pages 297–310. Springer, 2021.
- 44 K. Jansen, K. M. Klein, M. Maack, and M. Rau. Empowering the configuration-IP: New PTAS results for scheduling with setup times. *Mathematical Programming*, 195(1-2):367–401, 2022. doi:10.1007/S10107-021-01694-3.
- 45 K. Jansen, F. Land, and K. Land. Bounding the running time of algorithms for scheduling and packing problems. *SIAM Journal on Discrete Mathematics*, 30(1):343–366, 2016. doi:10.1137/140952636.
- 46 K. Jansen, A. Lassota, M. Maack, and T.s Pikiés. Total completion time minimization for scheduling with incompatibility cliques. In *Proceedings of the 31st International Conference on Automated Planning and Scheduling (ICAPS)*, pages 192–200, 2021.
- 47 K. Jansen, A. Lassota, and L. Rohwedder. Near-linear time algorithm for n-fold ILPs via color coding. *SIAM Journal on Discrete Mathematics*, 34(4):2282–2299, 2020. doi:10.1137/19M1303873.
- 48 K. Jansen and M. Mastrolilli. Scheduling unrelated parallel machines: Linear programming strikes back. university of kiel. Technical report, Technical Report 1004, 2010.
- 49 K. Jansen and L. Rohwedder. On integer programming, discrepancy, and convolution. *Mathematics of Operations Research*, 48(3):1481–1495, 2023. doi:10.1287/MOOR.2022.1308.
- 50 C. Jin. An improved FPTAS for 0-1 knapsack. In *Proceedings of the 46th International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 132 of *LIPICs*, pages 76:1–76:14, 2019. doi:10.4230/LIPICs.ICALP.2019.76.
- 51 H. W. Lenstra Jr. Integer programming with a fixed number of variables. *Mathematics of Operations Research*, 8(4):538–548, 1983. doi:10.1287/MOOR.8.4.538.
- 52 R. Kannan. Minkowski’s convex body theorem and integer programming. *Mathematics of Operations Research*, 12(3):415–440, 1987. doi:10.1287/MOOR.12.3.415.
- 53 R. M. Karp. Reducibility among combinatorial problems. In *Complexity of Computer Computations*, pages 85–103. Springer, 1972.
- 54 K. M. Klein. About the complexity of two-stage stochastic IPs. *Mathematical Programming*, pages 1–19, 2021.
- 55 D. Knop and M. Koutecký. Scheduling meets n-fold integer programming. *Journal of Scheduling*, 21(5):493–503, 2018. doi:10.1007/S10951-017-0550-0.
- 56 D. Knop, M. Koutecký, A. Levin, M. Mnich, and S. Onn. Multitype integer monoid optimization and applications. *arXiv preprint arXiv:1909.07326*, 2019.
- 57 D. Knop, M. Koutecký, A. Levin, M. Mnich, and S. Onn. High-multiplicity n-fold IP via configuration LP. *Mathematical Programming*, 200(1):199–227, 2023. doi:10.1007/S10107-022-01882-9.
- 58 D. Knop, M. Koutecký, and M. Mnich. Combinatorial n-fold integer programming and applications. *Mathematical Programming*, 184(1):1–34, 2020. doi:10.1007/S10107-019-01402-2.
- 59 D. Knop, M. Pilipczuk, and M. Wrochna. Tight complexity lower bounds for integer linear programming with few constraints. *ACM Transactions on Computation Theory*, 12(3):1–19, 2020. doi:10.1145/3397484.

- 60 M. Koutecký, A. Levin, and S. Onn. A parameterized strongly polynomial algorithm for block structured integer programs. In *Proceedings of the 45th International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 107 of *LIPIcs*, pages 85:1–85:14, 2018. doi:10.4230/LIPIcs.ICALP.2018.85.
- 61 A. Kulik and H. Shachnai. There is no EPTAS for two-dimensional knapsack. *Information Processing Letters*, 110(16):707–710, 2010. doi:10.1016/J.IPL.2010.05.031.
- 62 A. Lassota and K. Ligthart. Parameterized algorithms for matching integer programs with additional rows and columns. In *Proceedings of the 52nd International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 334 of *LIPIcs*, pages 112:1–112:18, 2025. doi:10.4230/LIPIcs.ICALP.2025.112.
- 63 K. Lee, J. Y. T. Leung, Z. Jia, W. Li, M. L. Pinedo, and B. M. T. Lin. Fast approximation algorithms for bi-criteria scheduling with machine assignment costs. *European Journal of Operational Research*, 238(1):54–64, 2014. doi:10.1016/J.EJOR.2014.03.026.
- 64 W. Li, X. Chen, M. Y. Kovalyov, and J. Ou. Budgeted identical parallel machine scheduling problem to minimize makespan. *Unpublished manuscript*, 2025.
- 65 W. Li and J. Ou. Approximation algorithms for scheduling parallel machines with an energy constraint in green manufacturing. *European Journal of Operational Research*, 314(3):882–893, 2024. doi:10.1016/J.EJOR.2023.11.008.
- 66 X. Mao. $(1 - \varepsilon)$ -approximation of knapsack in nearly quadratic time. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing (STOC)*, pages 295–306, 2024.
- 67 M. Mucha, K. Węgrzycki, and M. Włodarczyk. A subquadratic approximation scheme for partition. In *Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 70–88. SIAM, 2019.
- 68 T. Oertel, J. Paat, and R. Weismantel. A colorful Steinitz Lemma with application to block-structured integer programs. *Mathematical Programming*, 204(1):677–702, 2024. doi:10.1007/S10107-023-01971-3.
- 69 S. Onn. Theory and applications of n-fold integer programming. In *Mixed Integer Nonlinear Programming*, pages 559–593. Springer, 2011.
- 70 C. H. Papadimitriou. On the complexity of integer programming. *Journal of the ACM*, 28(4):765–768, 1981. doi:10.1145/322276.322287.
- 71 R. Ravi and A. Sinha. Hedging uncertainty: Approximation algorithms for stochastic optimization problems. *Mathematical Programming*, 108:97–114, 2006. doi:10.1007/S10107-005-0673-5.
- 72 V. Reis and T. Rothvoss. The subspace flatness conjecture and faster integer programming. In *Proceedings of the 64th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 974–988. IEEE, 2023.
- 73 D. B. Shmoys and C. Swamy. Stochastic optimization is (almost) as easy as deterministic optimization. In *Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 228–237. IEEE, 2004.
- 74 D. B. Shmoys and É. Tardos. An approximation algorithm for the generalized assignment problem. *Mathematical Programming*, 62(1):461–474, 1993. doi:10.1007/BF01585178.
- 75 C. Swamy and D. B. Shmoys. Sampling-based approximation algorithms for multi-stage stochastic optimization. In *Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 357–366. IEEE, 2005.