Parameterized Algorithms for MILPs with Small Treedepth

Cornelius Brand,[1] **Martin Koutecký,**[1] **Sebastian Ordyniak,**[2]
[1] Computer Science Institute, Charles University, Prague, Czech Republic,
[2] School of Computing, University of Leeds, United Kingdom,
{cbrand, koutecky}@iuuk.mff.cuni.cz, s.ordyniak@leeds.ac.uk

Abstract

Solving (mixed) integer (linear) programs, (M)I(L)Ps for short, is a fundamental optimisation task with a wide range of applications in artificial intelligence and computer science in general. While hard in general, recent years have brought about vast progress for solving structurally restricted, (non-mixed) ILPs: n-fold, tree-fold, 2-stage stochastic and multi-stage stochastic programs admit efficient algorithms, and all of these special cases are subsumed by the class of ILPs of small treedepth.

In this paper, we extend this line of work to the mixed case, by showing an algorithm solving MILP in time $f(a, d)\,\mathrm{poly}(n)$, where a is the largest coefficient of the constraint matrix, d is its treedepth, and n is the number of variables.

This is enabled by proving bounds on the denominators (fractionality) of the vertices of bounded-treedepth (non-integer) linear programs. We do so by carefully analysing the inverses of invertible sub-matrices of the constraint matrix. This allows us to afford scaling up the mixed program to the integer grid, and applying the known methods for integer programs.

We then trace the limiting boundary of our "bounded fractionality" approach both in terms of going beyond MILP (by allowing non-linear objectives) as well as its usefulness for generalising other important known tractable classes of ILP. On the positive side, we show that our result can be generalised from MILP to MIP with piece-wise linear separable convex objectives with integer breakpoints. On the negative side, we show that going even slightly beyond such objectives or considering other natural related tractable classes of ILP leads to unbounded fractionality.

Finally, we show that restricting the structure of only the integral variables in the constraint matrix does not yield tractable special cases.

Introduction

Integer Linear Programming (ILP) is a fundamental hard problem as well as a widely used and highly successful framework for solving difficult computational problems in AI, e.g., problems related to planning (van den Briel, Vossen, and Kambhampati 2005; Vossen et al. 1999), vehicle routing (Toth and Vigo 2001), process scheduling (Floudas and Lin 2005), packing (Lodi, Martello, and Monaci 2002), and network hub location (Alumur and Kara 2008) that can often be solved efficiently using a translation to ILP. This naturally motivates the search for tractable classes for ILP. In the '80s, Lenstra and Kannan (Kannan 1987; Lenstra 1983) and Papadimitriou (Papadimitriou 1981) have shown that the classes of ILPs with few variables or few constraints and small coefficients, respectively, are polynomially solvable. A line of research going back almost 20 years (Hemmecke, Onn, and Romanchuk 2013; Chen and Marx 2018; Eisenbrand, Hunkenschröder, and Klein 2018; Aschenbrenner and Hemmecke 2007; Hemmecke, Köppe, and Weismantel 2014; Ganian, Ordyniak, and Ramanujan 2017; Ganian and Ordyniak 2018; Dvorák et al. 2017) has recently culminated with the discovery of another tractable class of ILPs (Eisenbrand et al. 2019; Koutecký, Levin, and Onn 2018), namely ILPs with small treedepth and coefficients. The obtained results already found various algorithmic applications in areas such as scheduling (Knop and Koutecký 2018; Chen et al. 2017; Jansen et al. 2018), stringology and social choice (Knop, Koutecký, and Mnich 2017a,b), and the travelling salesman problem (Chen and Marx 2018).

The language of "special tractable cases" has been developed in the theory of parameterized complexity (Cygan et al. 2015). We say that a problem is *fixed-parameter tractable* (FPT) parameterized by k if it has an algorithm solving every instance I in time $f(k)\,\mathrm{poly}(|I|)$ for some computable function f, and we call this an FPT *algorithm*. Say that the height of a rooted forest is its largest root-leaf distance. A graph $G = (V, E)$ has treedepth d if d is the smallest height of a rooted forest $F = (V, E')$ in which each edge of G is between an ancestor-descendant pair in F, and we write $\mathrm{td}(G) = d$. The *primal graph* $G_P(A)$ of a matrix $A \in \mathbb{R}^{m \times n}$ has a vertex for each column of A, and two vertices are connected if an index $k \in [m] = \{1, \ldots, m\}$ exists such that both columns are non-zero in row k. The *dual graph* $G_D(A)$ is defined as $G_D(A) := G_P(A^\intercal)$. Define the *primal treedepth of* A to be $\mathrm{td}_P(A) = \mathrm{td}(G_P(A))$, and analogously $\mathrm{td}_D(A) = \mathrm{td}(G_D(A))$. The recent results state that there is an algorithm solving ILP in time $f(\|A\|_\infty, \min\{\mathrm{td}_P(A), \mathrm{td}_D(A)\})\,\mathrm{poly}(n)$, hence ILP is FPT parameterized by $\|A\|_\infty$ and $\min\{\mathrm{td}_P(A), \mathrm{td}_D(A)\}$. Besides this class, other parameterizations of ILP have been successfully employed to show tractability results, such as bounding the treewidth of the primal graph and the largest variable domain (Jansen and Kratsch 2015), the treewidth of the incidence graph and the largest solution prefix sum (Ganian, Ordyniak, and Ramanujan 2017), or the signed clique-width of the incidence graph (Eiben et al. 2018).

It is therefore natural to ask whether these tractability results can be generalised to more general settings than ILP. In this paper we ask this question for Mixed ILP (MILP), where both integer and non-integer variables are allowed:

$$\min \{\mathbf{cx} \mid A\mathbf{x} = \mathbf{b}, \mathbf{l} \leq \mathbf{x} \leq \mathbf{u}, \mathbf{x} \in \mathbb{Z}^z \times \mathbb{Q}^q\} , \qquad (1)$$

with $A \in \mathbb{Z}^{m \times z+q}$, $\mathbf{l}, \mathbf{u}, \mathbf{c} \in \mathbb{Z}^{z+q}$ and $\mathbf{b} \in \mathbb{Z}^m$.

MILP is a prominent modelling tool widely used in practice. For example, Bixby (Bixby 2002) says in his famous analysis of LP solver speed-ups, *"[I]nteger programming, and most particularly the mixed-integer variant, is the dominant application of linear programming in practice."* Already Lenstra has shown that MILP with few integer variables is polynomially solvable, naturally extending his result on ILPs with few variables. Analogously, we seek to extend the recent tractability results from ILP to MILP, most importantly for the parameterization by treedepth and largest coefficient. Our main result is as follows:

Theorem 1. *MILP is* FPT *parameterized by* $\|A\|_\infty$ *and* $\min\{\mathrm{td}_P(A), \mathrm{td}_D(A)\}$.

We note that our result also extends to the inequality form of MILP with constraints of the form $A\mathbf{x} \leq \mathbf{b}$ by the fact that introducing slack variables does not increase treedepth too much (Eisenbrand et al. 2019, Lemma 56).

The proof goes by reducing an MILP instance to an ILP instance whose parameters do not increase too much, and then applying the existing algorithms for ILP. A key technical result concerns the *fractionality* of an MILP instance, which is the minimum of the maxima of the denominators in optimal solutions. For example, it is well-known that the natural LP for the VERTEX COVER problem has half-integral optima, that is, there exists an optimum with all values in $\{0, \frac{1}{2}, 1\}$.

The usual way to go about proving fractionality bounds is via Cramer's rule and a sufficiently good bound on the determinant. As witnessed by any proper integer multiple of the identity, determinants can grow large even for matrices of very benign structure. Instead, we need to analyse much more carefully the structure of the inverse of the appearing invertible sub-matrices, allowing us to show:

Theorem 2. *A MILP instance with a constraint matrix A has an optimal solution $\mathbf{x}$ whose largest denominator is bounded by $(\|A\|_\infty)^{d!}(d!)^{d!/2}$, where $d = \min\{\mathrm{td}_P(A), \mathrm{td}_D(A)\}$.*

We are not aware of any prior work which lifts a positive result for ILP to a result for MILP in this way.

We also explore the limits of approaching the problem by bounding the fractionality of inverses: Other ILP classes with parameterized algorithms involve constraint matrices with small primal treewidth (Jansen and Kratsch 2015), small incidence treewidth (Ganian, Ordyniak, and Ramanujan 2017), small signed clique-width (Eiben et al. 2018) and 4-block n-fold matrices (Hemmecke, Köppe, and Weismantel 2014). Here, we obtain a negative answer: For each of these parameters, there exist families of MILP-instances with constant parameters, but unbounded fractionality. This is detailed in Lemma 18 below. The produced families also show that Theorem 2 is almost optimal:

Corollary 3. *There is a MILP instance with $\mathrm{td}_P(A), \mathrm{td}_D(A) = d$, $\|A\|_\infty = 2$, and fractionality 2^{2^d}.*

Compare this with our upper bound $2^{2^{d+\log d+\log\log d}}$. Next, we consider extending the positive result of Theorem 1 to separable convex functions, which is the regime considered in (Eisenbrand et al. 2019). We show that merely bounding the fractionality will unfortunately not suffice, which is detailed in Lemma 20 below. However, we show that for one important class of separable convex objectives, the fractionality does *not* increase, specifically: piece-wise linear functions with integer breakpoints. Let f be any separable convex function, and define f' to agree with f on integer points, and to be linear between them. In a sense, f' is an approximation of f which has a simpler structure. Using f' as a proxy for f is thus common in practice (Bazaraa, Sherali, and Shetty 2013; Lin et al. 2013). Moreover, functions of this form appear in applications of IPs with small treedepth (Knop, Koutecký, and Mnich 2017a; Bredereck et al. 2020).

Theorem 4. *MIP is* FPT *parameterized by* $\|A\|_\infty$ *and* $\min\{\mathrm{td}_P(A), \mathrm{td}_D(A)\}$ *if the objective function is piece-wise linear separable convex with integer breakpoints.*

By appropriate scaling, the integrality of breakpoints in the preceding theorem can be relaxed to requiring only breakpoints with fractionality bounded in the parameters.

Finally, we consider a different way to extend tractable ILP classes to MILP. Divide the constraint matrix A of an MILP instance in two parts corresponding to the integer and continuous variables as $A = (A_\mathbb{Z}\ A_\mathbb{Q})$. What structural restrictions have to be placed on $A_\mathbb{Z}$ and $A_\mathbb{Q}$ in order to obtain tractability of MILP? We show a general hardness result in this direction, which is made precise in Lemma 21. Note that the main reason for intractability is that we allow arbitrary interactions between the integer and the non-integer variables of the instance. Thus, Lemma 21 implies that this interaction between integral and fractional variables has to be restricted in some way in order to obtain a tractable fragment of MILP.

Related Work

We have already mentioned related work on structural parameterizations of ILP. The closest work to ours was done by Hemmecke (Hemmecke 2003) in 2003 when he studied a mixed-integer test set related to the Graver basis, which is the engine behind all recent progress on ILPs of small treedepth. It is unclear how to apply his approach, however, because it requires bounding the norm of elements of the mixed-integer test set, where the bound obtained by (a strengthening of) (Hemmecke 2003, Lemma 19),(Hemmecke 2001, Lemma 2.7.2), is polynomial in n, too much to obtain an FPT algorithm. Kotnyek (Kotnyek 2002) characterised k-integral matrices, i.e., matrices whose solutions have fractionality bounded by k, however it is unclear how his characterisation could be used to show Theorem 2, so we take a different route. Lenstra (Lenstra 1983) showed how to solve MILPs with few integer variables using the fact that a projection of a polytope is again a polytope; applying this approach to our case would require us to show that if P is a polytope described by inequalities with small treedepth, then a projection of P also has an inequality description of small treedepth. This is unclear. In a vein somewhat similar to our bounded-fractionality approach, ideas related to half-integrality have recently led to improved FPT algorithms (Iwata, Wahlstrom, and Yoshida 2016; Iwata, Yamaguchi, and Yoshida 2018; Guillemot 2011), some of which have been experimentally evaluated (Pilipczuk and Ziobro 2018). More fundamentally, half-integrality of two-commodity flow (Hu 1963; Karzanov 1998) and VERTEX COVER (Nemhauser and Trotter 1974) has been known and made use of for half a century.

Preliminaries

We consider zero a natural number, i.e., $0 \in \mathbb{N}$. We write vectors in boldface (e.g., $\mathbf{x}, \mathbf{y}$) and their entries in normal font (e.g., the i-th entry of $\mathbf{x}$ is x_i). For positive integers $m \leq n$ we set $[m, n] := \{m, \ldots, n\}$ and $[n] := [1, n]$. The following Proposition is well-known, but crucial.

Proposition 5. *Let $A \in \mathbb{Z}^{n \times n}$ be a full rank square matrix. Then,* $\mathrm{frac}(A^{-1}) \leq (\|A\|_\infty)^n n^{n/2}$.

Proof. Let $\mathrm{adj}(A)$ be the adjugate matrix of A (which has certain subdeterminants of A as entries.) Cramer's rule states that $A^{-1} = \frac{1}{\det(A)} \cdot \mathrm{adj}(A)$ holds. Hadamard's bound on $|\det(A)|$ can be derived from the fact that $|\det(A)$ is the volume of the parallelepiped spanned by the columns A_i of A. This in turn has the product of their lengths as an upper bound, attained precisely on pairwise orthogonal columns. Therefore, $|\det(A)| \leq \prod_{i=1}^{n} ||A_i||_2$, whence the Proposition easily follows. □

The proof makes it clear what makes it necessary to go beyond Cramer's rule: $|\det(A)|$ might be prohibitively large, while cancellations in $\mathrm{adj}(A)/\det(A)$ may lead to low fractionality in A^{-1} nonetheless.

Reducing MILP to ILP

Assume that an MILP instance is given and that some optimum $\mathbf{x} = (\mathbf{x}_\mathbb{Z}, \mathbf{x}_\mathbb{Q})$ exists whose set of denominators is D, and we know $M = \max D$. Recall $\mathrm{lcm}(D)$ is the least common multiple of the elements of D, and $\mathrm{lcm}(D) \leq M! =: \tilde{M}$. Then $\mathrm{lcm}(D)\mathbf{x}_\mathbb{Q}$ is an integral vector. Our idea here is to restrict our search among all optima of (1) to search among those optima with small fractionality, that is, with small denominators. Consider the *integralized MILP* instance:

$$\min\{(\tilde{M}\mathbf{c}_\mathbb{Z}\ \mathbf{c}_\mathbb{Q})\mathbf{z} \colon \mathbf{z} \in \mathbb{Z}^{z+q}, (\tilde{M} \cdot A_\mathbb{Z}\ A_\mathbb{Q})\mathbf{z} = \tilde{M} \cdot \mathbf{b}, \\ (\mathbf{l}_\mathbb{Z}, \tilde{M}\mathbf{l}_\mathbb{Q}) \leq (\mathbf{z}_\mathbb{Z}, \mathbf{z}_\mathbb{Q}) \leq (\mathbf{u}_\mathbb{Z}\tilde{M}\mathbf{u}_\mathbb{Q})\} \tag{2}$$

We claim that the optimum of (1) can be recovered from the optimum of (2):

Lemma 6. *Let M be the fractionality of* (1) *and $(\mathbf{z}_\mathbb{Z}\ \mathbf{z}_\mathbb{Q}) \in \mathbb{Z}^{z+q}$ be an optimum of* (2)*. Then $\mathbf{x} = (\mathbf{z}_\mathbb{Z}\ \frac{1}{\tilde{M}}\mathbf{z}_\mathbb{Q})$ is an optimum of* (1).

Proof. It is clear that there is a bijection between solutions $\mathbf{x}$ of (1) where $\mathbf{x}_\mathbb{Q}$ has all entries with a denominator $\tilde{M}$ and solutions $\mathbf{z}$ of (2). The optimality of $\mathbf{x}$ then follows from M being the fractionality of (1) and $M!$ always being divisible by $\mathrm{lcm}(D)$. □

The Graphs of A and Treedepth

We assume that $G_P(A)$ and $G_D(A)$ are connected, otherwise A has (up to row and column permutations) a block diagonal structure and solving (1) amounts to solving smaller (1) instances (for each block) independently.

Definition 7 (Treedepth)**.** *The* closure $\mathrm{cl}(F)$ *of a rooted tree F is the graph obtained from F by making every vertex adjacent to all of its ancestors. The* height *of a tree F denoted $\mathrm{ht}(F)$ is the maximum number of vertices on any root-leaf path. We denote by $\mathrm{dt}_F(v)$ the* depth *of vertex v in F, i.e., the number of vertices on the path from v to the root of F. A* td-decomposition of G *is a tree F such that $G \subseteq \mathrm{cl}(F)$. The* treedepth $\mathrm{td}(G)$ *of a connected graph G is the minimum height of its* td*-decompositions.*

To facilitate the analysis of our results we use two parameters called topological height (introduced by Eisenbrand et al. (Eisenbrand et al. 2019)) and topological length:

Definition 8 (Topological height and Topological length)**.** *A vertex of a rooted tree F is* degenerate *if it has exactly one child, and* non-degenerate *otherwise (i.e., if it is a leaf or has at least two children). The* topological height of F*, denoted $\mathrm{th}(F)$, is the maximum number of non-degenerate vertices on any root-leaf path in F. The* topological length of F*, denoted $\mathrm{tl}(F)$, is the maximum number of consecutive degenerate vertices on any root-leaf path in F. Clearly,* $\mathrm{th}(F), \mathrm{tl}(F) \leq \mathrm{ht}(F)$.

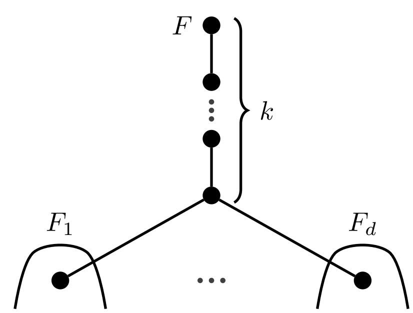

Figure 1: The treedepth decomposition F of $G_P(A)$ for the situation in Lemma 9.

We also need a lemma from (Eisenbrand et al. 2019); refer also to Figure 1 for an illustration.

Lemma 9 (Primal Decomposition (Eisenbrand et al. 2019, Lemma 19))**.** *Let $A \in \mathbb{Z}^{m \times n}$, $G_P(A)$, and a* td*-decomposition F of $G_P(A)$ be given, where $n, m \geq 1$. Then there exists an algorithm computing in time $\mathcal{O}(n)$ a decom-*

position of A

$$A = \begin{pmatrix} \bar{A}_1 & A_1 & & \\ \vdots & & \ddots & \\ \bar{A}_d & & & A_d \end{pmatrix}, \qquad \text{(block-structure)}$$

and td-*decompositions* $F_1, \ldots, F_d$ *of* $G_P(A_1), \ldots, G_P(A_d)$, *respectively, where* $d \in \mathbb{N}$, $\bar{A}_i \in \mathbb{Z}^{m_i \times k}$, $A_i \in \mathbb{Z}^{m_i \times n_i}$, $\mathrm{th}(F_i) \leq \mathrm{th}(F) - 1$, $\mathrm{ht}(F_i) \leq \mathrm{ht}(F) - k$, $k \leq \mathrm{tl}(F)$, *for* $i \in [d]$, $n_1, \ldots, n_d, m_1, \ldots, m_d \in \mathbb{N}$.

Fractionality of Bounded-Treedepth Matrices

This section is devoted to a proof of our main tractability result stated in Theorem 1, i.e., showing that MILP (like ILP) is fixed-parameter tractable parameterized by $\|A\|_\infty$ and $d = \min\{\mathrm{td}_P(A), \mathrm{td}_D(A)\}$. The main ingredient for the proof is Theorem 2 providing a bound on the fractionality of an optimal solution for MILP:

Theorem 2. *A MILP instance with a constraint matrix* A *has an optimal solution* $\mathbf{x}$ *whose largest denominator (fractionality) is bounded by* $(\|A\|_\infty)^{d!}(d!)^{d!/2}$.

We start by observing that the fractionality of an optimal solution of a MILP instance can be obtained from the fractionality of the inverse of some full rank square sub-matrix of the non-integer part of the constraint matrix A. Consider any optimal solution $(\mathbf{x}^*_\mathbb{Z}, \mathbf{x}^*_\mathbb{Q})$ of (1). The fractional part $\mathbf{x}^*_\mathbb{Q}$ is necessarily an optimal solution of the *linear* program

$$\begin{aligned} \min\{\mathbf{c}\mathbf{x}_\mathbb{Q} :\ & A_\mathbb{Q}\mathbf{x}_\mathbb{Q} = \mathbf{b} - A_\mathbb{Z}\mathbf{x}^*_\mathbb{Z}, \\ & \mathbf{l}_\mathbb{Q} \leq \mathbf{x}_\mathbb{Q} \leq \mathbf{u}_\mathbb{Q}, \mathbf{x}_\mathbb{Q} \in \mathbb{Q}^q\} \ . \end{aligned} \tag{3}$$

To bound the fractionality of (1), it therefore suffices to consider the fractionality of (3), and we shall hence assume that $A = A_\mathbb{Q}$.

Let us now recall some basic facts about vertices of polytopes adapted to the specifics of our situation. Consider a vertex of the polytope described by the solutions of the system of

$$A\mathbf{x} = \mathbf{b}, \mathbf{l} \leq \mathbf{x} \leq \mathbf{u}\,, \tag{4}$$

with $A, \mathbf{b}, \mathbf{x}, \mathbf{l}, \mathbf{u}$ as usual. Let $\mathbf{x}$ be any solution of (4). Being a vertex means satisfying n linearly independent constraints with equality. Without loss of generality (Eisenbrand et al. 2019, Proposition 4), A has full rank.

Since these first m equations necessarily hold for any solution $\mathbf{x}$, we have m linearly independent constraints satisfied, and there remain $n - m$ of the in total $2n$ upper and lower bounds to be satisfied. Without loss of generality, we may assume that it is indeed the first $n - m$ lower bound constraints that are met with equality, that is, $x_1 = l_1, \ldots, x_{n-m} = l_{n-m}$ holds. Let

$$\begin{aligned} \mathbf{x}_N &= (x_1, \ldots, x_{n-m}) \in \mathbb{Q}^{n-m}, \\ \mathbf{x}_B &= (x_{n-m+1}, \ldots, x_n) \in \mathbb{Q}^m\,, \end{aligned}$$

and partition accordingly the n columns of A as $A = (A_N\ A_B)$. Letting $\mathbf{b}' = \mathbf{b} - A_N\mathbf{x}_N$, the solution $\mathbf{x} = (\mathbf{x}_N, \mathbf{x}_B)$ satisfies

$$A_B\mathbf{x}_B = \mathbf{b}'. \tag{5}$$

Observe that $A_B \in \mathbb{Z}^{m \times m}$ is a square matrix with trivial kernel (that is, $A\mathbf{x} = \mathbf{0}$ only for $\mathbf{x} = \mathbf{0}$), thus invertible. Therefore, $\mathbf{x}_B = A_B^{-1}\mathbf{b}'$. (Otherwise, there is a direction $\mathbf{y}$ in the kernel such that both $\mathbf{x} + \epsilon\mathbf{y}$ and $\mathbf{x} - \epsilon\mathbf{y}$ are feasible, hence $\mathbf{x}$ was not a vertex.) Hence, in order to bound the fractionality of the vertex $\mathbf{x}$, it is enough to bound the fractionalities of the entries of A_B^{-1}. Therefore, to bound the fractionality of (1), it is sufficient to bound the fractionality of the inverse of any full rank square sub-matrix of the constraint matrix A. We will denote with $\mathrm{frac}(A)$ the *fractionality* of A, meaning the maximum denominator appearing over all entries, represented as fractions in lowest terms, of A. We will start by showing Theorem 2 for the case of primal treedepth, i.e., taking into account the discussion thus far (together with the fact that the treedepth of any sub-matrix of A is bounded by the treedepth of A) it is sufficient to show that:

Lemma 10. *Let* A *be a square matrix with full rank having a* td*-decomposition* F *of* $G_P(A)$. *Then,* $\mathrm{frac}(A^{-1})$ *is at most* $(\|A\|_\infty)^b b^{b/2}$, *where* $b = \min\{\mathrm{tl}(F)^{\mathrm{th}(F)+1}(\mathrm{th}(F)!), \mathrm{ht}(F)!\}$.

Remark 11. *Note that to show the bound stated in Theorem 2, it is sufficient to show the lemma for* $b = (\mathrm{ht}(F)!)$. *However, the bound given in Lemma 10 allows us to obtain better bounds for important special cases. For instance, for the case of 2-stage stochastic and* n*-fold ILP, we obtain that* $\mathrm{frac}(A^{-1}) \leq (\|A\|_\infty)^{2t^3}(2t^3)^{t^3}$ *since* $\mathrm{th}(F) = 2$ *and* $t = \mathrm{tl}(F)$ *is the block size.*

The main idea for the proof of Lemma 10 is to show that the matrix A contains a small sub-matrix A' with at most b columns and rows such that the fractionality of A^{-1} is at most the fractionality of $(A')^{-1}$, which can be bounded using Proposition 5. Towards showing this, we will employ a pruning procedure that works along the td-decomposition F of $G_P(A)$ in a bottom-up manner. The crucial ingredient of this procedure is given in Lemma 13 that in essence allows us to remove all but at most $\mathrm{dt}_F(v)$ many children (together with the columns and rows induced by the variables contained in the sub-trees below those children) of any non-degenerate vertex v of F. The following lemma shows a general property for the fractionality of the inverse of a matrix that makes this pruning step possible.

Lemma 12. *Let* $A \in \mathbb{Z}^{n \times n}$ *be a square matrix with full rank of the form* $\begin{pmatrix} B & 0 \\ R & A_D \end{pmatrix}$, *where* A_D *is a block diagonal matrix. Then, there is a block* A_B *in* A_D *such that* $\mathrm{frac}(A^{-1}) \leq \mathrm{frac}(A_R^{-1})$, *where* A_R *is obtained from* A *after removing all columns and rows from* A *that are in* A_D *but not in* A_B.

Proof. Note that both B and A_D are full rank square matrices because A_D is a square matrix and A is a full rank square matrix. By elementary matrix calculus, the inverse of A is given by $\begin{pmatrix} B^{-1} & 0 \\ R' & A_D^{-1} \end{pmatrix}$, where $R' = -A_D^{-1} \cdot R \cdot B^{-1}$.

Let e be an entry of A^{-1} with the maximum fractionality (among all entries in A^{-1}). If e is contained in B^{-1}, then setting A_B to an arbitrary block of A_D satisfies the claim of the lemma. If e is in A_D^{-1}, then setting A_B to be the block in A_D containing e satisfies the lemma. This is because A_D is block diagonal, and therefore the inverse of A_D is the block diagonal matrix of the inverses of the blocks. Finally, if e is contained in R', then because $R' = -A_D^{-1} \cdot R \cdot B^{-1}$, the entry e of R' is obtained by multiplying a row r of A_D^{-1} with a column of $R \cdot B^{-1}$. Therefore and because R has only integer entries, setting A_B to be the block of A_D having a non-zero entry at row r satisfies the claim of the lemma. $\square$

For a set of variables V, the sub-matrix of A *induced* on V contains all columns that correspond to a variable in V projected onto all rows of A that have a non-zero entry in at least one column in V.

Lemma 13. *Let $A \in \mathbb{Z}^{n\times n}$ be a square matrix with full rank having a* td*-decomposition F of $G_P(A)$, let v be a non-degenerate vertex of F and let C_v be the set of all children of v in F. Then there is a set C of at most $\mathrm{dt}_F(v)$ children of v in F such that the sub-matrix A_P of A obtained after removing all rows and columns in the sub-matrix of A induced on the set of all variables occurring in any sub-tree of F rooted at a child in $C_v \setminus C$, satisfies:*

- *A_P is a square matrix with full rank,*
- $\mathrm{frac}(A^{-1}) \leq \mathrm{frac}(A_P^{-1})$.

Proof. Let A_v be the sub-matrix of A induced on all variables occurring in the sub-tree of F rooted at v. Then A is of the form $\begin{pmatrix} B & 0 \\ R & A_v \end{pmatrix}$, since all rows not in A_v only have zero entries at all columns in A_v. Let r be the number of non-zero columns in R. Note that $r \leq \mathrm{dt}_F(v) - 1$ and because of Lemma 9, we obtain that A_v is of the form (block-structure), with $d = |C|$, and where $\bar{A}_j$ only contains the column corresponding to the variable v. Consider a block A_j with dimensions $m_j \times n_j$. Since A has full rank, $m_j \geq n_j$. Otherwise, the columns of A_j would not be linearly independent in A. Because A has full rank, we also obtain that $r + 1 + \sum_{j=1}^{|C|} n_j = \sum_{j=1}^{|C|} m_j$. Therefore, the number r' of different values for j such that $m_j > n_j$ is at most $r + 1$. W.l.o.g., we can assume that the first r' inequalities are strict and consequently A_v has the form $\begin{pmatrix} B' & 0 \\ R' & A_D \end{pmatrix}$, where A_D is a block diagonal square matrix (consisting of the blocks $A_{r'+1}, \ldots, A_{|C|}$) and B' consists only of the blocks $A_1, \ldots, A_{r'}$. Note that A now has the form $\begin{pmatrix} B & 0 \\ R & A_D \end{pmatrix}$ and satisfies the conditions in Lemma 12. Let A_k be the block of A_D, whose existence is ensured by Lemma 12. We claim that setting C to the children corresponding to the blocks $A_1, \ldots, A_{r'}, A_k$ satisfies the statement of the lemma. Indeed, $|C| \leq r' + 1 \leq \mathrm{dt}_F(v)$. Moreover, A_P is a square matrix with full rank because so is A and the removed blocks A_j are squares. Finally, $\mathrm{frac}(A^{-1}) \leq \mathrm{frac}(A_P^{-1})$ by Lemma 12. $\square$

The following lemma now shows how to apply the reduction given in Lemma 13 along the td-decomposition F, to obtain a sub-matrix of A with at most b columns and rows.

Lemma 14. *Let $A \in \mathbb{Z}^{n\times n}$ be a square matrix with full rank having a* td*-decomposition F of $G_P(A)$. Then there exists a sub-matrix A_P of A having at most $b = \min\{\mathrm{tl}(F)^{\mathrm{th}(F)+1}(\mathrm{th}(F)!), \mathrm{ht}(F)!\}$ columns and rows such that $\mathrm{frac}(A^{-1}) \leq \mathrm{frac}(A_P^{-1})$.*

Proof. Note that Lemma 13 allows us to reduce the size of A while not decreasing the fractionality of its inverse as long as F contains a non-degenerate vertex v with more than $\mathrm{dt}_F(v)$ children. To see this let A_P be the sub-matrix of A obtained after applying the lemma for some non-degenerate vertex v of F. Then A_P together with the td-decomposition obtained from F after removing the sub-trees rooted by a child in $C_v \setminus C$ again satisfy the conditions in the statement of the lemma and moreover $\mathrm{frac}(A^{-1}) \leq \mathrm{frac}(A_P^{-1})$. Let A_P be the sub-matrix obtained from A after applying the reduction rule given by Lemma 13 exhaustively and let F_P be the td-decomposition of $G_P(A_P)$. Then $\mathrm{frac}(A^{-1}) \leq \mathrm{frac}(A_P^{-1})$ and moreover every vertex v in F_P has at most $\mathrm{dt}_F(v)$ children, which implies that F_P has at most $b = \min\{\mathrm{tl}(F)^{\mathrm{th}(F)+1}(\mathrm{th}(F)!), \mathrm{ht}(F)!\}$ vertices. Therefore, A_P has at most b columns (and rows) and satisfies the statement of the lemma. $\square$

We are now ready to show Lemma 10.

Proof of Lemma 10. Let A_P be the sub-matrix of A, whose existence is ensured by Lemma 14. Because $\mathrm{frac}(A^{-1}) \leq \mathrm{frac}(A_P^{-1})$, it suffices to provide the bound for $\mathrm{frac}(A_P^{-1})$. Recall that A_P has at most b columns and rows. Therefore, by Proposition 5, the fractionality of the inverse of A_P is at most $(\|A\|_\infty)^b b^{b/2}$, as required. $\square$

The following corollary shows that the fractionality can be bounded in the same manner in terms of the treedepth of the dual graph.

Corollary 15. *Let A be a square matrix with full rank having a* td*-decomposition F of $G_D(A)$. Then, $\mathrm{frac}(A^{-1})$ is at most $(\|A\|_\infty)^b b^{b/2}$, where $b = \min\{\mathrm{tl}(F)^{\mathrm{th}(F)+1}(\mathrm{th}(F)!), \mathrm{ht}(F)!\}$.*

Proof. Because $G_P(A^\intercal) = G_D(A)$, we obtain that F is a td-decomposition of $G_P(A^\intercal)$. Therefore, Lemma 10 implies that $\mathrm{frac}((A^\intercal)^{-1})$ is at most $(\|A\|_\infty)^b b^{b/2}$. The corollary now follows because $(A^{-1})^\intercal = (A^\intercal)^{-1}$. $\square$

Theorem 2 now follows immediately from Lemma 10 and Corollary 15, which allows us to conclude with the proof of our main tractability result of this section.

Proof of Theorem 1. Theorem 2 gives us an exact bound M' on the largest coefficient of the (2) instance, and it is clear that the structure of non-zeroes (hence the primal and dual graphs) of the constraint matrix of (2) is identical to that of A.

Hence, by Lemma 6, (1) can be solved by solving (2), which can be done (by the results of (Eisenbrand et al. 2019)) in FPT time parameterized by $\|A\|_\infty$ and $\min\{\mathrm{td}_P(A), \mathrm{td}_D(A)\}$). (To be precise, we need to solve (2) for every $1 \le \tilde{M} \le M'$.) ☐

Piece-wise Linear Separable Convex Objectives

A generalisation of (1) to non-linear objectives is

$$\min\{f(\mathbf{x}) \mid A\mathbf{x} = \mathbf{b}, \mathbf{l} \le \mathbf{x} \le \mathbf{u}, \mathbf{x} \in \mathbb{Z}^z \times \mathbb{Q}^q\}\ , \quad (6)$$

and here we focus on the case when f is separable convex, meaning $f(\mathbf{x}) = \sum_{i=1}^n f_i(x_i)$ with $f_i : \mathbb{R} \to \mathbb{R}$ univariate convex for each $i \in [n]$. Moreover, we assume that f is piece-wise linear with breakpoints at integer points, i.e., for every $a \in \mathbb{R}$ and $i \in [n]$, $f_i(a) = \{a\} f_i(\lfloor a \rfloor) + (1 - \{a\}) f_i(\lceil a \rceil)$, where $\{a\} = a - \lfloor a \rfloor$.

We adapt a variable transformation of Hochbaum and Shantikumar (Hochbaum and Shantikumar 1990) to show that (6) admits a linearization that retains the fractionality of the original linear instance. This transformation was originally used to show that integer separable convex minimisation can be reduced to integer linear minimisation when A has small sub-determinants, but our use differs in three aspects: our variables are mixed integer, the matrix A may have large sub-determinants, and most importantly, we only use it to obtain a fractionality bound; we never need to solve the newly constructed instance. Specifically, we will transform an input (6) into a (1) whose parameters we define next:

$$\min\left\{\mathbf{c}\mathbf{y} \mid \hat{A}\mathbf{y} = \hat{\mathbf{b}},\, \mathbf{0} \le \mathbf{y} \le \mathbf{1},\, \mathbf{y} \in \mathbb{Z}^{\hat{z}} \times \mathbb{Q}^{\hat{q}}\right\} \quad (7)$$

The hatted data are obtained as follows: For each $i \in [z+q]$, replace the variable x_i with $u_i - l_i$ variables y_i^j, $j \in [u_i - l_i]$. Hence, in (7), and the number of integer variables is $\hat{z} = \sum_{i=1}^{z} u_i - l_i$, the number of continuous variables is $\hat{q} = \sum_{i=z+1}^{z+q} u_i - l_i$. We define the column of $\hat{A}$ corresponding to the variable y_i^j.The lower and upper bound for all variables is 0 and 1, respectively. Let the right-hand side be $\hat{\mathbf{b}} = \mathbf{b} - A\mathbf{l} = \sum_{i=1}^{z+q} A_i l_i$. Finally, the coefficient in the objective function $\mathbf{c}$ for variable y_i^j is the slope of f_i between points $l_i + (j-1)$ and $l_i + j$. Specifically, $c_i^j = f_i(l_i + j) - f_i(l_i + (j-1))$. Define a mapping $\varphi : \mathbb{Z}^z \times \mathbb{Q}^q \to \mathbb{Z}^{\hat{z}} \times \mathbb{Q}^{\hat{q}}$, as follows: given $\mathbf{x} \in \mathbb{Z}^z \times \mathbb{Q}^q$, $\varphi(\mathbf{x}) = \mathbf{y}$, where for each $i \in [z+q]$, $j \in [u_i - l_i]$, $y_i^j = \max\{0, \min\{1, x_i - l_i - (j-1)\}\}$.

Lemma 16. *1. A vector* $\mathbf{x}$ *is feasible in* (6) *iff* $\varphi(\mathbf{x})$ *is feasible in* (7).

2. If $\mathbf{l} \le \mathbf{x} \le \mathbf{u}$, *then* $f(\mathbf{x}) = \mathbf{c}\varphi(\mathbf{x}) + \sum_{i=1}^{z+q} f_i(l_i)$.

3. Let $\mathbf{x}^*$ *be an optimum of* (6). *Then* $\varphi(\mathbf{x}^*)$ *is an optimum of* (7).

The proof amounts to careful checking of the construction, and is deferred to the full version.

Lemma 17. *Every square sub-matrix* A' *of* $\hat{A}$ *of full rank has* $\mathrm{td}_P(A') \le \mathrm{td}_P(A)$ *and* $\mathrm{td}_D(A') \le \mathrm{td}_D(A)$.

Proof. For A' to have full rank, it cannot contain duplicate columns. Hence, A' is also a square sub-matrix of A, a case in which we have already shown the claim to hold. ☐

Proof of Theorem 4. By this lemma, the fractionality M of (7) is bounded by $\mathrm{frac}(A)$, and by Lemma 16 and the definition of φ, $\mathrm{frac}(A)$ is also a fractionality bound on (1) when f is separable convex piece-wise linear with integer breakpoints. Let $\hat{f}$ be defined component-wise from f as follows: for $i \in [1, z]$, let $\hat{f}_i = f_i$, and for $i \in [z+1, z+q]$, let $\hat{f}_i(x_i) = f_i(x_i/M)$. Then, to solve (1) in this regime, it is enough to optimise $\hat{f}$ over (2). ☐

Limits of the "Bounded Fractionality" Approach

In this section, we show the limits of our "bounded fractionality" approach. We start by showing its limits for various important known tractable classes of ILP, i.e., the class of small primal treewidth and domain (Jansen and Kratsch 2015), small incidence treewidth and largest solution prefix sum (Ganian, Ordyniak, and Ramanujan 2017), small signed clique-width of the incidence graph (Eiben et al. 2018), and the class of 4-block n-fold matrices (Hemmecke, Köppe, and Weismantel 2014). We show that all these classes exhibit unbounded fractionality.

Lemma 18. *For every* $n \in \mathbb{N}$, *there are MILP instances* I_1 *and* I_2 *with constraint matrices* A_1 *and* A_2, *such that* A_1 *has constant primal, dual, and incidence treewidth and signed incidence clique-width and* $\|A_1\|_\infty = 2$, *and* A_2 *is 4-block* n*-fold with all blocks being just* (1), *and the fractionality is* $2^{\Omega(n)}$ *for* I_1 *and* $\Omega(n)$ *for* I_2.

Proof. Consider the $n \times n$ matrix

$$A_1 = \begin{pmatrix} 2 & -1 & 0 & \cdots & 0 \\ 0 & 2 & -1 & \cdots & 0 \\ \vdots & & \ddots & & \vdots \\ 0 & 0 & \cdots & & 2 \end{pmatrix},$$

It is easy to verify that the matrix B with $B_{ij} = 2^{i-j-1}$ for $i \le j$ and $B_{ij} = 0$ otherwise is the inverse of A_1. Moreover, the primal, dual, incidence treewidth of A_1 is at most 1, the signed incidence clique-width of A_1 is at most 2, and $\|A_1\|_\infty = 2$.

It is again easy to verify that below are A_2 and its inverse, both $n \times n$, with $n' = n - 2$, and A_2 is a 4-block n-fold matrix with all blocks of size 1:

$$A_2 = \begin{pmatrix} 1 & 1 & 1 & \cdots & 1 \\ 1 & 1 & 0 & \cdots & 0 \\ 1 & 0 & 1 & \cdots & 0 \\ \vdots & & & \ddots & \\ 1 & 0 & 0 & \cdots & 1 \end{pmatrix},$$

$$A_2^{-1} = \begin{pmatrix} -\frac{1}{n'} & \frac{1}{n'} & \frac{1}{n'} & \cdots & \frac{1}{n'} \\ \frac{1}{n'} & \frac{n'-1}{n} & -\frac{1}{n'} & \cdots & -\frac{1}{n'} \\ \frac{1}{n'} & -\frac{1}{n'} & \frac{n'-1}{n} & \cdots & -\frac{1}{n'} \\ \vdots & & & \ddots & \\ \frac{1}{n'} & -\frac{1}{n'} & -\frac{1}{n'} & \cdots & \frac{n'-1}{n} \end{pmatrix}.$$

Because for each vertex $\mathbf{x}$ of a polyhedron there exists an objective vector $\mathbf{c}$ such that (1) is uniquely optimal in $\mathbf{x}$, and the fact that we have demonstrated inverses with high fractionality, there must exist vertices of high fractionality and corresponding objectives, which give the desired instances I_1 and I_2. □

Remark 19. *The $\Omega(n)$ fractionality lower bound in part 2 of Lemma 18 may be seen as mild given that for 4-block n-fold we would seek an algorithm running in time $n^{f(k)}$, for f some function and k largest block size, and that (the more permissive) n-fold IP problem has such an algorithm even when its entries are polynomial in n. However, this is not true for the 2-stage stochastic IP problem, which is NP-hard with polynomially bounded coefficients already with constant-size blocks (Dvorák et al. 2017). Because 4-block n-fold IP is at least as hard as 2-stage stochastic IP, the bounded fractionality approach cannot work for giving an $n^{f(k)}$ algorithm for 4-block n-fold MILP.*

We now show the limits of our approach for generalising our results from MILP to MIP for certain types of separable convex functions.

Lemma 20. *There are MIP instances with the following properties:*

1. *$A = (1 \cdots 1)$, $b = 1$, $f(\mathbf{x}) = \sum_i (x_i)^2$, $\mathrm{td}_D(A) = 1$, fractionality n,*
2. *dimension 1, no constraints, $f(x) = (x - \frac{1}{k})^2$, fractionality k,*
3. *dimension 1, no equality constraints, $0 \le x \le 1$, $f(x) = x^3 + 2x^2 - x$ univariate cubic convex, unbounded fractionality (minimum is $\frac{\sqrt{7}}{3} - \frac{2}{3}$).*

Proof. All instances have unique optima, and it is straightforward to verify that in part 1 of the Lemma, it is the point $\mathbf{x} = (\frac{1}{n}, \ldots, \frac{1}{n})$, in part 2 it is $x = \frac{1}{k}$, for any k, and in part 3, the minimum is irrational $x = \frac{\sqrt{7}}{3} - \frac{2}{3}$, hence fractionality is unbounded. The objective $f(x) = x^3 + 2x^2 - x$ is not convex on $\mathbb{R}$, but it is between 0 and 1. □

The Limits of Tractability for Structured MILPs

It is well-known that MILP is fixed-parameter tractable parameterized by the number of integer variables. It is therefore natural to ask, whether for our Theorem 1 it could be sufficient to only put restrictions on the integer part of the instance. Here, we show that this is not the case. We show hardness for the feasibility version of MILP, which is deciding the nonemptiness of the set $\{\mathbf{x} \in \mathbb{Z}^z \times \mathbb{Q}^q \mid A\mathbf{x} = \mathbf{b}, \mathbf{l} \le \mathbf{x} \le \mathbf{u}\}$.

Lemma 21. *Let $\mathcal{C}$ be a class of ILP instances for which the feasibility problem is NP-hard. Then there exists a class of MILP instances $\mathcal{C}'$ whose feasibility problem is NP-hard and whose constraint matrix is $A = \begin{pmatrix} \mathbf{0} & A_{\mathbb{Q}} \\ I & -I \end{pmatrix}$, where I is the identity matrix and $A_{\mathbb{Q}}$ is a constraint matrix of an instance from $\mathcal{C}$.*

The proof is deferred to the full version of the paper.

Remark 22. *It is an interesting question for future work whether we can generalise our results for MILP if we put additional restrictions on the interactions between integer and non-integer variables. A similar approach has recently been explored for generalising the tractability result for ILP based on primal treedepth to MILP (Ganian, Ordyniak, and Ramanujan 2017) using a hybrid decompositional parameter called torso-width.*

Open Problems

We close with three open problems motivating future research. First, what is the complexity of general MIP for matrices with bounded primal and dual treedepth? Our Lemma 20 shows that a different approach is needed. Second, is 4-block n-fold MILP in XP? At first sight, it may seem that to get an XP algorithm, it should suffice to bound the fractionality by $\mathrm{poly}(n)$ (and nothing better is possible by Lemma 18). However, the current XP algorithm for the pure integer case depends *exponentially* on the largest coefficient of the constraint matrix, so solving (2) would be too slow. Third, Lemma 21 suggests that new tractable fragments of MILP may be characterized by having bounded interaction between the integer and continuous variables. Hence, we ask: what is the complexity of MILP where $A_{\mathbb{Z}}$ comes from an ILP tractable fragment, $A_{\mathbb{Q}}$ is arbitrary, and the number of rows which are nonzero in both the integer and continuous variables is small? If this is hard, what restraints need to be placed on $A_{\mathbb{Q}}$ to obtain a tractable fragment?

Acknowledgements

Cornelius Brand was supported by OP RDE project No. CZ.02.2.69/0.0/0.0/18_053/0016976 International mobility of research, technical and administrative staff at the Charles University.

Martin Koutecký was partially supported by Charles University project UNCE/SCI/004 and by the project 19-27871X of GA ČR. Sebastian Ordyniak acknowledges support from the Engineering and Physical Sciences Research Council (EPSRC, project EP/V00252X/1).

References

Alumur, S. A.; and Kara, B. Y. 2008. Network hub location problems: The state of the art. *European Journal of Operational Research* 190(1): 1–21.

Aschenbrenner, M.; and Hemmecke, R. 2007. Finiteness theorems in stochastic integer programming. *Foundations of Computational Mathematics* 7(2): 183–227.

Bazaraa, M. S.; Sherali, H. D.; and Shetty, C. M. 2013. *Nonlinear programming: theory and algorithms*. John Wiley & Sons.

Bixby, R. E. 2002. Solving real-world linear programs: A decade and more of progress. *Operations research* 50(1): 3–15.

Bredereck, R.; Faliszewski, P.; Niedermeier, R.; Skowron, P.; and Talmon, N. 2020. Mixed integer programming with convex/concave constraints: Fixed-parameter tractability and applications to multicovering and voting. *Theor. Comput. Sci.* 814: 86–105. doi:10.1016/j.tcs.2020.01.017. URL https://doi.org/10.1016/j.tcs.2020.01.017.

Chen, L.; and Marx, D. 2018. Covering a tree with rooted subtrees - parameterized and approximation algorithms. In Czumaj, A., ed., *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, LA, USA, January 7-10, 2018*, 2801–2820. SIAM.

Chen, L.; Marx, D.; Ye, D.; and Zhang, G. 2017. Parameterized and Approximation Results for Scheduling with a Low Rank Processing Time Matrix. In Vollmer, H.; and Vallée, B., eds., *34th Symposium on Theoretical Aspects of Computer Science, STACS 2017, March 8-11, 2017, Hannover, Germany*, volume 66 of *LIPIcs*, 22:1–22:14. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik.

Cygan, M.; Fomin, F. V.; Kowalik, L.; Lokshtanov, D.; Marx, D.; Pilipczuk, M.; Pilipczuk, M.; and Saurabh, S. 2015. *Parameterized Algorithms*. Springer.

Dvorák, P.; Eiben, E.; Ganian, R.; Knop, D.; and Ordyniak, S. 2017. Solving Integer Linear Programs with a Small Number of Global Variables and Constraints. In Sierra, C., ed., *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19-25, 2017*, 607–613. ijcai.org.

Eiben, E.; Ganian, R.; Knop, D.; and Ordyniak, S. 2018. Unary Integer Linear Programming with Structural Restrictions. In Lang, J., ed., *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden*, 1284–1290. ijcai.org.

Eisenbrand, F.; Hunkenschröder, C.; Klein, K.; Koutecký, M.; Levin, A.; and Onn, S. 2019. An Algorithmic Theory of Integer Programming. *CoRR* abs/1904.01361.

Eisenbrand, F.; Hunkenschröder, C.; and Klein, K.-M. 2018. Faster Algorithms for Integer Programs with Block Structure. In *45th International Colloquium on Automata, Languages, and Programming (ICALP 2018)*, 49:1–49:13.

Floudas, C.; and Lin, X. 2005. Mixed Integer Linear Programming in Process Scheduling: Modeling, Algorithms, and Applications. *Annals of Operations Research* 139(1): 131–162. ISSN 0254-5330.

Ganian, R.; and Ordyniak, S. 2018. The complexity landscape of decompositional parameters for ILP. *Artif. Intell.* 257: 61–71.

Ganian, R.; Ordyniak, S.; and Ramanujan, M. S. 2017. Going Beyond Primal Treewidth for (M)ILP. In Singh, S. P.; and Markovitch, S., eds., *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA*, 815–821. AAAI Press.

Guillemot, S. 2011. FPT algorithms for path-transversal and cycle-transversal problems. *Discrete Optimization* 8(1): 61–71.

Hemmecke, R. 2001. *On the decomposition of test sets*. Ph.D. thesis, Universität Duisburg.

Hemmecke, R. 2003. On the positive sum property and the computation of Graver test sets. *Math. Program.* 96(2): 247–269.

Hemmecke, R.; Köppe, M.; and Weismantel, R. 2014. Graver basis and proximity techniques for block-structured separable convex integer minimization problems. *Mathematical Programming* 145(1-2, Ser. A): 1–18. ISSN 0025-5610.

Hemmecke, R.; Onn, S.; and Romanchuk, L. 2013. N-fold integer programming in cubic time. *Mathematical Programming* 1–17.

Hochbaum, D. S.; and Shantikumar, J. G. 1990. Convex Separable Optimization Is Not Much Harder than Linear Optimization. *J. ACM* 37(4): 843–862.

Hu, T. C. 1963. Multi-Commodity Network Flows. *Operations Research* 11(3): 344–360.

Iwata, Y.; Wahlstrom, M.; and Yoshida, Y. 2016. Half-integrality, LP-branching, and FPT algorithms. *SIAM Journal on Computing* 45(4): 1377–1411.

Iwata, Y.; Yamaguchi, Y.; and Yoshida, Y. 2018. 0/1/All CSPs, Half-Integral A-Path Packing, and Linear-Time FPT Algorithms. In *59th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2018, Paris, France, October 7-9, 2018*, 462–473.

Jansen, B. M. P.; and Kratsch, S. 2015. A Structural Approach to Kernels for ILPs: Treewidth and Total Unimodularity. In Bansal, N.; and Finocchi, I., eds., *Algorithms - ESA 2015 - 23rd Annual European Symposium, Patras, Greece, September 14-16, 2015, Proceedings*, volume 9294 of *Lecture Notes in Computer Science*, 779–791. Springer.

Jansen, K.; Klein, K.; Maack, M.; and Rau, M. 2018. Empowering the Configuration-IP - New PTAS Results for Scheduling with Setups Times. In Blum, A., ed., *10th Innovations in Theoretical Computer Science Conference, ITCS 2019, January 10-12, 2019, San Diego, California, USA*, volume 124 of *LIPIcs*, 44:1–44:19. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik.

Kannan, R. 1987. Minkowski's convex body theorem and integer programming. *Mathematics of Operations Research* 12(3): 415–440. ISSN 0364-765X.

Karzanov, A. V. 1998. Minimum 0-extensions of graph metrics. *European Journal of Combinatorics* 19(1): 71–101.

Knop, D.; and Koutecký, M. 2018. Scheduling meets n-fold integer programming. *J. Scheduling* 21(5): 493–503.

Knop, D.; Koutecký, M.; and Mnich, M. 2017a. Combinatorial n-fold Integer Programming and Applications. In Pruhs, K.; and Sohler, C., eds., *25th Annual European Symposium on Algorithms, ESA 2017, September 4-6, 2017, Vienna, Austria*, volume 87 of *LIPIcs*, 54:1–54:14. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik.

Knop, D.; Koutecký, M.; and Mnich, M. 2017b. Voting and Bribing in Single-Exponential Time. In Vollmer, H.; and Vallée, B., eds., *34th Symposium on Theoretical Aspects of Computer Science, STACS 2017, March 8-11, 2017, Hannover, Germany*, volume 66 of *LIPIcs*, 46:1–46:14. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik.

Kotnyek, B. 2002. *A generalization of totally unimodular and network matrices.* Ph.D. thesis, London School of Economics and Political Science (United Kingdom).

Koutecký, M.; Levin, A.; and Onn, S. 2018. A Parameterized Strongly Polynomial Algorithm for Block Structured Integer Programs. In Chatzigiannakis, I.; Kaklamanis, C.; Marx, D.; and Sannella, D., eds., *45th International Colloquium on Automata, Languages, and Programming, ICALP 2018, July 9-13, 2018, Prague, Czech Republic*, volume 107 of *LIPIcs*, 85:1–85:14. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik.

Lenstra, Jr., H. W. 1983. Integer programming with a fixed number of variables. *Mathematics of Operations Research* 8(4): 538–548.

Lin, M.-H.; Carlsson, J. G.; Ge, D.; Shi, J.; and Tsai, J.-F. 2013. A review of piecewise linearization methods. *Mathematical problems in Engineering* 2013.

Lodi, A.; Martello, S.; and Monaci, M. 2002. Two-dimensional packing problems: A survey. *European Journal of Operational Research* 141(2): 241–252.

Nemhauser, G. L.; and Trotter, L. E. 1974. Properties of vertex packing and independence system polyhedra. *Mathematical Programming* 6(1): 48–61.

Papadimitriou, C. H. 1981. On the complexity of integer programming. *J. ACM* 28(4): 765–768.

Pilipczuk, M.; and Ziobro, M. 2018. Experimental Evaluation of Parameterized Algorithms for Graph Separation Problems: Half-Integral Relaxations and Matroid-based Kernelization. *arXiv preprint arXiv:1811.07779* .

Toth, P.; and Vigo, D., eds. 2001. *The Vehicle Routing Problem*. Philadelphia, PA, USA: Society for Industrial and Applied Mathematics. ISBN 0-89871-498-2.

van den Briel, M.; Vossen, T.; and Kambhampati, S. 2005. Reviving Integer Programming Approaches for AI Planning: A Branch-and-Cut Framework. In *Automated Planning and Scheduling, ICAPS 2005. Proceedings*, 310–319. AAAI.

Vossen, T.; Ball, M. O.; Lotem, A.; and Nau, D. S. 1999. On the Use of Integer Programming Models in AI Planning. In *International Joint Conference on Artificial Intelligence, IJCAI 99. Proceedings*, 304–309. Morgan Kaufmann.