Optimizing Sparsity over Lattices and Semigroups

Iskander Aliev[1], Gennadiy Averkov[2(✉)], Jesús A. De Loera[3], and Timm Oertel[1]

[1] Cardiff University, Cardiff, UK
[2] Brandenburg University of Technology Cottbus-Senftenberg, Senftenberg, Germany
averkov@b-tu.de
[3] University of California, Davis, USA

Abstract. Motivated by problems in optimization we study the *sparsity* of the solutions to systems of linear Diophantine equations and linear integer programs, i.e., the number of non-zero entries of a solution, which is often referred to as the ℓ_0-norm. Our main results are improved bounds on the ℓ_0-norm of sparse solutions to systems $A\boldsymbol{x} = \boldsymbol{b}$, where $A \in \mathbb{Z}^{m \times n}$, $\boldsymbol{b} \in \mathbb{Z}^m$ and $\boldsymbol{x}$ is either a general integer vector (lattice case) or a non-negative integer vector (semigroup case). In the lattice case and certain scenarios of the semigroup case, we give polynomial time algorithms for computing solutions with ℓ_0-norm satisfying the obtained bounds.

1 Introduction

This paper discusses the problem of finding sparse solutions to systems of linear Diophantine equations and integer linear programs. We investigate the ℓ_0-norm $\|\boldsymbol{x}\|_0 := |\{i : x_i \neq 0\}|$, a function widely used in the theory of *compressed sensing* [6,9], which measures the sparsity of a given vector $\boldsymbol{x} = (x_1, \ldots, x_n)^\top \in \mathbb{R}^n$ (it is clear that the ℓ_0-norm is actually not a norm).

Sparsity is a topic of interest in several areas of optimization. The ℓ_0-norm minimization problem over reals is central in the theory of the classical compressed sensing, where a linear programming relaxation provides a guaranteed approximation [8,9]. Support minimization for solutions to Diophantine equations is relevant for the theory of compressed sensing for discrete-valued signals [11,12,17]. There is still little understanding of discrete signals in the compressed sensing paradigm, despite the fact that there are many applications in which the signal is known to have discrete-valued entries, for instance, in wireless communication [22] and the theory of error-correcting codes [7]. Sparsity was also investigated in integer optimization [1,10,20], where many combinatorial optimization problems have useful interpretations as sparse semigroup problems. For example, the edge-coloring problem can be seen as a problem in the semigroup generated by matchings of the graph [18]. Our results provide natural out-of-the-box sparsity bounds for problems with linear constraints and integer variables in a general form.

D. Bienstock and G. Zambelli (Eds.): IPCO 2020, LNCS 12125, pp. 40–51, 2020.
https://doi.org/10.1007/978-3-030-45771-6_4

1.1 Lattices: Sparse Solutions of Linear Diophantine Systems

Each integer matrix $A \in \mathbb{Z}^{m \times n}$ determines the lattice $\mathcal{L}(A) := \{A\boldsymbol{x} : \boldsymbol{x} \in \mathbb{Z}^n\}$ generated by the columns of A. By an easy reduction via row transformations, we may assume without loss of generality that the rank of A is m.

Let $[n] := \{1, \ldots, n\}$ and let $\binom{[n]}{m}$ be the set of all m-element subsets of $[n]$. For $\gamma \subseteq [n]$, consider the $m \times |\gamma|$ submatrix A_γ of A with columns indexed by γ. One can easily prove that the determinant of $\mathcal{L}(A)$ is equal to

$$\gcd(A) := \gcd\left\{\det(A_\gamma) : \gamma \in \binom{[n]}{m}\right\}.$$

Since $\mathcal{L}(A_\gamma)$ is the lattice spanned by the columns of A indexed by γ, it is a sublattice of $\mathcal{L}(A)$. We first deal with a natural question: *Can the description of a given lattice $\mathcal{L}(A)$ in terms of A be made sparser by passing from A to A_γ with γ having a smaller cardinality than n and satisfying $\mathcal{L}(A) = \mathcal{L}(A_\gamma)$?* That is, we want to discard some of the columns of A and generate $\mathcal{L}(A)$ by $|\gamma|$ columns with $|\gamma|$ being possibly small.

For stating our results, we need several number-theoretic functions. Given $z \in \mathbb{Z}_{>0}$, consider the prime factorization $z = p_1^{s_1} \cdots p_k^{s_k}$ with pairwise distinct prime factors $p_1, \ldots, p_k$ and their multiplicities $s_1, \ldots, s_k \in \mathbb{Z}_{>0}$. Then the number of prime factors $\sum_{i=1}^k s_i$ counting the multiplicities is denoted by $\Omega(z)$. Furthermore, we introduce $\Omega_m(z) := \sum_{i=1}^k \min\{s_i, m\}$. That is, by introducing m we set a threshold to account for multiplicities. In the case $m = 1$ we thus have $\omega(z) := \Omega_1(z) = k$, which is the number of prime factors in z, not taking the multiplicities into account. The functions Ω and ω are called *prime Ω-function and prime ω-function*, respectively, in number theory [15]. We call Ω_m the *truncated prime Ω-function.*

Theorem 1. *Let $A \in \mathbb{Z}^{m \times n}$, with $m \le n$, and let $\tau \in \binom{[n]}{m}$ be such that the matrix A_τ is non-singular. Then the equality $\mathcal{L}(A) = \mathcal{L}(A_\gamma)$ holds for some γ satisfying $\tau \subseteq \gamma \subseteq [n]$ and*

$$|\gamma| \le m + \Omega_m\left(\frac{|\det(A_\tau)|}{\gcd(A)}\right). \tag{1}$$

Given A and τ, the set γ can be computed in polynomial time.

One can easily see that $\omega(z) \le \Omega_m(z) \le \Omega(z) \le \log_2(z)$ for every $z \in \mathbb{Z}_{>0}$. The estimate using $\log_2(z)$ gives a first impression on the quality of the bound (1). It turns out, however, that $\Omega_m(z)$ is much smaller on the average. Results in number theory [15, §22.10] show that the average values $\frac{1}{z}(\omega(1) + \cdots + \omega(z))$ and $\frac{1}{z}(\Omega(1) + \cdots + \Omega(z))$ are of order $\log\log z$, as $z \to \infty$.

As an immediate consequence of Theorem 1 we obtain

Corollary 2. *Consider the linear Diophantine system*

$$A\boldsymbol{x} = \boldsymbol{b},\ \boldsymbol{x} \in \mathbb{Z}^n \tag{2}$$

with $A \in \mathbb{Z}^{m\times n}$*,* $\boldsymbol{b} \in \mathbb{Z}^m$ *and* $m \le n$*. Let* $\tau \in \binom{[n]}{m}$ *be such that the* $m \times m$ *matrix* A_τ *is non-singular. If* (2) *is feasible, then* (2) *has a solution* $\boldsymbol{x}$ *satisfying the sparsity bound*

$$\|\boldsymbol{x}\|_0 \le m + \Omega_m\left(\frac{|\det(A_\tau)|}{\gcd(A)}\right).$$

Under the above assumptions, for given $A, \boldsymbol{b}$ *and* τ*, such a sparse solution can be computed in polynomial time.*

From the optimization perspective, Corollary 2 deals with the problem

$$\min\{\|\boldsymbol{x}\|_0 \, : \, A\boldsymbol{x} = \boldsymbol{b}, \ \boldsymbol{x} \in \mathbb{Z}^n\}$$

of minimization of the ℓ_0-norm over the affine lattice $\{\boldsymbol{x} \in \mathbb{Z}^n \, : \, A\boldsymbol{x} = \boldsymbol{b}\}$.

1.2 Semigroups: Sparse Solutions in Integer Programming

Consider next the standard form of the feasibility constraints of integer linear programming

$$A\boldsymbol{x} = \boldsymbol{b}, \ \boldsymbol{x} \in \mathbb{Z}^n_{\ge 0}. \tag{3}$$

For a given matrix A, the set of all $\boldsymbol{b}$ such that (3) is feasible, is the *semigroup* $\mathcal{S}g(A) = \{A\boldsymbol{x} : \boldsymbol{x} \in \mathbb{Z}^n_{\ge 0}\}$ generated by the columns of A.

If (3) has a solution, i.e., $\boldsymbol{b} \in \mathcal{S}g(A)$, *how sparse can such a solution be?* In other words, we are interested in the ℓ_0-norm minimization problem

$$\min\left\{\|\boldsymbol{x}\|_0 \, : \, A\boldsymbol{x} = \boldsymbol{b}, \ \boldsymbol{x} \in \mathbb{Z}^n_{\ge 0}\right\}. \tag{4}$$

It is clear that Problem (4) is NP-hard, because deciding the feasibility of (3) [23, §18.2] or even solving the relaxation of (4) with the condition $\boldsymbol{x} \in \mathbb{Z}^n_{\ge 0}$ replaced by $\boldsymbol{x} \in \mathbb{R}^n$ [19] is NP-hard.

Taking the NP-hardness of Problem (4) into account, our aim is to *estimate* the optimal value of (4) under the assumption that this problem is feasible. In [2, Theorem 1.1 (i)] (see also [1, Theorem 1]), it was shown that for any $\boldsymbol{b} \in \mathcal{S}g(A)$, there exists a $\boldsymbol{x} \in \mathbb{Z}^n$, such that $A\boldsymbol{x} = \boldsymbol{b}$ and

$$\|\boldsymbol{x}\|_0 \le m + \left\lfloor \log_2\left(\frac{\sqrt{\det(AA^\top)}}{\gcd(A)}\right)\right\rfloor. \tag{5}$$

In [1, Theorem 2], it was shown that (5) cannot be improved significantly, but nevertheless we show here how to improve it in some special cases. As a consequence of Theorem 1 we obtain the following.

Corollary 3. *Let* $A \in \mathbb{Z}^{m\times n}$ *be a matrix whose columns positively span* $\mathbb{R}^m$ *and let* $\boldsymbol{b} \in \mathbb{Z}^m$*. Then* $\mathcal{L}(A) = \mathcal{S}g(A)$*. Furthermore, if* $\boldsymbol{b} \in \mathcal{L}(A)$*, and* $\tau \in \binom{[n]}{m}$ *is a set, for which the matrix* A_τ *is non-singular, then there is a solution* $\boldsymbol{x}$ *of the integer-programming feasibility problem* $A\boldsymbol{x} = \boldsymbol{b}, \boldsymbol{x} \in \mathbb{Z}^m_{\ge 0}$ *that satisfies the sparsity bound*

$$\|\boldsymbol{x}\|_0 \le 2m + \Omega_m\left(\frac{|\det(A_\tau)|}{\gcd(A)}\right). \tag{6}$$

Under the above assumptions, for given $A, \boldsymbol{b}$ and τ, such a sparse solution $\boldsymbol{x}$ can be computed in polynomial time.

Note that for a fixed m, (6) is usually much tighter than (5), because the function $\Omega_m(z)$ is bounded from above by the logarithmic function $\log_2(z)$ and is much smaller than $\log_2(z)$ on the average. Furthermore, $|\det(A_\tau)| \leq \sqrt{\det(AA^\top)}$ in view of the Cauchy-Binet formula.

We take a closer look at the case $m = 1$ of a single equation and tighten the given bounds in this case. That is, we consider the *knapsack feasibility problem*

$$\boldsymbol{a}^\top \boldsymbol{x} = b, \ \boldsymbol{x} \in \mathbb{Z}^n_{\geq 0}, \tag{7}$$

where $\boldsymbol{a} \in \mathbb{Z}^n$ and $b \in \mathbb{Z}$. Without loss of generality we can assume that all components of the vector $\boldsymbol{a}$ are not equal to zero. It follows from (5) that a feasible problem (7) has a solution $\boldsymbol{x}$ with

$$\|\boldsymbol{x}\|_0 \leq 1 + \left\lfloor \log_2 \left(\frac{\|\boldsymbol{a}\|_2}{\gcd(\boldsymbol{a})} \right) \right\rfloor . \tag{8}$$

If all components of $\boldsymbol{a}$ have the same sign, without loss of generality we can assume $\boldsymbol{a} \in \mathbb{Z}^n_{>0}$. In this setting, Theorem 1.2 in [2] strengthens the bound (8) by replacing the ℓ_2-norm of the vector $\boldsymbol{a}$ with the ℓ_∞-norm. It was conjectured in [2, page 247] that a bound $\|\boldsymbol{x}\|_0 \leq c + \lfloor \log_2 (\|\boldsymbol{a}\|_\infty / \gcd(\boldsymbol{a})) \rfloor$ with an absolute constant c holds for an *arbitrary* $\boldsymbol{a} \in \mathbb{Z}^n$. We obtain the following result, which covers the case that has not been settled so far and yields a confirmation of this conjecture.

Corollary 4. *Let $\boldsymbol{a} = (a_1, \ldots, a_n)^\top \in (\mathbb{Z} \setminus \{0\})^n$ be a vector that contains both positive and negative components. If the knapsack feasibility problem $\boldsymbol{a}^\top \boldsymbol{x} = b,\ \boldsymbol{x} \in \mathbb{Z}^n_{\geq 0}$ has a solution, then there is a solution $\boldsymbol{x}$ satisfying the sparsity bound*

$$\|\boldsymbol{x}\|_0 \leq 2 + \min \left\{ \omega \left(\frac{|a_i|}{\gcd(\boldsymbol{a})} \right) \ : \ i \in [n] \right\} .$$

Under the above assumptions, for given $\boldsymbol{a}$ and b, such a sparse solution $\boldsymbol{x}$ can be computed in polynomial time.

Our next contribution is that, given additional structure on A, we can improve on [2, Theorem 1.1 (i)], which in turn also gives an improvement on [2, Theorem 1.2]. For $\boldsymbol{a}_1, \ldots, \boldsymbol{a}_n \in \mathbb{R}^m$, we denote by $\operatorname{cone}(\boldsymbol{a}_1, \ldots, \boldsymbol{a}_n)$ the convex conic hull of the set $\{\boldsymbol{a}_1, \ldots, \boldsymbol{a}_n\}$. Now assume the matrix $A = (\boldsymbol{a}_1, \ldots, \boldsymbol{a}_n) \in \mathbb{Z}^{m \times n}$ with columns $\boldsymbol{a}_i$ satisfies the following conditions:

$$\boldsymbol{a}_1, \ldots, \boldsymbol{a}_n \in \mathbb{Z}^m \setminus \{\boldsymbol{0}\}, \tag{9}$$

$$\operatorname{cone}(\boldsymbol{a}_1, \ldots, \boldsymbol{a}_n) \text{ is an } m\text{-dimensional pointed cone,} \tag{10}$$

$$\operatorname{cone}(\boldsymbol{a}_1) \text{ is an extreme ray of } \operatorname{cone}(\boldsymbol{a}_1, \ldots, \boldsymbol{a}_n). \tag{11}$$

Note that the previously best sparsity bound for the general case of the integer-programming feasibility problem is (5). Using the Cauchy-Binet formula, (5) can be written as

$$\|\boldsymbol{x}\|_0 \leq m + \log_2 \frac{\sqrt{\sum_{I \in \binom{[n]}{m}} \det(A_I)^2}}{\gcd(A)}.$$

The following theorem improves this bound in the *"pointed cone case"* by removing a fraction of m/n of terms in the sum under the square root.

Theorem 5. *Let* $A = (\boldsymbol{a}_1, \ldots, \boldsymbol{a}_n) \in \mathbb{Z}^{m \times n}$ *satisfy* (9)–(11) *and, for* $\boldsymbol{b} \in \mathbb{Z}^m$, *consider the integer-programming feasibility problem*

$$A\boldsymbol{x} = \boldsymbol{b},\ \boldsymbol{x} \in \mathbb{Z}^n_{\geq 0}. \tag{12}$$

If (12) *is feasible, then there is a feasible solution* $\boldsymbol{x}$ *satisfying the sparsity bound*

$$\|\boldsymbol{x}\|_0 \leq m + \left\lfloor \log_2 \frac{q(A)}{\gcd(A)} \right\rfloor,$$

where

$$q(A) := \sqrt{\sum_{I \in \binom{[n]}{m} : 1 \in I} \det(A_I)^2}.$$

We omit the proof of this result due to the page limit for the IPCO proceedings. Instead we focus on the particularly interesting case $m = 1$. In this case, assumption (10) is equivalent to $\boldsymbol{a} \in \mathbb{Z}^n_{>0} \cup \mathbb{Z}^n_{<0}$. Without loss of generality, one can assume $\boldsymbol{a} \in \mathbb{Z}^n_{>0}$.

Theorem 6. *Let* $\boldsymbol{a} = (a_1, \ldots, a_n)^\top \in \mathbb{Z}^n_{>0}$ *and* $b \in \mathbb{Z}_{\geq 0}$. *If the knapsack feasibility problem* $\boldsymbol{a}^\top \boldsymbol{x} = b,\ \boldsymbol{x} \in \mathbb{Z}^n_{\geq 0}$ *has a solution, there is a solution* $\boldsymbol{x}$ *satisfying the sparsity bound*

$$\|\boldsymbol{x}\|_0 \leq 1 + \left\lfloor \log_2 \left(\frac{\min\{a_1, \ldots, a_n\}}{\gcd(\boldsymbol{a})} \right) \right\rfloor.$$

When dealing with bounds for sparsity it would be interesting to understand *the worst case scenario among all members of the semigroup*, which is described by the function

$$\mathrm{ICR}(A) = \max_{\boldsymbol{b} \in \mathcal{S}g(A)} \min\{\|\boldsymbol{x}\|_0 \,:\, A\boldsymbol{x} = \boldsymbol{b},\ \boldsymbol{x} \in \mathbb{Z}^n_{\geq 0}\}. \tag{13}$$

We call $\mathrm{ICR}(A)$ the *integer Carathéodory rank* in resemblance to the classical problem of finding the integer Carathéodory number for Hilbert bases [24]. Above results for the problem $A\boldsymbol{x} = \boldsymbol{b},\ \boldsymbol{x} \in \mathbb{Z}^n_{\geq 0}$ can be phrased as upper bounds on $\mathrm{ICR}(A)$. We are interested in the complexity of computing $\mathrm{ICR}(A)$. The first question is: *can the integer Carathéodory rank of a matrix A be computed at all?* After all, remember that the semigroup has infinitely many elements and, despite the fact that $\mathrm{ICR}(A)$ is a finite number, a direct usage of (13) would result into the determination of the sparsest representation $A\boldsymbol{x} = \boldsymbol{b}$ for all of the infinitely many elements $\boldsymbol{b}$ of $\mathcal{S}g(A)$. It turns out that $\mathrm{ICR}(A)$ is computable, as the inequality $\mathrm{ICR}(A) \le k$ can be expressed as the formula $\forall \boldsymbol{x} \in \mathbb{Z}^n_{\geq 0} \, \exists \boldsymbol{y} \in \mathbb{Z}^n_{\geq 0} : (A\boldsymbol{x} = A\boldsymbol{y}) \wedge (\|\boldsymbol{y}\|_0 \le k)$ in *Presburger arithmetic* [14]. Beyond this fact, the complexity status of computing $\mathrm{ICR}(A)$ is largely open, even when A is just one row:

Problem 7. *Given the input $\boldsymbol{a} = (a_1, \dots, a_n)^\top \in \mathbb{Z}^n$, is it NP-hard to compute* $\mathrm{ICR}(\boldsymbol{a}^\top)$*?*

The *Frobenius number* $\max \mathbb{Z}_{\geq 0} \setminus \mathcal{S}g(\boldsymbol{a}^\top)$, defined under the assumptions $\boldsymbol{a} \in \mathbb{Z}^n_{>0}$ and $\gcd(\boldsymbol{a}) = 1$, is yet another value associated to $\mathcal{S}g(\boldsymbol{a}^\top)$. The Frobenius number can be computed in polynomial time when n is fixed [5,16] but is NP-hard to compute when n is not fixed [21]. It seems that there might be a connection between computing the Frobenius number and $\mathrm{ICR}(\boldsymbol{a}^\top)$.

2 Proofs of Theorem 1 and its consequences

The proof of Theorem 1 relies on the theory of finite Abelian groups. We write Abelian groups additively. An Abelian group G is said to be a *direct sum* of its finitely many subgroups $G_1, \dots, G_m$, which is written as $G = \bigoplus_{i=1}^m G_i$, if every element $x \in G$ has a unique representation as $x = x_1 + \cdots + x_m$ with $x_i \in G_i$ for each $i \in [m]$. A *primary cyclic group* is a non-zero finite cyclic group whose order is a power of a prime number. We use G/H to denote the quotient of G modulo its subgroup H.

The fundamental theorem of finite Abelian groups states that every finite Abelian group G has a *primary decomposition*, which is essentially unique. This means, G is decomposable into a direct sum of its primary cyclic groups and that this decomposition is unique up to automorphisms of G. We denote by $\kappa(G)$ the number of direct summands in the primary decomposition of G.

For a subset S of a finite Abelian group G, we denote by $\langle S \rangle$ the subgroup of G generated by S. We call a subset S of G *non-redundant* if the subgroups $\langle T \rangle$ generated by proper subsets T of S are properly contained in $\langle S \rangle$. In other words, S is non-redundant if $\langle S \setminus \{x\} \rangle$ is a proper subgroup of $\langle S \rangle$ for every $x \in S$. The following result can be found in [13, Lemma A.6].

Theorem 8. *Let G be a finite Abelian group. Then the maximum cardinality of a non-redundant subset S of G is equal to $\kappa(G)$.*

We will also need the following lemmas, proved in the Appendix.

Lemma 1. *Let G be a finite Abelian group representable as a direct sum $G = \bigoplus_{j=1}^m G_j$ of $m \in \mathbb{Z}_{>0}$ cyclic groups. Then $\kappa(G) \le \Omega_m(|G|)$.*

Lemma 2. *Let Λ be a sublattice of $\mathbb{Z}^m$ of rank $m \in \mathbb{Z}^m_{>0}$. Then $G = \mathbb{Z}^m/\Lambda$ is a finite Abelian group of order* $\det(\Lambda)$ *that can be represented as a direct sum of at most m cyclic groups.*

Proof (Theorem 1). Let $\boldsymbol{a}_1, \ldots, \boldsymbol{a}_n$ be the columns of A. Without loss of generality, let $\tau = [m]$. We use the notation $B := A_\tau$.

Reduction to the case $\gcd(A) = 1$*.* For a non-singular square matrix M, the columns of $M^{-1}A$ are representations of the columns of A in the basis of columns of M. In particular, for a matrix M whose columns form a basis of $\mathcal{L}(A)$, the matrix $M^{-1}A$ is integral and the $m \times m$ minors of $M^{-1}A$ are the respective $m \times m$ minors of A divided by $\det(M) = \gcd(A)$. Thus, replacing A by $M^{-1}A$, we pass from $\mathcal{L}(A)$ to $\mathcal{L}(M^{-1}A) = \left\{M^{-1}z : z \in \mathcal{L}(A)\right\}$, which corresponds to a change of a coordinate system in $\mathbb{R}^m$ and ensures that $\gcd(A) = 1$.

Sparsity bound (1)*.* The matrix B gives rise to the lattice $\Lambda := \mathcal{L}(B)$ of rank m, while Λ determines the finite Abelian group $\mathbb{Z}^m/\Lambda$.

Consider the canonical homomorphism $\phi : \mathbb{Z}^m \to \mathbb{Z}^m/\Lambda$, sending an element of $\mathbb{Z}^m$ to its coset modulo Λ. Since $\gcd(A) = 1$, we have $\mathcal{L}(A) = \mathbb{Z}^m$, which implies $\langle T \rangle = \mathbb{Z}^m/\Lambda$ for $T := \{\phi(\boldsymbol{a}_{m+1}), \ldots, \phi(\boldsymbol{a}_n)\}$. For every non-redundant subset S of T, we have

$$\begin{aligned} |S| &\leq \kappa(\mathbb{Z}^m/\Lambda) && \text{(by Theorem 8)} \\ &\leq \Omega_m(|\det(A_\tau)|) && \text{(by Lemmas 1 and 2).} \end{aligned}$$

Fixing a set $I \subseteq \{m+1, \ldots, n\}$ that satisfies $|I| = |S|$ and $S = \{\phi(\boldsymbol{a}_i) : i \in I\}$, we reformulate $\langle S \rangle = \mathbb{Z}^m/\Lambda$ as $\mathbb{Z}^m = \mathcal{L}(A_I) + \Lambda = \mathcal{L}(A_I) + \mathcal{L}(A_\tau) = \mathcal{L}(A_{I \cup \tau})$. Thus, (1) holds for $\gamma = I \cup \tau$.

Construction of γ *in polynomial time.* The matrix M used in the reduction to the case $\gcd(A) = 1$ can be constructed in polynomial time: one can obtain M from the Hermite Normal Form of A (with respect to the column transformations) by discarding zero columns. For the determination of γ, the set I that defines the non-redundant subset $S = \{\phi(\boldsymbol{a}_i) : i \in I\}$ of $\mathbb{Z}^m/\Lambda$ needs to be determined. Start with $I = \{m+1, \ldots, n\}$ and iteratively check if some of the elements $\phi(\boldsymbol{a}_i) \in \mathbb{Z}^m/\Lambda$, where $i \in I$, is in the group generated by the remaining elements. Suppose $j \in I$ and we want to check if $\phi(\boldsymbol{a}_j)$ is in the group generated by all $\phi(\boldsymbol{a}_i)$ with $i \in I \setminus \{j\}$. Since $\Lambda = \mathcal{L}(A_\tau)$, this is equivalent to checking $\boldsymbol{a}_j \in \mathcal{L}(A_{I\setminus\{j\}\cup\tau})$ and is thus reduced to solving a system of linear Diophantine equations with the left-hand side matrix $A_{I\setminus\{j\}\cup\tau}$ and the right-hand side vector $\boldsymbol{a}_j$. Thus, carrying the above procedure for every $j \in I$ and removing j from I whenever $\boldsymbol{a}_j \in \mathcal{L}(A_{I\setminus\{j\}\cup\tau})$, we eventually arrive at a set I that determines a non-redundant subset S of $\mathbb{Z}^m/\Lambda$. This is done by solving at most $n - m$ linear Diophantine systems in total, where the matrix of each system is a sub-matrix of A and the right-hand vector of the system is a column of A. □

Remark 1 (Optimality of the bounds). For a given $\Delta \in \mathbb{Z}_{\geq 2}$ let us consider matrices $A \in \mathbb{Z}^{m \times n}$ with $\Delta = |\det(A_\tau)|/\gcd(A)$. We construct a matrix A that shows the optimality of the bound (1). As in the proof of Theorem 1, we assume $\tau = [m]$ and use the notation $B = A_\tau$. Consider the prime factorization $\Delta = p_1^{n_1} \cdots p_s^{n_s}$. We will fix the matrix B to be a diagonal matrix with diagonal entries $d_1, \ldots, d_m \in \mathbb{Z}_{>0}$ so that $\det(B) = d_1 \cdots d_m = \Delta$.

The diagonal entries are defined by distributing the prime factors of Δ among the diagonal entries of B. If the multiplicity n_i of the prime p_i is less than m, we introduce p_i as a factor of multiplicity 1 in n_i of the m diagonal entries of B. If the multiplicity n_i is at least m, we are able distribute the factors p_i among *all* of the diagonal entries of B so that each diagonal entry contains the factor p_i with multiplicity at least 1.

The group $\mathbb{Z}^m/\Lambda = \mathbb{Z}^m/\mathcal{L}(B)$ is a direct sum of m cyclic groups $G_1, \ldots, G_m$ of orders $d_1, \ldots, d_m$, respectively. By the Chinese Remainder Theorem, these cyclic groups can be further decomposed into the direct sum of primary cyclic groups. By our construction, the prime factor p_i of the multiplicity $n_i < m$ generates a cyclic direct summand of order p_i in n_i of the subgroups $G_1, \ldots, G_m$. If $n_i \geq m$, then each of the groups $G_1, \ldots, G_m$ has a direct summand, which is a non-trivial cyclic group whose order is a power of p_i. Summarizing, we see that the decomposition of $\mathbb{Z}^m/\Lambda$ into primary cyclic groups contains n_i summands of order p_i, when $n_i < m$, and m summands, whose order is a power of p_i, when $n_i \geq m$. The total number of summands is thus $\sum_{i=1}^{s} \min\{m, n_i\} = \Omega_m(\Delta)$.

Now, fix $n = m + \Omega_m(\Delta)$ and choose columns $\boldsymbol{a}_{m+1}, \ldots, \boldsymbol{a}_n$ so that $\phi(\boldsymbol{a}_{m+1}), \ldots, \phi(\boldsymbol{a}_n)$ generate all direct summands in the decomposition of $\mathbb{Z}^m/\Lambda$ into primary cyclic groups. With this choice, $\phi(\boldsymbol{a}_{m+1}), \ldots, \phi(\boldsymbol{a}_n)$ generate $\mathbb{Z}^m/\Lambda$, which means that $\mathcal{L}(A) = \mathbb{Z}^m$ and implies $\gcd(A) = 1$. On the other hand, any proper subset $\{\phi(\boldsymbol{a}_{m+1}), \ldots, \phi(\boldsymbol{a}_n)\}$ generates a proper subgroup of $\mathbb{Z}^m/\Lambda$, as some of the direct summands in the decomposition of $\mathbb{Z}^m/\Lambda$ into primary cyclic groups will be missing. This means $\mathcal{L}(A_{[m]\cup I}) \subsetneq \mathbb{Z}^m$ for every $I \subsetneq \{m+1, \ldots, n\}$.

Proof (Corollary 2). Feasiblity of (2) can be expressed as $\boldsymbol{b} \in \mathcal{L}(A)$. Choose γ from the assertion of Theorem 1. One has $\boldsymbol{b} \in \mathcal{L}(A) = \mathcal{L}(A_\gamma)$ and so there exists a solution $\boldsymbol{x}$ of (2) whose support is a subset of γ. This sparse solution $\boldsymbol{x}$ can be computed by solving the Diophantine system with the left-hand side matrix A_γ and the right-hand side vector $\boldsymbol{b}$.

Proof (Corollary 3). Assume that the Diophantine system $A\boldsymbol{x} = \boldsymbol{b}$, $\boldsymbol{x} \in \mathbb{Z}^n$ has a solution. It suffices to show that, in this case, the integer-programming feasibility problem $A\boldsymbol{x} = \boldsymbol{b}$, $\boldsymbol{x} \in \mathbb{Z}^n_{\geq 0}$ has a solution, too, and that one can find a solution of the desired sparsity to the integer-programming feasibility problem in polynomial time.

One can determine γ as in Theorem 1 in polynomial time. Using γ, we can determine a solution $\boldsymbol{x}^* = (x_1^*, \ldots, x_n^*)^\top \in \mathbb{Z}^n$ of the Diophantine system $A\boldsymbol{x} = b$, $\boldsymbol{x} \in \mathbb{Z}^n$ satisfying $x_i^* = 0$ for $i \in [n] \setminus \gamma$ in polynomial time, as described in the proof of Corollary 2.

Let $\boldsymbol{a}_1, \ldots, \boldsymbol{a}_n$ be the columns of A. Since the matrix A_τ is non-singular, the m vectors $\boldsymbol{a}_i$, where $i \in \tau$, together with the vector $\boldsymbol{v} = -\sum_{i\in\tau} \boldsymbol{a}_i$ positively span $\mathbb{R}^n$. Since all columns of A positive span $\mathbb{R}^n$, the conic version of the Carathéodory theorem implies the existence of a set $\beta \subseteq [m]$ with $|\beta| \leq m$, such that $\boldsymbol{v}$ is in the conic hull of $\{\boldsymbol{a}_i \,:\, i \in \beta\}$. Consequently, the set $\{\boldsymbol{a}_i \,:\, i \in \beta \cup \tau\}$ and by this also the larger set $\{\boldsymbol{a}_i \,:\, i \in \beta \cup \gamma\}$ positively span $\mathbb{R}^m$. Let $I = \beta \cup \gamma$. By construction, $|I| \leq |\beta| + |\gamma| \leq m + |\gamma|$.

Since the vectors $\boldsymbol{a}_i$ with $i \in I$ positively span $\mathbb{R}^m$, there exist a choice of rational coefficients $\lambda_i > 0$ $(i \in I)$ with $\sum_{i\in I} \lambda_i \boldsymbol{a}_i = 0$. After rescaling we can assume $\lambda_i \in \mathbb{Z}_{>0}$. Define $\boldsymbol{x}' = (x'_1, \ldots, x'_n)^\top \in \mathbb{Z}^n_{\geq 0}$ by setting $x'_i = \lambda_i$ for $i \in I$ and $x'_i = 0$ otherwise. The vector $\boldsymbol{x}'$ is a solution of $A\boldsymbol{x} = \boldsymbol{0}$. Choosing $N \in \mathbb{Z}_{>0}$ large enough, we can ensure that the vector $\boldsymbol{x}^* + N\boldsymbol{x}'$ has non-negative components. Hence, $\boldsymbol{x} = \boldsymbol{x}^* + N\boldsymbol{x}'$ is a solution of the system $A\boldsymbol{x} = \boldsymbol{b},\ \boldsymbol{x} \in \mathbb{Z}^n_{\geq 0}$ satisfying the desired sparsity estimate. The coefficients λ_i and the number N can be computed in polynomial time.

Proof (Corollary 4). The assertion follows by applying Corollary 3 for $m = 1$ and all $\tau = \{i\}$ with $i \in [n]$.

3 Proof of Theorem 6

Lemma 3. *Let $a_1, \ldots, a_t \in \mathbb{Z}_{>0}$, where $t \in \mathbb{Z}_{>0}$. If $t > 1 + \log_2(a_1)$, then the system*

$$\begin{aligned} &y_1 a_1 + \cdots + y_t a_t = 0, \\ &y_1 \in \mathbb{Z}_{\geq 0},\ y_2, \ldots, y_t \in \{-1, 0, 1\}. \end{aligned}$$

in the unknowns $y_1, \ldots, y_t$ has a solution that is not identically equal to zero.

Proof. The proof is inspired by the approach in [3, §3.1] (used in a different context) that suggests to reformulate the underlying equation over integers as two strict inequalities and then use Minkowski's first theorem [4, Ch. VII, Sect. 3] from the geometry of numbers. Consider the convex set $Y \subseteq \mathbb{R}^t$ defined by $2t$ strict linear inequalities

$$\begin{aligned} -1 <& y_1 a_1 + \cdots + y_t a_t < 1, \\ -2 <& y_i < 2 \text{ for all } i \in \{2, \ldots, t\}. \end{aligned}$$

Clearly, the set Y is the interior of a hyper-parallelepiped and can also be described as $Y = \{\boldsymbol{y} \in \mathbb{R}^t : \|M\boldsymbol{y}\|_\infty < 1\}$, where M is the upper triangular matrix

$$M = \begin{pmatrix} a_1 & a_2 & \cdots & a_t \\ & 1/2 & & \\ & & \ddots & \\ & & & 1/2 \end{pmatrix}.$$

It is easy to see that the t-dimensional volume $\mathrm{vol}(Y)$ of Y is

$$\mathrm{vol}(Y) = \mathrm{vol}(M^{-1}[-1, 1]^t) = \frac{1}{\det(M)} 2^t = \frac{4^t}{2a_1}.$$

The assumption $t > 1 + \log_2(a_1)$ implies that the volume of Y is strictly larger than 2^t. Thus, by Minkowski's first theorem, the set Y contains a non-zero integer vector $\boldsymbol{y} = (y_1, \ldots, y_t)^\top \in \mathbb{Z}^t$. Without loss of generality we can assume that $y_1 \geq 0$ (if the latter is not true, one can replace $\boldsymbol{y}$ by $-\boldsymbol{y}$). The vector $\boldsymbol{y}$ is a desired solution from the assertion of the lemma. □

Proof (Theorem 6). Without loss of generality we can assume that $\gcd(\boldsymbol{a}) = 1$. In fact, if b is divisible by $\gcd(\boldsymbol{a})$ we can convert $\boldsymbol{a}^\top \boldsymbol{x} = b$ to $\overline{\boldsymbol{a}}^\top \boldsymbol{x} = \overline{b}$ with $\overline{\boldsymbol{a}} = \frac{\boldsymbol{a}}{\gcd(\boldsymbol{a})}$ and $\overline{b} = \frac{b}{\gcd(\boldsymbol{a})}$, and, if b is not divisible by $\gcd(\boldsymbol{a})$, the knapsack feasibility problem $\boldsymbol{a}^\top \boldsymbol{x} = b,\ \boldsymbol{x} \in \mathbb{Z}^n_{\geq 0}$ has no solution.

Without loss of generality, let $a_1 = \min\{a_1, \dots, a_n\}$. We need to show the existence of solution of the knapsack feasibility problem satisfying $\|\boldsymbol{x}\|_0 \leq 1 + \log_2(a_1)$.

Choose a solution $\boldsymbol{x} = (x_1, \dots, x_n)^\top$ of the knapsack feasibility problem with the property that the number of indices $i \in \{2, \dots, n\}$ for which $x_i \neq 0$ is minimized. Without loss of generality we can assume that, for some $t \in \{2, \dots, n\}$ one has $x_2 > 0, \dots, x_t > 0, x_{t+1} = \dots = x_n = 0$. Lemma 3 implies $t \leq 1 + \log_2(a_1)$. In fact, if the latter was not true, then a solution $\boldsymbol{y} \in \mathbb{R}^t$ of the system in Lemma 3 could be extended to a solution $\boldsymbol{y} \in \mathbb{R}^n$ by appending zero components. It is clear that some of the components $y_2, \dots, y_t$ are negative, because $a_2 > 0, \dots, a_t > 0$. It then turns out that, for an appropriate choice of $k \in \mathbb{Z}_{\geq 0}$, the vector $\boldsymbol{x}' = (x'_1, \dots, x'_n)^\top = \boldsymbol{x} + k\boldsymbol{y}$ is a solution of the same knapsack feasibility problem satisfying $x'_1 \geq 0, \dots, x'_t \geq 0,\ x'_{t+1} = \dots = x'_n = 0$ and $x'_i = 0$ for at least one $i \in \{2, \dots, t\}$. Indeed, one can choose k to be the minimum among all a_i with $i \in \{2, \dots, t\}$ and $y_i = -1$.

The existence of $\boldsymbol{x}'$ with at most $t - 1$ non-zero components x'_i with $i \in \{2, \dots, n\}$ contradicts the choice of $\boldsymbol{x}$ and yields the assertion. □

Acknowledgements. The second author is supported by the DFG (German Research Foundation) within the project number 413995221. The third author acknowledges partial support from NSF grant 1818969.

A Appendix

Proof (Lemma 1). Consider the prime factorization $|G| = p_1^{n_1} \cdots p_s^{n_s}$. Then $|G_j| = p_1^{n_{i,j}} \cdots p_s^{n_{i,j}}$ with $0 \leq n_{i,j} \leq n_i$ and, by the Chinese Remainder Theorem, the cyclic group G_j can be represented as $G_j = \bigoplus_{i=1}^s G_{i,j}$, where $G_{i,j}$ is a cyclic group of order $p_i^{n_{i,j}}$. Consequently, $G = \bigoplus_{i=1}^s \bigoplus_{j=1}^m G_{i,j}$. This is a decomposition of G into a direct sum of primary cyclic groups and, possibly, some trivial summands $G_{i,j}$ equal to $\{0\}$. We can count the non-trivial direct summands whose order is a power of p_i, for a given $i \in [s]$. There is at most one summand like this for each of the groups G_j. So, there are at most m non-trivial summands in the decomposition whose order is a power of p_i. On the other hand, the direct sum of all non-trivial summands whose order is a power of p_i is a group of order $p_i^{n_{i,1} + \dots + n_{i,s}} = p_i^{n_i}$ so that the total number of such summands is not larger than n_i, as every summand contributes the factor at least p_i to the power $p_i^{n_i}$. This shows that the total number of non-zero summands in the decomposition of G is at most $\sum_{i=1}^s \min\{m, n_i\} = \Omega_m(|G|)$. □

Proof (Lemma 2). The proof relies on the relationship of finite Abelian groups and lattices, see [23, §4.4]. Fix a matrix $M \in \mathbb{Z}^{m \times m}$ whose columns form a basis of Λ. Then $|\det(M)| = \det(\Lambda)$. There exist unimodular matrices $U \in \mathbb{Z}^{m \times m}$

and $V \in \mathbb{Z}^{m \times m}$ such that $D := UMV$ is diagonal matrix with positive integer diagonal entries. For example, one can choose D to be the Smith Normal Form of M [23, §4.4]. Let $d_1, \ldots, d_m \in \mathbb{Z}_{>0}$ be the diagonal entries of D. Since U and V are unimodular, $d_1 \cdots d_m = \det(D) = \det(\Lambda)$.

We introduce the quotient group $G' := \mathbb{Z}^m/\Lambda' = (\mathbb{Z}/d_1\mathbb{Z}) \times \cdots \times (\mathbb{Z}/d_m\mathbb{Z})$ with respect to the lattice $\Lambda' := \mathcal{L}(D) = (d_1\mathbb{Z}) \times \cdots \times (d_m\mathbb{Z})$. The order of G' is $d_1 \cdots d_m = \det(D) = \det(\Lambda)$ and G' is a direct sum of at most m cyclic groups, as every $d_i > 1$ determines a non-trivial direct summand.

To conclude the proof, it suffices to show that G' is isomorphic to G. To see this, note that $\Lambda' = \mathcal{L}(D) = \mathcal{L}(UMV) = \mathcal{L}(UM) = \{Uz : z \in \Lambda\}$. Thus, the map $z \mapsto Uz$ is an automorphism of $\mathbb{Z}^m$ and an isomorphism from Λ to Λ'. Thus, $z \mapsto Uz$ induces an isomorphism from the group $G = \mathbb{Z}^m/\Lambda$ to the group $G' = \mathbb{Z}^m/\Lambda'$. □

References

1. Aliev, I., De Loera, J.A., Eisenbrand, F., Oertel, T., Weismantel, R.: The support of integer optimal solutions. SIAM J. Optim. **28**(3), 2152–2157 (2018)
2. Aliev, I., De Loera, J.A., Oertel, T., O'Neill, C.: Sparse solutions of linear diophantine equations. SIAM J. Appl. Algebra Geom. **1**(1), 239–253 (2017)
3. Averkov, G.: On the size of lattice simplices with a single interior lattice point. SIAM J. Discrete Math. **26**(2), 515–526 (2012)
4. Barvinok, A.I.: A Course in Convexity. Graduate Studies in Mathematics, vol. 54. American Mathematical Society, Providence, RI (2002)
5. Barvinok, A.I., Woods, K.: Short rational generating functions for lattice point problems. J. AMS **16**(4), 957–979 (2003)
6. Boche, H., Calderbank, R., Kutyniok, G., Vybíral, J.: A survey of compressed sensing. In: Boche, H., Calderbank, R., Kutyniok, G., Vybíral, J. (eds.) Compressed Sensing and its Applications. ANHA, pp. 1–39. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-16042-9_1
7. Candès, E., Rudelson, M., Tao, T., Vershynin, R.: Error correction via linear programming. In: 46th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2005), pp. 668–681 (2005)
8. Candès, E.J., Romberg, J.K., Tao, T.: Stable signal recovery from incomplete and inaccurate measurements. Comm. Pure Appl. Math. **59**(8), 1207–1223 (2006)
9. Candès, E.J., Tao, T.: Decoding by linear programming. IEEE Trans. Inform. Theory **51**(12), 4203–4215 (2005)
10. Eisenbrand, F., Shmonin, G.: Carathéodory bounds for integer cones. Oper. Res. Lett. **34**(5), 564–568 (2006)
11. Flinth, A., Kutyniok, G.: PROMP: a sparse recovery approach to lattice-valued signals. Appl. Comput. Harmon. Anal. **45**(3), 668–708 (2018)
12. Fukshansky, L., Needell, D., Sudakov, B.: An algebraic perspective on integer sparse recovery. Appl. Math. Comput. **340**, 31–42 (2019)
13. Geroldinger, A., Halter-Koch, F.: Non-Unique Factorizations: Algebraic Combinatorial and Analytic Theory. Pure and Applied Mathematics. Chapman and Hall/CRC, Boca Raton (2006)
14. Haase, C.: A survival guide to Presburger arithmetic. ACM SIGLOG News **5**(3), 67–82 (2018)

15. Hardy, G.H., Wright, E.M., Heath-Brown, R., Silverman, J.: An Introduction to the Theory of Numbers. Oxford Mathematics. OUP, Oxford (2008)
16. Kannan, R.: Lattice translates of a polytope and the frobenius problem. Combinatorica **12**(2), 161–177 (1992)
17. Konyagin, S.V.: On the recovery of an integer vector from linear measurements. Mat. Zametki **104**(6), 863–871 (2018)
18. Lovász, L.: Matching structure and the matching lattice. J. Comb. Theory Ser. B **43**(2), 187–222 (1987)
19. Natarajan, B.K.: Sparse approximate solutions to linear systems. SIAM J. Comput. **24**(2), 227–234 (1995)
20. Oertel, T., Paat, J., Weismantel, R.: Sparsity of integer solutions in the average case. In: Lodi, A., Nagarajan, V. (eds.) IPCO 2019. LNCS, vol. 11480, pp. 341–353. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-17953-3_26
21. Ramírez-Alfonsín, J.L.: Complexity of the Frobenius problem. Combinatorica **16**(1), 143–147 (1996)
22. Rossi, M., Haimovich, A.M., Eldar, Y.C.: Spatial compressive sensing for MIMO radar. IEEE Trans. Signal Process. **62**(2), 419–430 (2014)
23. Schrijver, A.: Theory of Linear and Integer Programming. Wiley-Interscience Series in Discrete Mathematics. John Wiley & Sons Ltd., Chichester (1986). A Wiley-Interscience Publication
24. Sebö, A.: Hilbert bases, Carathéodory's Theorem and combinatorial optimization. In: Proceedings of the 1st Integer Programming and Combinatorial Optimization Conference, Waterloo, Ont., Canada, pp. 431–455. University of Waterloo Press (1990)