

Computing a Basis for an Integer Lattice: A Special Case

Haomin Li

Cheriton School of Computer Science
University of Waterloo
h439li@uwaterloo.ca

Arne Storjohann

Cheriton School of Computer Science
University of Waterloo
astorjoh@uwaterloo.ca

ABSTRACT

Consider an integer matrix $A \in \mathbb{Z}^{n \times (n-1)}$ that has full column rank $n - 1$. The set of all \mathbb{Z} -linear combinations of the rows of A generates a lattice, denoted by $\mathcal{L}(A)$. An algorithm is given that computes a matrix $C \in \mathbb{Z}^{(n-1) \times n}$ such that $B := CA \in \mathbb{Z}^{(n-1) \times (n-1)}$ is a basis for A , that is, B has full row rank $n - 1$ and $\mathcal{L}(B) = \mathcal{L}(A)$. The matrix C will satisfy $\log \|C\| \leq 4 \log n + 2 \log \|A\|$ (where $\|\cdot\|$ denotes the maximum entry in absolute value) and will have at most $n(1 + \log_2 n)$ nonzero entries. The cost of the algorithm is about the same as that required to multiply together two square integer matrices of dimension n that have magnitude of entries bounded by $\|A\|$.

CCS CONCEPTS

• **Computing methodologies** → **Linear algebra algorithms; Exact arithmetic algorithms.**

KEYWORDS

integer matrix, lattice basis, exact arithmetic algorithm

ACM Reference Format:

Haomin Li and Arne Storjohann. 2022. Computing a Basis for an Integer Lattice: A Special Case. In *Proceedings of the 2022 International Symposium on Symbolic and Algebraic Computation (ISSAC '22)*, July 4–7, 2022, Villeneuve-d'Ascq, France. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3476446.3536184>

1 INTRODUCTION

The set of all \mathbb{Z} -linear combinations of the rows of an integer matrix A is a (row) lattice, denoted by $\mathcal{L}(A)$. We consider the problem of computing a basis for $\mathcal{L}(A)$. The most general version of the problem takes as input an $A \in \mathbb{Z}^{n \times m}$, and asks as output a $B \in \mathbb{Z}^{r \times m}$ such that $\mathcal{L}(B) = \mathcal{L}(A)$, where r is the rank of A . Such a B is a basis for $\mathcal{L}(A)$. A constraint is that we want the bitlength of entries in the basis to be not much larger than those of A . Formally, we require that $\log \|B\| \in O(\log n + \log \|A\|)$, where $\|\cdot\|$ for a matrix or vector denotes the largest entry in absolute value.

Before continuing, we mention that $\mathcal{L}(A)$ will always denote the lattice generated by the rows of A , and, for brevity, we will write “ B is a basis for A ” to mean “ B is a basis for $\mathcal{L}(A)$.”

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](https://permissions.acm.org).
ISSAC '22, July 4–7, 2022, Villeneuve-d'Ascq, France

© 2022 Association for Computing Machinery.
ACM ISBN 978-1-4503-8688-3/22/07...\$15.00
<https://doi.org/10.1145/3476446.3536184>

One of our motivations for studying this problem is to extend fast algorithms for computing normal forms of integer matrices to input matrices with arbitrary shape and rank profile. For example, the fastest algorithm for computing the Smith form of an integer matrix by Birmpilis et al. [2] requires as input a square nonsingular input matrix.

By first producing a basis $B \in \mathbb{Z}^{r \times m}$ for $A \in \mathbb{Z}^{n \times m}$, then producing a basis $\tilde{B} \in \mathbb{Z}^{r \times r}$ for $\text{Transpose}(B) \in \mathbb{Z}^{m \times r}$, we can apply the Smith form algorithm with the nonsingular input \tilde{B} to obtain the Smith form of A . We refer to Li and Nguyen [6] for other applications of computing a lattice basis with entries not much larger than those of the input matrix, and for a survey of previous approaches for solving this problem.

A classical case of the problem is $n = 2$ and $m = 1$. Computing a basis for

$$A = \begin{bmatrix} a_1 \\ a_2 \end{bmatrix} \in \mathbb{Z}^{2 \times 1}$$

corresponds to computing $g = \gcd(a_1, a_2)$. The extended gcd problem asks also for a lattice generator $C \in \mathbb{Z}^{1 \times 2}$ such that

$$CA = \begin{bmatrix} c_1 & c_2 \end{bmatrix} \begin{bmatrix} a_1 \\ a_2 \end{bmatrix} = \begin{bmatrix} g \end{bmatrix}.$$

We will use a cost function B such that the extended gcd problem can be solved in $O(B(\log \|A\|))$ bit operations, where $\|\cdot\|$ denotes the maximal entry in absolute value.

Another case is n arbitrary and $m = 1$. Given $A \in \mathbb{Z}^{n \times 1}$, find $g = \gcd(A)$, the gcd of all entries in A . [8, Corollary 6.5] gives an $O(n B(\log \|A\|))$ bit operations algorithm that solves the vector extended gcd problem, producing a $C \in \mathbb{Z}^{1 \times n}$ that satisfies $\|C\| \leq \|A\|$.

In this paper, we consider the special case $m = n - 1 = r$, that is, the input is a full column rank matrix $A \in \mathbb{Z}^{n \times (n-1)}$. To the best of our knowledge, the fastest previous algorithm to compute a basis B with $\log \|B\| \in O(\log n + \log \|A\|)$ is that of Li and Nguyen [6, Theorem 3.8], which solves the problem in the general case, that is, for n , m and r arbitrary. Applied to the special case $m = n - 1 = r$ that we consider here, their algorithm has cost bounded by $O(n^\omega B(nd))$ bit operations, where $d = \log n + \log \|A\|$ and ω is the exponent of matrix multiplication. In this paper, we give an algorithm that solves the problem in $O(n^\omega B(d)(\log n)^2)$ bit operations. We remark that this cost estimate matches that of the Smith form algorithm of Birmpilis et al. [2] and thus our technique can be used for the application mentioned earlier, at least for the special case $m = n - 1 = r$.

Instead of producing B directly, we first compute a matrix $C \in \mathbb{Z}^{(n-1) \times n}$ and then set $B := CA$. We show that $\log \|C\| \leq 4 \log n + 2 \log \|A\|$ and that C will have at most $n(1 + \log_2 n)$ nonzero entries, independent of $\|A\|$. To the best of our knowledge, we are not aware

that the existence of such a C with only $O(n \log n)$ nonzero entries was previously known.

The sparsity of C can be useful in some situations. Suppose, in addition to the full column rank input matrix $A \in \mathbb{Z}^{n \times (n-1)}$, we have additional linearly dependent columns $\bar{A} \in \mathbb{Z}^{n \times m}$. If $C \in \mathbb{Z}^{(n-1) \times n}$ is such that CA is a basis for A , then a basis for the rank $n-1$ augmented matrix $\begin{bmatrix} A & \bar{A} \end{bmatrix} \in \mathbb{Z}^{n \times ((n-1)+m)}$ is given by $C \begin{bmatrix} A & \bar{A} \end{bmatrix}$. The additional cost to compute $C\bar{A}$ is only $O(nmM(d) \log n)$ bit operations, where M is a cost function for integer multiplication, and $d = \log n + \log \|\bar{A}\|$.

We now give an outline of our approach to compute C . Our idea is to find a vector $\mathbf{v} \in \mathbb{Z}^{n \times 1}$ such that the augmented matrix

$$\begin{bmatrix} A & \mathbf{v} \end{bmatrix} \in \mathbb{Z}^{n \times n} \quad (1)$$

is nonsingular. Then we compute $C \in \mathbb{Z}^{(n-1) \times n}$ to be a left kernel basis of \mathbf{v} that satisfies $\|C\| \leq \|\mathbf{v}\|^2$. Of course, the augmentation vector \mathbf{v} must be special in order for this to work. We show that if $\mathcal{L}(\begin{bmatrix} A & \mathbf{v} \end{bmatrix})$ contains the last row of I_n , then a left kernel basis $C \in \mathbb{Z}^{(n-1) \times n}$ of \mathbf{v} has the property that CA is a basis for A . Our algorithm has three main computational steps:

- (1) Compute the left kernel basis $\mathbf{w} \in \mathbb{Z}^{1 \times n}$ of A .
- (2) Compute an augmentation vector $\mathbf{v} \in \mathbb{Z}^{n \times 1}$ from A and \mathbf{w} .
- (3) Compute a left kernel basis $C \in \mathbb{Z}^{(n-1) \times n}$ of \mathbf{v} .

The rest of this paper is organised as follows. Section 2 gives basic mathematical background. Section 3 establishes, via a more general result, the sufficient condition on the augmentation vector \mathbf{v} as described above. Section 4 adapts from the literature some computational tools related to deterministic system solving. Section 5 uses these computational tools to obtain deterministic algorithms for the kernel basis problem (step 1). Section 6 gives the algorithm to produce the augmentation vector \mathbf{v} (step 2). Section 7 gives the algorithm to compute a left kernel bases of \mathbf{v} (step 3) and also states the main result of the paper.

Cost model

Following von zur Gathen and Gerhard [3, Section 8.3], cost estimates are given using a function $M(d)$ that bounds the number of bit operations required to multiply two integers bounded in magnitude by 2^d . We use $B(d)$ to bound the cost of integer gcd-related computations such as the extended euclidean algorithm. We can always take $B(d) = O(M(d) \log d)$. If $M(d) \in \Omega(d^{1+\epsilon})$ for some $\epsilon > 0$, then $B(d) \in O(M(d))$.

As usual, we assume that M is superlinear and subquadratic. We also assume that $M(ab) \in O(M(a)M(b))$ for $a, b \geq 1$. We assume that $\omega > 2$, and to simplify cost estimates, we make the common assumption that $M(d) \in O(d^{\omega-1})$. The assumptions stated in this paragraph apply also to B .

2 MATHEMATICAL BACKGROUND

An important notion for us is that of a primitive integer matrix.

DEFINITION 1 (PRIMITIVE MATRIX). *A matrix $A \in \mathbb{Z}^{n \times m}$ is primitive if one of the following conditions hold.*

- (i) $(n > m)$ A has full column rank m and I_m is a basis for A .

- (ii) $(n < m)$ A has full row rank n and I_n is a basis for $\text{Transpose}(A)$.
- (iii) $(n = m)$ A is unimodular (that is, $\det A = \pm 1$).

DEFINITION 2 (LEFT KERNEL BASIS OF AN INTEGER MATRIX). *A left kernel basis of an $A \in \mathbb{Z}^{n \times m}$ with rank of r is a matrix $K \in \mathbb{Z}^{(n-r) \times n}$ such that*

- (i) K has rank $n-r$,
- (ii) KA is the $(n-r) \times m$ zero matrix, and
- (iii) K is primitive.

We remark that a matrix K that satisfies only properties (i) and (ii) of Definition 2 is a nullspace basis of A over \mathbb{Q} : property (ii) ensures that the set of all \mathbb{Q} -linear combinations of the rows of K generates $\{\mathbf{v} \in \mathbb{Q}^{1 \times n} \mid \mathbf{v}A = \mathbf{0}_{1 \times m}\}$. The additional property (iii) in Definition 2 ensures that K satisfies $\mathcal{L}(K) = \{\mathbf{v} \in \mathbb{Z}^{1 \times n} \mid \mathbf{v}A = \mathbf{0}_{1 \times m}\}$.

LEMMA 3. *Let $A \in \mathbb{Z}^{n \times m}$ have full column rank. If $B \in \mathbb{Z}^{m \times m}$ is a basis for A , then AB^{-1} is integral and primitive.*

PROOF. Since B is a basis for A , B is nonsingular and we have

$$UA = \begin{bmatrix} B \\ \hline \end{bmatrix} \quad (2)$$

for some unimodular matrix $U \in \mathbb{Z}^{n \times n}$. Postmultiplying both sides of (2) by B^{-1} gives

$$UAB^{-1} = \begin{bmatrix} I_m \\ \hline \end{bmatrix}. \quad (3)$$

Premultiplying both sides of equation (3) by U^{-1} gives

$$AB^{-1} = U^{-1} \begin{bmatrix} I_m \\ \hline \end{bmatrix}. \quad (4)$$

Since U is unimodular, U^{-1} is integral. The right side of equation (4) is integral. Hence, AB^{-1} is integral. It follows from equation (3) that AB^{-1} is primitive. \square

LEMMA 4. *Let $A \in \mathbb{Z}^{n \times m}$ of rank r be given. If $C \in \mathbb{Z}^{r \times n}$ is such that CA is a basis for A , and $K \in \mathbb{Z}^{(n-r) \times n}$ is a left kernel basis of A , then*

$$\begin{bmatrix} C \\ K \end{bmatrix} \in \mathbb{Z}^{n \times n} \quad (5)$$

is unimodular.

PROOF. Since K is primitive, it follows from Definition 1.(ii) that there exists a unimodular matrix U such that $KU = \begin{bmatrix} & \\ & I_{n-r} \end{bmatrix}$. Let

$$\begin{bmatrix} C \\ K \end{bmatrix} U = \begin{bmatrix} V & * \\ \hline & I_{n-r} \end{bmatrix} \quad (6)$$

and

$$U^{-1}A = \begin{bmatrix} A_1 \\ \hline A_2 \end{bmatrix}, \quad (7)$$

where $A_1 \in \mathbb{Z}^{r \times m}$. Multiplying the right hand side of (6) by that of (7) gives

$$\begin{bmatrix} V & * \\ \hline & I_{n-r} \end{bmatrix} \begin{bmatrix} A_1 \\ \hline A_2 \end{bmatrix} = \begin{bmatrix} CA \\ \hline \end{bmatrix}. \quad (8)$$

We conclude from (8) that A_2 is the zero matrix and then from (7) that A_1 is a basis for A . Since CA is also a basis for A , the matrix V must be unimodular. The result now follows from (6). \square

3 LATTICE GENERATOR VIA KERNEL BASIS

Suppose $A \in \mathbb{Z}^{n \times (n-1)}$ and $\mathbf{v} \in \mathbb{Z}^{n \times 1}$ are such that $[A \mid \mathbf{v}]$ is nonsingular. In this section, we establish a sufficient condition on \mathbf{v} to ensure that a left kernel basis $C \in \mathbb{Z}^{(n-1) \times n}$ of \mathbf{v} has the property that CA is a basis for A .

We rely on the following theorem that applies only to those nonsingular matrices $A \in \mathbb{Z}^{n \times n}$ for which there exists a unimodular decoupling matrix $U \in \mathbb{Z}^{n \times n}$ such that

$$UA = \left[\begin{array}{c|c} *1 & \\ \hline & *2 \end{array} \right].$$

Note that for such an A , we have $\mathcal{L}(A) = \{ [\mathcal{L}(*1) \mid \mathcal{L}(*2)] \}$.

THEOREM 5. *Let $A = [A_1 \mid A_2] \in \mathbb{Z}^{n \times n}$ be nonsingular, where $A_1 \in \mathbb{Z}^{n \times m_1}$ and $A_2 \in \mathbb{Z}^{n \times m_2}$. For any $C_1 \in \mathbb{Z}^{m_1 \times n}$ and $C_2 \in \mathbb{Z}^{m_2 \times n}$, we have*

- (a) $\left[\begin{array}{c} C_1 \\ \hline C_2 \end{array} \right] [A_1 \mid A_2] = \left[\begin{array}{c|c} C_1 A_1 & \\ \hline & C_2 A_2 \end{array} \right] \in \mathbb{Z}^{n \times n}$, with
 (b) $C_1 A_1$ a basis for A_1 , and
 (c) $C_2 A_2$ a basis for A_2

if and only if

- (i) $\left[\begin{array}{c} C_1 \\ \hline C_2 \end{array} \right] \in \mathbb{Z}^{n \times n}$ is unimodular,
 (ii) C_1 is a kernel basis for A_2 , and
 (iii) C_2 is a kernel basis for A_1 .

PROOF. (Only if): Assume that (a), (b) and (c) hold. Since $A = [A_1 \mid A_2]$ is nonsingular, A_i has full column rank ($i = 1, 2$). From (b) and (c), we have that $C_i A_i \in \mathbb{Z}^{m_i \times m_i}$ also has full column rank ($i = 1, 2$). It follows that $C_i A_i$ is nonsingular ($i = 1, 2$).

Postmultiplying both sides of the equation in (a) by

$$\left[\begin{array}{c|c} (C_1 A_1)^{-1} & \\ \hline & (C_2 A_2)^{-1} \end{array} \right]$$

gives

$$\left[\begin{array}{c} C_1 \\ \hline C_2 \end{array} \right] [A_1 (C_1 A_1)^{-1} \mid A_2 (C_2 A_2)^{-1}] = \left[\begin{array}{c|c} I_{m_1} & \\ \hline & I_{m_2} \end{array} \right]. \quad (9)$$

Since $(C_i A_i)$ is a basis for A_i ($i = 1, 2$), it follows from Lemma 3 that

$$U := [A_1 (C_1 A_1)^{-1} \mid A_2 (C_2 A_2)^{-1}]$$

is an integer matrix. Therefore,

$$\det \left(\left[\begin{array}{c} C_1 \\ \hline C_2 \end{array} \right] \right) = \pm 1$$

and property (i) holds. It remains to show that (ii) and (iii) hold. Since U is an integer matrix, it follows from (9) that C_1 satisfies the properties required by Definition 2 to be a left kernel of A_2 . In particular, C_1 is primitive and $C_1 A_2 = \mathbf{0}_{m_1 \times m_2}$. Similarly, C_2 is a left kernel of A_1 .

(If): Assume that (i), (ii) and (iii) hold. It follows from (ii) and (iii) that (a) holds. It now follows from property (i) that

$$\mathcal{L}([A_1 \mid A_2]) = \mathcal{L} \left(\left[\begin{array}{c|c} C_1 A_1 & \\ \hline & C_2 A_2 \end{array} \right] \right).$$

Properties (b) and (c) now follow by noting that A_i and $C_i A_i$ have the same column dimension ($i = 1, 2$). \square

The following corollary identifies a particular class of matrices for which Theorem 5 applies.

COROLLARY 6. *Let $A = [A_1 \mid A_2] \in \mathbb{Z}^{n \times n}$ be nonsingular, where $A_1 \in \mathbb{Z}^{n \times m_1}$ and $A_2 \in \mathbb{Z}^{n \times m_2}$. If $\mathcal{L}(A)$ contains the last m_2 rows of I_n , then any left kernel basis $C \in \mathbb{Z}^{m_1 \times n}$ of A_2 has the property that CA_1 is a basis for A_1 .*

PROOF. Since $\mathcal{L}(A)$ contains the last m_2 rows of I_n , there exists a unimodular matrix $U \in \mathbb{Z}^{n \times n}$ such that

$$UA = \left[\begin{array}{c|c} * & \\ \hline & I_{m_2} \end{array} \right].$$

Decompose U as

$$U = \left[\begin{array}{c} C_1 \\ \hline C_2 \end{array} \right],$$

where $C_1 \in \mathbb{Z}^{m_1 \times n}$ and $C_2 \in \mathbb{Z}^{m_2 \times n}$. Then the equation in property (a) of Theorem 5 holds, and from the unimodularity of U , it follows that C_1 is a left kernel basis for A_2 , and C_2 is a left kernel basis for A_1 . We are thus exactly in the situation of Theorem 5, and can conclude that $C_1 A_1$ is a basis for A_1 . Since any two left kernel bases of A_2 are left equivalent, we can replace C_1 with C in U and the matrix remains unimodular. \square

4 COMPUTATIONAL TOOLS

The main computational tool we need is an algorithm for nonsingular linear system solving. A deterministic algorithm has recently been given by Birmpilis et al. [1, Section 11], but was analyzed under a more restrictive cost model than we are using in this paper. We restate the result here.

THEOREM 7 (SYSTEM SOLVING). *There exists an algorithm that takes as input a nonsingular matrix $A \in \mathbb{Z}^{n \times n}$ and a vector $\mathbf{b} \in \mathbb{Z}^{n \times 1}$, and returns as output $A^{-1} \mathbf{b} \in \mathbb{Q}^{n \times 1}$, together with the minimal $e \in \mathbb{Z}_{>0}$ such that $eA^{-1} \mathbf{b}$ is integral. Let $d = \log n + \log \|A\|$. If $\log \|\mathbf{b}\| \in O(nd)$, the running time of the algorithm is $O(n^\omega B(d)(\log n)^2)$ bit operations.*

PROOF. [1, Section 11] gives an algorithm that computes an integer vector $\mathbf{x} \in \mathbb{Z}^{n \times 1}$ together with an $f \in \mathbb{Z}_{\geq 0}$ such that $2^f A^{-1} \mathbf{b}$ can be recovered from \mathbf{x} using rational number reconstruction. The cost bound of [1, Theorem 22] is

$$O(n^\omega M(d)(\log n + \log \log \|A\|)(\log n)). \quad (10)$$

The cost model used in [1] does not make the assumptions that $\omega > 2$ and $M(d) \in O(d^{\omega-1})$. If we do make these assumptions, and we replace the subroutine in [1, Section 6] with the integer analogue of the algorithm supporting [4, Theorem 7], the $\log \log \|A\|$ term in (10) can be avoided, giving the cost bound $O(n^\omega M(d)(\log n)^2)$ to recover \mathbf{x} .

To recover $A^{-1} \mathbf{b}$ from \mathbf{x} requires rational number reconstruction. By [3, Section 5.10], the cost of computing e and $eA^{-1} \mathbf{b}$ is

$O(n B(nd))$ bit operations; this simplifies to $O(n^\omega B(d))$ using the assumptions $B(nd) \in O(B(n)B(d))$ and $B(n) \in O(n^{\omega-1} B(d))$. \square

The algorithm supporting Theorem 7 was based on computing a 2-massager for the input matrix. Although originally defined for nonsingular matrices, the notion of a 2-massager extends directly to matrices of full column rank.

DEFINITION 8 (2-SMITH FORM). Let $A \in \mathbb{Z}^{n \times m}$ have full column rank m . The 2-Smith form of A is the matrix $\text{diag}(2^{e_1}, 2^{e_2}, \dots, 2^{e_m})$ with 2^{e_i} the largest power of two which divides the invariant factor i of the Smith form of A over \mathbb{Z} , $1 \leq i \leq m$.

DEFINITION 9. Let $A \in \mathbb{Z}^{n \times m}$ have full column rank m . A 2-massager for A is a triple of matrices (P, S, M) from $\mathbb{Z}^{m \times m}$ such that,

- P is a permutation,
- $S = \text{diag}(s_1, \dots, s_m)$ is the 2-Smith form of A ,
- M is unit upper triangular, and
- $APMS^{-1}$ is integral with $\text{Rem}(APMS^{-1}, 2) \in \mathbb{Z}/(2)^{n \times m}$ having full column rank over $\mathbb{Z}/(2)$.

(P, S, M) is a reduced 2-massager if entries in column i of M are from $[0, \dots, s_i - 1]$, $1 \leq i \leq n$.

EXAMPLE 10. A 2-massager for the the matrix

$$A := \begin{bmatrix} 92 & 27 & 99 \\ -124 & 32 & 116 \\ 556 & -42 & -150 \\ 176 & 249 & 77 \\ -8 & 195 & -121 \end{bmatrix} \in \mathbb{Z}^{5 \times 3}$$

is given by

$$P, S, M := \begin{bmatrix} & & 1 \\ 1 & & \\ & 1 & \end{bmatrix}, \begin{bmatrix} 1 & & \\ & 4 & \\ & & 16 \end{bmatrix}, \begin{bmatrix} 1 & 3 & 5 \\ & 1 & 15 \\ & & 1 \end{bmatrix}.$$

The massaged matrix

$$B := APMS^{-1} = \begin{bmatrix} 27 & 45 & 107 \\ 32 & 53 & 111 \\ -42 & -69 & -119 \\ 249 & 206 & 161 \\ 195 & 116 & -53 \end{bmatrix}$$

has full column rank over $\mathbb{Z}/(2)$.

Birmpilis et al. [1, Section 10] give an algorithm to compute a 2-massager of a nonsingular matrix. The algorithm extends naturally to an $n \times m$ matrix of full column rank m .

THEOREM 11 (2-MASSAGER COMPUTATION). There exists an algorithm that takes as input a full column rank $A \in \mathbb{Z}^{n \times m}$, and returns as output a reduced 2-massager for A together with the massaged matrix $B := APMS^{-1}$. The cost of the algorithm is $O(nm^{\omega-1} M(d)(\log m)^2)$ bit operations, where $d = \log m + \log \|A\|$.

PROOF. Similar to the proof of Theorem 7, if we replace the subroutine in [1, Section 6] with the integer analogue of the algorithm supporting [4, Theorem 7], then [1, Theorem 21] establishes the theorem in the case $n = m$, that is, for a square nonsingular $m \times m$ input matrix, a 2-massager can be computed in

$$O(m^\omega M(d)(\log m)^2)$$

bit operations. The only required modifications now to handle case of a full column rank rectangular $n \times m$ input matrix is to incorporate blocking into steps 1–4 of the algorithm supporting [1, Theorem 19].

Step 1 computes an LQUP decomposition of an $m \times m$ matrix over $\mathbb{Z}/(2)$ at a cost of $O(n^\omega)$ bit operations. In the rectangular case this costs $O(nm^{\omega-1})$ bit operations.

Step 2 computes a single nonsingular system solution of dimension m . In the rectangular case, step 2 can be accomplished by computing $\lceil n/m \rceil$ such linear systems, thus replacing the m^ω term in the cost estimate for the square case by $nm^{\omega-1}$.

Steps 3 and 4 compute a triangular Smith form of a square matrix over ring. In the rectangular case, these steps are accomplished using the integer analogue of [4, Theorem 7] followed up by the triangularization of a matrix that, compared to the square case, has row dimension now bounded by n . The cost increases similarly as for step 2. \square

One additional algorithm that we need is a fast solution to the vector extended gcd problem [8, Corollary 6.5].

Algorithm 1: VectorGcdex(w)

Input : $w \in \mathbb{Z}^{n \times 1}$

Output : $(e, g) \in (\mathbb{Z}^{1 \times n}, \mathbb{Z})$, satisfying

- $ew = g$, with $g = \text{gcd}(w)$,
- $\|e\| \leq \|w\|$, and
- the number of nonzero entries in e is bounded by $1 + \log_2 \|w\|$.

Runtime: $O(n B(\log \|w\|))$ bit operations

5 KERNEL BASIS OF NULLITY ONE

We give an algorithm to compute a left kernel basis $w \in \mathbb{Z}^{1 \times n}$ of a full column rank input matrix $A \in \mathbb{Z}^{n \times (n-1)}$. In this special case, where the nullity of A is one, the kernel basis w is unique (up to sign) and can be computed using nonsingular linear system solving. If $Q \in \mathbb{Z}^{n \times n}$ is a permutation such that

$$QA = \begin{bmatrix} \bar{A} \\ a \end{bmatrix} \in \mathbb{Z}^{n \times (n-1)} \quad (11)$$

with $\bar{A} \in \mathbb{Z}^{(n-1) \times (n-1)}$ nonsingular, then a left nullspace of QA over \mathbb{Q} is

$$\left[a\bar{A}^{-1} \mid -1 \right] \in \mathbb{Q}^{1 \times n}.$$

If $e \in \mathbb{Z}_{>0}$ is minimal such that $ea\bar{A}^{-1}$ is integral, then

$$w = Q \left[ea\bar{A}^{-1} \mid -e \right] \in \mathbb{Z}^{1 \times n} \quad (12)$$

is primitive and is thus a kernel basis of A (Definition 2).

THEOREM 12 (KERNEL OF NULLITY ONE). A left kernel basis of $w \in \mathbb{Z}^{1 \times n}$ of a full column rank $A \in \mathbb{Z}^{n \times (n-1)}$ can be computed in $O(n^\omega B(d)(\log n)^2)$ bit operations, where $d = \log n + \log \|A\|$.

PROOF. To find a suitable Q , compute a 2-massager (P, M, S) of A together with the massaged matrix $B = APMS^{-1}$. The massaged matrix B has full column rank modulo 2. Working modulo 2, use

the LSP decomposition of Ibarra et al. [5, Theorem 3.2] to compute the row rank profile $[i_1, i_2, \dots, i_{n-1}]$ of \mathbf{B} over $\mathbb{Z}/(2)$. Since \mathbf{PMS}^{-1} is nonsingular, the submatrix of \mathbf{A} comprised of rows i_1, i_2, \dots, i_{n-1} is also nonsingular. Define \mathbf{Q} accordingly.

Compute the minimal $e \in \mathbb{Z}_{>0}$ such that $e\mathbf{a}\bar{\mathbf{A}}^{-1}$ is integral, and compute $e\mathbf{a}\bar{\mathbf{A}}^{-1}$ itself. Return \mathbf{w} as in (12). The claimed cost bound follows from Theorems 11 (2-massager computation) and 7 (System solving). \square

The following corollary follows from Hadamard's inequality [3, Theorem 16.6] and Cramer's rule [3, Theorem 25.6].

COROLLARY 13. *The kernel basis produced by Theorem 12 satisfies $\log \|\mathbf{w}\| \in O(nd)$.*

6 AUGMENTATION VECTOR

Storjohann [9, Section 13] describes a Las Vegas randomized algorithm `DetReduction` that takes as input a nonsingular integer input matrix, and replaces the last column to obtain a new matrix whose lattice contains the last row of \mathbf{I} . Algorithm `DetReduction` exploits the fact that the input matrix is nonsingular. Here, we adapt the approach to obtain the following result. In the proof, we focus mainly on showing how to avoid the requirement that the input matrix be nonsingular, and how to avoid randomization by using the subroutines developed in previous sections.

THEOREM 14 (AUGMENTATION VECTOR). *There exists an algorithm that takes as input a full column rank integer matrix $\mathbf{A} \in \mathbb{Z}^{n \times (n-1)}$, and returns as output an integer vector $\mathbf{v} \in \mathbb{Z}^{n \times 1}$ such that*

- (1) $\mathcal{L}(\left[\begin{array}{c|c} \mathbf{A} & \mathbf{v} \end{array} \right])$ contains the last row of \mathbf{I}_n , and
- (2) $\|\mathbf{v}\| \leq n^2 \|\mathbf{A}\|$.

The cost of the algorithm is $O(n^\omega B(d) (\log n)^2)$ bit operations, where $d = \log n + \log \|\mathbf{A}\|$.

PROOF. Compute a kernel basis $\mathbf{w} \in \mathbb{Z}^{1 \times n}$ of \mathbf{A} . Let $\mathbf{P} \in \mathbb{Z}^{n \times n}$ be a permutation matrix such that $\mathbf{w}\mathbf{P}$ has last entry of maximal magnitude. Then

$$(\mathbf{w}\mathbf{P})(\mathbf{P}^{-1}\mathbf{A}) = \left[\begin{array}{c|c} \bar{\mathbf{w}} & w_n \end{array} \right] \left[\begin{array}{c} \bar{\mathbf{A}} \\ \mathbf{a} \end{array} \right] = \mathbf{0}_{1 \times (n-1)}$$

where $|w_n| = \|\mathbf{w}\|$ and $\bar{\mathbf{A}} \in \mathbb{Z}^{(n-1) \times (n-1)}$ is nonsingular since $w_n \neq 0$. Going forward, we will assume, without loss of generality, that $\mathbf{P} = \mathbf{I}_n$.

Use Algorithm 1 (`VectorGcdex`) to compute a vector $\mathbf{b} \in \mathbb{Z}^{n \times 1}$ such that

$$\mathbf{w}\mathbf{b} = \left[\begin{array}{c|c} \bar{\mathbf{w}} & w_n \end{array} \right] \left[\begin{array}{c} \bar{\mathbf{b}} \\ b_n \end{array} \right] = 1.$$

Then

$$\left[\begin{array}{c|c} \bar{\mathbf{w}} & w_n \end{array} \right] \left[\begin{array}{c|c} \bar{\mathbf{A}} & \bar{\mathbf{b}} \\ \mathbf{a} & b_n \end{array} \right] = \left[\begin{array}{c|c} & \\ & 1 \end{array} \right]. \quad (13)$$

It follows from (13) that $\mathcal{L}(\left[\begin{array}{c|c} \mathbf{A} & \mathbf{b} \end{array} \right])$ contains the last row of \mathbf{I}_n , and thus \mathbf{b} is candidate for our solution vector. However, the entries of \mathbf{b} can be too large.

We now show how to adjust \mathbf{b} to obtain a reduced solution \mathbf{v} . Postmultiplying both sides of equation (13) by a unimodular matrix

$$\left[\begin{array}{c|c} \mathbf{I}_{n-1} & -\bar{\mathbf{y}} \\ \hline & 1 \end{array} \right] \in \mathbb{Z}^{n \times n}$$

gives

$$\left[\begin{array}{c|c} \bar{\mathbf{w}} & w_n \end{array} \right] \left[\begin{array}{c|c} \bar{\mathbf{A}} & \bar{\mathbf{b}} - \bar{\mathbf{A}}\bar{\mathbf{y}} \\ \hline \mathbf{a} & b_n - \mathbf{a}\bar{\mathbf{y}} \end{array} \right] = \left[\begin{array}{c|c} & \\ & 1 \end{array} \right]. \quad (14)$$

Let $\mathbf{y} \in \mathbb{Q}^{(n-1) \times 1}$ be the solution to the linear system $\bar{\mathbf{A}}\mathbf{y} = \bar{\mathbf{b}}$. Let $\bar{\mathbf{y}} \in \mathbb{Z}^{(n-1) \times n}$ be the integral part of \mathbf{y} , that is, $\bar{\mathbf{y}}$ is the unique integer vector such that $\|\mathbf{y} - \bar{\mathbf{y}}\| < 1$. Then $\|\bar{\mathbf{b}} - \bar{\mathbf{A}}\bar{\mathbf{y}}\| \leq n\|\mathbf{A}\|$. We choose our solution vector to be $\mathbf{v} := \bar{\mathbf{b}} - \bar{\mathbf{A}}\bar{\mathbf{y}}$. It remains to show that the last entry $v_n = b_n - \mathbf{a}\bar{\mathbf{y}}$ of \mathbf{v} has magnitude bounded by $n^2\|\mathbf{A}\|$. From equation (14), we have $\bar{\mathbf{w}}\bar{\mathbf{v}} + w_n v_n = 1$. It follows that

$$|v_n| = \left\| \frac{\bar{\mathbf{w}}\bar{\mathbf{v}}}{w_n} - \frac{1}{w_n} \right\| \leq \left\| \frac{\bar{\mathbf{w}}\bar{\mathbf{v}}}{w_n} \right\| + \left| \frac{1}{w_n} \right|.$$

Since $\bar{\mathbf{w}}$ has column dimension $n-1$, $\|\bar{\mathbf{w}}\| \leq |w_n| = \|\mathbf{w}\|$, and $\|\bar{\mathbf{v}}\| \leq n\|\mathbf{A}\|$, we have

$$|v_n| \leq (n-1)\|\bar{\mathbf{v}}\| + 1 \leq (n-1)n\|\mathbf{A}\| + 1 \leq n^2\|\mathbf{A}\|.$$

Corollary 13 gives the bound $\log \|\mathbf{w}\| \in O(nd)$. From the specification of Algorithm 1 (`VectorGcdex`), we have $\|\mathbf{b}\| \leq \|\mathbf{w}\|$, so $\log \|\mathbf{b}\| \in O(nd)$ also. The claimed running time bound now follows from Theorem 12 (`Kernel of nullity one`), the running time of Algorithm 1 (`VectorGcdex`), and Theorem 7 (`System solving`). \square

EXAMPLE 15. *The input matrix*

$$\mathbf{A} = \begin{bmatrix} 231 & 303 & -118 & 16 \\ 344 & 202 & -389 & 163 \\ 185 & 190 & -80 & 136 \\ 263 & 189 & 196 & 66 \\ 259 & 157 & 131 & 136 \end{bmatrix} \in \mathbb{Z}^{5 \times 4}$$

has left kernel

$$\mathbf{w} = \left[-3330033 \quad 825718 \quad 4373666 \quad 7018431 \quad -8377548 \right].$$

`VectorGcdex` computes a $\mathbf{b} \in \mathbb{Z}^{5 \times 1}$ such that

$$\left[\begin{array}{c|c} \bar{\mathbf{A}} & \bar{\mathbf{b}} \\ \hline \mathbf{a} & b_5 \end{array} \right] = \left[\begin{array}{cccc|c} 231 & 303 & -118 & 16 & -2022969 \\ 344 & 202 & -389 & 163 & 0 \\ 185 & 190 & -80 & 136 & 2 \\ 263 & 189 & 196 & 66 & 0 \\ 259 & 157 & 131 & 136 & 804121 \end{array} \right]$$

with $\mathcal{L}(\left[\begin{array}{c|c} \mathbf{A} & \mathbf{b} \end{array} \right])$ containing the last row of \mathbf{I}_5 . Decompose $\mathbf{y} = \bar{\mathbf{A}}^{-1}\bar{\mathbf{b}} = \bar{\mathbf{y}} + \mathbf{r}$ where $\bar{\mathbf{y}} \in \mathbb{Z}^{4 \times 1}$ and $\mathbf{r} \in \mathbb{Q}^{4 \times 1}$ with $\|\mathbf{r}\| < 1$:

$$\mathbf{y} = \left[\begin{array}{c} \frac{219432205277}{94247415} \\ -\frac{10934282147281}{1319463810} \\ \frac{896384538211}{527785524} \\ \frac{709438547431}{75397932} \end{array} \right] = \bar{\mathbf{y}} + \mathbf{r} = \left[\begin{array}{c} 2328 \\ -8286 \\ 1698 \\ 9409 \end{array} \right] + \left[\begin{array}{c} \frac{24223157}{94247415} \\ -\frac{1205017621}{1319463810} \\ \frac{204718459}{527785524} \\ \frac{19405243}{75397932} \end{array} \right].$$

Our augmentation vector is given by

$$\mathbf{v} = \mathbf{b} - A\bar{\mathbf{y}} = \begin{bmatrix} -259 \\ -205 \\ -122 \\ -12 \\ 9 \end{bmatrix}.$$

7 GENERATING A LATTICE BASIS

We begin by giving an algorithm to compute a left kernel basis of a vector $\mathbf{v} \in \mathbb{Z}^{n \times 1}$. The algorithm we give is closely based on a more general algorithm [8, Chapter 6.2] that computes a left kernel basis for an $n \times m$ integer matrix. The variation we give here for the special case $m = 1$ is simpler and allows us to obtain a tighter explicit bound on the number of nonzero entries in the kernel.

Algorithm 2: SparseKernelBasis(\mathbf{v}, n)

Input : $\mathbf{v} \in \mathbb{Z}^{n \times 1}$ with all entries nonzero
Output : $\mathbf{K} \in \mathbb{Z}^{(n-1) \times n}$, a left kernel basis of \mathbf{v} satisfying

- $\|\mathbf{K}\| \leq \|\mathbf{v}\|^2$, and
- the number of nonzero entries in \mathbf{K} is bounded by $n(1 + \log_2 n)$.

Runtime: $O(n \text{ B}(\log \|\mathbf{v}\|) \log n)$ bit operations

```

1 if  $n = 1$  then
2   return the  $0 \times 1$  matrix
3  $n_1 := \lfloor n/2 \rfloor$ 
4  $n_2 := n - n_1$ 
5 Decompose  $\mathbf{v} = \begin{bmatrix} \mathbf{v}_1 \\ \mathbf{v}_2 \end{bmatrix}$  where  $\mathbf{v}_1 \in \mathbb{Z}^{n_1 \times 1}$  and  $\mathbf{v}_2 \in \mathbb{Z}^{n_2 \times 1}$ 
6  $\mathbf{K}_1 := \text{SparseKernelBasis}(\mathbf{v}_1, n_1)$ 
7  $\mathbf{K}_2 := \text{SparseKernelBasis}(\mathbf{v}_2, n_2)$ 
8  $\mathbf{e}_1, g_1 := \text{VectorGcdex}(\mathbf{v}_1, n_1)$ 
9  $\mathbf{e}_2, g_2 := \text{VectorGcdex}(\mathbf{v}_2, n_2)$ 
10  $g := \text{gcd}(g_1, g_2)$ 

```

11 **return** $\begin{bmatrix} \begin{array}{c|c} -(g_2/g) \mathbf{e}_1 & (g_1/g) \mathbf{e}_2 \\ \hline \mathbf{K}_1 & \\ \hline & \mathbf{K}_2 \end{array} \end{bmatrix}$

We will prove correctness of Algorithm SparseKernelBasis using three lemmas. The first lemma shows that the output is indeed a kernel basis. The second lemma establishes that the output satisfies the two stated conditions. The third lemma gives an analysis of the running time.

LEMMA 16. SparseKernelBasis outputs a kernel basis for \mathbf{v} .

PROOF. The algorithm uses a divide and conquer approach. The output of the base case ($n = 1$) is correct since $\mathbf{v} \in \mathbb{Z}^{1 \times 1}$ is nonzero by assumption.

For the recursive case ($n > 1$), we have subproblems of size n_1 and n_2 , with $n = n_1 + n_2$. By Lemma 4, the matrix

$$U_i = \begin{bmatrix} \mathbf{e}_i \\ \mathbf{K}_i \end{bmatrix} \in \mathbb{Z}^{n_i \times n_i}$$

is unimodular ($i = 1, 2$). The matrix $U := \text{diag}(U_1, U_2) \in \mathbb{Z}^{n \times n}$ is thus unimodular, and by construction we have

$$\begin{bmatrix} U_1 & & \\ & U_2 & \\ & & \end{bmatrix} \begin{bmatrix} \mathbf{v}_1 \\ \mathbf{v}_2 \end{bmatrix} = \begin{bmatrix} g_1 \\ g_2 \end{bmatrix} \in \mathbb{Z}^{n \times 1}, \quad (15)$$

where \mathbf{c} contains only two nonzero entries, g_1 at index 1 and g_2 at index $n_1 + 1$.

Note that the vector $\begin{bmatrix} -g_2/g & g_1/g \end{bmatrix} \in \mathbb{Z}^{1 \times 2}$ is primitive and thus can be extended to a unimodular matrix

$$\begin{bmatrix} *1 & *2 \\ -g_2/g & g_1/g \end{bmatrix} \in \mathbb{Z}^{2 \times 2}$$

that satisfies

$$\begin{bmatrix} *1 & *2 \\ -g_2/g & g_1/g \end{bmatrix} \begin{bmatrix} g_1 \\ g_2 \end{bmatrix} = \begin{bmatrix} g \\ \end{bmatrix}.$$

Let E be equal to I_n except with principal $(n_1+1) \times (n_1+1)$ submatrix equals to

$$\begin{bmatrix} *1 & & & *2 \\ & 1 & & \\ & & \ddots & \\ & & & 1 \\ -g_2/g & & & g_1/g \end{bmatrix} \in \mathbb{Z}^{(n_1+1) \times (n_1+1)}.$$

Then E is unimodular with

$$\begin{bmatrix} *1 & & & *2 \\ & \ddots & & \\ & & 1 & \\ -g_2/g & & & g_1/g \\ & & & \ddots & \\ & & & & 1 \end{bmatrix} \begin{bmatrix} \mathbf{c} \\ g_1 \\ \hline g_2 \end{bmatrix} = \begin{bmatrix} \bar{\mathbf{c}} \\ g \end{bmatrix} \quad (16)$$

By (15) and (16), $(E\mathbf{U})\mathbf{v} = \bar{\mathbf{c}}$. Since the last $n - 1$ entries of $\bar{\mathbf{c}}$ are zeros, and $E\mathbf{U}$ is unimodular, the last $n - 1$ rows of $E\mathbf{U}$ is a kernel basis of \mathbf{v} . The algorithm outputs the last $n - 1$ rows of $E\mathbf{U}$, except with row 1 and row $n_1 + 1$ interchanged. \square

LEMMA 17. The output $\|\mathbf{K}\|$ of SparseKernelBasis satisfies:

- (i) $\|\mathbf{K}\| \leq \|\mathbf{v}\|^2$, and
- (ii) the number of nonzero entries in \mathbf{K} is bounded by $n(1 + \log_2 n)$.

PROOF. The algorithm splits the problem into two subproblems with input $\mathbf{v}_1 \in \mathbb{Z}^{\lfloor n/2 \rfloor \times 1}$, and $\mathbf{v}_2 \in \mathbb{Z}^{\lfloor n/2 \rfloor \times 1}$. The kernel basis \mathbf{K} for \mathbf{v} is comprised of the kernel basis \mathbf{K}_i of \mathbf{v}_i ($i = 1, 2$), and one extra row

$$\mathbf{e} := \begin{bmatrix} -(g_2/g) \mathbf{e}_1 & | & (g_1/g) \mathbf{e}_2 \end{bmatrix} \in \mathbb{Z}^{1 \times n},$$

which we focus on to establish properties (i) and (ii).

First consider property (i). Since $g_i = \gcd(\mathbf{v}_i)$, we have $|g_i/g| \leq \|\mathbf{v}\|$ ($i = 1, 2$). According to the output specification of Algorithm 1 (VectorGcdex), we have $\|\mathbf{e}_i\| \leq \|\mathbf{v}_i\| \leq \|\mathbf{v}\|$ ($i = 1, 2$). It follows that $\|\mathbf{e}\| \leq \|\mathbf{v}\|^2$. Induction on n (base cases $n = 1, 2$) now shows that property (i) holds.

Now consider property (ii). Since the extra row \mathbf{e} has at most n nonzero entries, an upper bound on the number of nonzero entries in \mathbf{K} satisfies the recurrence

$$T(n) = T\left(\left\lfloor \frac{n}{2} \right\rfloor\right) + T\left(\left\lceil \frac{n}{2} \right\rceil\right) + n, \quad T(1) = 0.$$

From Sloane and Inc. [7], we obtain that

$$T(n) = n(\lfloor \log_2 n \rfloor + 3) - 2^{\lceil \log_2 n \rceil + 1}.$$

Dropping the floor and ceiling and simplifying yields the claimed upper bound. \square

LEMMA 18. *The running time of SparseKernelBasis is bounded by $O(n \mathbf{B}(\log \|\mathbf{v}\|) \log n)$ bit operations.*

PROOF. The nonrecursive work is dominated by the calls to Algorithm 1 (VectorGcdex) which has cost $O(n \mathbf{B}(\log \|\mathbf{v}\|))$. The running time of the algorithm thus satisfies the recurrence

$$T(n) = T\left(\left\lfloor \frac{n}{2} \right\rfloor\right) + T\left(\left\lceil \frac{n}{2} \right\rceil\right) + O(n \mathbf{B}(\log \|\mathbf{v}\|)),$$

which solves to the running time bound stated in the lemma. \square

THEOREM 19 (SPARSE KERNEL BASIS). *There exists an algorithm that takes as input a nonzero vector $\mathbf{v} \in \mathbb{Z}^{n \times 1}$, and returns as output a left kernel basis $\mathbf{K} \in \mathbb{Z}^{(n-1) \times n}$ for \mathbf{v} . The cost of the algorithm is $O(n \mathbf{B}(\log \|\mathbf{v}\|) \log n)$ bit operations. The output will satisfy $\|\mathbf{K}\| \leq \|\mathbf{v}\|^2$. The number of nonzero entries in \mathbf{K} will be bounded by $n(1 + \log_2 n)$.*

PROOF. Suppose \mathbf{v} has m nonzero entries. Let \mathbf{P} be an $n \times n$ permutation matrix such that

$$\mathbf{P}\mathbf{v} = \begin{bmatrix} \bar{\mathbf{v}} \\ \hline \end{bmatrix} \in \mathbb{Z}^{n \times 1}$$

with all entries in $\bar{\mathbf{v}} \in \mathbb{Z}^{m \times 1}$ nonzero. Use Algorithm SparseKernelBasis to compute a kernel basis $\bar{\mathbf{K}}$ for $\bar{\mathbf{v}}$. Then

$$\begin{bmatrix} \bar{\mathbf{K}} & \hline \hline \mathbf{I}_{n-m} \end{bmatrix} \mathbf{P}$$

is a kernel basis for \mathbf{v} . The result now follows from Lemmas 16–18. \square

We can now state the main result of the paper.

THEOREM 20. *There exists an algorithm that takes as input a full column rank $\mathbf{A} \in \mathbb{Z}^{n \times (n-1)}$, and returns as output a $\mathbf{C} \in \mathbb{Z}^{(n-1) \times n}$ such that $\mathbf{C}\mathbf{A}$ is a basis for \mathbf{A} . The output will satisfy $\|\mathbf{C}\| \leq n^4 \|\mathbf{A}\|^2$ and the number of nonzero entries in \mathbf{C} will be bounded by $n(1 + \log_2 n)$. The cost of the algorithm is $O(n^\omega \mathbf{B}(d) (\log n)^2)$ bit operations, where $d = \log n + \log \|\mathbf{A}\|$.*

PROOF. Use the algorithm supporting Theorem 14 (Augmentation vector) to compute a $\mathbf{v} \in \mathbb{Z}^{n \times 1}$ such that $\mathcal{L}(\begin{bmatrix} \mathbf{A} & \mathbf{v} \end{bmatrix})$ contains the last row of \mathbf{I}_n . Use the algorithm supporting Theorem 19 (Sparse kernel basis) to compute a left kernel \mathbf{C} of \mathbf{v} . Return \mathbf{C} . Correctness follows from Corollary 6.

By Theorem 14, \mathbf{v} satisfies $\|\mathbf{v}\| \leq n^2 \|\mathbf{A}\|$. The claimed bound for $\|\mathbf{C}\|$ and the bound on the number of nonzero entries in \mathbf{C} follow from Theorem 19. The runtime bound follows from Theorems 14 and 19. \square

8 CONCLUSIONS

Given a full column rank $\mathbf{A} \in \mathbb{Z}^{n \times (n-1)}$, we have given a deterministic algorithm that produces a $\mathbf{C} \in \mathbb{Z}^{(n-1) \times n}$ such that $\mathbf{C}\mathbf{A}$ is a basis for \mathbf{A} . Note that any such \mathbf{C} must have at least $n - 1$ nonzero entries, otherwise it would be rank deficient. We show that the \mathbf{C} produced by our algorithm has at most $n(1 + \log_2 n)$ nonzero entries, a logarithmic factor away from the lower bound.

A natural direction for further research is to extend the approach to a matrix $\mathbf{A} \in \mathbb{Z}^{n \times (n-k)}$ of full column rank $n - k$. Using the idea in the first part of the proof of Theorem 12, we may assume, up to permuting the rows of \mathbf{A} , that we can write \mathbf{A} as

$$\mathbf{A} = \begin{bmatrix} \bar{\mathbf{A}} \\ \mathbf{a}_1 \\ \vdots \\ \mathbf{a}_k \end{bmatrix}$$

where $\bar{\mathbf{A}}$ is nonsingular. Then k applications of the algorithm supporting Theorem 20 can be used to produce a \mathbf{C} such that $\mathbf{C}\mathbf{A}$ is a basis for \mathbf{A} .

Initialize $\mathbf{S}_0 = \bar{\mathbf{A}}$. For $i = 1, 2, \dots, k$ in succession, let

$$\bar{\mathbf{S}}_i := \begin{bmatrix} \mathbf{S}_{i-1} \\ \hline \mathbf{a}_i \end{bmatrix}$$

and use Theorem 20 to find a $\bar{\mathbf{C}}_i$ such that $\mathbf{S}_i := \bar{\mathbf{C}}_i \bar{\mathbf{S}}_i$ is a basis for $\bar{\mathbf{S}}_i$. To combine the $\bar{\mathbf{C}}_i$, set

$$\mathbf{C}_i = \left[\begin{array}{c|c} \bar{\mathbf{C}}_i & \hline \hline \mathbf{I}_{k-i} \end{array} \right] \in \mathbb{Z}^{(n-i) \times (n-i+1)}$$

for $1 \leq i \leq k$. The final transforming matrix is then given by

$$\mathbf{C} = \mathbf{C}_k \mathbf{C}_{k-1} \dots \mathbf{C}_1.$$

If $k = O(1)$, then the \mathbf{C} produced will still satisfy $\log \|\mathbf{C}\| \in O(\log n + \log \|\mathbf{A}\|)$, but is no longer guaranteed to be sparse. For non-constant k , further ideas will be required.

REFERENCES

- [1] S. Birmpilis, G. Labahn, and A. Storjohann. Deterministic reduction of integer nonsingular linear system solving to matrix multiplication. In *Proc. Int'l. Symp. on Symbolic and Algebraic Computation: ISSAC'19*, page 58–65, New York, NY, USA, 2019. ACM.
- [2] S. Birmpilis, G. Labahn, and A. Storjohann. A Las Vegas algorithm for computing the Smith form of a nonsingular integer matrix. In *Proc. Int'l. Symp. on Symbolic and Algebraic Computation: ISSAC'20*, page 38–45, New York, NY, USA, 2020. ACM.
- [3] J. von zur Gathen and J. Gerhard. *Modern Computer Algebra*. Cambridge University Press, 3rd edition, 2013.

- [4] S. Gupta, S. Sarkar, A. Storjohann, and J. Valeriote. Triangular x -basis decompositions and derandomization of linear algebra algorithms over $K[x]$. *Journal of Symbolic Computation*, 47(4), 2012. Festschrift for the 60th Birthday of Joachim von zur Gathen.
- [5] O. Ibarra, S. Moran, and R. Hui. A generalization of the fast LUP matrix decomposition algorithm and applications. *Journal of Algorithms*, 3:45–56, 1982.
- [6] J. Li and P. Q. Nguyen. Computing a lattice basis revisited. In *Proc. Int'l. Symp. on Symbolic and Algebraic Computation: ISSAC'19*, page 275–282, New York, NY, USA, 2019. ACM.
- [7] N. J. A. Sloane and The OEIS Foundation Inc. Entry a033156 in the on-line encyclopedia of integer sequences, 2020. URL <http://oeis.org/A033156>.
- [8] A. Storjohann. *Algorithms for Matrix Canonical Forms*. PhD thesis, Swiss Federal Institute of Technology, ETH–Zurich, 2000.
- [9] A. Storjohann. The shifted number system for fast linear algebra on integer matrices. *Journal of Complexity*, 21(4):609–650, 2005. Festschrift for the 70th Birthday of Arnold Schönhage.