

Smith Normal Form over Local Rings and Related Problems

by

Mustafa Elsheikh

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Doctor of Philosophy
in
Computer Science

Waterloo, Ontario, Canada, 2017

© Mustafa Elsheikh 2017

Examining Committee Membership

The following served on the Examining Committee for this thesis. The decision of the Examining Committee is by majority vote.

External Examiner	Jean-Guillaume Dumas Professeur, Mathématiques Appliquées Université Grenoble Alpes, France
Supervisor	Mark Giesbrecht Professor, David R. Cheriton School of Computer Science University of Waterloo, Canada
Internal Member	George Labahn Professor, David R. Cheriton School of Computer Science University of Waterloo, Canada
Internal Member	Arne Storjohann Associate Professor, David R. Cheriton School of Computer Science University of Waterloo, Canada
Internal-external Member	Kevin Hare Professor, Department of Pure Mathematics University of Waterloo, Canada

Author's Declaration

This thesis consists of material all of which I authored or co-authored: see Statement of Contributions included in the thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Statement of Contributions

I am the sole author of Chapters 1, 4, and 7. Chapter 2 is based on an article co-authored with Mark Giesbrecht, B. David Saunders, and Andy Novocin. Chapter 3 is based on collaboration with Mark Giesbrecht, and B. David Saunders. Chapter 5 is based on an article co-authored with Mark Giesbrecht. Chapter 6 is based on collaboration with Mark Giesbrecht, and Andy Novocin.

Abstract

The Smith normal form is a diagonalization of matrices with many applications in diophantine analysis, graph theory, system control theory, simplicial homology, and more recently, in topological analysis of big data. Efficient computation of Smith normal form is a well-studied area for matrices with integer and polynomial entries. Existing successful algorithms typically rely on elimination for dense matrices and iterative Krylov space methods for sparse matrices.

Our interest lies in computing Smith normal form for sparse matrices over local rings, where traditional iterative methods face challenges due to the lack of unique minimal polynomials. We explore different approaches to tackling this problem for two local rings: the integers modulo a prime power, and the polynomials modulo a power of an irreducible polynomial. Over local polynomial rings, we find success in linearization into larger dimension matrices over the base field. Effectively we transform the problem of computing the Smith normal form into a small number of rank problems over the base field. The latter problem has existing efficient algorithms for sparse and dense matrices.

The problem is harder over local integer rings. We take the approach of hybrid sparse-dense algorithms. We also tackle a restricted version of the problem where we detect only the first non-trivial invariant factor. We also give an algorithm to find the first few invariant factors using iterative rank-1 updates. This method becomes dense when applied to finding all the invariant factors.

We digress slightly into the related problem of preconditioning. We show that linear-time preconditioners are suitable for computing Smith normal form, and computing nullspace samples. For the latter problem we design an algorithm for computing uniform samples from the nullspace.

On a separate track, we focus on the properties of the Smith normal form decomposition. We relate the invariant factors to the eigenvalues. Our ultimate goal is to extend the applications of numerical algorithms for computing eigenvalues to computing the invariant factors of symbolic matrices.

Acknowledgements

I would like to thank Mark Giesbrecht for being a mentor, a collaborator, and a good friend.

I would like to thank my family for their constant love and support. Thank you Inas.

Table of Contents

1	Introduction	1
1.1	Smith Normal Form	1
1.2	The Black-Box Model	4
1.3	Probabilistic Algorithms	9
1.4	Thesis Outline	10
2	Sparse Smith Normal Form	12
2.1	Introduction	12
2.2	Previous Work	14
2.2.1	Finding the Last Invariant Factor	14
2.2.2	The Valence Method	15
2.2.3	Using the Characteristic Polynomial	17
2.3	Linearization of Polynomial Matrices	18
2.3.1	A Black-Box for the Embedding	22
2.3.2	The Algorithm	22
2.4	Smith Normal Form over $\mathbb{Z}/p^e\mathbb{Z}$	24
2.4.1	Nullspace Method	24
2.4.2	Probabilistic Dimension Reduction	25

2.4.3	Working with Small Primes	28
2.4.4	The Algorithm	28
2.5	Detecting Non-Trivial Smith Normal Form	30
3	Preconditioning	34
3.1	Introduction	34
3.2	Nullspace Sampling	36
3.2.1	Transpose Preconditioners	37
3.2.2	Avoiding Extension Fields	41
3.3	Smith Normal Form	43
3.3.1	Preconditioning and Determinantal Divisors	44
3.4	Conclusion	48
4	Rank Reduction	49
4.1	Introduction	49
4.2	Invariant Factors and Rank-1 Updates	52
4.3	Block Reduction	59
4.4	Nullspace Sampling	64
5	The Eigenvalues and the Invariant Factors	68
5.1	Introduction	68
5.1.1	The p -adic Correspondence	72
5.1.2	Previous Work	72
5.2	Establishing p -Correspondence	74
5.3	Density of p -Characterized Matrices	78
5.3.1	Density at Large Primes	82
5.4	Density at Small Primes	89

6	Ranks of Remainder Matrices	91
6.1	Introduction	91
6.2	Quotient and Remainder Matrices	92
6.2.1	Remainder of Rank-1 Matrices	93
6.2.2	A Note on Latin Squares	96
6.2.3	Rank Theorem	98
6.3	p -Adic Matrices	100
6.3.1	Binary Code Matrices	100
6.3.2	Non-Symmetric Matrices	104
6.3.3	Odd Primes	105
7	Conclusion and Future Work	107
	References	111

Chapter 1

Introduction

In this thesis we will study the problem of computing the Smith normal form over local rings, and some related problems regarding computing nullspace vectors, rank properties and eigenvalues properties. We are primarily concerned with the case of sparse matrices, but we will occasionally consider dense matrices as well. We start this chapter by defining several concepts central to our discussion. We finish by giving an overview of the chapters of this thesis.

1.1 Smith Normal Form

Let F be a field, and let A be an $n \times n$ matrix over F . We use $\det(A)$, $\text{rank}(A)$, $\text{charpoly}(A)$, $\text{im}(A)$, $\text{ker}(A)$ to denote the determinant, rank, characteristic polynomial, image, and right kernel of A , respectively. We use $\text{minpoly}(A)$ to denote the minimal polynomial of A , that is, the lowest degree non-zero monic polynomial $f \in F[x]$ such that $f(A) = 0$. We use $\{\lambda_1, \dots, \lambda_n\}$ to denote the eigenvalues of A .

Two notions of matrix transformation are notable: similarity transformation and equivalence transformation. We say that A and B are similar if there exists an invertible matrix W such that $A = W^{-1}BW$. On the other hand, we say that A and B are equivalent if there exists two invertible matrices P, Q such that $A = PBQ$. When working over a principal ideal ring R , we require that P, Q be unimodular, i.e., have a determinant which is unit

in the ring. Some matrix invariants, such as rank (when suitably defined over the ring), are preserved under both transformations. Other invariants, such as the characteristic polynomial (and hence) the eigenvalues, are preserved under the similarity transformation only. This will have important implications later when we discuss the preconditioning operations.

Consider a matrix A over a field. The eigenvalues of A will typically lie in an extension field \mathbb{K} (that is, the splitting field of the characteristic polynomial of A). It is well-known that over \mathbb{K} , the matrix A can be brought to the Jordan form using a similarity transformation $J = WAW^{-1}$. The matrix J is a block-diagonal matrix given by

$$J = \begin{bmatrix} J_1 & & & \\ & J_2 & & \\ & & \ddots & \\ & & & J_\ell \end{bmatrix},$$

where each J_i is called the Jordan block, which is a $k_i \times k_i$ matrix

$$J_i = \begin{bmatrix} \lambda_i & 1 & & \\ & \ddots & \ddots & \\ & & \ddots & 1 \\ & & & \lambda_i \end{bmatrix}.$$

The number of Jordan blocks associated with an eigenvalue and the dimension of each block is determined by the geometric and algebraic multiplicities of that eigenvalue. We note that $\text{minpoly}(A)$ is always a factor of $\text{charpoly}(A)$. However, as we will discuss later, if all the Jordan blocks corresponding to the zero eigenvalue have size at most 1, then $\text{charpoly}(A)$ and $\text{minpoly}(A)$ differ by at least a power of x , which will be useful in computing properties of sparse matrices such as rank.

The Jordan form and the eigenvalues describe a canonical representation of matrices under similarity transformations. On the other hand, the Smith normal form describes a canonical representation of matrices under equivalence transformations. Suppose that A has entries from a ring \mathbb{R} which we will require to be a principal ideal ring (PIR). By a

PIR we mean a ring in which every ideal is principal. There exist two unimodular matrices U, V over \mathbf{R} such that $A = USV$ where

$$S = \text{diag}(s_1, \dots, s_r, 0, \dots, 0),$$

and $s_i \mid s_{i+1}$ for all $1 \leq i \leq r-1$. The notion of rank over this ring is defined by the Smith normal form. That is, r is the rank of A over \mathbf{R} .

Definition 1.1. *The matrix S is called the Smith normal form of A , and the diagonal elements are called the invariant factors of A .*

The invariant factors are unique up to multiplication by units, but the transformation matrices are not necessarily unique. The existence and uniqueness of the Smith normal form was first proven by [Smith, 1861] for matrices over \mathbb{Z} and principal ideal domains. Kaplansky [Kaplansky, 1949] extended the notion of the Smith normal form to principal ideal rings.

The Smith normal form has found many applications in diophantine analysis [Chou and Collins, 1982], integer programming [Hu, 1969], combinatorics [Wallis et al., 1972], determining the structure of Abelian groups [Newman, 1972], class groups [Hafner and McCurley, 1989], system theory [Kailath, 1980], and in the study of symplectic spaces [Chandler et al., 2010].

The invariant factors are typically defined as the diagonal elements of S after the unimodular diagonalization of A . Alternatively, over a principal ideal domain, the invariant factors can be defined explicitly as follows. For $i \in [1, n]$, let \mathcal{C}_i^n denote the set of all i -tuples of integers of the form $t = (t_1, \dots, t_i)$ where $1 \leq t_1 < \dots < t_i \leq n$. For $\sigma, \tau \in \mathcal{C}_i^n$, let $A \binom{\sigma}{\tau}$ denote the determinant of the $i \times i$ submatrix of A selected by the rows $\sigma_1, \dots, \sigma_i$ and the columns τ_1, \dots, τ_i . This is the *minor* of A selected by σ and τ . Finally, let Δ_i denote i th *determinantal divisor* of A , that is,

$$\Delta_i = \text{gcd} \left\{ A \binom{\sigma}{\tau} : \sigma, \tau \in \mathcal{C}_i^n \right\},$$

the greatest common divisor of all $i \times i$ minors of A . Then the invariant factors are given by $s_1 = \Delta_1$ and $s_i = \Delta_i / \Delta_{i-1}$ for $i \in [2, r]$.

Our main concern will be computing the invariant factors of matrices over local rings. In what follows let R be a principal ideal domain, e be a positive integer, and π be a generator of a maximal ideal in R . By *local ring* we mean a principal ideal ring which has a unique maximal ideal. The local rings we study in this thesis are of the form $R/(\pi^e)$.

Over $R/(\pi^e)$ every ideal is generated by a power of π (where the exponent is an integer between 0 and $e - 1$). In this setting, S is given by

$$S = \text{diag}(\underbrace{1, \dots, 1}_{r_0}, \underbrace{\pi, \dots, \pi}_{r_1}, \dots, \underbrace{\pi^{e-1}, \dots, \pi^{e-1}}_{r_{e-1}}, \underbrace{0, \dots, 0}_{r_e}),$$

where $r_0 + r_1 + \dots + r_{e-1} = r$, and $r + r_e = n$.

It is known, see for example [Gerstein, 1977, Corollary 1], that the Smith normal form of any matrix A over R factors into the product $S = \prod S_{\pi_k}^{e_k}$, where each $S_{\pi_k}^{e_k}$ is the Smith normal form of the image of A in the ring $R/(\pi_k^{e_k})$. The product ranges over all irreducible factors of $s_r = \prod \pi_k^{e_k}$. In this context, some authors refer to S as the *global* Smith normal form and S_{π}^e is often called the *local Smith normal form* at π . Similarly, the invariant factors over $R/(\pi_k^{e_k})$ are called the *local invariant factors*. This local-global approach is often used in practice to compute the Smith normal form of integer matrices [Dumas et al., 2001, Lübeck, 2002], and polynomial matrices [Wilkening and Yu, 2011].

Concrete examples of $R/(\pi^e)$ are $\mathbb{Z}/p^e\mathbb{Z}$ where p is a prime, and $F[x]/(f^e)$ where f is an irreducible polynomial. These two rings capture the localization of Smith normal form of integer and polynomial matrices at a factor of the determinant (or a factor of the largest invariant factor). We often assume that p, e (or f, e) are given. In general, they can be found by computing the largest invariant factor, s_r , and then computing its factorization [Eberly et al., 2000, Dumas et al., 2001].

1.2 The Black-Box Model

The complexity analysis of any algorithm should take into account the cost of arithmetic operations in the underlying ring or field. When working over a finite field or a finite ring, the cost of the arithmetic operations is usually considered constant. In this case, we report

the *algebraic complexity* of the algorithm. However when the field or ring is infinite, then the size of the expressions can grow with the number of operations performed. In this case the *bit complexity* is an appropriate measure. When pertaining to matrix algorithms, the complexity is reported in terms of the matrix dimension n , the rank r , and the size of the largest entry in the input matrix (typically denoted by $\log \|A\|$).

It is often convenient to discuss the complexity without the logarithmic factors. We use the “soft-O” notation. We say that $f \in \tilde{O}(g)$ if $f \in O(g \log^c(g))$ for some positive constant c .

When designing algorithms for linear algebra we desire that the complexity be optimal in terms of space and time. A clear distinction has to be made between sparse matrices and dense matrices. In this context a matrix is sparse when the number of non-zero elements is much smaller than n^2 . A matrix is considered dense otherwise.

The notion of optimal algorithms for dense matrices is often linked to the complexity of matrix multiplication. This is because fundamental algorithms, such as Gaussian elimination, are known to have runtime in the order of matrix multiplication [Bunch and Hopcroft, 1974]. Over a field, if we can multiply two $n \times n$ matrices in $O(n^\omega)$ field operations, then we can perform Gaussian elimination and linear system solution, determinant, rank and other useful quantities in $O(n^\omega)$ field operations. The exponent ω has been progressively improved from 3 (using the naive method) to 2.8074 [Strassen, 1969] and beyond. The best known value for ω is currently 2.3728639 [Le Gall, 2014].

On the other hand, applying elimination on sparse matrices can be challenging. It is often the case that the sparse input matrix has a sub-quadratic number of non-zero elements, which we often denote by μ . The goal for designing efficient algorithms for sparse matrices is to have algorithms which are sensitive to the input size (i.e., to μ and $\log \|A\|$) and require no more than quadratic time and linear space. This is the case for most matrix problems over finite fields. As we will see in this thesis, we typically incur additional low degree factors in the size of the local ring.

When dealing with sparse matrices we would like to preserve the sparsity of the input matrix. If we use elimination, then adding two columns or rows could turn a zero entry into

a non-zero entry, i.e., we introduce a *fill-in*. Fill-in can easily increase the time and space requirements of the algorithm. It is known that the order of choosing pivots will impact the amount of fill-in introduced by Gaussian elimination. However, finding optimal pivoting is NP-complete [Yannakakis, 1981]. Several methods have been developed to reorder the input matrix such that the fill-in is minimized. Heuristic methods, such as the folklore Markowitz method, can be effective in practice (see [Dumas and Villard, 2002] for experiments with several reordering algorithms). Notably, nested dissection and graph-based methods have been successful in bounding the fill-in to $O(n \log n)$ non-zero entries [George, 1973], [Lipton et al., 1979], and have been applied with success to general fields [Alon and Yuster, 2010]. However, these methods are applicable to certain classes of matrices. In particular, they require certain classes of the underlying graphs such as planar graphs or bounded genus graphs. It is currently not known how to transform arbitrary input matrices to satisfy the properties required by these algorithms.

A successful approach to computing with sparse matrices is to treat the matrix as a *black-box*. In this model, direct manipulations of the matrix entries are not allowed. Instead, we are only allowed to compute matrix-vector products, which is often called *application* of the black-box.

Definition 1.2. *Let A be an $n \times n$ matrix over a ring \mathbb{R} . A black-box for A with cost μ is a function $\mathbb{R}^n \rightarrow \mathbb{R}^n$ which requires μ operations in \mathbb{R} to compute $Av \in \mathbb{R}^n$ for any vector $v \in \mathbb{R}^n$.*

When this is not ambiguous, we shall use the same letter (e.g., A) to denote the matrix A , and the black-box for A . The black-box approach is not particularly limiting. Given two black-boxes A and B , we can add, multiply, and compose black-boxes since $(A \pm B)v = Av \pm Bv$, and $ABv = A(Bv)$. We can also use repeated applications to compute the i th power of a black-box, i.e., $A^i v$ in i steps. We can transpose a black-box either by explicit construction, or by using Tellegen’s theorem [Tellegen, 1952, Penfield Jr. et al., 1970, Bostan et al., 2003] to get a transpose black-box at the same cost. We can augment two black-boxes, pad a black-box with zeros, or construct a black-box for the leading submatrix by means of constant number of applications, and padding/trimming

the input and output vectors. Finally, we can evaluate any polynomial f of degree d (i.e., $v \mapsto f(A)v$) using Horner's rule at the cost of $O(d\mu + dn)$ operations and space $O(n)$ elements.

This model is also useful when dealing with structured matrices such as Toeplitz matrices or circulant matrices, where matrix-vector multiplication can be computed efficiently using fast Fourier transform or similar techniques which costs less than n^2 operations.

The complexity of black-box algorithms is thus expressed in terms of the number of matrix-vector products used, and any additional arithmetic operations. That is, it is expressed in μ , n , and, when applicable, $\log \|A\|$. Space requirements are desired to be linear, i.e., kept to the storage of a few vectors.

There has been great success in applying black-box methods over finite and arbitrary fields, starting with Wiedemann's algorithm [Wiedemann, 1986], where the cost of many linear algebra problems has been reduced to computing a linear number of matrix-vector products. Wiedemann's algorithm has been further analyzed and enhanced by [Kaltofen and Saunders, 1991] and many others. It has been generalized to block projections by [Coppersmith, 1994] and [Kaltofen, 1995]. Alternative algorithms based on Lanczos' method and other Krylov subspace methods were developed by [Lambert, 1996], [Eberly and Kaltofen, 1997], [Eberly, 2004], and [Hovinen and Eberly, 2005] in scalar and block settings. It was shown that Lanczos' algorithm has the same cost as Wiedemann's algorithm. An ultimate goal is to add Smith normal form over local rings to the list of problems efficiently solvable by black-box methods.

The key idea of Wiedemann's algorithms is reducing matrix problems to computing the minimal polynomial of the matrix. For a matrix A of size $n \times n$ over a field F , and two vectors u, v of size n , the sequence

$$\{u^T v, u^T A v, \dots, u^T A^{2n-1} v\},$$

is linearly-recurrent and has a minimal generating polynomial $f \in F[x]$. In general we have $f(x) \mid \text{minpoly}(A)$. If u, v are chosen uniformly at random from F^n then with probability at least $1 - n/|F|$, we have $f(x) = \text{minpoly}(A)$. The cost of computing $\text{minpoly}(A)$ is

$2n\mu$ to compute the iterates $u^T A^i v$ and an additional $O(n^2)$ to compute $f(x)$ using the Berlekamp-Massey algorithm or similar Padé approximation methods. The storage is $O(n)$ elements in F .

Once we compute $\text{minpoly}(A) = x^d + f_{d-1}x^{d-1} + \dots + f_0$, we can solve many matrix problems. For example, the solution to the linear system $Ax = b$ is readily available as $A^{-1}b = (-1/f_0)(A^{d-1}b + f_{d-1}A^{d-2}b + \dots + f_2Ab + f_1b)$. To compute quantities such as determinant and rank, the common approach is to find a pre- and post-multiplier matrices $P, Q \in F^{n \times n}$ and construct $B = PAQ$ such that the minimal polynomial of B encodes the desired quantity. For example, the rank of A would equal the degree of $\text{minpoly}(B)$. This technique is called *preconditioning*. To avoid introducing an additional overhead, the matrices P, Q are often diagonal or structured matrices, and hence admit a fast black-box construction. To make this method general for any input matrix, P and Q are typically random and chosen from prespecified distribution. For a comprehensive list of preconditioning for various matrix problems see [Chen et al., 2002, §2]. We will discuss preconditioning in Chapter 3.

Computations over \mathbb{Z} and \mathbb{Q} and $F[x]$ can suffer from expression swell. To control the size of the intermediate coefficients, the minimal polynomial can be computed modulo a collection of primes followed by use of the Chinese Remainder Theorem. The cost increases by the number of primes required to reconstruct the minimal polynomial. Naively, the number of primes is $O(n)$ using Hadamard’s determinantal bound [von zur Gathen and Gerhard, 2003, Geddes et al., 1992]. Sharper bounds can be used; for example, [Dumas et al., 2000, §3] use Oval of Cassini bound [Brauer, 1946] to replace n with the degree of the minimal polynomial. The Chinese Remainder approach also has the practical advantage of parallelism. Alternatively, we can use Hensel lifting [Dixon, 1982] which also requires $O(n)$ iterations modulo a randomly chosen prime.

The degree of the minimal polynomial can be smaller than n . Hence the computation of the minimal generating polynomial of the sequence $\{u^T A^i v\}$ can benefit from the *early termination* technique of [Lobo, 1995, Kaltofen et al., 2000, Eberly, 2003]. The early termination heuristic stops the iterative computation when the minimal polynomial remains the

same after few iterations. The complexity of Wiedemann’s algorithm becomes $O(\tilde{r}\mu + r^2)$.

We will often assume that an appropriate choice of black-box algorithms is made, noting that there is considerable difference in their effectiveness in practice and over various ground fields.

Wiedemann-based methods have been successful in reducing the cost of many linear algebra problems for sparse matrices using the black-box model over fields. However, linearly-recurrent sequences over local rings do not have a unique minimal generator and therefore these algorithms fail to work over local rings. We propose algorithms that try to avoid direct minimal polynomial computations over the local rings.

1.3 Probabilistic Algorithms

The aforementioned algorithms often require making random choices of vectors or matrices. The success of these algorithms (i.e., the event that a desired property holds) is probabilistic and often relies on the following lemma.

Fact 1.1 (Schwartz-Zippel lemma [Zippel, 1979, Schwartz, 1980, Demillo and Lipton, 1978]). *Let F be a field (or an integral domain) and S be a finite subset of F . Let $f \in F[x_1, \dots, x_n]$ be a non-zero polynomial of total degree d . If v_1, \dots, v_n are chosen independently and uniformly at random from S , then $\Pr[f(v_1, \dots, v_n) \neq 0] \geq 1 - d/|S|$.*

There are two types of probabilistic algorithms. Monte Carlo algorithms always terminate and return correct results with controllably high probability. If the output is correct with probability at least $1/2$, then the success probability can be amplified to $1 - \epsilon$ for any small $\epsilon \in [0, 1]$ as follows. Repeat the algorithm $O(\log \frac{1}{\epsilon})$ times and apply a majority voting scheme on the outputs, then the success probability is amplified to at least $1 - \epsilon$ by the Chernoff bound [Motwani and Raghavan, 1995].

If we can verify the correctness of the output, then the algorithm is randomized of the Las Vegas type – always correct, probably terminates. In this case we have to repeat the algorithm until the output is verified to be correct. The reported complexity is the expected

runtime. For example, if we are computing a nullspace vector v using a randomized Las Vegas algorithm, then we can repeat the algorithm with different random choices until $Av = 0$.

1.4 Thesis Outline

This thesis is motivated by the problem of computing Smith normal forms of sparse matrices over local rings. In the course of studying this problem, we also encounter and study a few related and interesting problems. We will focus on the following problems:

1. In Chapter 2 we study the problem of computing the Smith Normal Form for sparse matrices over $\mathbb{Z}/p^e\mathbb{Z}$ and $\mathbb{F}[x]/(f^e)$. We will give two algorithms towards this end. We will also give an algorithm for finding the first non-trivial invariant factor for a black-box matrix. For $\mathbb{F}[x]/(f^e)$ the algorithms depend on a number of tools, such as sparse matrix rank computation over finite fields, for which the best-known efficient algorithms are probabilistic.
2. In Chapter 3 we will discuss the problem of preconditioning. We will extend the application of well-known linear time preconditioners to the problems of computing nullspace vectors and computing a Smith normal form.
3. In Chapter 4 we will discuss a new approach to computing the Smith normal form. A simple elimination process based on rank-1 updates will be discussed. Interestingly, this will give us an improved nullspace sampling algorithm.
4. In Chapter 5 we take a different approach to understanding the invariant factors over local rings. We consider their relationship with other matrix invariants. A new characterization of the invariant factors in terms of the eigenvalues will be presented. We use p -adic valuations as a measure of size. Density estimates will be given for cases when the p -adic valuation of the eigenvalues coincide with the p -adic valuation of the invariant factors.

5. In Chapter 6 we will focus on the Smith normal form decomposition $A = USV$. We will consider two computational notions that capture the local representation of matrices, and the carry digits that occur in the computations: the base- p expansion of matrix entries and the action of the remainder operator on matrices.

Chapter 2

Sparse Smith Normal Form

In this chapter we present two algorithms for computing the Smith normal form of sparse matrices over local rings. The two rings under study are localization of the polynomial ring in one variable and localization of the integers. In designing the algorithms we will take into account the sparsity of the input matrix. The results of this chapter appeared in [Elsheikh et al., 2012].

2.1 Introduction

We are primarily concerned with computing the Smith normal form of sparse matrices over local principal ideal rings of the form $\mathbb{R}/(\pi^e)$ where \mathbb{R} is a principal ideal domain, e is a positive integer, and π is a generator of a maximal ideal in \mathbb{R} . The Smith normal form of any matrix over $\mathbb{R}/(\pi^e)$ has powers of π on its diagonal. In particular, we use the following notation to count the multiplicities of the invariant factors:

$$S = \text{diag}(\underbrace{1, \dots, 1}_{r_0}, \underbrace{\pi, \dots, \pi}_{r_1}, \dots, \underbrace{\pi^{e-1}, \dots, \pi^{e-1}}_{r_{e-1}}, \underbrace{0, \dots, 0}_{r_e}),$$

where $r_0 + r_1 + \dots + r_{e-1} = r$, and $r + r_e = n$.

Existing approaches to computing the Smith normal form differ between sparse and dense matrices. Table 2.1 summarizes the time complexities of various existing algorithms,

Table 2.1: Time complexity of computing the Smith normal form.

Algorithm	Complexity	Ring	Randomness	Sparsity
[Storjohann, 2000]	n^ω	PIR	Deterministic	Dense
[Storjohann, 2000]	$n^{\omega+1}$	\mathbb{Z}	Deterministic	Dense
[Eberly et al., 2000]	$n^{2+\omega/2}$	\mathbb{Z}	Monte Carlo	Dense
[Dumas et al., 2001]	n^3	$\mathbb{Z}/p^e\mathbb{Z}$	Deterministic	Dense
[Storjohann and Labahn, 1997]	$n^3 \deg(A)(\deg(A) + n^2)$	$\mathbb{Z}[x]$	Las Vegas	Dense
[Storjohann, 2003]	$n^\omega \deg(A)$	$\mathbb{F}[x]$	Las Vegas	Dense
[Zhou et al., 2015]	$n^\omega \text{avg}(\deg(A))$	$\mathbb{F}[x]$	Las Vegas	Dense
[Kaltofen and Villard, 2005]	$n^{2.69726263}$	\mathbb{Z}	Monte Carlo	Dense
[Giesbrecht, 2001]	$n^2\mu + n^3$	\mathbb{Z}	Monte Carlo	Sparse
[Dumas et al., 2001]	$n\mu \deg(\text{minpoly})$	\mathbb{Z}	Monte Carlo	Sparse
[Eberly et al., 2007]	$n^{1.579}\mu + n^{2.579}$	\mathbb{Z}	Monte Carlo	Sparse

while omitting logarithmic factors in n and the size of the maximal entry, $\log \|A\|$. The presented complexity is counted in terms of bit operations for \mathbb{Z} , field operations for $\mathbb{F}[x]$, and ring operations for PIRs and $\mathbb{Z}/p^e\mathbb{Z}$. The cost of matrix-vector multiplication is denoted by μ . Complexity statements were simplified by assuming square matrices of order n and replacing the rank factors with n . See the cited references for refined complexity statements.

Notably, for dense matrices, elimination offers optimal or near-optimal complexity in terms of matrix multiplication time. For sparse matrices (where the number of non-zero elements $\mu \ll n^2$), elimination generally introduces fill-in and hence prohibitive storage requirements.

Existing non-elimination based algorithms are essentially cubic. However, blocking techniques are used to achieve sub-cubic complexity. Effectively these algorithms extract the invariant factors from the minimal (or characteristic) polynomial of the matrix [Giesbrecht, 2001], [Eberly et al., 2007] or its largest coefficient [Dumas et al., 2001]. While the minimal polynomial itself can be computed in quadratic time over a field, the extra factors

in the complexity correspond to recovering the full precision of the invariant factors which can be as large as n bits.

The problem of computing the Smith normal form of sparse matrices over local rings presents its own challenges. One could simply carry out the computations over the global ring \mathbf{R} , then reduce the results modulo π^e . However computations over a ring \mathbf{R} (e.g., over $\mathbf{F}[x]$ or \mathbb{Z}) suffer from coefficient growth which is not clearly necessary in the local ring where the precision is bounded by e . For example, over $\mathbb{Z}/32\mathbb{Z}$ the invariant factors are bounded by 32, and thus one might hope to perform all computations modulo 32, and not with integers larger than 32.

We attempt to solve this problem by designing algorithms which do not suffer from fill-in or expression swell. In doing so, we pursue black-box algorithms. However we will occasionally fall back to dense algorithms.

2.2 Previous Work

In this section we review some of the relevant previous work, and highlight some of the key ideas in their algorithms.

2.2.1 Finding the Last Invariant Factor

In some applications it suffices to compute the last invariant factor only. For dense polynomial matrices, the deterministic algorithm by [Zhou et al., 2015] finds the last invariant factor over $\mathbf{F}[x]$ for any field \mathbf{F} in $O(n^\omega s)$ where s is a bound on the average column degree. We will present the approach of [Eberly et al., 2000]. While their results are focused on dense matrices, the core idea of finding the last invariant factor using rational solution is useful for sparse matrices as we shall see.

The algorithm employs an earlier idea by [Pan, 1988, Abbott et al., 1999] to find the largest invariant factor of a matrix by solving a random linear system. To compute the rest of the invariant factors, a random perturbation scheme is used to compute the k th

invariant factor for any $k \in [1, n]$. Combined with binary search, this method can compute the Smith normal form of any integer matrix in $O^\sim(n^{3.5})$ bit operations [Eberly et al., 2000, Theorem 4.2]. In fact, the complexity is sensitive to the size of the determinant.

More concretely, we are given a non-singular integer matrix A of size $n \times n$. To compute s_n , the largest invariant factor of A , we select a random integer vector b , and solve the system $Ax = b$ over \mathbb{Q} . Then s_n can be inferred from the LCM of the denominators of the entries of x . The process is repeated to get a provably good success probability. The cost is dominated by the time to solve the linear system over \mathbb{Q} , which is $O^\sim(n^3)$ bit operations if one uses p -adic lifting.

To compute the $(n - k)$ th invariant factor of A , or s_{n-k} , we construct two random integer matrices U and V of rank k , and apply the perturbation $B = A + UV$. It is shown that the n -th invariant factor of B is related to the k -th invariant factor of A . In particular, the GCD of $s_n(A)$ and $s_n(A + UV)$ is $s_{n-k}(A)$. Again, $s_n(A + UV)$ can be computed as above by solving a random linear system over the rationals.

A naive application of this method requires $n - 1$ random perturbations to compute $s_{n-1}, s_{n-2}, \dots, s_1$ which would result in quartic complexity. However, the authors show that the number of *distinct* invariant factors is bounded by $(\log \det A)^{0.5}$, which is on the order of $n^{0.5}$ (omitting the factor in $\|A\|$ for brevity). Binary search can be used to construct $n^{0.5}$ random perturbations of A and get all the invariant factors. The overall complexity is reduced to $O^\sim(n^{3.5})$. In order to apply this binary search method to sparse matrices, the dense perturbation matrix U and V should be replaced with suitable sparse matrices.

2.2.2 The Valence Method

The comprehensive work of [Dumas et al., 2000, Dumas et al., 2001, Dumas et al., 2003] present one of the best and most practical approaches to computing the Smith normal form of sparse integer matrices. Let A be an integer matrix of size $n \times n$ and rank r . Let d be the degree of the minimal polynomial of A . Let μ be cost of applying A to a vector. Let e be the largest exponent in the factorization of the largest invariant factor. The key steps are the following.

1. Compute the *valence* (the trailing coefficient of the characteristic polynomial) over \mathbb{Z} using Chinese remaindering. Typically this requires $O(n)$ primes due to the often pessimistic Hadamard determinantal bound. Instead, the ovals of Cassini provide a sharper bound, and hence only $O(d)$ primes are required. In order to maintain a low degree minimal polynomial, we can use non-preconditioned matrix AA^T or $A^T A$. This step requires $O(\tilde{d}^2 \mu)$ bit operations.
2. For each prime dividing the valence, compute the local Smith normal form using elimination. The algorithm requires $O(rn^2)$ operations. This algorithm is practical when r is much smaller than n .
3. If only the last invariant factor at prime p is required, then a clever algorithm is presented which costs $O(n\mu e^2)$ bit operations. The key idea is to precondition A into $B = p^e I + qA$ where q is an arbitrary prime. The matrix B is equivalent to A modulo p^e . However B is a non-singular diagonal matrix modulo q which implies a very fast q -adic lifting. This allows for removing a factor of n from the lifting complexity. To maintain the factor n saving, one should not lift the entire n components of the solution vector. Instead, a preconditioning matrix is used such that the first entry in the solution vector reveals the largest invariant factor. These ideas are combined to give a $O(\tilde{n}\mu e^2)$ bit complexity to compute the last invariant factor at p .

We note that this cannot be efficiently extended to compute all the invariant factors. Suppose we use a binary search method similar to [Eberly, 2000]. This will introduce dense preconditioners and μ will become n^2 . So the overall complexity would be at best (replace μ with n^2): $\sqrt{n} \times ne^2 n^2 = n^{3.5} e^2$.

The overall cost of computing the Smith normal form depends on the underlying approach. If local elimination is used, then the bit complexity is $O(\tilde{d}rn^2 + d\mu)$. If we only compute the last invariant factor for each relevant prime, then the bit complexity is $O(\tilde{d}n\mu e)$.

The valence method is very practical when the degree of the minimal polynomial d is much smaller than the rank, as was demonstrated by the authors in the case of simplicial

homology matrices.

2.2.3 Using the Characteristic Polynomial

The valence method used only the trailing coefficient of the characteristic polynomial. The work of [Giesbrecht, 1995], [Giesbrecht, 1996], [Giesbrecht, 2001] utilizes all the coefficients of the characteristic polynomial. This method works on both dense and sparse matrices, but we are more concerned with sparse matrices here. The key idea in this algorithm is reducing the computation of the determinantal divisors (and hence the invariant factors) to computing the coefficients of the characteristic polynomial as follows.

- For a given prime p , if the power of p dividing the leading $k \times k$ minor equals the power of p dividing its k th determinantal divisor, then the $(n - k)$ th coefficients of the characteristic polynomial is a multiple of the k th determinantal divisor.
- The algorithm tries to precondition the input matrix such that its characteristic polynomial and its minimal polynomial differ only by a power of x . Furthermore, all the leading minors of the preconditioned matrix should satisfy the condition above.
- Using a Toeplitz-based construction, random preconditioning can be achieved without introducing fill-in, and without introducing significant overhead in the matrix-vector application.
- Some extraneous primes will appear in the coefficients of the characteristic polynomial of the preconditioned matrix. By repeating the preconditioning and the characteristic polynomial computation, we can take the GCD of the respective coefficients. This will probabilistically remove the extraneous primes. Special attention is paid to the success probability of this step.

The algorithm uses $O(n^2 \log \|A\|)$ black-box calls modulo a prime p , and an additional $O(n^3 \log^2 \|A\|)$ bit operations. Assume the black-box costs $\mu \in O(n)$ operations. Then the bit complexity is cubic, or $O^\sim(n^{3+\epsilon} \log \|A\| + n^3 \log^2 \|A\|)$.

2.3 Linearization of Polynomial Matrices

We now present our first algorithm. Let \mathbf{F} be a field, $f \in \mathbf{F}[x]$ be an irreducible polynomial of degree d , and $e > 1$ be an integer. The ring $\mathbf{L} = \mathbf{F}[x]/(f^e)$ is a local ring and all its ideals are of the form $f^i\mathbf{L}$ for $0 \leq i < e$. Let $A \in \mathbf{L}^{n \times n}$ be a matrix over \mathbf{L} , whose Smith normal form is given by

$$S = \text{diag}(\underbrace{1, \dots, 1}_{r_0}, \underbrace{f, \dots, f}_{r_1}, \dots, \underbrace{f^{e-1}, \dots, f^{e-1}}_{r_{e-1}}, 0, \dots, 0). \quad (2.1)$$

Our goal is to compute the multiplicities, $\{r_0, r_1, \dots, r_{e-1}\}$, efficiently when A is sparse or given by a black-box. We assume that f and e are known a priori. In practice, we are given a matrix over $\mathbf{F}[x]$. We can compute its largest invariant factor using existing methods, we can then factor it to get each irreducible power in the factorization.

The approach we take is to embed the ring $\mathbf{L}^{n \times n}$ in the ring $\mathbf{F}^{nde \times nde}$ and reduce the computation of the Smith normal form to finding ranks of matrices in the base field \mathbf{F} , where known fast algorithms can be used.

This approach is known as *linearization* and has been used in the context of computing Hermite normal form [Kaltofen et al., 1987]. We first describe the classical embedding of \mathbf{L} into $\mathbf{F}^{de \times de}$. We then show that the multiplicities r_i 's for matrices over \mathbf{L} are revealed by their images over \mathbf{F} .

First, define the map $\varphi_e : \mathbf{L} \rightarrow \mathbf{F}^{de \times de}$, which maps polynomials into de by de matrices, as follows. Suppose $f^e = a_0 + a_1x + \dots + a_{de-1}x^{de-1} + x^{de}$, whose companion matrix is

$$C_{f^e} = \begin{bmatrix} 0 & 0 & \cdots & -a_0 \\ 1 & \ddots & & -a_1 \\ \vdots & \ddots & \ddots & \vdots \\ 0 & \cdots & 1 & -a_{de-1} \end{bmatrix}.$$

Define $\varphi_e(x) = C_{f^e}$, and $\varphi_e(x^i) = \varphi_e(x)^i$. By linearity, extend φ_e to all elements of $g = g_0 + g_1x + \dots + g_{de-1}x^{de-1} \in \mathbf{L}$ such that $\varphi_e(g) \in \mathbf{F}^{de \times de}$ is given by:

$$\varphi_e(g) = g(C_{f^e}) = g_0I + g_1C_{f^e} + g_2C_{f^e}^2 + \dots + g_{de-1}C_{f^e}^{de-1}.$$

We now show that φ_e is an isomorphism between the polynomials in \mathbf{L} and the subset of matrices given by $\mathbf{F}[C_{f^e}]$.

Lemma 2.1. *The map φ_e is a ring isomorphism between \mathbf{L} and $\mathbf{F}[C_{f^e}]$.*

Proof. We have $\varphi_e(1) = I$. Also, φ_e is a ring homomorphism because for any two polynomials $g, h \in \mathbf{L}$ we have

$$\begin{aligned}\varphi_e(g + h) &= \varphi_e\left(\sum_i (g_i + h_i)x^i\right) = \sum_i (g_i + h_i)C_{f^e}^i \\ &= \sum_i g_i C_{f^e}^i + \sum_i h_i C_{f^e}^i = \varphi_e(h) + \varphi_e(g),\end{aligned}$$

and

$$\begin{aligned}\varphi_e(g \cdot h) &= \varphi_e\left(\sum_i g_i h_{de-i-1} x^i\right) = \sum_i g_i h_{de-i-1} C_{f^e}^i \\ &= \left(\sum_i g_i C_{f^e}^i\right) \cdot \left(\sum_j h_j C_{f^e}^j\right) = \varphi_e(h) \cdot \varphi_e(g).\end{aligned}$$

For all $g \in \mathbf{L}$ there exists a corresponding element in $\mathbf{F}[C_{f^e}]$ given by $\phi(g)$. Conversely, any matrix $G \in \mathbf{F}[C_{f^e}]$ can be written as a linear combination $\sum_i g_i C_{f^e}^i$ where $i < de$ because C_{f^e} is a companion matrix of a polynomial of degree de . So $\varphi_e^{-1}(G)$ can be given by a polynomial in \mathbf{L} whose coefficients are g_i 's.

To show that the inverse is unique we let $g_1, g_2 \in \mathbf{L}$ be such that $h = g_1 - g_2 \neq 0$ and $\varphi_e(g_1) = \varphi_e(g_2)$. Then using the ring homomorphism properties we get $\varphi_e(g_1) - \varphi_e(g_2) = 0$, or $\varphi_e(g_1 - g_2) = \varphi_e(h)$. We get $\varphi_e(\sum h_i x^i) = \sum h_i C_{f^e}^i = 0$. But this implies that the minimal polynomial of C_{f^e} divides h . This is a contradiction because $\deg h \leq \max\{\deg g_1, \deg g_2\} < de$, which can not be divisible by a polynomial whose degree is de because $\deg \text{minpoly } C_{f^e} = de$. Therefore φ_e is a bijection between the sets \mathbf{L} and $\mathbf{F}[C_{f^e}]$. \square

The embedding φ_e has useful rank properties.

Lemma 2.2. $\text{rank}(\varphi_e(f^i)) = d(e - i)$ for all $0 \leq i \leq e$.

Proof. If $i = 0$, then $\text{rank}(\varphi_e(f^0)) = \text{rank}(I_{de \times de}) = de$. If we associate the elements of \mathbb{L} with vectors in \mathbb{F}^{de} , then for all $i > 0$ the matrix $f^i(C_{f^e})$ acts on \mathbb{F}^{de} as multiplication by $f^i \bmod f^e$. Thus its nullspace is generated by the images of polynomials in $f^{e-i}\mathbb{L}$. Any element $h \in f^{e-i}\mathbb{L}$ can be defined by choosing $g \in \mathbb{L}$ and forming the product $f^{e-i}g \pmod{f^e}$. Now write $g = Qf^i + R$ using Euclidean division, where degree of R is less than di . We get $h = Qf^i \cdot f^{e-i} + Rf^{e-i} \pmod{f^e} = Rf^{e-1}$. So h is completely specified by the di coefficients of R , and $f^{e-i}\mathbb{L}$ as a vector space has dimension di . Therefore $\text{rank}(f^i(C_{f^e})) = de - di$. \square

We extend the map φ_e to $n \times n$ matrices over \mathbb{L} using element-wise application. For any $A \in \mathbb{L}^{n \times n}$, $\varphi_e(A)$ is an $nde \times nde$ matrix over \mathbb{F} , where every entry $a_{i,j}$ of A is replaced by the $de \times de$ block $\varphi_e(a_{i,j})$. Applying φ_e to (2.1), we get

$$\begin{aligned} \varphi_e(S) = \text{diag}(\underbrace{\varphi_e(1), \dots, \varphi_e(1)}_{r_0}, \underbrace{\varphi_e(f), \dots, \varphi_e(f)}_{r_1}, \dots, \\ \underbrace{\varphi_e(f^{e-1}), \dots, \varphi_e(f^{e-1})}_{r_{e-1}}, 0, \dots, 0) \in \mathbb{F}^{nde \times nde}. \end{aligned} \quad (2.2)$$

Since $\varphi_e(S)$ is a block diagonal matrix, we get the following additive rank:

$$\text{rank}(\varphi_e(S)) = \sum_{i=0}^{e-1} r_i \text{rank}(\varphi_e(f^i)) = \sum_{i=0}^{e-1} r_i d(e - i), \quad (2.3)$$

where the latter equality uses Lemma 2.2.

Lemma 2.3. *If $U \in \mathbb{L}^{n \times n}$ is invertible, then $\varphi_e(U) \in \mathbb{F}^{nde \times nde}$ is invertible.*

Proof. If U is invertible, then there exists a $W \in \mathbb{L}^{n \times n}$ such that $UW = I$. Now apply φ_e to both sides and use the linearity of φ_e to get $\varphi_e(UW) = \varphi_e(U)\varphi_e(W) = \varphi_e(I) = I_{nde}$. So $\varphi_e(U)$ is invertible and its inverse is given by $\varphi_e(W)$. \square

We do not intend to map the Smith normal form $A = USV$ over \mathbb{L} to the Smith normal form $\varphi_e(A) = \varphi_e(U)\varphi_e(S)\varphi_e(V)$ over \mathbb{F} . This is because over \mathbb{F} there is no useful notion of the Smith normal form. Any matrix of rank r over \mathbb{F} , will have a trivial Smith normal form of r ones. Instead, we map the multiplicities in the invariant factors of A over \mathbb{L} to the rank of $\varphi_e(A)$ over \mathbb{F} .

Theorem 2.1. *Let $A \in \mathbb{L}^{n \times n}$ have the Smith normal form*

$$\text{diag}(\underbrace{1, \dots, 1}_{r_0}, \underbrace{f, \dots, f}_{r_1}, \dots, \underbrace{f^{e-1}, \dots, f^{e-1}}_{r_{e-1}}, 0, \dots, 0),$$

then $\text{rank}(\varphi_e(A)) = der_0 + d(e-1)r_1 + \dots + dr_{e-1}$.

Proof. There exist unimodular matrices $U, V \in \mathbb{L}^{n \times n}$ such that $UAV = S$. By the isomorphism of φ_e , we have $\varphi_e(U)\varphi_e(A)\varphi_e(V) = \varphi_e(S)$. By Lemma 2.3, $\varphi_e(U), \varphi_e(V)$ are invertible and thus $\text{rank}(\varphi_e(A)) = \text{rank}(\varphi_e(S)) = der_0 + d(e-1)r_1 + \dots + dr_{e-1}$ using (2.3). \square

As a consequence of Theorem 2.1 we have the following corollary.

Corollary 2.1. *Let $\rho_{\ell-1}$ denote $\text{rank}(\varphi_e(A \bmod f^\ell))$, where $1 \leq \ell \leq e$. Then*

$$\begin{bmatrix} d & 0 & \cdots & 0 \\ 2d & d & \cdots & 0 \\ \vdots & \ddots & \ddots & \vdots \\ ed & \cdots & 2d & d \end{bmatrix} \begin{bmatrix} r_0 \\ r_1 \\ \vdots \\ r_{e-1} \end{bmatrix} = \begin{bmatrix} \rho_0 \\ \rho_1 \\ \vdots \\ \rho_{e-1} \end{bmatrix}. \quad (2.4)$$

Proof. For all $1 \leq \ell \leq e$, we have $\text{rank}(\varphi_\ell(A \bmod f^\ell)) = d\ell r_0 + d(\ell-1)r_1 + \dots + dr_{\ell-1}$. The statement follows immediately from substituting ℓ with $1, 2, \dots, e-1$. \square

The system (2.4) can be solved in $O(e)$ operations over \mathbb{Z} since

$$\begin{bmatrix} 1 & & & \\ 2 & 1 & & \\ \vdots & \ddots & \ddots & \\ e & \cdots & 2 & 1 \end{bmatrix}^{-1} = \begin{bmatrix} 1 & & & \\ -2 & \ddots & & \\ 1 & \ddots & \ddots & \\ & \ddots & \ddots & \ddots \\ & & 1 & -2 & 1 \end{bmatrix}.$$

Next we consider how to efficiently compute $\{\rho_0, \rho_1, \dots, \rho_{e-1}\}$ for a given black-box matrix.

2.3.1 A Black-Box for the Embedding

Given a black-box for $A \in \mathbb{L}^{n \times n}$ we can construct a black-box for $\varphi_\ell(A \bmod f^\ell)$, for any $\ell \leq e$, at not much higher cost. We assume that the black-box for $A \in \mathbb{L}^{n \times n}$ costs μ operations in \mathbb{F} . If each column and row of A has at least one non-zero entry, i.e., A has at least n non-zero entries, then $\mu \geq nde$ because each entry of \mathbb{L} is a polynomial with de coefficients.

We show how to efficiently perform black-box computations under φ_e transformations. Let $M(d)$ denote the cost of multiplying two polynomials of degree at most d .

Lemma 2.4. *Suppose we are given a black-box for $A \in \mathbb{L}^{n \times n}$, where $\mathbb{L} = \mathbb{F}[x]/(f^e)$ as above. Let $\ell \in \{1, \dots, e\}$ and $\hat{v} \in \mathbb{F}^{d\ell n}$ with unique pre-image $v \in \mathbb{F}[x]/(f^\ell)$. Then we can compute $\hat{w} = \varphi_\ell(A \bmod f^\ell)\hat{v} \in \mathbb{F}^{d\ell n}$ using $O(\mu + nM(de))$ operations in \mathbb{F} .*

Proof. Assume that $\hat{v} \in \mathbb{F}^{d\ell n}$ is labelled as:

$$\hat{v} = (\hat{v}_{1,0}, \dots, \hat{v}_{1,d\ell-1}, \hat{v}_{2,0}, \dots, \hat{v}_{2,d\ell-1}, \dots, \hat{v}_{n,0}, \dots, \hat{v}_{n,d\ell-1}).$$

Construct the vector $v = (v_1, \dots, v_n) \in \mathbb{L}^n$, where $v_i = \sum_{0 \leq j < d\ell} \hat{v}_{i,j} x^j \in \mathbb{F}[x]$. Now, compute $w = Av \bmod f^\ell \in \mathbb{L}^n$ using μ operations for the black-box evaluation plus $O(nM(de))$ operations in \mathbb{F} for the n modular reduction. Let $w = (w_1, \dots, w_n)$ where $w_i = \sum_{0 \leq j < d\ell} \hat{w}_{i,j} x^j \in \mathbb{F}[x]$. Then \hat{w} is given by

$$\hat{w} = (\hat{w}_{1,0}, \dots, \hat{w}_{1,d\ell-1}, \dots, \hat{w}_{n,0}, \dots, \hat{w}_{n,d\ell-1}).$$

□

2.3.2 The Algorithm

The algorithm for computing the Smith normal form of a matrix $A \in \mathbb{L}^{n \times n}$ given by a black-box is now straightforward. Using Theorem 2.1 and Lemma 2.4 we reduce the computation of ρ_i 's in (2.4) to computing ranks of matrices over the ground field \mathbb{F} , which can be accomplished using existing efficient black-box algorithms over fields, for example based on Wiedemann's algorithm.

Algorithm 2.1. Given a black-box for $A \in \mathbb{L}^{n \times n}$, return r_0, \dots, r_{e-1} such that r_i is the multiplicity of f^i in the Smith normal Form of A , and the multiplicity of zero is $n - \sum_i r_i$.

1. For all $\ell \in \{1, \dots, e\}$, invoke a black-box rank algorithm on the black-box for $\varphi_\ell(A \bmod f^\ell) : \mathbb{F}^{d\ell n} \rightarrow \mathbb{F}^{d\ell n}$. Let $\rho_{\ell-1} = \text{rank}(\varphi_\ell(A \bmod f^\ell))$.
2. Solve (2.4) for r_0, \dots, r_{e-1} .
3. Return r_0, \dots, r_{e-1} .

Theorem 2.2. *Algorithm 2.1 is correct, and requires $\tilde{O}(\mu de^2 n)$ operations in \mathbb{F} . The space requirement of the algorithm is $O(\text{den})$ elements in \mathbb{F} .*

Proof. The correctness follows from the results and discussions in this section. We analyze the time and space complexity of step (1), which dominates the cost. It requires $O(de^2 n)$ black-box evaluations, and storage for $O(\text{den})$ elements in \mathbb{F} . \square

If A has a linear number of entries then $\mu \in \tilde{O}(\text{den})$ using fast polynomial arithmetic. The complexity in this case is $\tilde{O}(d^2 e^3 n^2)$. We expect any algorithm in this setup to cost at least $\tilde{O}(\text{den}^2)$ operations in \mathbb{F} where the term $\tilde{O}(de)$ accounts for the polynomial arithmetic with de coefficients over \mathbb{F} . Our algorithm is a factor of de^2 away from this complexity.

As a future work, one could hope to reduce this complexity by factor of e . For example, we can reuse the same iterate vectors in Algorithm 2.1 for the all values of e . However, in this case the success probability of the rank computation will depend on e . On the other hand, the factor de is a strong artifact of the linearization method. We do not currently know how to remove this factor.

2.4 Smith Normal Form over $\mathbb{Z}/p^e\mathbb{Z}$

In this section we will focus on the problem of computing the Smith normal form for matrices over the local ring $\mathbb{Z}/p^e\mathbb{Z}$. In some of the applications, the input matrices have few non-trivial factors, i.e., most of the invariant factors are 1's or 0's. For example, in the homology computations in [Babson et al., 1999, Björner and Welker, 1999] large boundary matrices (e.g. 135135 by 270270 matrix) will have only 220 threes, and the rest of the invariant factors are ones and zeroes. The algorithm we present in this section addresses that case.

For a prime p and an exponent $e \in \mathbb{Z}_{>1}$, let φ be the natural projection $\mathbb{Z}/p^e\mathbb{Z} \rightarrow \mathbb{Z}/p\mathbb{Z}$, which extends to vectors and matrices by element-wise application. Note that $x \in \mathbb{Z}/p^e\mathbb{Z}$ is a unit if and only if $\varphi(x) \neq 0$. Likewise, $A \in (\mathbb{Z}/p^e\mathbb{Z})^{n \times n}$ is unimodular if and only if $\varphi(A)$ is unimodular.

2.4.1 Nullspace Method

Let us introduce the approach by way of a sketched example. Suppose A is a matrix over $\mathbb{Z}/p^5\mathbb{Z}$ of size 100×100 . Let \sim denote the unimodular equivalence of two matrices. Assume

$$A \sim \text{diag}(1, 1, \dots, 1, p, p, p, p^3, p^4, 0, 0, \dots, 0),$$

with 45 ones and 50 zeroes. First, a reduction in dimension allows us to reduce A to an $\ell \times \ell$ matrix $\rho(A)$ having the same nonzero invariant factors, where ℓ is the rank, or slightly larger. We illustrate with $\ell = 52$:

$$\rho(A) \sim S = \text{diag}(1, 1, \dots, 1, p, p, p, p^3, p^4, 0, 0).$$

Over $\mathbb{Z}/p\mathbb{Z}$, the nullspace basis N' of $\varphi(\rho(A))$ (N' has 7 columns) is unimodularly equivalent to the nullspace of S . Let E' be the last 7 columns of the 52×52 identity matrix. Let E and N be arbitrary embeddings of E' and N' in $(\mathbb{Z}/p^5\mathbb{Z})^{52 \times 7}$ such that $\varphi(E) = E', \varphi(N) = N'$. Then $\rho(A)N$ and SE are multiples of p and

$$\rho(A)N \sim SE = \text{diag}(p, p, p, p^3, p^4, 0, 0).$$

In summary, the algorithm is to apply a reduction in dimension to dispose of zeroes, compute nullspace basis N to dispose of ones, and determine the nontrivial invariants by computing the Smith normal form of AN using dense methods. AN is an $n \times k$ matrix, where k is the number of nontrivial invariants or a slightly larger bound.

Reduction in dimension is a frequent tool and has been used for the Smith normal form computation, for example, in [Dumas et al., 2001]. However, their computation proceeds without disposing of the unit invariant factors. Thus the time complexities below, otherwise similar to theirs, differ in that we replace a rank factor ℓ by the number of nontrivial invariants, k .

2.4.2 Probabilistic Dimension Reduction

Let $A \in (\mathbb{Z}/p^e\mathbb{Z})^{n \times n}$, for which we have a fast black-box. Let A have the Smith normal form $\text{diag}(s_1, \dots, s_r, 0, \dots, 0) \in (\mathbb{Z}/p^e\mathbb{Z})^{n \times n}$. Our goal in this section is to construct $\rho(A)$. Given A and $\ell \in \{1, \dots, n\}$, we construct a black-box of similar cost for $B \in (\mathbb{Z}/p^e\mathbb{Z})^{\ell \times \ell}$ which has the Smith normal form $\text{diag}(s_1, \dots, s_\ell)$, i.e., with the first ℓ invariant factors of A .

Recall that \mathcal{C}_k^n denote the the set of k -tuples of distinct elements (in increasing order) of $\{1, \dots, n\}$. For $\sigma, \tau \in \mathcal{C}_k^n$, $B(\sigma, \tau)$ is the (σ, τ) minor of B . We use script letters, e.g. $\mathfrak{D}, \mathfrak{T}$, to denote matrices with indeterminate entries.

We use techniques similar to that derived in [Giesbrecht, 2001] with *scaled Toeplitz matrix preconditioners*. For a set of indeterminates $\Lambda = \{v_i, w_i, y_i\}$, let $\mathfrak{D}_1 = \text{diag}(v_1, \dots, v_n)$, $\mathfrak{D}_2 = \text{diag}(w_1, \dots, w_n)$, and \mathfrak{T} be a generic Toeplitz matrix given by

$$\mathfrak{T} = \begin{bmatrix} y_n & y_{n-1} & \cdots & y_1 \\ \vdots & y_n & \ddots & \vdots \\ y_{n-2} & & \ddots & y_{n-1} \\ y_{2n-1} & y_{2n-2} & \cdots & y_n \end{bmatrix}. \quad (2.5)$$

Lemma 2.5. *Let $\mathfrak{B} = \mathfrak{D}_1 \mathfrak{T} \mathfrak{D}_2$ be as in (2.5). Let $k \in \{1, \dots, n\}$ and $\sigma = (\sigma_1, \dots, \sigma_k)$, $\tau = (\tau_1, \dots, \tau_k) \in \mathcal{C}_k^n$. Then*

(i) $\mathfrak{Z}\left(\begin{smallmatrix} \sigma \\ \tau \end{smallmatrix}\right) \in \mathbb{Z}[\Lambda]$ has content 1;

(ii) $\mathfrak{B}\left(\begin{smallmatrix} \sigma \\ \tau \end{smallmatrix}\right) = v_{\sigma_1} \cdots v_{\sigma_k} w_{\tau_1} \cdots w_{\tau_k} \mathfrak{Z}\left(\begin{smallmatrix} \sigma \\ \tau \end{smallmatrix}\right)$.

Proof. Part (i) is from [Giesbrecht, 2001, Lemma 1.3]. For part (ii), we use the Cauchy-Binet formula to get

$$\mathfrak{B}\left(\begin{smallmatrix} \sigma \\ \tau \end{smallmatrix}\right) = \mathfrak{D}_1\left(\begin{smallmatrix} \sigma \\ \sigma \end{smallmatrix}\right) \mathfrak{Z}\left(\begin{smallmatrix} \sigma \\ \tau \end{smallmatrix}\right) \mathfrak{D}_2\left(\begin{smallmatrix} \tau \\ \tau \end{smallmatrix}\right) = v_{\sigma_1} \cdots v_{\sigma_k} w_{\tau_1} \cdots w_{\tau_k} \mathfrak{Z}\left(\begin{smallmatrix} \sigma \\ \tau \end{smallmatrix}\right).$$

□

Note that $\mathfrak{B}\left(\begin{smallmatrix} \sigma \\ \tau \end{smallmatrix}\right)$ uniquely identifies which minor of \mathfrak{B} was selected. By applying two preconditioner matrices $\mathfrak{B}_1, \mathfrak{B}_2$ to A , we can get a matrix whose leading minors are related to the determinantal divisors of A .

Lemma 2.6. *Let $A \in \mathbb{Z}^{n \times n}$, and $\mathfrak{B}_1, \mathfrak{B}_2$ be $n \times n$ matrices of distinct indeterminates from a set Λ , of the form (2.5), and let $\mathfrak{A} = \mathfrak{B}_1 A \mathfrak{B}_2$. Then for all $1 \leq k \leq n$, the content of $\psi_k = \mathfrak{A}\left(\begin{smallmatrix} 1 \dots k \\ 1 \dots k \end{smallmatrix}\right) \in \mathbb{Z}[\Lambda]$ equals Δ_k , the k th determinantal divisor of A .*

Proof. By the Cauchy-Binet formula we have

$$\mathfrak{A}\left(\begin{smallmatrix} 1 \dots k \\ 1 \dots k \end{smallmatrix}\right) = \sum_{\sigma, \tau \in C_k^n} \mathfrak{B}_1\left(\begin{smallmatrix} 1 \dots k \\ \sigma \end{smallmatrix}\right) A\left(\begin{smallmatrix} \sigma \\ \tau \end{smallmatrix}\right) \mathfrak{B}_2\left(\begin{smallmatrix} \tau \\ 1 \dots k \end{smallmatrix}\right).$$

Thus $\mathfrak{A}\left(\begin{smallmatrix} 1 \dots k \\ 1 \dots k \end{smallmatrix}\right)$ is a sum of polynomials of content 1, with distinct indeterminates, one for each $k \times k$ minor of A , times the value of that minor. Hence it must have content equal to the GCD of all $k \times k$ minors of A , which is equal to the k th determinantal divisor of A . □

Working with symbolic matrices is expensive. Instead, we use randomization to get a preconditioned matrix with high probability.

Theorem 2.3. *Let $A \in \mathbb{Z}^{n \times n}$, $p \geq 6n^2\xi$ be a prime, and $\xi \geq 2$. Let $B_1, B_2 \in \mathbb{Z}^{n \times n}$ be formed by a random assignment of variables in $\mathfrak{B}_1, \mathfrak{B}_2$ in (2.5) respectively, where choices are made uniformly from $L = \{0, \dots, 6n^2\xi - 1\}$, and $\widehat{A} = B_1 A B_2$. Then with probability at least $1 - 1/\xi$, for all $1 \leq k \leq n$, the order of p in Δ_k equals the order of p in $\widehat{A}\left(\begin{smallmatrix} 1 \dots k \\ 1 \dots k \end{smallmatrix}\right)$.*

Proof. Let ψ_k be as in Lemma 2.6, which has content equal to the k th determinantal divisor Δ_k of A . The total degree of a $k \times k$ minor of a Toeplitz matrix of indeterminates is at most k . Then $\deg \psi_k \leq 6k \leq 6n$. Now substitute random values for the variables in Λ as described, and apply the Schwartz-Zippel Lemma. We get ψ_k/Δ_k is a polynomial in the entries of the matrices B_1, B_2 and

$$\Pr[(\psi_k/\Delta_k) \not\equiv 0 \pmod{p}] \geq 1 - \frac{6n}{6n^2\xi}.$$

This is the probability that the order of p in Δ_k equals the order of p in the leading $k \times k$ minor of \widehat{A} . The probability that this happens jointly for all $1 \leq k \leq n$ is at least $(1 - 1/(n\xi))^n$. The lemma statement holds because $(1 - 1/(n\xi))^n \geq 1 - 1/\xi$ using Bernoulli's inequality [Carothers, 2000]. \square

The following result states that the dimension reduction to size $\ell \times \ell$ preserves the first ℓ invariant factors.

Corollary 2.2. *Let $p \geq 6n^2\xi$ be prime, $\xi > 1$, and $e \geq 1$. Suppose $A \in (\mathbb{Z}/p^e\mathbb{Z})^{n \times n}$ has the Smith normal form $\text{diag}(s_1, \dots, s_n) \in (\mathbb{Z}/p^e\mathbb{Z})^{n \times n}$. Let $B_1, B_2 \in (\mathbb{Z}/p^e\mathbb{Z})^{n \times n}$ be formed by a random assignments of variables in $\mathfrak{B}_1, \mathfrak{B}_2$ in (2.5) respectively, where choices are made uniformly from $L = \{0, \dots, 6n^2\xi - 1\} \pmod{p^e}$. Let $\widehat{A} = B_1AB_2 \in (\mathbb{Z}/p^e\mathbb{Z})^{n \times n}$. For all $1 \leq k \leq n$ let \widehat{A}_k be the leading $k \times k$ submatrix of \widehat{A} . Then with probability at least $1 - 1/\xi$, for all $k \in \{1, \dots, n\}$, the Smith normal form of \widehat{A}_k is $\text{diag}(s_1, \dots, s_k) \in (\mathbb{Z}/p^e\mathbb{Z})^{k \times k}$.*

Proof. The Smith normal form of A equals the Smith normal form of any $\widetilde{A} \in \mathbb{Z}^{n \times n}$ with $\widetilde{A} \equiv A \pmod{p^e}$, reduced modulo p^e . Thus, Theorem 2.3 implies that the order of p in the k th determinantal divisor of A equals the order of p in the leading $k \times k$ minor of \widetilde{A} , for all k , with probability at least $1 - 1/\xi$. This implies that $\widehat{A}_k = \widetilde{A}_k \pmod{p^e}$ will have the Smith normal form (s_1, \dots, s_k) for all $1 \leq k \leq n$ where \widetilde{A}_k is the leading $k \times k$ minor of \widetilde{A} , since $\Delta_k = s_1 \cdots s_k$ for $1 \leq k \leq n$. \square

Computationally, if we know that m is an upper bound on that rank of A , then we can work with truncated random scaled Toeplitz matrices $B_1 \in (\mathbb{Z}/p^e\mathbb{Z})^{m \times n}$ and $B_2 \in (\mathbb{Z}/p^e\mathbb{Z})^{n \times m}$. Then Corollary 2.2 implies that $\widehat{A} = B_1AB_2 \in (\mathbb{Z}/p^e\mathbb{Z})^{m \times m}$ has the same

non-zero invariant factors as A . This upper bound can be efficiently obtained by computing the rank of A modulo a small set of randomly chosen primes.

2.4.3 Working with Small Primes

Corollary 2.2 requires that $p \geq 6n^2\xi$. For smaller primes the algorithm may well work, but this appears much more difficult to prove. The following method can be used to remedy this.

We construct the *Galois ring* extension $\text{GR}(p^e, d) = \mathbb{Z}[x]/(p^e, f)$, where $d = \lceil \log_p(6n^2\xi) \rceil$ and $f \in \mathbb{Z}[x]$ is a polynomial of degree at least d , and the image of f is irreducible over $\mathbb{Z}/p\mathbb{Z}$. The ring $\text{GR}(p^e, d)$ is a principal ideal ring and all its non-trivial ideals are generated by powers of p [McDonald, 1974, §XVI]. Smith normal form is defined for matrices over $\text{GR}(p^e, d)$ [McDonald, 1974, exercise XVI.1, XVI.2]. Analogues of Theorem 2.3 and Corollary 2.2 over $\text{GR}(p^e, d)$ can be proven similarly. When choosing random elements from $\text{GR}(p^e, d)$ we are choosing random polynomials of degree less than d . Since $\text{GR}(p^e, d)$ contains the finite field $\text{GF}(p^d)$, the Schwartz-Zippel Lemma can be applied.

Corollary 2.3. *Let p be prime, $e \geq 1$, $\xi \geq 1$ and $d = \lceil \log_p(6n^2\xi) \rceil$. Let $\text{GR}(p^e, d) = \mathbb{Z}[x]/(p^e, f)$ for $f \in \mathbb{Z}[x]$ of degree d which is irreducible modulo p . Suppose $A \in (\mathbb{Z}/p^e\mathbb{Z})^{n \times n}$ has the Smith normal form $\text{diag}(s_1, \dots, s_n) \in (\mathbb{Z}/p^e\mathbb{Z})^{n \times n}$. Let $B_1, B_2 \in \text{GR}(p, d)^{n \times n}$ be formed by random assignments of the indeterminates in $\mathfrak{B}_1, \mathfrak{B}_2$ in (2.5) respectively, where the random choices are made uniformly from $L = \{\sum_{0 \leq i < d} \alpha_i x^i : \alpha_i \in [0, p)\} \bmod p^e$. Let $\hat{A} = B_1 A B_2 \in \text{GR}(p^e, d)^{n \times n}$, and for $1 \leq k \leq n$ let \hat{A}_k be the leading $k \times k$ submatrix of \hat{A} . Then with probability at least $1 - 1/\xi$, for all $k \in \{1, \dots, n\}$, the Smith normal form of \hat{A}_k is $\text{diag}(s_1, \dots, s_k) \in (\mathbb{Z}/p^e\mathbb{Z})^{k \times k}$.*

2.4.4 The Algorithm

After reducing the dimension to a value at or near the number of nonzero invariant factors, the following algorithm is applied. We will discuss the value of ℓ below.

Algorithm 2.2. Given a black-box for $B \in (\mathbb{Z}/p^e\mathbb{Z})^{n \times n}$, and a bound ℓ for the number of nonzero invariant factors, compute the invariant factors of B .

1. Construct A as the $\ell \times \ell$ dimension reduction of B .
2. Let $r_0 = \text{rank}(\varphi(A))$ over $\mathbb{Z}/p\mathbb{Z}$. Let $k = \ell - r_0$.
3. Compute $N' \in (\mathbb{Z}/p^e\mathbb{Z})^{\ell \times k}$, a lifting to $\mathbb{Z}/p^e\mathbb{Z}$ of a right nullspace basis of $\varphi(A)$ over $\mathbb{Z}/p\mathbb{Z}$.
4. Let $N = AN' \in (\mathbb{Z}/p^e\mathbb{Z})^{\ell \times k}$. This involves k matrix vector products with A .
5. Compute the Smith normal form of N over $\mathbb{Z}/p^e\mathbb{Z}$ by dense methods:

$$\text{diag}(\underbrace{p, \dots, p}_{r_1}, \underbrace{p^2, \dots, p^2}_{r_2}, \dots, \underbrace{p^{e-1}, \dots, p^{e-1}}_{r_{e-1}}, \underbrace{0, \dots, 0}_{r_e}).$$

6. Return r_0, \dots, r_{e-1} .

We will analyze the algorithm holding e and p constant. Considering them as parameters would introduce a factor of $O(e \log(p))$ in the complexity. Let the cost of matrix-vector product by B be μ . Since we are holding e and p constant, this is the same cost for application to vectors in $(\mathbb{Z}/p\mathbb{Z})^n$ and in $(\mathbb{Z}/p^e\mathbb{Z})^n$.

Theorem 2.4. *Algorithm 2.2 is a correct Monte Carlo algorithm. The time complexity is $O(\ell k(k^{\omega-2} + \mu))$ operations in $\mathbb{Z}/p^e\mathbb{Z}$, where k is the number of non-trivial (neither 0 nor 1) invariant factors, and ℓ is the reduced dimension. The memory requirement is $O(k\ell)$.*

Proof. Step 1: Toeplitz matrices may be applied to vectors via polynomial multiplication, so the cost of the black-box for A is $O(M(n) + \mu)$. But $M(n)$ is $O(n)$ and $\mu \geq n$. Thus the black-box cost of A is $O(\mu)$.

Step 2: The rank over $\mathbb{Z}/p\mathbb{Z}$ can be computed by a black-box method in $O((\ell\mu) \log(\xi))$ to achieve probability of error less than $1/\xi$ [Wiedemann, 1986]. Memory requirement is $O(1)$ vectors in $(\mathbb{Z}/p\mathbb{Z})^\ell$.

Step 3: Let $k = \ell - r_0$ denote the nullity of A modulo p . By black-box methods, k random samples of the nullspace will yield a nullspace basis N' . Oversampling can be done and column echelon form computation is used to reduce to a basis of k columns if need be. The cost is $O(k(\ell\mu))$ operations in $\mathbb{Z}/p\mathbb{Z}$, and $O(k\ell)$ space [Chen et al., 2002].

Step 4: The cost of applying A to N' is $O(k\mu)$.

Step 5: Any nullspace basis for S over $\mathbb{Z}/p\mathbb{Z}$ is of the form EW' , where E is the last $\ell - r_0$ columns of the identity matrix, and W' is a $k \times k$ unimodular matrix. Then $0 = AN' = USVN = USEW'$ over $\mathbb{Z}/p\mathbb{Z}$, for some unimodular W' . This lifts to a factorization $AN = USEW$ over $\mathbb{Z}/p^e\mathbb{Z}$ with U, W being unimodular. Thus AN has Smith normal form SE . The Smith normal form of AN can be computed by elimination using $O(\ell^2 k^{\omega-2})$ operations in $\mathbb{Z}/p^e\mathbb{Z}$ [Storjohann, 2000]. \square

Since $\omega \leq 3$, $k \leq n$, and $\mu \geq n$, we can assume that $\mu > k^{\omega-2}$. The complexity is then dominated by $O(\ell k \mu)$. This algorithm is useful when there are only few non-trivial invariant factors, i.e., when $\ell \ll n$.

The value of ℓ can be inferred by running the algorithm multiple times with $\ell = 2, 4, 8, \dots$, and stopping when the resulting Smith normal form remains unchanged. There will be at most $\lceil \log r \rceil \leq \lceil \log n \rceil$ steps. The cost of the algorithm will increase by a $\log n$ factor, which does not change the statement of the result.

2.5 Detecting Non-Trivial Smith Normal Form

In this section we present an algorithm to detect whether a black-box matrix has a non-trivial Smith normal form. Let A be a matrix over $\mathbb{Z}/p^e\mathbb{Z}$ whose Smith normal form is $\text{diag}(s_1, \dots, s_n)$.

Definition 2.1. *The first non-trivial invariant factor of A is the non-zero s_i with the smallest index $i \in [1, n]$ such that $p \mid s_i$.*

The p -adic lifting technique of [Dixon, 1982] is widely used to compute solutions of linear systems. It works by adding a p -adic digit to the solution after every iteration. We

show that lifting can be used to detect the first non-trivial invariant factor when used to solve $Ax = 0$ modulo p^e .

Algorithm 2.3. Given an $n \times n$ matrix A over $\mathbb{Z}/p^e\mathbb{Z}$, return a vector $v \in (\mathbb{Z}/p^e\mathbb{Z})^n$ in the nullspace of A or “FAIL at iteration i ”.

1. Let $A_p = A \bmod p$.
2. Let $r = 0 \in (\mathbb{Z}/p^e\mathbb{Z})^n$.
3. For $i = 0$ to $e - 1$ do:
 - (a) If $r \neq 0$: solve $A_p v_i = r$ over $\mathbb{Z}/p\mathbb{Z}$.
 - (b) If $r = 0$: set v_i to a random non-zero vector in the nullspace of A_p over $\mathbb{Z}/p\mathbb{Z}$.
 - (c) If $A(v_0 + pv_1 + \dots + p^i v_i) \neq 0 \bmod p^{i+1}$ then return “FAIL at iteration i ”.
 - (d) Let $r := (r - Av_i)/p$.
4. Return $v_0 + pv_1 + \dots + p^{e-1}v_{e-1}$.

Sampling the nullspace of A_p over $\mathbb{Z}/p\mathbb{Z}$ can be done using black-box methods at the same cost of solving a linear system over $\mathbb{Z}/p\mathbb{Z}$, for example, using Algorithm *random-LINSOLVE0* of [Chen et al., 2002].

We show that if A has a non-trivial invariant factor then Algorithm 2.3 will fail to find any nullspace vector. In particular, if the first non-trivial invariant factor is p^k then step 3 will return FAIL at iteration k . We first start with a simple case of A being in the Smith normal form.

Lemma 2.7. Let A be an $n \times n$ matrix over $\mathbb{Z}/p^e\mathbb{Z}$ given by

$$A = \begin{bmatrix} I_{\ell \times \ell} & & \\ & p^k I_{m \times m} & \\ & & p^{k+\tau} I_{t \times t} \end{bmatrix},$$

where $\tau > 0$ and $t \geq 0$. If we invoke Algorithm 2.3 on A , then it will fail at iteration $i = k$ with probability at least $1 - (1/p)^m$.

Proof. We have

$$A = \begin{bmatrix} I_{\ell \times \ell} & & \\ & 0_{m \times m} & \\ & & 0_{t \times t} \end{bmatrix} \pmod{p^k}, \text{ and } A = \begin{bmatrix} I_{\ell \times \ell} & & \\ & p^k I_{m \times m} & \\ & & 0_{t \times t} \end{bmatrix} \pmod{p^{k+1}}.$$

Let w be the resulting vector after iteration $i = k - 1$ such that $Aw = 0 \pmod{p^k}$. Then we must have

$$w = (0, \dots, 0, a_1, \dots, a_m, *, \dots, *),$$

where a_1, \dots, a_m can be any values from $[0, p^k)$ and the entries denoted by $*$ are arbitrary values which we do not need to specify for the sake of this argument. Now let v be the resulting vector from lifting w after iteration $i = k$ such that $Av = 0 \pmod{p^{k+1}}$. Since v only modifies w by adding a p^k digit, we have

$$v = (0, \dots, 0, a_1 + p^k b_1, \dots, a_m + p^k b_m, *, \dots, *).$$

From $Av = 0 \pmod{p^{k+1}}$ we have $p^k(a_j + p^k b_j) = 0 \pmod{p^{k+1}}$ for all $j \in [1, m]$. So we must have $a_j = 0 \pmod{p}$ for all $j \in [1, m]$. Now suppose that in computing w , we have $a_j \neq 0 \pmod{p}$ for any $j \in [1, m]$, then there exists no v such that $Av = 0 \pmod{p^{k+1}}$. In this case, step 3(c) will report failure at iteration k . What is the probability that any $a_j \neq 0 \pmod{p}$?

The values of $a_j \pmod{p}$ digits are computed at iteration $i = 0$ for which the residue vector r is zero. Step 3(b) will perform random sampling from the nullspace of A over $\mathbb{Z}/p\mathbb{Z}$. If the nullspace sampling is uniform, then with probability at least $1 - (1/p)^m$, one of the entries a_j will have $a_j \neq 0 \pmod{p}$. Step 3(c) will then fail with the stated probability. \square

The probability is the same when A is not in the Smith normal form.

Lemma 2.8. *Let A be an $n \times n$ matrix over $\mathbb{Z}/p^e\mathbb{Z}$ which has the Smith normal form*

$$S = \begin{bmatrix} I_{\ell \times \ell} & & \\ & p^k I_{m \times m} & \\ & & p^{k+\tau} I_{t \times t} \end{bmatrix},$$

where $\tau > 0$ and $t \geq 0$. If we invoke Algorithm 2.3 on A , then it will fail at iteration $i = k$ with probability at least $1 - (1/p)^m$.

Proof. Let $A = USV$ where U, V are unimodular. In step $i = 0$ of the algorithm, we select a vector $x \pmod{p}$ uniformly at random from $\ker(A)$. This is equivalent to selecting a random vector v uniformly at random from $\ker(S)$ where $v = Vx$. We can apply the same argument of Lemma 2.7 that a “bad” selection of $v \pmod{p}$ (i.e., a selection that will cause the algorithm to report failure at Step k) will occur with probability at least $1 - 1/p^m$, which is the same probability for the random selection of x because V is a bijection between the two kernels as given by $\ker(S) = V \ker(A)$. \square

Algorithm 2.4. Given a black-box for $A \in (\mathbb{Z}/p^e\mathbb{Z})^{n \times n}$, return the first non-trivial invariant factor, or TRIVIAL if the Smith normal form of A is trivial.

1. Run Algorithm 2.3 on A as input.
2. If the algorithm returned FAIL at iteration k , then return p^k .
3. If the algorithm successfully terminated after e iterations, then return TRIVIAL.

Theorem 2.5. Algorithm 2.4 is a correct Monte Carlo algorithm with success probability at least $1 - 1/p^{r_k}$, where r_k is the multiplicity of the invariant factor p^k . It requires $O(en\mu)$ operations in $\mathbb{Z}/p^e\mathbb{Z}$ and a space of $O(n)$ elements in $\mathbb{Z}/p^e\mathbb{Z}$, where μ is the cost of the black-box for A over $\mathbb{Z}/p^e\mathbb{Z}$.

Proof. Correctness follows from the lemmas above. The cost is dominated by step 3 of Algorithm 2.3. There are at most e iterations of this step. Each iteration costs $O(n\mu)$ operations and $O(n)$ space using Wiedemann-based methods over $\mathbb{Z}/p\mathbb{Z}$. \square

The success probability is at least $1/2$ for all $p \geq 2$ and $r_k \geq 1$. The smallest value occurs when $p = 2$ (recall that r_k is unknown). In this case we can repeat the algorithm and apply majority voting to get a controllably high success probability.

Chapter 3

Preconditioning

The success of iterative algorithms for sparse linear algebra is usually tied to preconditioner matrices. These matrices are typically structured matrices with random entries from a specified distribution. Special attention is paid to the additional cost introduced by multiplying preconditioner matrices with vectors. In this chapter, we extend the application of faster known preconditioners to the problems of nullspace sampling, and computing Smith normal form.

3.1 Introduction

Wiedemann-based methods often apply structured pre- and post-multiplier matrices to the input matrix, so that the minimal polynomial of the resulting “preconditioned” matrix encodes useful information about the original matrix such as rank, determinant, etc. The preconditioner matrices, however, increase the overall cost of the algorithm. It is desirable to have preconditioners which are sparse and can be applied to vectors in a linear number of operations. Several preconditioners have been proposed with cost varying from linear to quasilinear. Preconditioners are often structured or diagonal matrices with random entries, and their success is probabilistic and relative to the field size. Over small fields, few preconditioners work directly without the need to construct field extensions, which adds a logarithmic factor to the complexity. Block Wiedemann’s algorithm can work well

over small fields without the use of preconditioning given that the blocking size is greater than the number of blocks in the Frobenius normal form of the input matrix [Villard, 1997]. Preconditioning can ensure having a small number of Frobenius blocks [Eberly, 2004]. A comprehensive review of preconditioning is presented in [Chen et al., 2002].

In this chapter we address the application of faster (though already known) preconditioners to the problem of sampling from the nullspace over large and small fields, and the problem of computing Smith normal form of integer matrices. The fastest preconditioners for nullspace and Smith normal form are based on random Toeplitz matrices [Kaltofen and Saunders, 1991], [Giesbrecht, 1995] which increase the cost of matrix-vector products by a factor of $O(n \log(n) \log \log(n))$, or $O(n)$ polynomial multiplications. We show how to replace Toeplitz matrices with scaled-transpose preconditioners of [Eberly and Kaltofen, 1997] in more circumstances, and hence reduce the overhead to linear time.

In §3.2, we present a new Monte Carlo algorithm for sampling uniformly from the nullspace of a black-box matrix, even if the field size is small. The algorithm can be applied with two different random preconditioners: the “transpose box” with diagonal scaling, and the so-called *Wiedemann-sparse* matrices introduced by [Wiedemann, 1986] and further studied by [Chen et al., 2002]. Our nullspace sampling algorithm with Wiedemann’s preconditioner is similar to the algorithm of [Eberly, 2004], and offers the same asymptotic cost. However, we rely on a simple computation of the minimal polynomial rather than computing a Frobenius decomposition.

In §3.3 we show that the scaled-transpose preconditioner $A \mapsto D_1 A^T D_2 A$ can replace Toeplitz preconditioners in existing Smith normal form algorithms [Giesbrecht, 1995, Giesbrecht, 1996, Giesbrecht, 2001]. This approach offers an alternative in practice since Toeplitz preconditioners typically involve (FFT-based) polynomial multiplications, while scaled-transpose preconditioners require scalar multiplications only.

3.2 Nullspace Sampling

Let A be an $n \times n$ singular matrix of rank r over a field F . Let $P, Q \in F^{n \times n}$ be preconditioners such that with high probability the preconditioned matrix $\tilde{A} = PAQ$ has a minimal polynomial $xg(x)$ where $g(0) \neq 0$. The following lemma shows that sampling uniformly at random from the nullspace of A can be reduced to computing $Qg(\tilde{A})w$ for a randomly chosen $w \in F^n$.

Lemma 3.1. *Let $A \in F^{n \times n}$ have rank $= r < n$. Let $P, Q \in F^{n \times n}$ be such that Q is invertible, $\ker(A) = Q \ker(PAQ)$, $\text{minpoly}(PAQ) = xg(x)$, and $g(0) \neq 0$. Then $\ker(A) = Q \text{im}(g(PAQ))$.*

Proof. Let $\tilde{A} = PAQ$. Then $\tilde{A}g(\tilde{A}) = 0$ and $g(\tilde{A}) \neq 0$ by minimality. So $\text{im}(g(\tilde{A})) \subseteq \ker(\tilde{A})$. For the inclusion \supseteq , let $v \in \ker(\tilde{A})$. Then $\tilde{A}^i v = 0$ for all $i \geq 1$. This implies that $g(\tilde{A})(v/g(0)) = v$, i.e., $v \in \text{im}(g(\tilde{A}))$. Thus $\ker(\tilde{A}) = \text{im}(g(\tilde{A}))$. The claim follows since $\ker(A) = Q \ker(\tilde{A})$. \square

We will state the algorithm, and discuss two random preconditioners which satisfy the conditions of Lemma 3.1. The cost is discussed after we make concrete choices for the preconditioners.

Algorithm 3.1. [Nullspace Sampling] Given a matrix $A \in F^{n \times n}$ over a field F , and an error bound $0 < \epsilon < 1$, the output is a random vector $v \in F^n$ sampled uniformly from $\ker(A)$ with probability of correctness at least $1 - \epsilon$.

1. Choose a preconditioner $P, Q \in F^{n \times n}$ such that Lemma 3.1 holds with probability at least $1 - \epsilon$. Let $\tilde{A} = PAQ$.
2. Compute $f(x) = \text{minpoly}(\tilde{A})$. Let $g(x) = f(x)/x$.
3. Choose a vector w uniformly at random from F^n .
4. Return $Qg(\tilde{A})w$.

Example 3.1. Let $A = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}$ whose minimal polynomial is x^2 . If we choose $P = A^T$ and $Q = I_{2 \times 2}$, then

$$\tilde{A} = PAQ = A^T A = \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix},$$

and $\text{minpoly}(\tilde{A}) = x(x-1) = xg(x)$. Let $w = \begin{bmatrix} w_1 \\ w_2 \end{bmatrix}$, and compute

$$v = Qg(\tilde{A})w = (\tilde{A} - I)w = \begin{bmatrix} -1 & 0 \\ 0 & 0 \end{bmatrix} w = \begin{bmatrix} -w_1 \\ 0 \end{bmatrix}.$$

Indeed $Av = 0$ for any choice of w_1 . □

3.2.1 Transpose Preconditioners

To get our first concrete instance of Algorithm 3.1, we use a transpose-based preconditioner. It has a linear cost for matrix-vector product but its success probability requires a “large” field size. When the field is small we will work over a field extension \mathbb{K} , and discuss how to uniformly project the nullspace vectors from \mathbb{K} to \mathbb{F} .

In what follows, let S be a finite subset of $\mathbb{F} \setminus \{0\}$. Let $\phi_1, \dots, \phi_n, \psi_1, \dots, \psi_n$ be a set of independent indeterminates. Let $\mathfrak{D}_1 = \text{diag}(\phi_1, \dots, \phi_n)$, $\mathfrak{D}_2 = \text{diag}(\psi_1, \dots, \psi_n)$, and $\mathfrak{A} = \mathfrak{D}_1 A^T \mathfrak{D}_2 A$. Let d_1, \dots, d_{2n} be random assignments of $\phi_1, \dots, \phi_n, \psi_1, \dots, \psi_n$, respectively, where d_j 's are chosen independently and uniformly at random from S . Let $D_1 = \text{diag}(d_1, \dots, d_n)$, $D_2 = \text{diag}(d_{n+1}, \dots, d_{2n})$ and $\tilde{A} = D_1 A^T D_2 A$. In other words, \tilde{A} , D_1 , D_2 are random evaluations of \mathfrak{A} , \mathfrak{D}_1 , and \mathfrak{D}_2 , respectively.

Lemma 3.2. *With probability at least $1 - 2n/|S|$ we have $\ker(\tilde{A}) = \ker(A)$.*

Proof. Clearly $\text{rank}(\tilde{A}) \leq \text{rank}(A)$. To prove $\text{rank}(\tilde{A}) = \text{rank}(A)$, it suffices to show the existence of an $r \times r$ non-zero minor of \tilde{A} . From $\text{rank}(A) = r$ we know that the largest non-zero minor of A is of size $r \times r$, and that at least one such minor exists. Call this

minor $A\binom{\tau'}{\sigma}$ where $\tau', \sigma \in \mathcal{C}_r^n$. Then

$$\mathfrak{A}\binom{\sigma}{\sigma} = \sum_{\tau \in \mathcal{C}_r^n} \mathfrak{D}_1\binom{\sigma}{\sigma} A^T\binom{\sigma}{\tau} \mathfrak{D}_2\binom{\tau}{\tau} A\binom{\tau}{\sigma} \quad (3.1)$$

$$= \sum_{\tau \in \mathcal{C}_r^n} A\binom{\tau}{\sigma}^2 \phi_{\sigma_1} \cdots \phi_{\sigma_r} \psi_{\tau_1} \cdots \psi_{\tau_r}, \quad (3.2)$$

which is a multivariate polynomial in $\mathbb{F}[\phi_1, \dots, \phi_n, \psi_1, \dots, \psi_n]$ of total degree at most $2r \leq 2n$. This polynomial is non-zero because $A\binom{\tau'}{\sigma} \neq 0$, and the monomials $\phi_{\sigma_1} \cdots \phi_{\sigma_r} \psi_{\tau_1} \cdots \psi_{\tau_r}$ are distinct for different σ, τ pairs, so no cancellations will happen among the summands. If we evaluate $\mathfrak{A}\binom{\sigma}{\sigma}$ at $\{d_1, \dots, d_{2n}\}$, we get that $\tilde{A}\binom{\sigma}{\sigma}$ is non-zero with probability at least $1 - 2n/|S|$ using the Schwartz-Zippel lemma. So \tilde{A} has a non-zero $r \times r$ minor, and hence has rank r , with the stated probability.

The inclusion $\ker(A) \subseteq \ker(\tilde{A})$ is straightforward. By rank arguments above and the rank-nullity theorem, we get $\text{nullity}(A) = \text{nullity}(\tilde{A})$ with the stated probability. But $\ker(A)$ is a subspace of $\ker(\tilde{A})$ and $\dim(\ker(A)) = \dim(\ker(\tilde{A}))$ so we must have that $\ker(A) = \ker(\tilde{A})$. \square

Lemma 3.3 (Theorem 4.5 of [Chen et al., 2002]). *With probability at least $1 - n/|S|$, $\text{minpoly}(\tilde{A}) = xg(x)$ and $g(0) \neq 0$.*

Thus the preconditioner $A \mapsto D_1 A^T D_2 A$ satisfies the conditions of Lemma 3.1. We can apply it to Algorithm 3.1.

Theorem 3.1. *Let A be an $n \times n$ singular matrix over \mathbb{F} . Let μ be the number of field operations required to multiply A by a vector. Then a sample from the nullspace of A can be computed using Algorithm 3.1. The output of the algorithm is correct with probability at least $1 - \epsilon$. The cost is $O(r\mu + r^2)$ operations in \mathbb{F} , and storage for $O(n)$ elements in \mathbb{F} .*

Proof. First we construct the preconditioner. Choose a finite set $S \subseteq \mathbb{F} \setminus \{0\}$ of size $|S| > 3n/\epsilon$. Choose d_1, \dots, d_{2n} independently and uniformly at random from S . Let $D_1 = \text{diag}(d_1, \dots, d_n)$, $D_2 = \text{diag}(d_{n+1}, \dots, d_{2n})$. Apply Algorithm 3.1 and choose $P = D_1 A^T D_2$, and $Q = I$ in Step 1 of the algorithm.

Correctness follows from Lemmas 3.1, 3.2, and 3.3 which hold with joint probability at least $(1 - n/|S|)(1 - 2n/|S|) \geq 1 - 3n/|S| = 1 - \epsilon$.

The cost of the matrix-vector multiplications $v \mapsto D_1v$, $v \mapsto D_2v$ is n operations in \mathbb{F} . The cost of applying $\tilde{A} = D_1A^TD_2A$ to a vector is $2\mu + 4n$. So the cost of computing $\text{minpoly}(\tilde{A})$ is $4r\mu + O(r^2)$ using Wiedemann’s algorithm, under the assumption that $r\mu$ dominates rn . The cost of computing $g(\tilde{A})w$ is $2r\mu + O(rn)$ using Horner’s rule [von zur Gathen and Gerhard, 2003, Geddes et al., 1992].

The storage of D_1, D_2 is $2n$ elements. Wiedemann’s algorithm, and Horner’s evaluation of $g(\tilde{A})w$, require $O(n)$ space. □

Working Over Small Fields

When $|\mathbb{F}| < 3n/\epsilon$, the success probability of Algorithm 3.1 diminishes. We call this case a “small field”. We can remedy this by working over an algebraic extension field \mathbb{K}/\mathbb{F} with at least $3n/\epsilon$ elements. All proofs in this section extend to a larger field. As before, let $M(d)$ denote the number of field operations required to multiply two polynomials of degree at most d . To build \mathbb{K} , set $[\mathbb{K} : \mathbb{F}] = e$ such that $\text{char}(\mathbb{F}) \nmid e$ and $|\mathbb{F}|^e > 3n/\epsilon$. Thus $e \in O(\log(n/\epsilon))$ holding $|\mathbb{F}|$ fixed. The algorithm of [Shoup, 1994] can be used to find an irreducible polynomial in $\mathbb{F}[x]$ of degree e using $O((e \log e + \log |\mathbb{F}|) M(e))$ operations in \mathbb{F} . This cost is dominated by other steps of Algorithm 3.1. Working over \mathbb{K} , the cost of nullspace sampling increases by a factor of $M(e)$, and the storage increases by a factor of e . Since A has entries in \mathbb{F} , the cost of applying A or A^T to a vector in \mathbb{K}^n increases by a factor of e rather than $M(e)$. We have proven the following result.

Lemma 3.4. *Working over an extension \mathbb{K}/\mathbb{F} as described above, Algorithm 3.1 is correct. The computed vector is in the nullspace of A with probability at least $1 - \epsilon$. The algorithm requires $O(r\mu \log(n/\epsilon) + r^2 M(\log(n/\epsilon)))$ operations in \mathbb{F} , and storage of $O(n \log(n/\epsilon))$ elements in \mathbb{F} .*

The computed nullspace vector in Lemma 3.4 might have components from $\mathbb{K} \setminus \mathbb{F}$. We show how to project nullspace vectors from \mathbb{K}^n to \mathbb{F}^n while preserving the uniform sampling.

Recall $\text{Tr} : \mathbb{K} \rightarrow \mathbb{F}$ is the field trace defined by $\text{Tr}(\alpha) = \sum_{i=0}^{e-1} \alpha^{q^i}$ where $e = [\mathbb{K} : \mathbb{F}]$ and $q = |\mathbb{F}|$. $\text{Tr}(\cdot)$ is a surjective \mathbb{F} -linear map [Lidl and Niederreiter, 1986]. If $\alpha \in \mathbb{F}$ then $\text{Tr}(\alpha) = e\alpha$. If $\text{char}(\mathbb{F}) \nmid e$, then $\text{Tr}(\alpha) \neq 0$ whenever $\alpha \neq 0$.

The map $\text{Tr}(\cdot)$ can be extended to vectors and matrices by component-wise application. By the normal basis theorem, there exists an element $\theta \in \mathbb{K}$ such that $\{\theta, \theta^q, \dots, \theta^{q^{e-1}}\}$ is a normal basis for \mathbb{K} as an \mathbb{F} -vector space. By the linear independence of the basis, we have $\text{Tr}(\theta^{q^j}) \neq 0$ for all $0 \leq j \leq e-1$.

Let $A \in \mathbb{F}^{n \times n}$, $v \in \mathbb{K}^n$ such that $Av = 0$. Then $\text{Tr}(Av) = \text{Tr}(A) \text{Tr}(v)$ by linearity. Also $\text{Tr}(A) = eA$ and $\text{Tr}(0) = 0$. So $A \text{Tr}(v) = 0$. Thus we can project a given nullspace vector $v \in \mathbb{K}^n$ to a nullspace vector $\text{Tr}(v) \in \mathbb{F}^n$. For a uniformly sampled nullspace vector $v \in \mathbb{K}^n$, we have to show that $\text{Tr}(v)$ is a uniform sample of the nullspace of A over \mathbb{F} .

Lemma 3.5. *Let \mathbb{K} be an algebraic extension of \mathbb{F} of degree e . If α is chosen uniformly at random from \mathbb{K} , then for all $\sigma \in \mathbb{F}$ we have $\Pr[\text{Tr}(\alpha) = \sigma] = 1/|\mathbb{F}|$.*

Proof. Let $\theta \in \mathbb{K}$ be a normal basis generator of \mathbb{K} . Then $\alpha = \sum_{i=0}^{e-1} \alpha_i \theta^{q^i}$ where α_i 's are chosen independently and uniformly at random from \mathbb{F} . By linearity of Tr , we have $\text{Tr}(\alpha) = \sum_{i=0}^{e-1} \alpha_i \text{Tr}(\theta^{q^i}) = \text{Tr}(\theta) \sum_{i=0}^{e-1} \alpha_i$. Now for any $\sigma \in \mathbb{F}$, we have $\Pr[\text{Tr}(\alpha) = \sigma] = \Pr[\text{Tr}(\theta) \sum_{i=0}^{e-1} \alpha_i = \sigma]$. Split the sum into disjoint events:

$$\sum_{\tau \in \mathbb{F}} \Pr[\text{Tr}(\theta) \sum_{i=1}^{e-1} \alpha_i = \tau] \Pr[\text{Tr}(\theta) \alpha_0 = \sigma - \tau].$$

But $\Pr[\text{Tr}(\theta) \alpha_0 = \sigma - \tau] = 1/|\mathbb{F}|$ because α_0 is chosen uniformly and independently at random. The sum becomes

$$|\mathbb{F}|^{-1} \sum_{\tau \in \mathbb{F}} \Pr[\text{Tr}(\theta) \sum_{i=1}^{e-1} \alpha_i = \tau] = 1/|\mathbb{F}|,$$

which is the uniform distribution on \mathbb{F} . □

Lemma 3.6. *Let $A \in \mathbb{F}^{n \times n}$ have rank r and nullity $\nu > 0$. Let \mathbb{K} be an algebraic extension of \mathbb{F} of degree e . Define the two subspaces $V_{\mathbb{F}} = \{u \in \mathbb{F}^n : Au = 0\}$ and $V_{\mathbb{K}} = \{v \in \mathbb{K}^n : Av = 0\}$; i.e., the kernel of A over \mathbb{F} and \mathbb{K} respectively. If $v \in \mathbb{K}^n$ is chosen uniformly at random from $V_{\mathbb{K}}$, then $\Pr[\text{Tr}(v) = u] = 1/|V_{\mathbb{F}}|$ for all $u \in V_{\mathbb{F}}$.*

Proof. We have $\dim V_{\mathbb{F}} = \dim V_{\mathbb{K}} = \nu$, but $|V_{\mathbb{F}}| < |V_{\mathbb{K}}|$. Trace is surjective, so it suffices to prove that Tr maps any uniformly sampled ν -dimensional vector $v \in V_{\mathbb{K}}$ into any vector $u \in V_{\mathbb{F}}$ with probability $1/|V_{\mathbb{F}}| = 1/|\mathbb{F}|^{\nu}$. Let $\theta \in \mathbb{K}$ be a normal basis generator of \mathbb{K} over \mathbb{F} . Let $\{b_1, \dots, b_{\nu}\}$ be any basis for $V_{\mathbb{F}}$. Thus it is also a basis for $V_{\mathbb{K}}$. Fix the following two bases: $\{\text{Tr}(\theta)b_1, \dots, \text{Tr}(\theta)b_{\nu}\}$ for $V_{\mathbb{F}}$, and $\{\theta b_1, \dots, \theta b_{\nu}\}$ for $V_{\mathbb{K}}$. In this setting, $\text{Tr} : V_{\mathbb{K}} \rightarrow V_{\mathbb{F}}$ maps each basis element θb_i of $V_{\mathbb{K}}$ to its corresponding unique basis elements $\text{Tr}(\theta)b_i$ of $V_{\mathbb{F}}$.

Let $v = (v_1, \dots, v_{\nu}) \in V_{\mathbb{K}}$ where v_i 's are chosen uniformly and independently at random from \mathbb{K} . For any $u = (u_1, \dots, u_{\nu}) \in V_{\mathbb{F}}$, we have $\Pr[\text{Tr}(v) = u] = \prod_{i=1}^{\nu} \Pr[\text{Tr}(v_i) = u_i]$. By Lemma 3.5 we have, $\Pr[\text{Tr}(v_i) = v_i] = 1/|\mathbb{F}|$. So $\Pr[\text{Tr}(v) = u] = 1/|\mathbb{F}|^{\nu}$ which is the probability for the uniform distribution on $V_{\mathbb{F}}$. \square

How much does the trace computation cost? Given $\alpha \in \mathbb{K}$, $\text{Tr}(\alpha)$ is the coefficient of x^{e-1} of the minimal polynomial of α in $\mathbb{F}[x]$. It can be computed in $O(e^2 + M(e)\sqrt{e})$ operations in \mathbb{F} [Shoup, 1999]. The cost of computing $\text{Tr}(v)$ for the n components of $v \in \mathbb{K}^n$ is dominated by other steps of Algorithm 3.1.

3.2.2 Avoiding Extension Fields

The second preconditioner we use in Algorithm 3.1 is a sparse binary matrix which was first introduced by [Wiedemann, 1986], and further studied by [Chen et al., 2002, Eberly, 2004]. It has a quasilinear cost for matrix-vector products. But its advantage is that the success probability is high even over small fields. Thus we can avoid constructing field extensions and save a $\log(n)$ factor in the cost. However, this factor is compensated by $\log^2(n)$ factor in applying the preconditioner to vectors. This tradeoff can be exploited in practice, by choosing between this preconditioner and the diagonal scaling.

Let \mathbb{F} be a finite field of size q . We construct two $n \times n$ matrices L, R with 0-1 entries sampled independently from the following distribution:

$$\Pr[L_{i,j} = 1] = \min\left(1 - \frac{1}{q}, \frac{\log(n)}{j}\right),$$

where $L_{i,j}$ is the entry i, j of L . The matrix R is constructed similarly. This construction gives us a sparse matrix with good preconditioning properties.

Lemma 3.7 ([Wiedemann, 1986] and §7 of [Chen et al., 2002]). *The matrices R and L are invertible with probability at least $1/4$. The expected number of non-zero entries in R, L is $O(n \log^2 n)$. If $A \in \mathbb{F}^{n \times n}$, then with probability at least $1 - 1/n$, the minimal polynomial of LAR is $xg(x)$ and $g(0) \neq 0$.*




Figure 3.1: The non-zero entries in a 500×500 matrix constructed using Wiedemann’s distribution.

The preconditioning $A \mapsto LAR$ satisfies Lemma 3.1. Note that the success probability does not depend on the field size and therefore it is suitable for small fields such as $\text{GF}(2)$. The cost for $v \mapsto Lv$ or $v \mapsto Rv$ is $O(n \log^2 n)$ operations in \mathbb{F} . The space requirement for L, R is $O(n \log^2 n)$.

Corollary 3.1. *Let A be an $n \times n$ matrix over \mathbb{F} . Let μ be the number of field operations required to multiply A by a vector. Let R, L be constructed as in Lemma 3.7. Call Algorithm 3.1 with $P = L, Q = R$ in Step 1. Then the output is a correct nullspace sample with probability at least $1 - 1/n$. The cost is $O(r\mu + rn \log^2 n)$ operations in \mathbb{F} . The space complexity is $O(n \log^2 n)$ elements in \mathbb{F} .*

Proof. Correctness follows from Lemma 3.7. The cost is dominated by the minpoly computation which is $2r\mu + 2rO(n \log^2 n) + O(r^2)$. Storage is dominated by L and R . \square

Solving Singular Linear Systems

It is worth noting that we can also (trivially) address the problem of uniform sampling from the solution space of singular linear systems. We are given a singular system $Ax = b$ over a field \mathbb{F} . Assume that there are no rows of the augmented matrix $[A \mid -b]$ which are identically zero. Invoke Algorithm 3.1, to get a random nullspace vector $x' = (x_1, \dots, x_n, x_{n+1}) \in \mathbb{F}^{n+1}$. If $x_{n+1} = 0$, then we repeat the resampling until $x_{n+1} \neq 0$. It is expected that we need a constant number of repetitions. To see why, fix a choice of P, Q and let $B = Qg(P[A \mid -b]Q)$. Now let w be a vector of symbolic entries and let $x' = Bw = (x_1, \dots, x_n, x_{n+1})$. Then x_{n+1} is a linear polynomial in entries of w . This polynomial can not be identically zero because we removed all rows of $[A \mid -b]$ that are entirely zero. Evaluate the entries of w using uniform random choices and apply the Schwartz-Zippel lemma. We get that $\Pr[x_{n+1} \neq 0] \geq 1 - 1/|\mathbb{F}| \geq 1/2$. This is a Bernoulli random variable and the expected number of retries until we get $x_{n+1} \neq 0$ is constant.

Finally, set $x = (x_1/x_{n+1}, \dots, x_n/x_{n+1})$ which is a uniform random solution of $Ax = b$.

3.3 Smith Normal Form

The goal of this section is not to present a new algorithm. Instead, we extend the applications of the transpose preconditioner into computing sparse Smith normal form. We replace Toeplitz preconditioners in existing sparse algorithms [Giesbrecht, 1995, Giesbrecht, 1996, Giesbrecht, 2001, Eberly et al., 2007] with more efficient diagonal preconditioners.

The main idea behind the cited algorithms is that the i th determinantal divisor (the GCD of all $i \times i$ minors) divides the $(n-i)$ th coefficient of the characteristic polynomial. Under suitable preconditioning this division is “maximal”, allowing the determinantal divisors, and hence the invariant factors to be extracted from the coefficients of the characteristic polynomial. A preconditioning is considered successful if the characteristic polynomial of the resulting matrix has the form $x^k f(x)$ where $f(x)$ is square free and $k > 0$. Therefore the coefficients of the characteristic polynomial can be recovered from the minimal polynomial, which is now easier to compute using methods such as Wiedemann’s because it

is square free. Toeplitz matrices were used to achieve this preconditioning in [Giesbrecht, 1996]. Here we show that scaled transpose matrices are suitable for the same task.

3.3.1 Preconditioning and Determinantal Divisors

Let $A \in \mathbb{Z}^{n \times n}$, $\text{rank}(A) = r < n$. Let S be a finite subset of $\mathbb{Z} \setminus \{0\}$. Let $\phi_1, \dots, \phi_n, \psi_1, \dots, \psi_n$ be a set of independent indeterminates. Let $\mathfrak{D}_1 = \text{diag}(\phi_1, \dots, \phi_n)$, $\mathfrak{D}_2 = \text{diag}(\psi_1, \dots, \psi_n)$, and $\mathfrak{A} = \mathfrak{D}_1 A^T \mathfrak{D}_2 A$. Let d_1, \dots, d_{2n} be random assignments of $\phi_1, \dots, \phi_n, \psi_1, \dots, \psi_n$, respectively, where d_j 's are chosen independently and uniformly at random from S . Let $D_1 = \text{diag}(d_1, \dots, d_n)$, $D_2 = \text{diag}(d_{n+1}, \dots, d_{2n})$ and $\tilde{A} = D_1 A^T D_2 A$. In other words, \tilde{A} , D_1 , and D_2 are random evaluations of \mathfrak{A} , \mathfrak{D}_1 , and \mathfrak{D}_2 , respectively.

The next lemma shows that the coefficients of the characteristic polynomial and the minimal polynomial of \tilde{A} agree up to a shift of at most x^{n-r} .

Lemma 3.8. *We have:*

1. *The characteristic polynomial of \mathfrak{A} is a product of a power of x and a square-free polynomial.*
2. *With probability at least $1 - 4n^2/|S|$, $\text{charpoly}(\tilde{A})$ is a product of a power of x^k and a square-free polynomial.*
3. *With probability at least $1 - 4n^2/|S|$, $\text{charpoly}(\tilde{A}) = x^{n-r}g$ where g is a polynomial such that $g(0) \neq 0$.*

Proof. For Part 1, we adopt the argument from Theorem 4.1 of [Chen et al., 2002]. Let $f = \text{charpoly}(\mathfrak{D}_1 A^T \mathfrak{D}_2 A) = x^n + f_1 x^{n-1} + \dots + f_{n-1} x + f_n$. It is known that $f_i = (-1)^i \sum \mathfrak{A}(\sigma)$ for all $\sigma \in \mathcal{C}_i^n$. Substitute (3.2) to get:

$$f_i = (-1)^i \sum_{\sigma, \tau \in \mathcal{C}_i^n} A \binom{\tau}{\sigma}^2 \phi_{\sigma_1} \cdots \phi_{\sigma_i} \psi_{\tau_1} \cdots \psi_{\tau_i}, \quad (3.3)$$

where $\phi_{\sigma_1} \cdots \phi_{\sigma_i} \psi_{\tau_1} \cdots \psi_{\tau_i}$ are distinct for different σ, τ pairs. So f_i is a non-zero homogeneous polynomial in $\mathbb{Z}[\phi_1, \dots, \phi_n, \psi_1, \dots, \psi_n]$ with degree 1 in each of ϕ_j 's, ψ_j 's, and total

degree $2i < 2n$. Now every factorization of f has a *linear degree* in ϕ_j 's and ψ_j 's. So the only possible repeated factors of f must be free of ϕ_j 's and ψ_j 's, i.e., factors of the form x^k . So f is square-free up to a power of x .

For Part 2, consider, f , the characteristic polynomial of \mathfrak{A} . From the previous part we know that $h = f/x^k$ is a square-free polynomial and therefore its discriminant, disc_h , is non-zero. Let $h = x^{n-k} + f_1x^{n-k-1} + \dots + f_{n-k-1}x + f_{n-k}$, where $f_{n-k} \neq 0$. Then disc_h is a non-zero polynomial in $\mathbb{Z}[\phi_1, \dots, \phi_n, \psi_1, \dots, \psi_n]$ whose total degree is $(2(n-k)-1)(2n) < 4n^2$. If we evaluate \mathfrak{A} and h at $\{d_1, \dots, d_{2n}\}$, and apply Schwartz-Zippel lemma to disc_h , we get $\text{disc}_h(d_1, \dots, d_{2n})$ is a non-zero integer with probability at least $1 - 4n^2/|S|$. Thus $\text{charpoly}(\tilde{A})/x^k$ is a square-free polynomial with the same probability.

Part 3 follows from the previous parts, Lemma 3.3, and by observing in (3.3) that f_i must be zero for all $i > r$. \square

The following lemma shows that the coefficients of $\text{charpoly}(\tilde{A})$ contain the prime power divisors of the Δ_i 's. This can be considered as a replacement for the diagonal Toeplitz preconditioners of [Giesbrecht, 2001, Section 1]. For a non-zero integer a , we use $\text{ord}_p(a)$ to denote the exact power of p dividing a . We will not use $\text{ord}_p(0)$ in this chapter.

Lemma 3.9. *Let*

$$f = \text{charpoly}(\mathfrak{A}) = f_n + f_{n-1}x + \dots + f_1x^{n-1} + x^n,$$

where $f_i \in \mathbb{Z}[\phi_1, \dots, \phi_n, \psi_1, \dots, \psi_n]$. For all $1 \leq i \leq r$ we have $\text{ord}_p(f_i) = 2 \text{ord}_p(\Delta_i)$.

Proof. If $\text{ord}_p(\Delta_i) = e$ then for some $\sigma, \tau \in \mathcal{C}_i^n$ we have $\text{ord}_p(A(\begin{smallmatrix} \tau \\ \sigma \end{smallmatrix})) = e$ and for all other minors the order is at least e . From (3.3), the coefficients of f_i are squares of $(i \times i)$ minors of A . Then $\text{ord}_p(f_i) = 2e$. \square

It follows that $\text{ord}_p(f_i(d_1, \dots, d_{2n})) \geq 2 \text{ord}_p(\Delta_i)$ where d_j 's are the random evaluations of ϕ_j 's and ψ_j 's. We now discuss the conditions under which equality is achieved.

Lemma 3.10. *Let $A \in \mathbb{Z}^{n \times n}$ be as above, with rank r and non-zero determinantal divisors $\Delta_1, \dots, \Delta_r \in \mathbb{Z}$, and let $f = \text{charpoly}(\mathfrak{A}) = f_n + f_{n-1}x + \dots + f_1x^{n-1} + x^n$, where $f_i \in$*

$\mathbb{Z}[\phi_1, \dots, \phi_n, \psi_1, \dots, \psi_n]$ for all $i \in [1, n]$. Let $\lambda \in \mathbb{Z}_{>0}$ and $S = \{1, \dots, \lambda\}$. Let d_1, \dots, d_{2n} be assignments of $\phi_1, \dots, \phi_n, \psi_1, \dots, \psi_n$, respectively, which are chosen independently and uniformly at random from S . Let $p > \lambda$ be a given prime. Then the probability that for all $i \in [1, r]$, we have $\text{ord}_p(f_i(d_1, \dots, d_{2n})) = 2 \text{ord}_p(\Delta_i)$, is at least $1 - 2nr/\lambda$.

Proof. From $\text{ord}_p(f_i) = 2 \text{ord}_p(\Delta_i)$ we conclude that $(f_i/\Delta_i^2) \in \mathbb{Z}[\phi_1, \dots, \phi_n, \psi_1, \dots, \psi_n]$ is non-zero modulo p . Moreover, it has degree $2i \leq 2n$ by (3.3).

In choosing elements uniformly from S , we are choosing distinct elements from $S \bmod p$ because $p > \lambda$. Then for any $i \in [1, r]$, the probability that $(f_i/\Delta_i^2)(d_1, \dots, d_{2n}) = 0$ is at most $2n/\lambda$ by the Schwartz-Zippel Lemma. The lemma statement follows by taking the joint probability for all $i \in [1, r]$. \square

Lemma 3.11. *Let $A \in \mathbb{Z}^{n \times n}$ have rank r , and suppose we choose d_1, \dots, d_{2n} from S as in Lemma 3.10. Then the number of distinct primes dividing $f_r(d_1, \dots, d_{2n})$ is less than $n(1 + 2 \log_2 n + 2 \log_2 \lambda + \log_2 \|A\|)$.*

Proof. The coefficient $f_r(d_1, \dots, d_{2n})$ is the sum of the symmetric $r \times r$ minors of $D_1 A^T D_2 A$ which has size $\|D_1 A^T D_2 A\| \leq n \lambda^2 \|A\|^2$. By Hadamard's bound, the $r \times r$ minors have absolute value at most $r^r (n \lambda^2 \|A\|)^r$, and there are at most $\binom{n}{r} < 2^n$ of them. The number of distinct prime factors is at most \log_2 of the product of these quantities, and the lemma immediately follows. \square

The following algorithm computes p -adic approximations to the determinantal divisors for large primes. Our focus is to introduce the linear-time preconditioning and hence we will not repeat the analysis of [Giesbrecht, 2001].

Algorithm 3.2. [Smith Normal Form - Large Primes] Given a matrix $A \in \mathbb{Z}^{n \times n}$ of rank r , compute $\delta_1, \dots, \delta_r \in \mathbb{Z}$ such that for all $i \in [1, r]$ we have $\text{ord}_p \delta_i = 2 \text{ord}_p \Delta_i$ for all large primes p with probability at least $8/9$.

1. Choose λ such that $\frac{1}{\lambda} 2n^3 (1 + 2 \log n + 2 \log \lambda + \log \|A\|) < 1/3$.

2. Let $S = \{1, 2, \dots, \lambda\}$.
3. For k from 1 to 2:
 - (a) Choose d_1, \dots, d_{2n} independently and uniformly at random from S .
 - (b) Let $D_1 = \text{diag}(d_1, \dots, d_n)$, $D_2 = \text{diag}(d_{n+1}, \dots, d_{2n})$, and $\tilde{A} = D_1 A^T D_2 A$.
 - (c) Compute $f^{(k)}(x) = \text{minpoly}(\tilde{A})$.
 - (d) Let $g^{(k)}(x) = f^{(k)}(x)/x = g_r^{(k)} + g_{r-1}^{(k)}x + \dots + x^r$.
4. For all $1 \leq i \leq r$, set $\delta_i = \text{gcd}(g_i^{(1)}, g_i^{(2)})$.
5. Return $\delta_1, \dots, \delta_r$.

Repeating the computations twice and computing coefficient-wise GCD of the resulting polynomials amplifies the probability to $1 - (1 - 2/3)^2 \geq 8/9$.

The choice of λ in the algorithm above satisfies Lemmas 3.10 and 3.11 for all relevant primes at once. So with the stated probability we can correctly compute $\text{ord}_p(\Delta_i)$ for all primes $p > \lambda$, where $\lambda \in O(n^3)$. This defines a notion of *large primes*: $p > \lambda$, and small primes otherwise.

For a small prime p , we can construct a ring extension similar to the construction we used in §2.4.3. Let e be an integer such that $p^e > \lambda$. Let $\Gamma_p \in \mathbb{Z}[x]$ be a polynomial of degree e whose image over $\mathbb{Z}/p\mathbb{Z}$ is irreducible. The ring $\mathbb{R}_p = \mathbb{Z}[x]/(\Gamma_p)$ contains a copy of $\text{GF}(p^e)$ with at least λ elements. We can apply the Schwartz-Zippel lemma over \mathbb{R}_p modulo p . Similar to Lemma 3.10, we can show that $p \nmid f_i/\Delta_i^2$ with high probability if the random choices for D_1, D_2 are taken as polynomials from a subset of \mathbb{R}_p . Thus the preconditioning $A \mapsto D_1 A^T D_2 A$ works for smaller primes.

In general, there are several small primes p_1, \dots, p_k which divide Δ_r . In [Giesbrecht, 2001, §2], all k primes are considered at once. The resulting ring extension is called a *rough* extension ring $\mathbb{R} = \mathbb{R}_{p_1} \oplus \dots \oplus \mathbb{R}_{p_k} = \mathbb{Z}[x]/(\Gamma)$, where $\Gamma = \Gamma_{p_1} \cdots \Gamma_{p_k} \in \mathbb{Z}[x]$. The preconditioning $A \mapsto D_1 A^T D_2 A$ works for this setup as well using the arguments above.

3.4 Conclusion

We extended the applications of linear-time preconditioners to random and uniform nullspace sampling, and computing Smith normal form of sparse and black-box matrices. For the case of small fields (for nullspace sampling) and small primes (for Smith normal form), we incurred a logarithmic penalty for working over extensions. In the nullspace case, we adopted Wiedemann's sparse preconditioners as an alternative to constructing field extensions, at (approximately) the same logarithmic overhead.

It remains an open question to find linear-time preconditioners, or apply known ones, to Wiedemann's algorithm over small fields (for the nullspace sampling) and with respect to small primes (for Smith normal form), without constructing extensions.

Chapter 4

Rank Reduction

This chapter presents a new technique for computing Smith normal form using rank-1 updates. The underlying technique is a rank reduction first formulated by Wedderburn in 1934. We will introduce rank reduction and then present two variants of our algorithm: iterative and block-iterative. Finally, we will present an application in computing nullspace vectors.

4.1 Introduction

Here we consider the problem of computing the Smith normal form of matrices with entries from a local ring. Throughout our presentation we will focus on $\mathbb{Z}/p^e\mathbb{Z}$ but the results are applicable to other local rings such as $\mathbb{F}[x]/(f^k)$.

As discussed in earlier chapters, existing approaches to computing Smith normal form rely on elimination [Storjohann, 2000], the characteristic polynomial [Dumas et al., 2001], [Giesbrecht, 2001], and random perturbations combined with solving random linear systems [Villard, 2000, Eberly et al., 2000]. Notably, the work of [Eberly et al., 2000] relies on modifying (or perturbing) the last k invariant factors. The approach we present here is in the realm of perturbing the invariant factors. Our perturbation schemes modify the first k invariant factors of a matrix using rank-1 and rank- k updates. For rank updates, we rely on the so-called Wedderburn rank reduction formula. Let A be an $n \times n$ matrix over a

field \mathbb{F} . Let $x, y \in \mathbb{F}^n$ be two vectors such that $w = y^T A x \in \mathbb{F}$ is non-zero. Then the rank-1 update given by

$$B := A - w^{-1} A x y^T A, \quad (4.1)$$

will result in $\text{rank}(B) = \text{rank}(A) - 1$. Wedderburn [Wedderburn, 1934, p. 69] first discovered this property for real and complex-valued matrices. The converse of the formula is also true [Householder, 1964]. Independently, Egerváry discovered this property and its converse, and used it in an iterative process to compute LU decompositions [Galántai, 2010]. Cline and Funderlic [Cline and Funderlic, 1979] have generalized these results from rank-1 updates to block updates as follows. Suppose $X, Y \in \mathbb{M}_{n \times k}(\mathbb{F})$ and $W = Y^T A X \in \mathbb{M}_k(\mathbb{F})$ such that W is invertible, and let

$$B := A - A X W^{-1} Y^T A. \quad (4.2)$$

Then $\text{rank}(B) = \text{rank}(A) - k$. One can then take the resulting matrix and re-apply a second iteration of rank reduction and so forth until the resulting matrix is zero.

The iterative nature of the rank-1 update lends itself to applications in numerical linear algebra where iterative methods are widely used. A comprehensive study of Wedderburn rank reduction formula can be found in [Chu et al., 1995]. They extend the applications of rank reduction to a general-purpose bi-conjugation process, and show that many matrix factorizations such as SVD, QR, Cholesky decomposition, Gram-Schmidt and Lanczos can be formulated in terms of this bi-conjugation process. The authors also show that Wedderburn rank reduction is related to the ABS method [Abaffy et al., 1984].

The work of Raboky and Amiri [Raboky and Amiri, 2013a, Raboky and Amiri, 2013b] is the only direct utilization of Wedderburn rank reduction that we know of in exact linear algebra. They develop a bi-conjugation process based on the Wedderburn rank formula and the ABS method to compute Smith normal form of integer matrices. Their algorithm is deterministic, and has a quadratic space complexity and hence it is only suitable for dense matrices.

The authors provide experimental results to show that their algorithm performs well in terms of space and time. We attempt at analyzing its time complexity. There are at

most n iterations in their algorithm. The key step in every iteration is finding integer vectors $t, w \in \mathbb{Z}^n$ such that $t^T A_i w = \gcd(A_i)$ where A_i is A 's perturbation at step i . Since A_i is completely known at step i , finding t, w reduces to solving a quadratic diophantine equation. The authors present an algorithm for solving this problem which is dominated by the time to compute an integer row basis for A_i . We note that this is bounded by the time to compute Hermite normal form. Thus the overall bit complexity is bounded by $O^\sim(n^{\omega+2})$. The algorithm has the advantage of explicitly constructing the transformation matrices.

Throughout this chapter we use the following notation. Let p^e be a prime power. All equalities in this chapter are over the ring $\mathbb{Z}/p^e\mathbb{Z}$ and hence should be understood as equivalences modulo p^e . For a matrix $M \in \mathbf{M}_n(\mathbb{Z}/p^e\mathbb{Z})$, let $\text{snf}M$ be the Smith normal form of M . Two matrices A, B are unimodularly equivalent if $\text{snf}A = \text{snf}B$ which we write as $A \sim B$. Let $\ker(M) \subseteq (\mathbb{Z}/p^e\mathbb{Z})^n$ denote the right kernel of M over $\mathbb{Z}/p^e\mathbb{Z}$, that is, the $\mathbb{Z}/p^e\mathbb{Z}$ module of all vectors v such that $Mv = 0$. With abuse of notation, we also use $\ker(M)$ to denote an $n \times b$ matrix whose b columns form a basis of the kernel. For a matrix M , and a set of vectors x_1, \dots, x_n (where $n \geq 1$), we use the notation $M \cup \{x_1, \dots, x_n\}$ to denote $\text{span}(M, x_1, \dots, x_n)$. Finally, $[M \ x]$ denotes the $n \times (n+1)$ matrix resulting from augmenting the column vector x to the matrix M . If p is a prime, M is a matrix, and v is a vector, then we use the notation $p \mid M$, and $p \mid v$ to mean $M \equiv 0 \pmod{p}$, and $v \equiv 0 \pmod{p}$, respectively.

We offer simple algorithms with polynomial time complexity. The algorithms are randomized and their success probabilities are controllably close to 1. Our algorithms only compute the invariant factors and cannot produce the unimodular transformation matrices. It is worth noting that this is sufficient for many applications in system theory [McMillan, 1952, Kailath, 1980], and in algebraic topology [Dumas et al., 2003], where the main interest is in the invariant factors, rather than the transformation matrices.

4.2 Invariant Factors and Rank-1 Updates

Let $A \in M_n(\mathbb{Z}/p^e\mathbb{Z})$ and let x, y be two non-zero vectors with entries from $\mathbb{Z}/p^e\mathbb{Z}$ with the condition that $y^T Ax$ is a unit in $\mathbb{Z}/p^e\mathbb{Z}$. We will discuss the existence of such vectors in a later section. Let

$$B = A - w^{-1}Axy^T A. \quad (4.3)$$

The following lemma shows the relationship between the kernels of A and B .

Lemma 4.1. *The sets $\ker(A)$ and $\ker(B)$ satisfy $\ker(B) = \ker(A) \cup \{x\}$.*

Proof. The proof follows [Wedderburn, 1934, §5.06]. The inclusion $\ker(A) \cup \{x\} \subseteq \ker(B)$ is straightforward. For the other direction, let v be any vector such that $Bv = 0$, and let $\alpha = w^{-1}y^T Av$. We have

$$Bv = Av - w^{-1}Axy^T Av = Av - \alpha Ax = A(v - \alpha x) = 0.$$

Then $v - \alpha x \in \ker(A)$. But $\alpha \in \mathbb{Z}/p^e\mathbb{Z}$, so $v \in \ker(A) \cup \{x\}$ over $\mathbb{Z}/p^e\mathbb{Z}$, and $\ker(B) \subseteq \ker(A) \cup \{x\}$. \square

Example 4.1. *Take A to be the 2×2 identity matrix over $\mathbb{Z}/4\mathbb{Z}$. If we choose $x = [3 \ 2]^T$ and $y = [1 \ 1]^T$, then $w = y^T Ax = [1 \ 1][3 \ 2]^T = 1$, and*

$$B = A - w^{-1}Axy^T A = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} - \begin{bmatrix} 3 & 3 \\ 2 & 2 \end{bmatrix} = \begin{bmatrix} -2 & -3 \\ -2 & -1 \end{bmatrix} = \begin{bmatrix} 2 & 1 \\ 2 & 3 \end{bmatrix}.$$

Indeed, x is now in the (right) kernel of B :

$$Bx = \begin{bmatrix} 2 & 1 \\ 2 & 3 \end{bmatrix} \begin{bmatrix} 3 \\ 2 \end{bmatrix} = 0.$$

Note that $\text{snf } A = \text{diag}(1, 1)$ and $\text{snf } B = \text{diag}(1, 4) = \text{diag}(1, 0)$ over $\mathbb{Z}/4\mathbb{Z}$. \square

In the rest of this section we will study the effects of the rank-1 reduction on the invariant factors.

Lemma 4.2. Let $M \in \mathbf{M}_n(\mathbb{Z}/p^e\mathbb{Z})$, $x \in (\mathbb{Z}/p^e\mathbb{Z})^n$, and $M = PDQ$ be the Smith decomposition of M where $P, Q \in \mathbf{GL}_n(\mathbb{Z}/p^e\mathbb{Z})$. Then $\ker(M) \sim \ker(D)$, and $\text{snf} \begin{bmatrix} M & x \end{bmatrix} = \text{snf} \begin{bmatrix} D & P^{-1}x \end{bmatrix}$.

Proof. If $M = PDQ$ then $\ker(D) = Q\ker(M)$ and so $\ker(D) \sim \ker(M)$ because Q is unimodular. For the second part, we have

$$\text{snf} \begin{bmatrix} M & x \end{bmatrix} = \text{snf} P^{-1} \begin{bmatrix} M & x \end{bmatrix} \begin{bmatrix} Q^{-1} & \\ & 1 \end{bmatrix} = \text{snf} \begin{bmatrix} D & P^{-1}x \end{bmatrix},$$

because Smith normal form is invariant under left and right multiplication by unimodular matrices. \square

The following lemma shows that under a certain choice of basis for $\ker(A)$, rank-1 reduction is equivalent to adding a unit vector \mathbf{e}_i to the kernel basis.

Lemma 4.3. Let the Smith normal form of $A = USV \in \mathbf{M}_n(\mathbb{Z}/p^e\mathbb{Z})$ be

$$S = \text{diag}(1, p^{e_2}, \dots, p^{e_n}),$$

where $0 \leq e_i \leq e$. Let $x \in (\mathbb{Z}/p^e\mathbb{Z})^n$ be such that $x \notin \ker(A)$ and $p \nmid Ax$. Then

$$[\ker(A) \quad x] \sim \text{diag}(1, p^{e-e_2}, \dots, p^{e-e_n}).$$

Proof. For the given S , we have $\ker(A)$ is generated by

$$V^{-1} \text{diag}(0, p^{e-e_2}, \dots, p^{e-e_n}).$$

Let $z = Vx$ for some vector z . Then $p \nmid Ax$ implies that the first entry of z must be a unit for otherwise $Ax = USz = (p^i, *, \dots, *)$ for some $i > 0$ which is impossible since $p \nmid Ax$. Now,

$$\begin{aligned} \text{snf}[\ker(A) \quad x] &= \text{snf} V[\ker(A) \quad x] \\ &= \begin{bmatrix} 0 & & & \text{unit} \\ & p^{e-e_2} & & * \\ & & \ddots & * \\ & & & p^{e-e_n} & * \end{bmatrix}. \end{aligned}$$

We do not need to specify the entries denoted by $*$. If we swap the first and last columns we get a lower triangular matrix. We can apply column and row operations to reduce the entries below the unit element to zero. After rearranging the entries in the canonical Smith normal form divisibility chain, we have $[\ker(A) \ x] \sim \text{diag}(1, p^{e-e_n}, \dots, p^{e-e_2})$. \square

As an example of Lemma 4.3, consider the following matrix over $\mathbb{Z}/4\mathbb{Z}$:

$$A = \begin{bmatrix} 1 & 0 \\ 0 & 2 \end{bmatrix}, \quad \ker(A) = \begin{bmatrix} 0 & 0 \\ 0 & 2 \end{bmatrix}.$$

If we choose $x = [1, 1]^T$ and $y = [1, 1]^T$, then the rank-1 reduction is

$$B = \begin{bmatrix} 2 & 2 \\ 2 & 2 \end{bmatrix}, \quad \ker(B) = \begin{bmatrix} 1 & 0 \\ 1 & 2 \end{bmatrix} \sim \begin{bmatrix} 1 & 0 \\ 0 & 2 \end{bmatrix}.$$

We see that adding x to the kernel of A has the effect of adding the vector \mathbf{e}_1 to the Smith normal form of $\ker(B)$.

In the above example the rank reduction modified the invariant factors of A from $\text{diag}(1, 2)$ to $\text{snf} B = \text{diag}(2, 0)$. The following result shows that rank reduction modifies the structure of the Smith normal form by decrementing the number of 1's and incrementing the number of 0's, while leaving all the other invariant factors intact.

Theorem 4.1. *Let the Smith normal form of $A \in \mathbf{M}_n(\mathbb{Z}/p^e\mathbb{Z})$ be $\text{diag}(1, p^{e_2}, \dots, p^{e_n})$. Let x, y be two vectors in $(\mathbb{Z}/p^e\mathbb{Z})^n$ such that $w = y^T A x$ is a unit in $\mathbb{Z}/p^e\mathbb{Z}$, and let $B = A - w^{-1} A x y^T A$. Then $\text{snf} B = \text{diag}(p^{e_2}, \dots, p^{e_n}, 0)$.*

Proof. If $p \nmid y^T A x$ then $A x \neq 0$ and $p \nmid A x$. Applying Lemma 4.2, Lemma 4.1, and Lemma 4.3 in order we get

$$\ker(\text{snf} B) \sim \ker(B) = [\ker(A) \ x] \sim \text{diag}(1, p^{e-e_n}, \dots, p^{e-e_2}).$$

But $\text{snf} B$ is a diagonal matrix, so it is trivial to infer from its kernel that

$$\text{snf} B \sim \text{diag}(0, p^{e_n}, \dots, p^{e_2}).$$

Now reorder the invariant factors of B to get the desired Smith normal form. \square

The iterative scheme for computing the invariant factors based on rank-1 reductions is now straightforward. Start with the input matrix A , and set s_1 to the GCD of all its elements. Replace A by A/s_1 such that Theorem 4.1 is applicable, and then apply the rank reduction to the resulting matrix. Then s_2/s_1 is the GCD of all elements of the resulting (rank-1 reduced) matrix. By iteratively applying this scheme, all the invariant factors of A will be discovered after at most $r \leq n$ iterations.

It is worth noting that this iterative process decomposes A as follows. As step i let A_i denote the current reduction of A . Let x_i, y_i denote the current choice of the reduction vectors x, y . Finally, let $u_i = A_i x_i$, $v_i^T = y_i^T A_i$, and $w_i = y_i^T A_i x_i$. Then

$$A = USW^{-1}V = \begin{bmatrix} | & | & & | \\ u_1 & u_2 & \cdots & u_r \\ | & | & & | \end{bmatrix} \begin{bmatrix} \frac{s_1}{w_1} & & & \\ & \frac{s_2}{w_2} & & \\ & & \ddots & \\ & & & \frac{s_r}{w_r} \end{bmatrix} \begin{bmatrix} - & v_1^T & - \\ - & v_2^T & - \\ & \vdots & \\ - & v_r^T & - \end{bmatrix}.$$

Note that U, V need not be unimodular.

We will now discuss a randomized approach for choosing x, y using a Bernoulli process without compromising the overall complexity of the algorithm. We will also present the deterministic approach in a subsequent discussion. The deterministic approach is obviously preferable. However, the randomized approach might be of use when one cannot efficiently inspect all the entries of A , e.g., when A is given by one or more black-boxes, or by a lengthy straight line program.

The next lemma shows that there are sufficiently many vectors x, y such that $p \nmid y^T Ax$.

Lemma 4.4. *Let x, y be n vectors whose entries are chosen uniformly at random from $\mathbb{Z}/p^e\mathbb{Z}$. If $s_1 = 1$, then $\Pr[p \nmid y^T Ax] \geq 1 - 2/p$.*

Proof. First treat the entries of x, y as algebraically independent symbols. Then

$$y^T Ax = \sum_{i=1}^n \sum_{j=1}^n a_{ij} x_i y_j \in (\mathbb{Z}/p^e\mathbb{Z})[x_1, \dots, x_n, y_1, \dots, y_n]. \quad (4.4)$$

But $s_1 = 1$ implies that at least one a_{ij} is a unit, and so $y^T Ax$ is a non-zero polynomial and has total degree exactly 2. In fact, $s_1 = 1$ implies that the image of $y^T Ax$ in $\mathbb{Z}/p\mathbb{Z}$ is also a non-zero polynomial of total degree equal to 2.

Now we choose the entries of x, y at random. When choosing values for x_j 's, y_i 's uniformly (and independently) at random from $[0, p^e)$, we are also choosing the values of $x_j \pmod p$ and $y_i \pmod p$ over $\mathbb{Z}/p\mathbb{Z}$ uniformly at random from $[0, p)$, since each element in $[0, p^e)$ can be written as a power series in p with coefficients from $[0, p)$. If we apply the Schwartz-Zippel lemma on the image of $y^T Ax$ over $\mathbb{Z}/p\mathbb{Z}$, we get $\Pr[p \mid y^T Ax] \leq \deg(y^T Ax)/|\mathbb{Z}/p\mathbb{Z}| = 2/p$. \square

Therefore a randomly chosen pair of vectors will satisfy the conditions of Theorem 4.1 with probability at least $1/2$ for all primes $p \geq 5$. To get a probability of at least $1/2$ when $p = 2$ or 3 , we sample x, y from a constant degree ring extension of $\mathbb{Z}/p^e\mathbb{Z}$, while retaining a good probability that w is a unit in the ring extension.

Let $\text{GR}(p^e, d)$ be the Galois ring extension of $\mathbb{Z}/p^e\mathbb{Z}$ whose degree is d . Over $\text{GR}(p^e, d)$, $w = y^T Ax$ is a unit if $p \nmid w$. To keep the presentation uniform for both $p = 2$ and $p = 3$, we choose the polynomial $\alpha^2 - \alpha - 1$ which is irreducible over both $\mathbb{Z}/2\mathbb{Z}$ and $\mathbb{Z}/3\mathbb{Z}$. Other irreducible polynomials can be independently chosen for either of the two rings. Then we use the extension $\text{GR}(p^e, 2) = \mathbb{Z}[\alpha]/(p^e, \alpha^2 - \alpha - 1)$ which has p^2 elements. Arithmetic operations over $\text{GR}(p^e, 2)$ require a constant number of operations over $\mathbb{Z}/p^e\mathbb{Z}$ using polynomial arithmetic in $(\mathbb{Z}/p^e\mathbb{Z})[x]$ with degrees at most 2. Hence working over this extension only introduces a constant overhead in the complexity of the algorithm. Finally, $\text{GR}(p^e, 2)$ contains an image of $\text{GF}(p^2)$ given by $\mathbb{Z}[\alpha]/(p, \alpha^2 - \alpha - 1)$ which will be useful in applying the Schwartz-Zippel lemma.

Lemma 4.5. *Let x, y be n vectors whose entries are chosen uniformly at random from $\text{GR}(p^e, 2)$. If $s_1 = 1$, then $\Pr[2 \nmid y^T Ax] \geq 1/2$.*

Proof. As in Lemma 4.4, $s_1 = 1$ implies that $y^T Ax$ is a non-zero polynomial of degree 2 over $\text{GF}(p^2)$. Let $c_1\alpha + c_0$ be an entry of x or y , which is chosen uniformly at random from $\text{GR}(p^e, 2)$. Then c_0, c_1 are chosen uniformly at random from $\mathbb{Z}/p^e\mathbb{Z}$ and $c_1\alpha + c_0 \pmod p$ is chosen uniformly at random from $\text{GF}(p^2)$. The rest of the proof is similar to Lemma 4.4, and we get

$$\Pr[y^T Ax \not\equiv 0 \pmod{p, \alpha^2 - \alpha - 1}] \geq 1 - \frac{2}{|\text{GF}(p^2)|} \geq 1/2,$$

using the Schwartz-Zippel lemma over $\text{GF}(p^2)$. □

We can now present an iterative algorithm for computing the Smith normal form over $\mathbb{Z}/p^e\mathbb{Z}$ using rank-1 updates and random choices for x, y .

Algorithm 4.1. Given a matrix $A = (a_{ij}) \in \text{M}_n(\mathbb{Z}/p^e\mathbb{Z})$, this algorithm returns A 's invariants factors: $s_1, \dots, s_n \in \mathbb{Z}/p^e\mathbb{Z}$.

1. Initialize $s_0 = 1$ and $s_1 = 0, \dots, s_n = 0$.
2. For $\ell = 1$ to n :
 - (a) Compute $\text{gcd}(A) = \text{gcd} \{a_{ij} : i, j \in [1, n]\}$.
 - (b) Set $s_\ell = s_{\ell-1} \cdot \text{gcd}(A)$.
 - (c) If $s_\ell = 0$ then **break**.
 - (d) Update A using $A := A / \text{gcd}(A)$.
 - (e) Find x, y :
 - i. Construct vectors x, y with entries sampled uniformly at random from $[0, p^e)$.
 - ii. Compute $w = y^T Ax$.
 - iii. Repeat steps 2(e)i, 2(e)ii until $p \nmid w$.
 - (f) Update A using $A := A - w^{-1}Ax y^T A \pmod{p^e}$.
3. Return s_1, \dots, s_n .

When $p = 2$ or 3 , all arithmetic operations will be done over $\text{GR}(p^e, 2)$. In particular, the random entries in step 2(e)i can be constructed as $a\alpha + b$ where $a, b \in [0, p^e)$ are chosen uniformly at random.

Theorem 4.2. For all primes $p \geq 2$, Algorithm 4.1 is a correct randomized Las Vegas algorithm. The expected cost of the algorithm is $O(n^3)$ operations over $\mathbb{Z}/p^e\mathbb{Z}$. The space complexity is $O(n^2)$ elements in $\mathbb{Z}/p^e\mathbb{Z}$.

Proof. First recall that scaling a matrix by a power of p also scales each invariant factor by the same power of p . The first run of step 2b correctly computes s_1 . Step 2d always scale A such that Theorem 4.1 is applicable. Suppose that the invariant factors of A are $[s_1, \dots, s_n]$. After the first run of step 2d, the algorithm transforms the invariant factors into $[1, s_2/s_1, \dots, s_n/s_1]$. After step 2f, the invariant factors become $[s_2/s_1, s_3/s_1, \dots, s_n/s_1, 0]$. The subsequent iteration then computes $\gcd(A) = s_2/s_1$. Hence $s_2 = s_1 \gcd(A)$. Scaling the matrix again by s_2/s_1 transforms the invariant factors into $[1, s_3/s_2, \dots, s_n/s_2, 0]$ and so on. Thus at iteration ℓ , the gcd computed by step 2a is equal to $s_\ell/s_{\ell-1}$. This loop stops when the ℓ th invariant factor is 0 and hence all subsequent invariant factors are zero, i.e., when ℓ is the rank of A . The algorithm is correct.

Each individual step is dominated by $O(n^2)$ operations, including 2f where the matrix multiplication can be factored into the outer product $(Ax)(y^T A)$.

Steps 2(e)i, 2(e)ii are essentially sampling a Bernoulli random variable where the success outcome is $y^T Ax \not\equiv 0 \pmod{p}$, which has probability at least $1/2$ by Lemma 4.4 (and Lemma 4.5). The expected number of trials until the success of a Bernoulli random variable is given by the inverse of the success probability, which is a constant. At each iteration we can verify the success of the random choice by testing that $y^T Ax$ is a unit. The algorithm is randomized Las Vegas.

The expected cost of the overall algorithm is $O(n^3)$. Assuming step 2f is performed out of place on A , the space complexity of the algorithm is $O(n^2)$ elements which are required to store x, y and the reduction of A at every iteration.

When $p = 2$ or 3 we work over $\text{GR}(p^e, 2)$. In this case, the complexity will only increase by a constant factor while maintaining good probability bounds (greater than 0.5). So the arguments of this proof hold for all primes. \square

Finally, the deterministic choice for x, y is rather simple. At every iteration, the first invariant factor of the current matrix A is 1 because we have already divided A by its gcd. Then there must be at least one entry a_{ij} such that $p \nmid a_{ij}$. If we let $x = \mathbf{e}_j$ and $y = \mathbf{e}_i$ then $w = y^T Ax = \mathbf{e}_i^T A \mathbf{e}_j = a_{ij}$, and we must have $p \nmid w$ for this choice. We get the following result.

Theorem 4.3. *In Algorithm 4.1, replace step 2e with: Set $x = \mathbf{e}_j$ and $y = \mathbf{e}_i$ where $i, j \in [1, n]$ are any pair of indices such as $p \nmid a_{ij}$. The resulting algorithm is deterministic, and costs $O(n^3)$ operations over $\mathbb{Z}/p^e\mathbb{Z}$. The space complexity is $O(n^2)$ elements in $\mathbb{Z}/p^e\mathbb{Z}$.*

This algorithm is similar to local elimination approach of [Dumas et al., 2001, algorithm LRE] which has cubic time complexity as well.

4.3 Block Reduction

The algorithm in the preceding section has cubic time complexity. To achieve subcubic complexity, we utilize fast matrix multiplication, and block rank reductions.

Let A be an $n \times n$ matrix over $\mathbb{Z}/p^e\mathbb{Z}$ and its Smith normal form be

$$\text{diag}(\underbrace{1, \dots, 1}_{r_0}, \underbrace{p, \dots, p}_{r_1}, \dots, \underbrace{p^{e-1}, \dots, p^{e-1}}_{r_{e-1}}, 0, \dots, 0).$$

Let $b \leq r_0$, and $X, Y \in \mathbf{M}_{n \times b}(\mathbb{Z}/p^e\mathbb{Z})$ such that $W = Y^T A X \in \mathbf{GL}_b(\mathbb{Z}/p^e\mathbb{Z})$. Let the columns of X be X_1, X_2, \dots, X_b . In what follows it is shown that the rank of

$$B := A - A X W^{-1} Y^T A,$$

is exactly $\text{rank}(A) - b$. The rank- b reduction decrements the number of ones in the Smith normal form by b while, at the same time, increments the number of zeroes by b .

First we establish that rank reduction adds the columns of X to the kernel of the new matrix.

Lemma 4.6. *We have $\ker(B) = \ker(A) \cup \{X_1, \dots, X_b\}$.*

Proof. If $Av = 0$ then $Bv = 0$. For all $i \in [1, b]$, we have

$$B X_i = B X \mathbf{e}_i = A X \mathbf{e}_i - A X W^{-1} Y^T A X \mathbf{e}_i = A X \mathbf{e}_i - A X \mathbf{e}_i = 0.$$

So $\ker(A) \cup \{X_1, \dots, X_b\} \subseteq \ker(B)$.

For the other direction of set inclusion, let $Bv = 0$ and $u = XW^{-1}Y^T Av$. Then

$$Av - AXW^{-1}Y^T Av = A(v - XW^{-1}Y^T Av) = A(v - Xu) = 0$$

or $v - Xu \in \ker(A)$. So $v \in \ker(A) \cup \{X_1, \dots, X_b\}$, or more generally, $\ker(B) \subseteq \ker(A) \cup \{X_1, \dots, X_b\}$. \square

Example 4.2. Let $A \in M_{4 \times 4}(\mathbb{Z}/8\mathbb{Z})$ be

$$A = \begin{bmatrix} 1 & & & \\ & 1 & & \\ & & 2 & \\ & & & 4 \end{bmatrix}.$$

If we choose X, Y as

$$X = \begin{bmatrix} 1 & 0 \\ 0 & 1 \\ 0 & 0 \\ 0 & 0 \end{bmatrix}, Y^T = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix}.$$

Then W is the 2×2 identity matrix, and

$$B = A - AXW^{-1}Y^T A = A - A \begin{bmatrix} 1 & & & \\ & 1 & & \\ & & 0 & \\ & & & 0 \end{bmatrix} A = \begin{bmatrix} 0 & & & \\ & 0 & & \\ & & 2 & \\ & & & 4 \end{bmatrix}.$$

Note $\text{rank}(B) = 2$, and both columns of X are nullspace vectors of B . \square

If $\ker(B) = [\ker(A) \ X]$, what is the Smith normal form of B ?

Theorem 4.4. If the Smith normal form of A is $\text{diag}(\underbrace{1, \dots, 1}_b, p^{e_{b+1}}, \dots, p^{e_n})$ where $0 \leq e_i \leq e$ for all $i \in [e_{b+1}, e_n]$. Then the Smith normal form of B is $\text{diag}(p^{e_{b+1}}, \dots, p^{e_n}, \underbrace{0, \dots, 0}_b)$.

Proof. We will examine the Smith normal form of $\ker(B)$. Let the Smith normal form of A be given by $A = USV$. We have $\ker(A) = V^{-1} \ker(S)$. Let $Z = VX$ for some $n \times b$

matrix Z . For any column X_i of X we have $p \nmid AX_i$ otherwise the i th column of W will be zero modulo p which is impossible since W is invertible modulo p . This implies $p \nmid SZ_i$, and this can only happen if the i th entry of Z_i is a unit because the upper $b \times b$ submatrix of S is the identity matrix. So the entries on the diagonal of Z are units.

We get $\ker(B) = [\ker(A) \ X] = V^{-1}[\ker(S) \ Z]$ which is unimodularly equivalent to

$$\left[\begin{array}{c|ccc} & \text{unit} & & * \\ & & \ddots & \\ & * & & \text{unit} \\ \hline p^{e-e_{b+1}} & & & \\ & \ddots & & \\ & & p^{e-e_n} & \\ \hline & & & * \end{array} \right].$$

If we apply column operations on the $b \times b$ rightmost top quadrant, we can reduce it to a lower triangular matrix. However we need to show that Z will reduce to a lower unit triangular matrix. Let C encode the desired column operations. Assume by way of contradiction that the top $b \times b$ submatrix of Z has rank less than b . Then for some i , the column $(ZC)_i$ is either 0 or divisible by p . Therefore $p \mid (SZC)_i$ and $p \mid U(SV^{-1}VC)_i = (AXC)_i$ which is impossible because AX has a full column rank, and a unimodular transformation C will not alter the rank. Therefore the assumption is wrong. The top $b \times b$ submatrix of Z has rank b

If we use the units on the diagonal to eliminate the lower part of the triangular matrix, we get

$$\left[\begin{array}{c|ccc} & \text{unit} & & 0 \\ & & \ddots & \\ & * & & \text{unit} \\ \hline p^{e-e_{b+1}} & & & \\ & \ddots & & \\ & & p^{e-e_n} & \\ \hline & & & * \end{array} \right] \sim \left[\begin{array}{c|ccc} & 0 & & I_{b \times b} \\ \hline p^{e-e_{b+1}} & & & \\ & \ddots & & \\ & & p^{e-e_n} & \\ \hline & & & 0 \end{array} \right].$$

So $\ker(B) \sim \text{diag}(I_{b \times b}, p^{e-e_{b+1}}, \dots, p^{e-e_n})$. Finally, $\ker(\text{snf} B) \sim \ker(B)$ so we can deduce that $\text{snf} B = \text{diag}(p^{e_{b+1}}, \dots, p^{e_n}, 0, \dots, 0)$. \square

In what follows we examine the construction of X, Y .

Lemma 4.7. *If the entries of X, Y are chosen uniformly at random from $[0, p^k)$ then the probability that $p \mid \det(W)$ is at most $2b/p$.*

Proof. First consider the case where the $2nb$ entries of X, Y are algebraically independent variables. Then by the Cauchy-Binet formula, $\det(W)$ is a polynomial of degree $2b$ in the entries of X, Y . This polynomial is not identically zero modulo p since there is at least one $b \times b$ minor of A which is not zero modulo p by the requirement that $b \leq r_0$. Now apply the Schwartz-Zippel lemma over $\mathbb{Z}/p\mathbb{Z}$ to get $\Pr[\det W^{-1} \equiv 0 \pmod{p}] \leq 2b/p$. \square

A single step of the randomized block rank reduction involves setting $b \leq r_0$, and selecting the entries of X, Y uniformly at random from $[0, p^e)$, computing $W^{-1} = (Y^T A X)^{-1}$, and then applying $A := A - A X W^{-1} Y^T A X$.

Lemma 4.8. *A single step of the randomized block rank reduction succeeds with probability at least $1 - 2b/p$ and costs $O(en^\omega)$ operations in $\mathbb{Z}/p^e\mathbb{Z}$.*

Proof. A single step succeeds if $\det W \not\equiv 0 \pmod{p}$, if the columns of X are linearly independent, and if the Smith normal form of X is $\text{diag}(1, \dots, 1)$. We will show that these conditions are redundant, so we do not need to consider the probability of these events independently.

The probability that W is invertible is at least $1 - 2b/p$ by Lemma 4.7. Using the Cauchy-Binet formula we have

$$\det W = \sum_{\sigma, \tau} Y^T \begin{pmatrix} 1..b \\ \sigma \end{pmatrix} A \begin{pmatrix} \sigma \\ \tau \end{pmatrix} X \begin{pmatrix} \tau \\ 1..b \end{pmatrix}.$$

If $p \nmid \det W$ then there exists at least one pair σ, τ such that the $b \times b$ minor $X \begin{pmatrix} \tau \\ 1..b \end{pmatrix}$ is a unit modulo p . This implies that the b th invariant factor of X is 1, and $\text{snf } X = \text{diag}(1, \dots, 1)$. The latter Smith normal form implies that the columns of X are linearly independent. Finally, if any column X_i of X is such that $A X_i = 0$, then the i th column of W will be zero, but this is impossible since W is invertible. Columns of X are not in the nullspace

of A . The rank reduction will reduce the rank by b and will succeed in converting b 1's to 0's in the Smith normal form of A .

The cost is dominated by the matrix multiplications, and the inversion of W which is $O(en^\omega)$ operations in $\mathbb{Z}/p^e\mathbb{Z}$ using e steps of lifting. \square

To obtain a good success probability of at least 0.5, we require that $2b/p \leq 1/2$, i.e., $p \geq 4n$. When the prime is small we can work over a Galois Ring extension $\text{GR}(p^e, d)$ with sufficient elements to maintain a good success probability.

Lemma 4.9. *A rank- b reduction step can be performed using $O(en^\omega \log^2 n)$ operations in $\mathbb{Z}/p^e\mathbb{Z}$, and will succeed with probability at least $1/2$.*

Proof. Choose a Galois Ring extension of degree $d \geq 3 + \log n$. Then the image of $\text{GR}(p^e, d) \bmod p$ will have $p^d \geq 4n$ elements. We can use [Shoup, 1994] to find a polynomial f of degree d which is irreducible over $\mathbb{Z}/p\mathbb{Z}$.

Now choose the entries of X, Y from $\text{GR}(p^e, d)$ and perform all arithmetic over this ring. As before, we apply the Schwartz-Zippel lemma modulo p to get

$$\Pr[\det W \equiv 0 \pmod{p, f}] \leq \frac{2b}{|\text{GF}(p^d)|} \leq \frac{2n}{4n} = \frac{1}{2}.$$

The arithmetic cost will increase by a factor equal to cost of performing polynomial arithmetic with degree at most d . That is, $\log^2 n$ if we use naive polynomial arithmetic. \square

Algorithm 4.2. Given a matrix $A = (a_{ij}) \in \mathbf{M}_n(\mathbb{Z}/p^e\mathbb{Z})$, this algorithm returns the invariant factors of A .

1. Let $s_0 = 1$, $s_i = 0$, and $r_i = 0$ for all $i \in [1, e]$.
2. For $\ell = 1$ to e :
 - (a) Let $\text{gcd}(A) = \text{gcd} \{a_{ij} : i, j \in [1, n]\}$.
 - (b) Let $s_\ell = s_{\ell-1} \cdot \text{gcd}(A)$.
 - (c) Let $A = A / \text{gcd}(A)$.

- (d) Let $r_\ell = \text{rank}(A)$ over $\mathbb{Z}/p\mathbb{Z}$.
 - (e) Construct two $n \times r_\ell$ matrices X, Y with uniform random entries from $[0, p^e)$, and compute $W = Y^T A X$. Repeat until W is non-singular over $\mathbb{Z}/p\mathbb{Z}$.
 - (f) Let $A = A - A X W^{-1} Y^T A$.
 - (g) If $A = 0$ then **break**.
3. Return $\underbrace{s_1, \dots, s_1}_{r_1}, \dots, \underbrace{s_\ell, \dots, s_\ell}_{r_\ell}, \underbrace{0, \dots, 0}_{n - (r_1 + \dots + r_\ell)}$.

Theorem 4.5. *Algorithm 4.2 is a correct Las Vegas algorithm. The expected cost is $O(e^2 n^\omega)$ operations in $\mathbb{Z}/p^e\mathbb{Z}$, and storage of $O(n^2)$ elements from $\mathbb{Z}/p^e\mathbb{Z}$.*

Proof. Correctness follows from the previous discussion. We can certify the success of step 2e by way of checking the invertibility of W . The success probability is at least $1/2$, and the expected number of trials is constant. So the algorithm is Las Vegas.

There are at most e iterations. The cost of each iteration is dominated by the cost of matrix multiplications and by inversion of W . Both can be done in $O(en^\omega)$ operations in $\mathbb{Z}/p^e\mathbb{Z}$. □

It is an interesting future work to derandomize this algorithm. In step 2d, we could compute the LDU decomposition of A over $\mathbb{Z}/p\mathbb{Z}$ and then use L, U with arbitrary lifting into $\mathbb{Z}/p^e\mathbb{Z}$ to construct X, Y . The complexity of the algorithm remains the same as above.

By contrast, [Storjohann, 2000, Proposition 7.16] gives a deterministic algorithm for computing Smith normal form over arbitrary principal ideal rings which requires $O(n^\omega \log r)$ ring operations (where r is the rank).

4.4 Nullspace Sampling

The block rank reduction formula gives rise to an interesting nullspace sampling algorithm over finite fields. Let A be an $n \times n$ matrix over a field F , and $\text{rank } A = r$. Let X, Y be

$n \times r$ matrices such that $W = Y^T A X \in \text{GL}_r(\mathbb{F})$. Then $A - A X W^{-1} Y^T A = 0$ because it has rank 0. If we let $K = I - X W^{-1} Y^T A$ then $AK = 0$. We can show that K gives a nullspace basis for A .

Lemma 4.10. $\text{im}K = \ker A$.

Proof. If $v \in \text{im}K$ then there exists a vector u such that $v = Ku$. We have $Av = AKu = 0$. So $\text{im}K \subseteq \ker A$. Conversely, if $Av = 0$ then $Kv = (I - XW^{-1}Y^T A)v = v - XW^{-1}Y^T Av = v$. So $\ker A \subseteq \text{im}K$. \square

We can then sample the nullspace of A by sampling the column space of K . We will use the same trace technique from Lemma 3.6 to ensure that the samples are over \mathbb{F}^n .

Algorithm 4.3. Given a matrix $A \in \text{M}_n(\mathbb{F})$, return a vector $v \in \mathbb{F}^n$ which is a uniform random sample from the nullspace of A .

1. Compute $r = \text{rank}(A)$.
2. If $|\mathbb{F}| \geq 4n$, then let $\mathbb{K} = \mathbb{F}$. If $|\mathbb{F}| < 4n$, then find an irreducible polynomial $f \in \mathbb{F}[x]$ of degree at least $\lceil \log_{|\mathbb{F}|} 4n \rceil$. Let $\mathbb{K} = \mathbb{F}[x]/(f)$. Let $\text{Tr} : \mathbb{K} \rightarrow \mathbb{F}$ be the field trace.
3. Construct two $n \times r$ matrices X, Y with uniform random entries from \mathbb{K} , and compute $W = Y^T A X$. Repeat until W is non-singular.
4. Construct a vector $w \in \mathbb{F}^n$ with entries chosen uniformly at random from \mathbb{F} .
5. Let $v = (I - XW^{-1}Y^T A)w$.
6. If v has components from $\mathbb{K} \setminus \mathbb{F}$ then let $v = \text{Tr}(v)$.
7. Return v .

Lemma 4.11. *Algorithm 4.3 is a correct Las Vegas algorithm. The expected cost is $O^\sim(n^\omega)$ operations in \mathbb{F} . The space complexity is $O^\sim(n^2)$ elements from \mathbb{F} .*

Proof. Arithmetic over the extension field introduces logarithmic factors in the complexity. The cost for computing $K = I - XW^{-1}Y^T A$ is $O(n^\omega)$ operations in \mathbb{F} , and $O(n^2)$ operations for computing the sample v . It follows from our results on randomized block rank reduction that this algorithm has a success probability of at least $1/2$. We can verify that rank reduction succeeded whenever W is invertible, and we can verify that v is a nullspace vector by computing Av . So the algorithm is Las Vegas. The expected number of steps is constant because the success probability is at least $1/2$. The storage cost is dominated by K which has quadratic elements from \mathbb{K} . \square

When A is sparse, using dense matrices X, Y and computing W^{-1} will introduce fill-in and therefore is not favourable. It remains open to find a sparse (or structured) choice for X, Y which gives a provably good success probability. We will outline a possible choice for X and Y , namely, using Krylov matrices.

Let A be a black-box matrix and let the cost of $v \mapsto Av$ be μ operations in \mathbb{F} . Let $x, y \in \mathbb{F}^n$ have entries chosen uniformly at random from \mathbb{F} . Construct X, Y as $n \times r$ Krylov matrices:

$$X = \begin{bmatrix} x & Ax & \cdots & A^{r-1}x \end{bmatrix}, \quad Y^T = \begin{bmatrix} y^T \\ y^T A \\ \vdots \\ y^T A^{r-1} \end{bmatrix}, \quad W = Y^T A X. \quad (4.5)$$

Then $W \in M_r(\mathbb{F})$ is a Hankel matrix:

$$W = \begin{bmatrix} y^T A x & y^T A^2 x & \cdots & y^T A^r x \\ y^T A^2 x & \ddots & & y^T A^{r+1} x \\ \vdots & & \ddots & \vdots \\ y^T A^r x & y^T A^{r+1} x & \cdots & y^T A^{2r-1} x \end{bmatrix}. \quad (4.6)$$

This construction is similar to the preconditioner matrices of [Eberly et al., 2007] in a scalar setting. The black-box for $K = I - XW^{-1}Y^T A$ can be constructed by lazily storing x, y and W^{-1} . If we precondition W properly then we can construct a black-box for its inverse in quasilinear time [Gohberg and Fel'dman, 1974, Labahn et al., 1990, Labahn and Shalom, 1992].

The cost of computing a random nullspace vector in this setup is 2μ operations to multiply a vector by A , and $2r\mu$ to multiply by X, Y^T , and $O(r)$ to multiply a vector by W^{-1} . Thus the asymptotic cost of computing each nullspace sample is $O(n\mu)$ operations in \mathbb{F} , and the space complexity is $O(n)$ elements from \mathbb{F} .

Chapter 5

The Eigenvalues and the Invariant Factors

In this chapter we study the relationship between the invariant factors of a matrix and its eigenvalues when viewed p -adically. Our motivation is both to understand the fundamental connection, and to design efficient algorithms to compute the Smith normal form of sparse integer matrices. Understanding the p -adic structure of the spectrum is a step towards reducing the computation of the Smith normal form to computing the characteristic polynomial or the ranks modulo prime powers, which would be arguably efficient.

Concretely, conditions are established under which the p -adic valuations of the invariant factors of an integer matrix are equal to the p -adic valuations of the eigenvalues. It is then shown that this correspondence is the typical case for “most” matrices. Density counts are given for when this property holds, as well as transformations to this typical case. The results of this chapter appeared in [Elsheikh and Giesbrecht, 2015].

5.1 Introduction

Let A be an $n \times n$ integer matrix whose rank is r , its invariant factors are s_1, \dots, s_r and its determinantal divisors are $\Delta_1, \dots, \Delta_r$. *A priori* the invariant factors of a matrix and the

eigenvalues of a matrix would seem to be rather different invariants. The former is related to the \mathbb{Z} -lattice structure of A and the latter to the geometry of the linear map. We show that, in fact, they are “usually” in one to one correspondence with respect to their p -adic valuations at a prime p . We demonstrate a simple sufficient condition under which this holds for any integer matrix, and provide bounds on the density of matrices for which it holds.

Throughout we will work with the p -adic numbers. There are a few equivalent ways to define the p -adic numbers. See [Koblitz, 1984] or [Gouvêa, 1997] for a full treatment on this subject. Hereby we give an explicit construction using p -adic expansions. Let p be a prime number. Any p -adic integer a can be uniquely written as $a = \sum_{i \geq N} a_i p^i$ where $N \geq 0$ and $a_i \in [0, p)$ [Gouvêa, 1997, Corollary 3.3.11]. The set \mathbb{Z}_p denotes the p -adic integers. A notable property of \mathbb{Z}_p is that it is a principal ideal ring and hence every matrix over this ring admits a Smith normal form. Any p -adic rational a can be uniquely written as $a = \sum_{i \geq N} a_i p^i$ where $a_i \in [0, p)$ and N is a possibly negative integer [Gouvêa, 1997, Corollary 3.3.12]. The set of p -adic rationals is denoted by \mathbb{Q}_p . It is easy to show that $\mathbb{Z} \subset \mathbb{Z}_p$ and $\mathbb{Q} \subset \mathbb{Q}_p$.

Let $v_p(a) \in \mathbb{N} \cup \{\infty\}$ be the p -adic order or p -adic valuation. For any $a \in \mathbb{Z}_p$, v_p is the number of times p divides a exactly, where $v_p(0)$ is taken to be ∞ . The valuation is extended to \mathbb{Q}_p by letting $v_p(a/b) = v_p(a) - v_p(b)$ for $a, b \in \mathbb{Z}_p$.

Many authors in computer algebra use \mathbb{Z}_p to denote the finite field with p elements. We remind the reader that we do not use this notation here since it is ambiguous. We reserve the symbol \mathbb{Z}_p for the p -adic integers and use $\mathbb{Z}/p\mathbb{Z}$ to denote the finite field with p elements.

The eigenvalues of an integer matrix A are the roots of its characteristic polynomial which has a natural image in $\mathbb{Z}_p[x]$ since \mathbb{Z}_p contains \mathbb{Z} . Thus, the eigenvalues of A can naturally be viewed as p -adic algebraic integers in a finite-degree algebraic extension field \mathbb{K}_p over \mathbb{Q}_p [Gouvêa, 1997, Proposition 5.4.5 (v)]. We make use of this fact because we do not view the eigenvalues as complex numbers, but rather as algebraic integers which allows us to relate their p -adic valuation to the powers of p dividing the invariant factors.

Example 5.1. Consider the matrix

$$A = \begin{bmatrix} 3 & -1 & 3 \\ 9 & -10 & 0 \\ 3 & 0 & 3 \end{bmatrix} = \overbrace{\begin{bmatrix} 1 & -1 & 0 \\ 1 & 0 & 1 \\ 0 & -1 & 0 \end{bmatrix}}^U \overbrace{\begin{bmatrix} 1 & & \\ & 3 & \\ & & 3^2 \end{bmatrix}}^S \overbrace{\begin{bmatrix} 0 & -1 & 0 \\ -1 & 0 & -1 \\ 1 & -1 & 0 \end{bmatrix}}^V,$$

for unimodular U, V and the Smith normal form S of A . Now consider the eigenvalues of A , which are roots of the characteristic polynomial

$$f = \det(xI - A) = x^3 + 4x^2 - 51x - 27 \in \mathbb{Z}[x].$$

We find it has three distinct roots in the 3-adics \mathbb{Z}_3 :

$$\lambda_1 = -1 - 3^3 - 3^4 - 3^5 - 3^6 - 3^8 + 3^9 + O(3^{10}),$$

$$\lambda_2 = -3 - 3^2 - 3^3 - 3^4 + 3^6 - 3^8 - 3^9 + O(3^{10}),$$

$$\lambda_3 = 3^2 - 3^3 - 3^5 + 3^6 - 3^8 + 3^9 + O(3^{10}).$$

In this example we see that $v_3(\lambda_1) = 0$, $v_3(\lambda_2) = 1$ and $v_3(\lambda_3) = 2$. We see that the diagonal entries of the Smith normal form have precisely the same p -adic valuations as the eigenvalues of A . □

In order to show the correspondence between the eigenvalues and the invariant factors, we need to extend the definition of the valuation v_p to the eigenvalues (more generally, to the elements of \mathbb{K}_p). If an element $a \in \mathbb{K}_p$ has a minimal polynomial $x^{d_a} + a_{d_a-1}x^{d_a-1} + \dots + a_0 \in \mathbb{Q}_p[x]$, then the valuation is uniquely given by $v_p(a) = (1/d_a)v_p(a_0)$. See [Koblitz, 1984, §3, pp. 66]. The image of the extended v_p is \mathbb{Q} , and its restriction to \mathbb{Q}_p agrees with the earlier definition of v_p on \mathbb{Q}_p . The valuation of a non-zero eigenvalue $v_p(\lambda_i)$ is independent of the choice of \mathbb{K}_p , since it only depends on the minimal polynomial of λ_i over \mathbb{Q}_p . In particular, the set of minimal polynomials of the non-zero eigenvalues is precisely the set of irreducible factors of the characteristic polynomial of the matrix over \mathbb{Q}_p regardless of the field extension. We have shown the following.

Lemma 5.1. Given an integer matrix and a prime p , $v_p(\lambda_1), \dots, v_p(\lambda_n)$ are invariants, and independent of the p -adic extension chosen to contain the eigenvalues.

In light of the above, we will treat integer matrices and their eigenvalues as being naturally embedded in \mathbb{Z}_p , \mathbb{Q}_p or \mathbb{K}_p as appropriate, under the p -adic valuation v_p .

It should be noted that the correspondence between the valuations of the eigenvalues and the invariant factors does not hold for all matrices.

Example 5.2. *Let*

$$A = \begin{bmatrix} 37 & 192 & 180 & 369 \\ 55 & 268 & 198 & 531 \\ 163 & 758 & 442 & 1539 \\ 198 & 908 & 486 & 1858 \end{bmatrix},$$

which has the following Smith normal form:

$$A = \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix} \begin{bmatrix} 1 & & & \\ & 2 & & \\ & & 2 & \\ & & & 4 \end{bmatrix} \begin{bmatrix} 163 & 758 & 442 & 1539 \\ 99 & 454 & 243 & 929 \\ -54 & -245 & -122 & -504 \\ -54 & -246 & -126 & -505 \end{bmatrix}.$$

The characteristic polynomial of A is

$$f = x^4 - 2605x^3 + 39504x^2 + 40952x + 16 \in \mathbb{Z}[x],$$

which factors over \mathbb{Q}_2 into

$$x + (1 + 2^2 + 2^3 + O(2^5)) \in \mathbb{Z}_2[x],$$

and the irreducible factor

$$x^3 + (2^3 + O(2^5))x^2 + (2^3 + 2^4 + O(2^5))x + (2^4 + O(2^5)) \in \mathbb{Z}_2[x].$$

Using a computer algebra system (or Newton slopes as in Fact 5.1 below) we find that the 2-adic valuations of the eigenvalues are $[0, 4/3, 4/3, 4/3]$. But the 2-adic valuations of the invariant factors of A are $[0, 1, 1, 2]$. The eigenvalues and the invariant factors are not in 1-1 correspondence w.r.t. their p -adic valuation. \square

In the remainder of this chapter we explore the conditions under which this correspondence between the p -adic valuation of the invariant factors and the eigenvalues occurs, and show that it is, in fact, the “typical” case, i.e., it holds for “most” matrices.

5.1.1 The p -adic Correspondence

We first define two important matrix properties for our purposes.

Definition 5.1. *Let $A \in \mathbb{Z}^{n \times n}$ be of rank r and p be any prime. Assume*

- (i) *The matrix A has the Smith normal form $S = \text{diag}(s_1, \dots, s_r, 0, \dots, 0)$ over \mathbb{Z} , so that $\Delta_i = s_1 \cdots s_i$ is the i th determinantal divisor of A , for all $1 \leq i \leq r$;*
- (ii) *The matrix A has non-zero eigenvalues (with multiplicity) $\lambda_1, \dots, \lambda_r$ in a finite-degree extension \mathbb{K}_p over \mathbb{Q}_p , and assume that $v_p(\lambda_1) \leq \dots \leq v_p(\lambda_r)$;*
- (iii) *The matrix A has characteristic polynomial $f = x^n + f_1 x^{n-1} + \dots + f_r x^{n-r} \in \mathbb{Z}[x]$ (note the reversed indexing).*

We say A is p -characterized if and only if $v_p(f_i) = v_p(\Delta_i)$ for all $i \in [1, r]$. We say A is p -correspondent if and only if $v_p(s_i) = v_p(\lambda_i)$ for all $i \in [1, r]$.

Note that if A is p -correspondent, then the valuations of the eigenvalues are non-negative integers (since $v_p(s_i) \geq 0$). Our main goal is to study the notion of p -correspondence; that is the relationship between the spectrum and the invariant factors. The notion of p -characterization is an auxiliary definition used throughout our proofs. In fact, we will discuss the caveats of our results regarding “small” primes, which are in large part an artifact of our proofs (which uses the Schwartz-Zippel lemma) and the notion of p -characterization.

We shall see that if A is p -characterized then A is p -correspondent. Of course, not all matrices are p -correspondent at any particular prime p , but it is generally possible to transform a matrix to a p -correspondent one.

In [§5.3](#) we establish that “most” matrices are p -correspondent. We will consider the density in each equivalence class defined by a given Smith normal form.

5.1.2 Previous Work

The work of [Newman and Thompson, 1991] studies matrices with algebraic integer entries. They study, among other things, links between the eigenvalues and the invariant factors.

In their setup, matrices have entries from a ring R which is not necessarily a PID. In this case many of the properties about the Smith normal form are not necessarily applicable. They overcome this by embedding the matrix in a ring extension such that the required properties hold. Let A be an $n \times n$ matrix over R . Let the eigenvalues of A be $\lambda_1, \dots, \lambda_n$ (in some extension) and the invariant factors of A be s_1, \dots, s_n . Then Theorem 6 of §8 states that for all $k \in [1, n]$: $s_1 \cdots s_k \mid \lambda_{i_1} \cdots \lambda_{i_k}$ where $I = \{i_1, \dots, i_k\}$ is any subset of $[1, n]$. In other words, the k th determinantal divisor, Δ_k , divides the products of any subset of size k of the eigenvalues. However, this does not show any correspondence between the individual eigenvalues and the invariant factors.

Rushanan [Rushanan, 1995] studies the relationship between the spectrum and the Smith normal form of non-singular integer matrices with integer eigenvalues. However his results are valid for any PID. We note the following theorem since it is the most relevant to our result. Given an integer matrix A define $G(A)$ to be the factor \mathbb{Z} -module $G(A) = \mathbb{Z}^n / \text{RowSpace}(A)$. The finite part of this module is given by the direct sum $\mathbb{Z}/s_1\mathbb{Z} \oplus \cdots \oplus \mathbb{Z}/s_r\mathbb{Z}$. Theorem 4 in [Rushanan, 1995] states that if $\lambda \in \mathbb{Z}$ is an eigenvalue of A with multiplicity m , then $(\mathbb{Z}/\lambda\mathbb{Z})^m$ is isomorphic to a subgroup of $G(A)$. This indeed addresses the association between an eigenvalue of multiplicity m and an invariant factor of the same multiplicity. However this result gives a divisibility relationship and it does not give conditions for when the equality of valuation is exact.

Finally, [Lorenzini, 2008] studies the Smith normal form of Laplacian matrices of graphs. For a graph G let L denote its Laplacian matrix. The group $\mathbb{Z}^n / \text{im}(L)$ can be computed using the Smith normal form of L since $\mathbb{Z}^n / \text{im}(L)$ is isomorphic to $\mathbb{Z}^r \oplus \mathbb{Z}/s_1\mathbb{Z} \oplus \cdots \oplus \mathbb{Z}/s_{n-1}\mathbb{Z}$. Let $\phi(G) = \mathbb{Z}/s_1\mathbb{Z} \oplus \cdots \oplus \mathbb{Z}/s_{n-1}\mathbb{Z}$ be the torsion part of the group $\mathbb{Z}^n / \text{im}(L)$. An interesting connection between the Smith normal form of L and the properties of G is that $|\phi(G)|$ is a graph invariant, namely, the number of spanning trees of G . Clearly this invariant is also given by $s_1 \cdots s_{n-1} = \Delta_{n-1}$ of the matrix L . The cutoff at $n - 1$ is due to the fact that Laplacian matrices have rank $n - 1$. The most relevant result of his work is the following. If λ is an eigenvalue of L with multiplicity m and $\nu = v_p(\lambda)$ for some prime p , then $\phi(G)$ contains a subgroup isomorphic to $(\mathbb{Z}/p^\nu\mathbb{Z})^m$. Thus this result gives a

correspondence between an eigenvalue λ of a given multiplicity and products of subsets of s_i forming a subgroup of order p^{ν_m} .

5.2 Establishing p -Correspondence

In this section we will prove that all p -characterized matrices are p -correspondent. First recall that the coefficients of the characteristic polynomial $f = x^n + \sum_{1 \leq i \leq n} f_i x^{n-i} \in \mathbb{Z}[x]$ of a matrix $A \in \mathbb{Z}^{n \times n}$ are related to the minors of A in the following manner.

Recall that $A \binom{\sigma}{\tau}$ is the *minor* of A selected by the sets of indices σ and τ . It is well-known that for all $i \in [1, n]$,

$$f_i = (-1)^i \sum_{\sigma \in \mathcal{C}_i^n} A \binom{\sigma}{\sigma}. \quad (5.1)$$

Since Δ_i divides all $i \times i$ minors, we have $\Delta_i \mid f_i$, i.e., $v_p(f_i) \geq v_p(\Delta_i)$. Moreover, if A has rank r we have $f_{r+1} = f_{r+2} = \dots = f_n = 0$.

We will use the so-called *Newton polygon* of the characteristic polynomial of A . Let f be the polynomial $x^n + \sum_{1 \leq i \leq n} f_i x^{n-i}$.

Definition 5.2. *The Newton polygon of f , denoted by $\text{NP}(f)$, is the lower convex hull of the following points in \mathbb{R}^2 : $\{(0, 0), (1, v_p(f_1)), \dots, (n, v_p(f_n))\}$.*

The polygon is represented by a list of points $(x_1, y_1), \dots, (x_k, y_k) \in \mathbb{R}^2$ with $x_1 < x_2 < \dots < x_k$. For each segment of $\text{NP}(f)$ connecting two adjacent points (x_{i-1}, y_{i-1}) and (x_i, y_i) , the *slope* of the segment is $m_i = (y_i - y_{i-1}) / (x_i - x_{i-1})$ and the *length* of the segment is the length of its projection onto the x -axis, taken as $\ell_i = x_i - x_{i-1}$. An important use of Newton polygon is the following.

Fact 5.1 ([Koblitz, 1984], §IV.3, Lemma 4). *Let $f = x^n + f_1 x^{n-1} + \dots + f_n \in \mathbb{Z}_p[x]$ and $f_n \neq 0$. Let the roots of f (counting multiplicity) be $\lambda_1, \dots, \lambda_n$ in an extension \mathbb{K}_p over \mathbb{Q}_p . If the Newton polygon of f has slopes m_1, \dots, m_k and lengths ℓ_1, \dots, ℓ_k as above, then for each $1 \leq j \leq k$, f has exactly ℓ_j roots $\lambda \in \mathbb{K}_p$ whose valuation is $v_p(\lambda) = m_j$.*

We now have all the tools to prove the following.

Theorem 5.1. *Let $A \in \mathbb{Z}^{n \times n}$ and p be a prime. If A is p -characterized then A is p -correspondent.*

Proof. Assume that A is p -characterized with rank r and characteristic polynomial $f = \sum_{0 \leq i \leq r} f_i x^{n-i} \in \mathbb{Z}[x]$, and A has the Smith normal form $S = \text{diag}(s_1, \dots, s_r, 0, \dots, 0) \in \mathbb{Z}^{n \times n}$. Also, assume that the p -adic valuations of the invariant factors s_1, \dots, s_r have multiplicities r_0, \dots, r_{e-1} as follows:

$$(v_p(s_1), \dots, v_p(s_r)) = (\underbrace{0, \dots, 0}_{r_0}, \underbrace{1, \dots, 1}_{r_1}, \dots, \underbrace{e-1, \dots, e-1}_{r_{e-1}}),$$

where $e = v_p(s_r) + 1$. Since A is p -characterized, by definition we have

$$v_p(f_i) = v_p(\Delta_i) = \sum_{1 \leq j \leq i} v_p(s_j),$$

for all $1 \leq i \leq r$. For notational convenience, define m_i as

$$m_i = v_p(\Delta_{r_0+r_1+\dots+r_i}) = r_1 + 2r_2 + \dots + i \cdot r_i.$$

Grouping the non-zero coefficients of f by their p -adic valuation we get

$$\begin{aligned} & (v_p(f_1), \dots, v_p(f_r)) \\ &= \left(\underbrace{0, \dots, 0}_{r_0}, \underbrace{1, 2, 3, \dots, r_1}_{r_1}, \underbrace{m_1 + 2, m_1 + 4, \dots, m_1 + 2r_2}_{r_2}, \right. \\ & \quad \left. \dots, \underbrace{m_{e-2} + (e-1), m_{e-2} + 2(e-1), \dots, m_{e-2} + r_{e-1}(e-1)}_{r_{e-1}} \right). \end{aligned}$$

$\text{NP}(f)$ is easily seen to consist of e segments, where segment i has slope i , and length r_i , for $0 \leq i < e$ (a segment i may have length 0 if $r_i = 0$). Thus, by Fact 5.1, f has r_i roots λ with $v_p(\lambda) = i$. This accounts for all the non-zero roots of f , since $r_0 + r_1 + \dots + r_{e-1} = \text{rank}(A)$. Since these roots are the non-zero eigenvalues of A , we immediately see that A is p -correspondent. \square

It should be noted that the converse of Theorem 5.1 is not necessarily true. The matrix in the following example is p -correspondent but not p -characterized.

Example 5.3. *The invariant factors of*

$$A = \begin{bmatrix} -20 & -2 & 81 & -388 \\ 18 & -6 & -84 & 375 \\ 7 & 34 & 3 & 41 \\ 13004 & -11695 & -64944 & 289315 \end{bmatrix},$$


are $[1, 3, 3, 9]$, and the 3-adic eigenvalues are:

$$2 + O(3), 2 \cdot 3 + O(3^3), 3 + O(3^2), 3^2 + O(3^3).$$

However, the 3-adic valuations of the determinantal divisors are $[0, 1, 2, 4]$ and the characteristic polynomial over $\mathbb{Z}_3[x]$ is:

$$x^4 + (1 + O(3))x^3 + (2 \cdot 3^2 + O(3^6))x^2 + (2 \cdot 3^2 + O(3^3))x + (3^4 + O(3)).$$

This is due to the fact that the Newton polygon of A is the convex hull of the segments defined by the coefficients of characteristic polynomial.



While the coefficients of the characteristic polynomial (points in left figure) do not correspond to the 3-adic valuations of the determinantal divisors, their lower convex cover (segments in right figure) corresponds to the 3-adic valuations of the invariant factors with slopes: 0, 1 (twice), and 2. □

We now prove two simple lemmas establishing p -correspondence under unimodular equivalence transformations and under similarity transformations.

Lemma 5.2. *Let $A \in \mathbb{Z}^{n \times n}$ and p be any prime. There exists an equivalence transformation $P, Q \in \mathrm{GL}_n(\mathbb{Z})$ such that PAQ is p -correspondent.*

Proof. Simply choose $P, Q \in \mathrm{GL}_n(\mathbb{Z})$ such that PAQ is in the Smith normal form $S = \mathrm{diag}(s_1, \dots, s_r, 0, \dots, 0)$. Then the eigenvalues of PAQ are s_1, \dots, s_r . \square

Lemma 5.3. *Let $A \in \mathbb{Z}^{n \times n}$ be non-singular, p be any prime. There exists a similarity transformation U with entries in an extension \mathbb{K}_p over \mathbb{Q}_p such that $U^{-1}AU$ is p -correspondent.*

Proof. Choose \mathbb{K}_p to be a splitting field of the minimal polynomial of A . It is well-known that any matrix over the splitting field of its characteristic polynomial (\mathbb{K}_p in our case) is similar to a matrix $J \in \mathbb{K}_p^{n \times n}$ in Jordan form [Meyer, 2000]. That is, there exists an invertible $W \in \mathbb{K}_p^{n \times n}$ such that $W^{-1}AW = \mathrm{diag}(J_1, \dots, J_\ell)$ where

$$J_i = \begin{bmatrix} \mu_i & 1 & & \\ & \ddots & \ddots & \\ & & \ddots & 1 \\ & & & \mu_i \end{bmatrix},$$

for some (not necessarily unique) eigenvalue $\mu_i \in \mathbb{K}_p$ of A , and J_i has dimensions $k_i \times k_i$. However, we can choose an alternative Jordan block \widehat{J}_i , similar to J_i , by applying the similarity transformation $\mathrm{diag}(1, 1/\mu_i, \dots, 1/\mu_i^{k_i-1})$ to J_i to get

$$\widehat{J}_i = \begin{bmatrix} \mu_i & \mu_i & & \\ & \ddots & \ddots & \\ & & \ddots & \mu_i \\ & & & \mu_i \end{bmatrix}.$$

The Smith normal form of \widehat{J}_i can be obtained as follows. Subtract the first column from the second column. Then subtract the second column from the third, and so forth. The resulting matrix is $\mathrm{diag}(\mu_i, \dots, \mu_i)$ which is in the Smith normal form when viewed as a matrix over the ring of algebraic integers \mathcal{O}_p .

Combining together the different Jordan blocks to form an alternative Jordan form \widehat{J} for A , we see that \widehat{J} is p -correspondent, and similar to A , as required. \square

If A is singular, Lemma 5.3 may not hold. Consider for example

$$A = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix},$$

whose only eigenvalue is zero, with multiplicity two. However, this matrix has rank one, and so one of the invariant factors must always be non-zero. This is also the case for any matrix similar to A .

5.3 Density of p -Characterized Matrices

In this section we show that most matrices which are unimodularly equivalent to a matrix $A \in \mathbb{Z}^{n \times n}$, are p -characterized (and hence p -correspondent) when p is large compared to n . The main tool is the following lemma.

Lemma 5.4. *Let $A \in \mathbb{Z}^{n \times n}$ have rank r . Let \mathcal{U}, \mathcal{V} be $n \times n$ matrices whose $2n^2$ entries are algebraically independent indeterminates u_{ij} and v_{ij} respectively. Let g_k be the coefficient of x^{n-k} in the characteristic polynomial of $B = \mathcal{U}A\mathcal{V}$. Then for all $k \in [1, r]$, g_k is a polynomial of total degree $2k$ and the content of g_k is Δ_k , the k th determinantal divisor of A .*

Proof. Assume throughout that $k \leq r$. Using the Cauchy-Binet formula,

$$\begin{aligned} g_k &= (-1)^k \sum_{\sigma \in \mathcal{C}_k^n} B \begin{pmatrix} \sigma \\ \sigma \end{pmatrix} = (-1)^k \sum_{\sigma, \tau, \omega \in \mathcal{C}_k^n} \mathcal{U} \begin{pmatrix} \sigma \\ \tau \end{pmatrix} A \begin{pmatrix} \tau \\ \omega \end{pmatrix} \mathcal{V} \begin{pmatrix} \omega \\ \sigma \end{pmatrix} \\ &= (-1)^k \sum_{\tau, \omega \in \mathcal{C}_k^n} A \begin{pmatrix} \tau \\ \omega \end{pmatrix} \Upsilon_{\tau, \omega}, \quad \text{where } \Upsilon_{\tau, \omega} = \sum_{\sigma \in \mathcal{C}_k^n} \mathcal{U} \begin{pmatrix} \sigma \\ \tau \end{pmatrix} \mathcal{V} \begin{pmatrix} \omega \\ \sigma \end{pmatrix}. \end{aligned} \tag{5.2}$$

We first show that $\Upsilon_{\tau, \omega}$ has content 1. By Leibniz's determinant expansion on the minor of \mathcal{U} selected by the first k rows, and the columns given by the indices in $\tau \in \mathcal{C}_k^n$, we have

$$\begin{aligned} \mathcal{U} \begin{pmatrix} (1, 2, \dots, k) \\ \tau \end{pmatrix} &= \sum_{\mu \in S_k} \text{sgn}(\mu) \prod_{1 \leq i \leq k} u_{\mu_i, \tau_i} \\ &= u_{1, \tau_1} u_{2, \tau_2} \cdots u_{k, \tau_k} + \sum_{\substack{\mu \in S_k \\ \mu \neq \text{id}}} \prod_{1 \leq i \leq k} \text{sgn}(\mu) u_{\mu_i, \tau_i}, \end{aligned}$$

where S_k is the symmetric group of permutations of k symbols, (μ_1, \dots, μ_k) is a permutation of $\{1, \dots, k\}$ and $\text{id} = (1, \dots, k)$ is the identity permutation. Similarly,

$$\mathcal{V}\left(\begin{matrix} \omega \\ (1, \dots, k) \end{matrix}\right) = v_{\omega_1,1} v_{\omega_2,2} \cdots v_{\omega_k,k} + \sum_{\substack{\mu \in S_k \\ \mu \neq \text{id}}} \prod_{1 \leq i \leq k} \text{sgn}(\mu) v_{\omega_i, \mu_i}.$$

We observe that $\mathcal{U}\left(\begin{matrix} (1,2,\dots,k) \\ \tau \end{matrix}\right)$ contains the distinguished monomial $u_{1,\tau_1} \cdots u_{k,\tau_k}$ which is not found in any of the remaining terms of the expansion of $\mathcal{U}\left(\begin{matrix} (1,\dots,k) \\ \tau \end{matrix}\right)$ and hence has coefficient 1 (since each permutation μ is distinct), and is not found in the expansion of $\mathcal{U}\left(\begin{matrix} \sigma' \\ \tau' \end{matrix}\right)$ for any other $\sigma', \tau' \in \mathcal{C}_k^n$ (since the variables in the term allow us to identify the subsets σ' and τ'). Similarly, $\mathcal{V}\left(\begin{matrix} \omega \\ (1,\dots,k) \end{matrix}\right)$ contains the distinguished monomial $v_{\omega_1,1} \cdots v_{\omega_k,k}$ with coefficient 1 which is not found in $\mathcal{V}\left(\begin{matrix} \omega' \\ \sigma' \end{matrix}\right)$ for any other $\omega', \sigma' \in \mathcal{C}_k^n$.

Thus, for every choice of τ, ω , the polynomial $\Upsilon_{\tau,\omega}$ has a monic distinguished term $u_{1,\tau_1} \cdots u_{k,\tau_k} v_{\omega_1,1} \cdots v_{\omega_k,k}$ not appearing in $\Upsilon_{\tau',\omega'}$ for any other $\tau', \omega' \in \mathcal{C}_k^n$. Thus $\Upsilon_{\tau,\omega}$ is non-zero, has degree $2k$, and has content 1.

It follows immediately that g_k has degree $2k$ and content which is the GCD of all $A\left(\begin{matrix} \tau \\ \omega \end{matrix}\right)$, which is precisely Δ_k . \square

A related result is found in [Giesbrecht, 2001, Theorem 1.4]. A similar technique is used in [Kaltofen and Saunders, 1991, Theorem 2], where a minor with symbolic entries is explicitly selected and shown to be lexicographically unique and hence the resulting polynomial, e.g. g_k , is shown to be non-zero.

The following lemma is used to count the number of matrices with a given property. While this result resembles the Schwartz-Zippel lemma [Zippel, 1979, Schwartz, 1980], similar statements can be traced to earlier literature, for example in [Kasami et al., 1968].

Lemma 5.5. *Let p be a prime, $\ell \geq 1$ be an integer, and $g \in \mathbb{Z}[x_1, \dots, x_n]$ be a non-zero polynomial of total degree k . Then the number of points $\alpha = (\alpha_1, \dots, \alpha_n) \in [0, \ell p)^n$ for which $g(\alpha) \equiv 0 \pmod{p}$ is at most $\ell^n k p^{n-1}$.*

Proof. As a shorthand, we call $\alpha \in \mathbb{Z}^n$ a p -root if $f(\alpha) \equiv 0 \pmod{p}$. For $\ell = 1$ the statement of the lemma becomes exactly Corollary 1 of [Schwartz, 1980]: the number of p -roots in $[0, p)^n$ is at most $k p^{n-1}$.

Now assume $\ell > 1$. Every p -root $b \in [0, \ell p)^n$ can be written with component-wise Euclidean division as $(b_1, \dots, b_n) = (\alpha_1 + r_1 p, \dots, \alpha_n + r_n p) = \alpha + (r_1 p, \dots, r_n p)$ where $r_i \in [0, \ell - 1)$ and $\alpha = (\alpha_1, \dots, \alpha_n) \in [0, p)^n$. Then α must be a p -root because $b \equiv \alpha \pmod{p}$. Conversely if $\alpha = (\alpha_1, \dots, \alpha_n) \in [0, p)^n$ is a p -root, then $(\alpha_1 + r_1 p, \dots, \alpha_n + r_n p) \in [0, \ell p)^n$ is a p -root for all the ℓ^n possible values of $(r_1, \dots, r_n) \in [0, \ell)^n$. Thus there are at most $\ell^n \cdot k p^{n-1}$ p -roots in $[0, \ell p)^n$. \square

Lemma 5.6. *Let $A \in \mathbb{Z}^{n \times n}$, $\epsilon > 0$, p a prime greater than $(n^2 + 3n)/\epsilon$, and N a non-zero integer divisible by p . The number of pairs of matrices (U, V) with entries from $[0, N)$ such that U and V are both non-singular modulo p , and that UAV is p -characterized, and hence p -correspondent, is at least $(1 - \epsilon)N^{2n^2}$.*

Proof. We show this count by associating each pair of matrices (U, V) with a point in $[0, N)^{2n^2}$ and then bounding the number of roots of a particular set of polynomials when evaluated in $[0, N)^{2n^2}$.

First consider the product $\mathcal{U}A\mathcal{V}$ where \mathcal{U}, \mathcal{V} have symbolic independent indeterminates u_{ij} and v_{ij} for all $i, j \in [1, n]$. Let the characteristic polynomial of $\mathcal{U}A\mathcal{V}$ be

$$g = x^n + g_1 x^{n-1} + \dots + g_k x^{n-k} + \dots + g_n,$$

Then each

$$\bar{g}_k = \frac{g_k}{\Delta_k(A)} \in \mathbb{Z}[u_{11}, u_{12}, \dots, v_{nn}],$$

is a polynomial in the entries of \mathcal{U}, \mathcal{V} with degree $2k$ and content 1 by Lemma 5.4.

Each pair of matrices U, V in the lemma statement defines a point in $[0, N)^{2n^2}$; the entries of U, V define the values for the $2n^2$ variables u_{ij} and v_{ij} . The coefficients of the characteristic polynomial of each matrix UAV is obtained by evaluating the polynomials g_k at the point in $[0, N)^{2n^2}$ defined by (U, V) . Then using Lemma 5.5, we have $\bar{g}_k \equiv 0 \pmod{p}$ in at most $(N/p)^{2n^2} \cdot 2kp^{2n^2-1} = N^{2n^2} \cdot 2k/p$ points.

The determinant of \mathcal{U} (resp. \mathcal{V}) is a polynomial of degree n in all of the $2n^2$ variables u_{ij} (resp. v_{ij}), and hence $\det U \equiv 0 \pmod{p}$ in at most $(N/p)^{2n^2} np^{2n^2-1} = N^{2n^2} n/p$ points in $[0, N)^{2n^2}$ by Lemma 5.5.

Thus the number of points in $[0, N)^{2n^2}$ for which $\det U \equiv 0 \pmod{p}$ or $\det V \equiv 0 \pmod{p}$, or that $\bar{g}_k \equiv 0 \pmod{p}$ for *some* $k \in [1, r]$ is at most

$$\frac{2nN^{2n^2}}{p} + \sum_{1 \leq k \leq r} \frac{2kN^{2n^2}}{p} = \frac{2nN^{2n^2}}{p} + \frac{r(r+1)N^{2n^2}}{p} \leq \frac{(n^2 + 3n)}{p} N^{2n^2} < \epsilon N^{2n^2}.$$

If all $\bar{g}_k \not\equiv 0 \pmod{p}$ for $k \in [1, r]$, then $v_p(\bar{g}_k) = 0$ and $v_p(g_k) = v_p(\Delta_k)$ for $k \in [1, r]$, so UAV is p -characterized, and hence p -correspondent. The number of pairs (U, V) for which this holds is then at least $N^{2n^2} - \epsilon N^{2n^2} = (1 - \epsilon)N^{2n^2}$. \square

Example 5.4. *Intuitively, Lemma 5.6 shows that most choices of the pairs (U, V) will result in UAV being p -correspondent. Consider the matrix:*

$$A = \begin{bmatrix} -48 & -83 & 91 & -497 \\ -407 & -666 & 637 & -3948 \\ 83 & 125 & -91 & 728 \\ -291 & -599 & 903 & -3717 \end{bmatrix}.$$

A is not p -correspondent since its invariant factors are $[1, 7, 7, 49]$ and its 7-adic eigenvalues are (using the Sage computer algebra system [Stein et al., 2014]):

$$\begin{aligned} &6 \cdot 7 + 7^2 + O(7^3), \\ &3 \cdot 7 + 3 \cdot 7^2 + O(7^3), \\ &1 \cdot 7 + 4 \cdot 7^2 + O(7^3), \\ &2 \cdot 7 + 3 \cdot 7^2 + O(7^3). \end{aligned}$$

Now consider a particular choice of $U, V \in \mathbb{Z}^{4 \times 4}$:

$$U = \begin{bmatrix} 6 & 1 & 0 & 20 \\ 1 & 1 & 1 & 0 \\ 1 & 1 & 1 & 2 \\ 1 & 3 & 0 & 1 \end{bmatrix}, \quad V = \begin{bmatrix} 1 & 1 & 1 & 17 \\ 0 & 0 & 3 & 2 \\ 0 & 5 & 1 & 3 \\ 1 & 0 & 9 & 56 \end{bmatrix},$$

and let

$$\tilde{A} = UAV = \begin{bmatrix} -87785 & 89700 & -758134 & -4630434 \\ -4089 & 2813 & -35060 & -213813 \\ -12105 & 11261 & -104336 & -636989 \\ -17618 & 12965 & -151217 & -922413 \end{bmatrix}.$$

Using Sage we can verify that $\det U \not\equiv 0 \pmod{7}$, $\det V \not\equiv 0 \pmod{7}$, that the invariant factors of \tilde{A} are $[1, 7, 7, 2^{10} \cdot 7^2 \cdot 17]$ and that the 7-adic valuations of the eigenvalues of \tilde{A} are $[0, 1, 1, 2]$. As expected from Lemma 5.6, \tilde{A} is p -correspondent. \square

5.3.1 Density at Large Primes

To establish the density of p -correspondent matrices, we consider the set \mathfrak{S}_S^m (defined below) of all matrices with a given Smith normal form S and integer entries from $[0, p^m)$, and show that most matrices in this set are p -characterized.

In our proofs we will embed integer matrices in the local ring $\mathbb{Z}/p^m\mathbb{Z}$ and study their local Smith normal form. If $\hat{A} \in \mathbb{Z}^{n \times n}$ is such that $\hat{A} \equiv A \pmod{p^m}$, and \hat{A} has integer Smith normal form $\text{diag}(s_1, \dots, s_{\hat{r}}, 0, \dots, 0) \in \mathbb{Z}^{n \times n}$ and A has the Smith normal form $\text{diag}(p^{e_1}, \dots, p^{e_r}, 0, \dots, 0) \in (\mathbb{Z}/p^m\mathbb{Z})^{n \times n}$ then $r \leq \hat{r}$ and $e_i = v_p(s_i)$ for $1 \leq i \leq r$.

The following lemma relates the construction UAV in Lemma 5.6 to integer matrices with prescribed p -adic valuations on their invariant factors.

For any integer a and any prime power p^m , we use $a \bmod p^m$ to denote the unique non-negative integer $r < p^m$ such that $a = qp^m + r$ for some integer q . We extend the “ $\bmod p^m$ ” operator to vectors and matrices using element-wise application. It is important to note that “ $\bmod p^m$ ” operator is not the same as the “ $\pmod{p^m}$ ” equivalence relation; for example, $(a + b) \bmod p^m \neq (a \bmod p^m) + (b \bmod p^m)$ in general.

Definition 5.3. Fix a prime p , positive integers m, n , and integers $0 \leq e_1 \leq e_2 \leq \dots \leq e_n$. Let $S = \text{diag}(p^{e_1}, \dots, p^{e_n}) \in \mathbb{Z}^{n \times n}$. Define $\mathfrak{S}_S^m \subseteq \mathbb{Z}^{n \times n}$ as the set of integer matrices with entries from $[0, p^m)$ whose Smith normal form $\text{diag}(s_1, \dots, s_n)$ satisfies $v_p(s_i) = e_i$ for all $i \in [1, n]$.

Lemma 5.7. Fix an integer n , a prime p , and integers $0 \leq e_1 \leq \dots \leq e_n$, and let $m > e_1 + \dots + e_n$. Let $S = \text{diag}(p^{e_1}, \dots, p^{e_n})$ and $\mathfrak{S}_S^m \subseteq \mathbb{Z}^{n \times n}$ as in Definition 5.3. Fix any $A \in \mathfrak{S}_S^m$. Let $L, R \in \mathbb{Z}^{n \times n}$ be any integer matrices satisfying $A = (LSR) \bmod p^m$. Then $v_p(\det L) = v_p(\det R) = 0$, and hence L, R are both invertible modulo p^m .

Proof. If $A = (LSR) \bmod p^m$ then there exists an integer matrix Q such that $A + p^m Q = LSR$. Taking the determinants of both sides, we have

$$\det(A + p^m Q) = \det(L) \det(S) \det(R).$$

Both sides are (products of) determinants, and hence polynomials in the matrix entries. Projecting modulo p^m we get

$$\det(A) \equiv \det(L) \det(S) \det(R) \pmod{p^m},$$

or equivalently

$$\det(A) + p^m q = \det(L) \det(S) \det(R),$$

for some $q \in \mathbb{Z}$.

Since $A \in \mathfrak{S}_S^m$ we know that $v_p(\det(A)) = v_p(\det(S))$, and moreover, $0 \leq v_p(\det(A)) < m$ by the conditions of the lemma. Thus $v_p(\det(A) + p^m q) = v_p(\det(A)) < m$, since the valuation, the number of times p divides $\det(A) + p^m q$, is unaffected by the second summand. Taking the valuation of both sides, we then have

$$v_p(\det(A) + p^m q) = v_p(\det(A)) = v_p(\det(L)) + v_p(\det(S)) + v_p(\det(R)).$$

Since $0 \leq v_p(\det A) = v_p(\det S) < m$, it must be the case that $v_p(\det(L)) = v_p(\det(R)) = 0$. □

Lemma 5.8. *Fix an integer n , a prime p , and integers $0 \leq e_1 \leq \dots \leq e_n$, and let $m > e_1 + \dots + e_n$. Let $S = \text{diag}(p^{e_1}, \dots, p^{e_n})$ and $\mathfrak{S}_S^m \subseteq \mathbb{Z}^{n \times n}$ as in Definition 5.3. Fix any $A \in \mathfrak{S}_S^m$. Define*

$$P_A = \{(L, R) : L, R \text{ have entries from } [0, p^m] \text{ and } A = (LSR) \bmod p^m\}.$$

Then $|P_A| = |\text{GL}_n(\mathbb{Z}/p^m\mathbb{Z})^2|/|\mathfrak{S}_S^m|$, independent of the choice of A .

Proof. We have chosen $[0, p^m]$ to represent $\mathbb{Z}/p^m\mathbb{Z}$, so any integer matrix from $[0, p^m]^{n \times n}$ has a unique image over $\mathbb{Z}/p^m\mathbb{Z}$ and vice versa. To keep track of the rings we are working

over, we use the subscript p^m to denote matrices over the ring $\mathbb{Z}/p^m\mathbb{Z}$. We first show that there is a bijection between P_A and

$$P'_A = \{(L_{p^m}, R_{p^m}) \in \mathrm{GL}_n(\mathbb{Z}/p^m\mathbb{Z})^2 : A_{p^m} \equiv L_{p^m} S_{p^m} R_{p^m} \pmod{p^m}\}.$$

If $(L, R) \in P_A$, and its image over $\mathbb{Z}/p^m\mathbb{Z}$ is (L_{p^m}, R_{p^m}) , then $(L_{p^m}, R_{p^m}) \in \mathrm{GL}_n(\mathbb{Z}/p^m\mathbb{Z})^2$ by Lemma 5.7. Also, $A = (LSR) \bmod p^m$ implies that $A + p^m Q = LSR$ for some integer matrix Q and so $A_{p^m} \equiv L_{p^m} S_{p^m} R_{p^m} \pmod{p^m}$. Thus $(L_{p^m}, R_{p^m}) \in P'_A$.

Conversely, let $(L_{p^m}, R_{p^m}) \in P'_A$ and their preimages be $L, R \in [0, p^m)^{n \times n}$. The equivalence $A_{p^m} \equiv L_{p^m} S_{p^m} R_{p^m} \pmod{p^m}$ implies

$$A + p^m Q_1 = (L + p^m Q_2)(S + p^m Q_3)(R + p^m Q_4),$$

for some integer matrices Q_1, Q_2, Q_3, Q_4 . This can be simplified to

$$A + p^m Q_5 = LSR,$$

for some integer matrix Q_5 . In other words,

$$A = (LSR) \bmod p^m,$$

and so $(L, R) \in P_A$. Thus there is a bijection between P_A and P'_A .

We now observe that the multiplicative group $\mathrm{GL}_n(\mathbb{Z}/p^m\mathbb{Z})^2$ acts on $(\mathbb{Z}/p^m\mathbb{Z})^{n \times n}$ via left and right multiplication: $(L_{p^m}, R_{p^m}) \in \mathrm{GL}_n(\mathbb{Z}/p^m\mathbb{Z})^2$ acts on $A_{p^m} \in (\mathbb{Z}/p^m\mathbb{Z})^{n \times n}$ to produce $L_{p^m} A_{p^m} R_{p^m} \in (\mathbb{Z}/p^m\mathbb{Z})^{n \times n}$. Then $\mathrm{orbit}(A_{p^m}) = \mathrm{orbit}(S_{p^m})$ under this group action since there exists at least one such L_{p^m}, R_{p^m} with $L_{p^m} A_{p^m} R_{p^m} \equiv S_{p^m} \pmod{p^m}$. Furthermore, the orbit of S_{p^m} corresponds to \mathfrak{S}_S^m : every matrix in \mathfrak{S}_S^m has a natural image over $\mathbb{Z}/p^m\mathbb{Z}$ which can be written as $L_{p^m} S_{p^m} R_{p^m} \pmod{p^m}$ for suitable choice of $L_{p^m}, R_{p^m} \in \mathrm{GL}_n(\mathbb{Z}/p^m\mathbb{Z})$, and conversely every matrix $L_{p^m} S_{p^m} R_{p^m} \pmod{p^m}$ corresponds to a preimage integer matrix in \mathfrak{S}_S^m . Therefore we know $|\mathrm{orbit}(S_{p^m})| = |\mathfrak{S}_S^m|$.

Let $\mathrm{stab}(S_{p^m})$ be the stabilizer of S_{p^m} defined as:

$$\{(L_{p^m}, R_{p^m}) : L_{p^m}, R_{p^m} \in (\mathbb{Z}/p^m\mathbb{Z})^{n \times n}, S_{p^m} \equiv L_{p^m} S_{p^m} R_{p^m} \pmod{p^m}\},$$

and let $A_{p^m} \equiv U_{p^m} S_{p^m} V_{p^m} \pmod{p^m}$ be the Smith decomposition of A_{p^m} , then every pair $(L_{p^m}, R_{p^m}) \in P'_A$ can be mapped to a pair $(U_{p^m}^{-1} L_{p^m}, R_{p^m} V_{p^m}^{-1}) \in \text{stab}(S_{p^m})$. Similarly, every pair $(L_{p^m}, R_{p^m}) \in \text{stab}(S_{p^m})$ can be mapped to a pair $(U_{p^m} L_{p^m}, R_{p^m} V_{p^m}) \in P'_A$. Thus $|P'_A| = |\text{stab}(S_{p^m})|$.

By the orbit-stabilizer theorem [Artin, 1991, Proposition 7.2], we have

$$|\text{orbit}(S_{p^m})| \cdot |\text{stab}(S_{p^m})| = |\text{GL}_n(\mathbb{Z}/p^m\mathbb{Z})^2|.$$

The lemma statement follows because $|\text{orbit}(S_{p^m})| = |\mathfrak{S}_S^m|$, and $|\text{stab}(S_{p^m})| = |P'_A| = |P_A|$. \square

Lemma 5.9. *Let $\phi \in \mathbb{Z}[x_1, \dots, x_\ell]$ be a non-zero polynomial and $a_1, \dots, a_\ell \in \mathbb{Z}$. Let p be a prime and $m \geq 1$ be an integer. Let $k = v_p(\phi(a_1, \dots, a_\ell))$ and $\bar{k} = v_p(\phi(a_1 \bmod p^m, \dots, a_\ell \bmod p^m))$. Then*

(i) *If $k < m$ then $\bar{k} = k$.*

(ii) *If $k \geq m$ then $\bar{k} \geq m$.*

(iii) *If $k = \infty$ then $\bar{k} \geq m$.*

Proof. Let $\phi(a_1, \dots, a_\ell) = p^k \alpha$ for some $\alpha \in \mathbb{Z}$ and $p \nmid \alpha$. For all $i \in [1, \ell]$, apply the Euclidean division to a_i and p^m to get $a_i = r_i + p^m q_i$ where $p^m \nmid q_i$ and $r_i = a_i \bmod p^m$. Then

$$\phi(r_1 + p^m q_1, \dots, r_\ell + p^m q_\ell) \equiv \phi(r_1, \dots, r_\ell) \pmod{p^m}.$$

(i) If $k < m$ then

$$\phi(r_1 + p^m q_1, \dots, r_\ell + p^m q_\ell) \equiv \phi(r_1, \dots, r_\ell) \equiv p^k \alpha \pmod{p^m},$$

and $\phi(r_1, \dots, r_\ell) = p^k \alpha + p^m u$ for some $u \in \mathbb{Z}$. Now $v_p(p^k \alpha + p^m u) = k$ since $p^m u$ has valuation at least $m > k$. So $\bar{k} = k$.

(ii) If $k \geq m$ then $\phi(r_1 + p^m q_1, \dots, r_\ell + p^m q_\ell) \equiv \phi(r_1, \dots, r_\ell) \equiv 0 \pmod{p^m}$, and $\phi(r_1, \dots, r_\ell) = p^{m+j} u_1$ for some $u_1 \in \mathbb{Z}$, $p \nmid u_1$ and some $j \geq 0$. Then $\bar{k} = m + j \geq m$.

(iii) If $k = \infty$ then $\phi(r_1 + p^m q_1, \dots, r_\ell + p^m q_\ell) = 0$, and $\phi(r_1, \dots, r_\ell) \equiv \phi(r_1 + p^m q_1, \dots, r_\ell + p^m q_\ell) \equiv 0 \pmod{p^m}$, which is similar to part (ii). \square

Lemma 5.10. Let $\phi_1, \dots, \phi_r \in \mathbb{Z}[x_1, \dots, x_\ell]$ be polynomials such that

$$v_p\left(\gcd\{\phi_1(a_1, \dots, a_\ell), \dots, \phi_r(a_1, \dots, a_\ell)\}\right) = k < m.$$

Then

$$v_p\left(\gcd\{\phi_1(a_1 \bmod p^m, \dots, a_\ell \bmod p^m), \dots, \phi_r(a_1 \bmod p^m, \dots, a_\ell \bmod p^m)\}\right) = k.$$

Proof. There exists an $i \in [1, r]$ such that $v_p(\phi_i(a_1, \dots, a_\ell)) = k$ whereas for all other $j \in [1, r] \setminus \{i\}$, we have $v_p(\phi_j(a_1, \dots, a_\ell)) \geq k$ (and possibly ∞). Then, by Lemma 5.9, $v_p(\phi_i(a_1 \bmod p^m, \dots, a_\ell \bmod p^m)) = k$ while for all j , $v_p(\phi_j(a_1 \bmod p^m, \dots, a_\ell \bmod p^m))$ is either k or higher than m (but not lower than k). Thus the valuation of the desired GCD is also k . \square

We now show that if A is non-singular, then the powers of p in the Smith normal form of A and $A \bmod p^m$ coincide when $m > v_p(\det A)$.

Lemma 5.11. Let $A \in \mathbb{Z}^{n \times n}$ be a non-singular matrix, $m > v_p(\det A)$ and $\bar{A} = A \bmod p^m$. Suppose the invariant factors of A and \bar{A} are s_1, \dots, s_n and $\bar{s}_1, \dots, \bar{s}_n$, respectively. Then $v_p(s_i) = v_p(\bar{s}_i)$ for $1 \leq i \leq n$.

Proof. Let Δ_i and $\bar{\Delta}_i$ be the i th determinantal divisors of A and \bar{A} respectively. We show equivalently that $v_p(\Delta_i) = v_p(\bar{\Delta}_i)$ for $1 \leq i \leq n$. Each Δ_i (resp. $\bar{\Delta}_i$) is the GCD of all $i \times i$ minors of A (resp. \bar{A}), where each such minor is a polynomial in the n^2 entries of A (resp. \bar{A}). Then by Lemma 5.10 we have $v_p(\Delta_i) = v_p(\bar{\Delta}_i)$ for all $i \in [1, n]$. \square

Lemma 5.12. Let $A \in \mathbb{Z}^{n \times n}$, $\det A \neq 0$ and $m > v_p(\det A)$. Let f_i^M denote the x^{n-i} coefficient of the characteristic polynomial of a matrix M . For all $i \in [1, n]$, if $v_p(f_i^A) = k < m$ then $v_p(f_i^{A \bmod p^m}) = k$.

Proof. Each f_i^A is the sum of all $i \times i$ symmetric minors of A , which is a polynomial in the entries of A . The claim then follows by Lemma 5.9. \square

We now apply the above lemmas to get the following.

Lemma 5.13. *Let A be a p -characterized non-singular matrix and let $m > v_p(\det A)$. Then $\bar{A} = A \bmod p^m$ is also p -characterized.*

Proof. Let Δ_i and $\bar{\Delta}_i$ be the i th determinantal divisors of A and \bar{A} respectively, for $1 \leq i \leq n$. If A is a p -characterized, then $v_p(f_i^A) = v_p(\Delta_i)$ for each $i \in [1, n]$. By Lemma 5.11 and Lemma 5.12, we have $v_p(\bar{\Delta}_i) = v_p(\Delta_i)$ and $v_p(f_i^{\bar{A}}) = v_p(f_i^A)$. So \bar{A} is p -characterized. \square

Example 5.5. *For a prime p consider the matrix A with its Smith normal form decomposition:*

$$A = \begin{bmatrix} p^3 + 1 & p \\ 2p^4 & p^2 \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ -p^4 & 1 \end{bmatrix} \begin{bmatrix} 1 & \\ & -p^2 + p^5 \end{bmatrix} \begin{bmatrix} 1 & p \\ -p^2 & -1 - p^3 \end{bmatrix}.$$

The characteristic polynomial of A is

$$f = x^2 - (1 + p^2 + p^3)x + p^2 - p^5.$$

Note that A is p -characterized. Now let $m = 3$ and consider $A \bmod p^m$ and its Smith normal form:

$$A \bmod p^3 = \begin{bmatrix} 1 & p \\ 0 & p^2 \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & \\ & p^2 \end{bmatrix} \begin{bmatrix} 1 & -p \\ 0 & 1 \end{bmatrix},$$

which has the characteristic polynomial

$$x^2 - (1 + p^2)x + p^2.$$

Thus $A \bmod p^3$ is p -characterized as well. \square

The following bound is a relatively well-known fact, but we prove it for completeness.

Lemma 5.14. $|\mathbf{M}_n(\mathbb{Z}/p^m\mathbb{Z})|/|\mathbf{GL}_n(\mathbb{Z}/p^m\mathbb{Z})| < 4$.

Proof. Any matrix $A \in \mathbf{M}_n(\mathbb{Z}/p^m\mathbb{Z})$ can be written as $A = A_0 + pA_1 + \dots + p^{m-1}A_{m-1}$ with A_i 's having entries from $[0, p)$. Then $A \in \mathbf{GL}_n(\mathbb{Z}/p^m\mathbb{Z})$ if and only if $A_0 \in \mathbf{GL}_n(\mathbb{Z}/p\mathbb{Z})$. There are $(p^{n^2})^{m-1}$ ways to construct the components A_1, \dots, A_{m-1} for each given $A_0 \in$

$\mathrm{GL}_n(\mathbb{Z}/p\mathbb{Z})$. So $|\mathrm{GL}_n(\mathbb{Z}/p^m\mathbb{Z})| = p^{(m-1)n^2} \cdot |\mathrm{GL}_n(\mathbb{Z}/p\mathbb{Z})|$. Next, recall the well-known density bound for non-singular matrices over finite fields:

$$\frac{|\mathrm{GL}_n(\mathbb{Z}/p\mathbb{Z})|}{p^{n^2}} = \left(1 - \frac{1}{p}\right) \left(1 - \frac{1}{p^2}\right) \cdots \left(1 - \frac{1}{p^n}\right) > 1/4.$$

Thus

$$\frac{|\mathrm{M}_n(\mathbb{Z}/p^m\mathbb{Z})|}{|\mathrm{GL}_n(\mathbb{Z}/p^m\mathbb{Z})|} = \frac{p^{mn^2}}{p^{(m-1)n^2} |\mathrm{GL}_n(\mathbb{Z}/p\mathbb{Z})|} = \frac{p^{n^2}}{|\mathrm{GL}_n(\mathbb{Z}/p\mathbb{Z})|} < 4.$$

□

We can now establish our main density result.

Theorem 5.2. *Let n be a positive integer, $\epsilon > 0$, and p be any prime greater than $16(n^2 + 3n)/\epsilon$. Fix a set of integers $0 \leq e_1 \leq e_2 \leq \cdots \leq e_n$ and let $m \geq e_1 + \cdots + e_n + 1$ and $S = \mathrm{diag}(p^{e_1}, \dots, p^{e_n}) \in \mathbb{Z}^{n \times n}$. Then the number of matrices in \mathfrak{S}_S^m which are p -characterized and hence p -correspondent is at least $(1 - \epsilon) \cdot |\mathfrak{S}_S^m|$.*

Proof. Let

$$P = \{(L, R) : L, R \in [0, p^m]^{n \times n}\}.$$

For any $A \in \mathfrak{S}_S^m$, let $P_A \subseteq P$ be as in Lemma 5.8:

$$P_A = \{(L, R) : L, R \text{ have entries from } [0, p^m] \text{ and } A = (LSR) \bmod p^m\}.$$

If at least one pair $(L, R) \in P_A$ is such that LSR is p -characterized, then A is p -characterized by Lemma 5.13 (recall $A = (LSR) \bmod p^m$ and $m \geq e_1 + \cdots + e_n + 1$ implies $m > v_p(\det A)$). On the other hand, if every pair $(L, R) \in P_A$ is such that LSR is not p -characterized then A can be either p -characterized or not (because the converse of Lemma 5.13 is not necessarily true; some non p -characterized matrices can become p -characterized after applying $\bmod p^m$). To derive an upper bound on the number of non p -characterized matrices in \mathfrak{S}_S^m , we allow the worst outcome: $A = (LSR) \bmod p^m$ is not p -characterized when LSR is not p -characterized for all pairs $(L, R) \in P_A$.

The number of sets, P_A , having every pair (L, R) with a non p -characterized product LSR , can be obtained as the ratio between the total number of pairs giving non

p -characterized products (which is at most $(\epsilon/16)|P|$ by Lemma 5.6) divided by the size of each P_A (which is $|\mathrm{GL}_n(\mathbb{Z}/p^m\mathbb{Z})^2|/|\mathfrak{S}_S^m|$ by Lemma 5.8). So the maximum number of matrices in \mathfrak{S}_S^m which are not p -characterized is

$$\frac{(\epsilon/16)|P|}{|P_A|} = \frac{(\epsilon/16)|\mathrm{M}_n(\mathbb{Z}/p^m\mathbb{Z})|^2}{|\mathrm{GL}_n(\mathbb{Z}/p^m\mathbb{Z})^2|/|\mathfrak{S}_S^m|} < \epsilon|\mathfrak{S}_S^m|,$$

where the inequality follows using Lemma 5.14.

Hence there are at least $(1 - \epsilon)|\mathfrak{S}_S^m|$ matrices in \mathfrak{S}_S^m which are p -characterized, and each one of those matrices is also p -correspondent by Theorem 5.1. \square

This result also implies that an integer matrix with entries sampled from uniformly at random from $[0, p^m)$ will be p -correspondent with probability at least $1 - \epsilon$.

5.4 Density at Small Primes

The density estimate of Theorem 5.2 is limited to large primes. We now report on experiments with small primes. For a given size n and a prime power p^m , we enumerate the set of all $n \times n$ matrices with entries from $[0, p^m)$ and $v_p(\text{determinant}) < m$. We then count the fraction of matrices which are p -correspondent.

Table 5.1 shows the density of p -characterized and p -correspondent non-singular matrices for small values of p, m, n . The fourth and fifth columns report the fraction (in percentage) of p -characterized and p -correspondent matrices among all $n \times n$ non-singular matrices with entries $[0, p^m)$ and whose determinant has p -adic valuation smaller than m . Recall from Example 5.3 that matrices can be p -correspondent but not necessarily p -characterized, thus the reported p -characterized density is lower than p -correspondent density.

The sixth column in the table reports the minimum percentage of p -characterized matrices among all the Smith normal forms. Given p^m and n , we consider the set of all $n \times n$ matrices with entries $[0, p^m)$ and $v_p(\text{determinant}) < m$. We partition these matrices by their Smith normal forms localized at p , where we only care about the powers of p in the invariant factors and treat the other prime powers as units. For example, when $p^m = 2^2$

Table 5.1: Density (in percentage) of p -characterized and p -correspondent matrices.

p	m	n	p -characterized	p -correspondent	min p -char.
2	1	2	56.25	81.25	33.33
	2	2	53.52	80.08	33.33
	3	2	53.34	80.00	33.33
	4	2	53.33	80.00	33.33
	1	3	29.10	71.29	18.75
	2	3	26.51	70.14	16.67
	1	4	15.61	66.67	6.667
3	1	2	67.90	90.12	62.50
	2	2	67.50	90.00	50.00
	3	2	67.50	90.00	50.00
	1	3	45.58	86.73	42.77
5	1	2	80.16	96.16	79.17
	2	2	80.13	96.15	78.96
7	1	2	85.76	98.00	85.42

and $n = 2$, we get the following (non-singular) Smith normal forms localized at 2:

$$\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \begin{bmatrix} 1 & 0 \\ 0 & 2 \end{bmatrix}.$$

We then count the fraction of p -characterized matrices in each partition and report the minimum percentage among all partitions.

Finally, the table shows that the density drops as n increases and as p decreases, which is consistent with the proofs for large primes. An open question is to prove similar density estimates for small primes, i.e., when p is small compared to n .

Chapter 6

Ranks of Remainder Matrices

This chapter presents two related results on ranks of integer matrices after applying element-wise division with remainder. This is a study towards understanding the interaction between the local invariant factors at a prime p and the p -adic expansion of matrices.

For a prime p and a matrix $A \in \mathbb{Z}^{n \times n}$, write A as $A = p(A \text{ quo } p) + (A \text{ rem } p)$ where the remainder and quotient operations are applied element-wise. Write the p -adic expansion of A as $A = A^{[0]} + pA^{[1]} + p^2A^{[2]} + \dots$ where each $A^{[i]} \in \mathbb{Z}^{n \times n}$ has entries between $[0, p-1]$. Upper bounds are proven for the integer ranks of $A \text{ rem } p$, and $A \text{ quo } p$. Also, upper bounds are proven for the finite field rank of $A^{[i]}$ for all $i \geq 0$ when $p = 2$, and a conjecture is presented for odd primes.

6.1 Introduction

Let p be a prime, and let A be an integer matrix whose Smith normal form is given by $A = USV$. Assume that $S = \text{diag}(1, \dots, 1, p, \dots, p, 0, \dots, 0)$; i.e., the only non-trivial invariant factor of A is p . If we write A , U , S , and V using their p -adic expansion, we get $(A_0 + pA_1) = (U_0 + pU_1)(S_0 + pS_1)(V_0 + pV_1)$, where $A_0 = U_0S_0V_0 \text{ mod } p$ and

$$A_1 = U_1S_0V_0 + U_0S_0V_1 + U_0S_1V_0, \tag{6.1}$$

where $\text{rank}(S_0) = r_0$, and $\text{rank}(S_1) = r_1$, which are the multiplicities of 1's and p 's in the Smith normal form, respectively. Furthermore, with appropriate preconditioning, $\text{rank}(A_0)$ is proportional to r_0 , and $\text{rank}(A_1)$ is proportional to $2r_0 + r_1$. Hence, this formulation leads to a belief that we can isolate A_0, A_1 , compute their ranks over $\mathbb{Z}/p\mathbb{Z}$, and discover multiplicities of the invariant factors. However, by closer inspection, equation (6.1) is in fact:

$$A_1 = U_1 S_0 V_0 + U_0 S_0 V_1 + U_0 S_1 V_0 + \overbrace{U_0 S_0 V_0 \text{ quo } p}^{\text{carry}}.$$

The extra term, $(U_0 S_0 V_0 \text{ quo } p)$, is introduced by the fact that arithmetic operations over \mathbb{Z} exhibit carries. These carries contribute to the overall ranks of matrix expressions. This leads to interesting questions about ranks of matrices under the remainder and quotient operators and the ranks of components of matrices when written p -adically.

6.2 Quotient and Remainder Matrices

For any integer n and any prime p , let $n \text{ rem } p$ and $n \text{ quo } p$ denote the remainder and quotient in the division $n = qp + r$. To ensure a unique representation, we choose the non-negative remainder $r \in [0, p)$. The operators $\text{rem } p$ and $\text{quo } p$ are naturally extended to vectors and matrices using element-wise application.

When convenient, we embed integer matrices in $\mathbb{Z}/p\mathbb{Z}$ using the natural element-wise projection $a \mapsto a \bmod p$. We use two notations for ranks. The integer rank is denoted by $\text{rank}(\cdot)$, while the rank over the finite field $\mathbb{Z}/p\mathbb{Z}$ is denoted by $\text{rank}_p(\cdot)$. Alternatively, if $r = \text{rank}(A)$ and the Smith normal form of A is $S = \text{diag}(s_1, \dots, s_r, 0, \dots, 0)$, then $\text{rank}_p(A) = r_0$ is the maximal index i such that $p \nmid s_i$. Finally, we use the notation $A_{*,j}$ for the j th column of $A \in \mathbb{Z}^{n \times n}$ and $a_{i,j}$ for the entry (i, j) of A .

The following is the main result of this section. Let A be an $n \times n$ matrix over \mathbb{Z} , $r = \text{rank}(A)$, $r_0 = \text{rank}_p(A)$, and assume $n > p^{r_0}$. Then

- (i) $\text{rank}(A \text{ rem } p) \leq (p^{r_0} - 1)(p + 1)/(2(p - 1))$;
- (ii) $\text{rank}(A \text{ quo } p) \leq r + (p^{r_0} - 1)(p + 1)/(2(p - 1))$.

We will first study the remainder problem for rank-1 matrices. Then we will generalize to arbitrary rank matrices.

6.2.1 Remainder of Rank-1 Matrices

Let p be any odd prime, $n \geq p$. Let $u \in \mathbb{Z}^n$ be any non-zero vector where the entries of $u \bmod p$ include $\{1, 2, \dots, p-1\}$.

Lemma 6.1. *The set of vectors $\{u \bmod p, (2u) \bmod p, \dots, ((p-1)u) \bmod p\}$ is linearly dependent and has rank $(p+1)/2$.*

First we prove this result for $n = p-1$. A generalization follows. Let $u = (1, 2, \dots, p-1) \in \mathbb{Z}^{(p-1)}$ and $M \in \mathbb{Z}^{(p-1) \times (p-1)}$ be the rank-1 matrix $M = uu^T$ and let $R = M \bmod p$.

Lemma 6.2. $\text{rank}(R) = (p+1)/2$.

Proof. Lemma 6.3 shows that $(p+1)/2$ is an upper bound on the rank and Lemma 6.5 shows that $(p+1)/2$ is a lower bound. \square

Lemma 6.3. $\text{rank}(R) \leq (p+1)/2$.

Proof. Let $1 \leq j \leq (p-1)/2$ and $1 \leq i \leq p-1$. Write $ij = qp + r$ where $0 \leq r < p$. Also $i, j < p$ implies $p \nmid i$ and $p \nmid j$, so $r \neq 0$. Then $i(p-j) = ip - ij = ip - qr - r = p(i-q-1) + (p-r)$ where $0 < (p-r) < p$. So $ij \bmod p + i(p-j) \bmod p = r + (p-r) = p$. But $R_{i,j} = ij \bmod p$, so for all $1 \leq i \leq (p-1)/2$ we have $R_{*,i} = (p, p, \dots, p)^T - R_{*,p-i}$. Thus there are $(p-1)/2$ linearly dependent columns, and no more than $(p+1)/2$ linearly independent columns. \square

To prove that $(p+1)/2$ is also a lower bound on the rank, it suffices (using Lemma 6.3) to consider the matrix B of size $(p-1) \times \frac{p+1}{2}$ which is formed by the first $(p-1)/2$ columns of R and the column $B_{*,(p+1)/2} = R_{*,(p+1)/2} + R_{*,(p-1)/2} = (p, \dots, p)^T$. The matrix B has

the following structure:

$$B = \begin{bmatrix} 1 & 2 & \cdots & \frac{p-1}{2} & p \\ 2 & 4 & \cdots & p-1 & p \\ 3 & 6 \text{ rem } p & \cdots & 3\frac{p-1}{2} \text{ rem } p & p \\ \vdots & \vdots & \ddots & \vdots & \\ (p-1) \text{ rem } p & 2(p-1) \text{ rem } p & \cdots & \frac{(p-1)^2}{2} \text{ rem } p & p \end{bmatrix}.$$

Lemma 6.4. *Either the right kernel of B is empty, or the first $(p-1)/2$ columns of B are linearly dependent.*

Proof. We will prove the statement by contradiction. Assume the contrary, that is, the right kernel of B is not empty and the first $(p-1)/2$ columns are linearly independent. Then there exists $(p+1)/2$ integers $c_1, \dots, c_{(p+1)/2}$ such that

$$c_1 B_{*,1} + c_2 B_{*,2} + \dots + c_{(p+1)/2} B_{*,(p+1)/2} = 0. \quad (6.2)$$

Apply this linear combination simultaneously to the first two rows of B to get

$$c_1 + 2c_2 + \dots + c_{(p-1)/2} (p-1)/2 + c_{(p+1)/2} p = 0, \quad (6.3)$$

$$2c_1 + 4c_2 + \dots + c_{(p-1)/2} (p-1) + c_{(p+1)/2} p = 0. \quad (6.4)$$

If we multiply $2 \times (6.3) - (6.4)$ we get $c_{(p+1)/2} = 0$. Substituting in (6.2), we get:

$$c_1 B_{*,1} + c_2 B_{*,2} + \dots + c_{(p-1)/2} B_{*,(p-1)/2} = 0. \quad (6.5)$$

But this contradicts the assumption that the first $(p-1)/2$ columns are linearly independent. The assumption is wrong, and the lemma statement holds. \square

Lemma 6.5. $(p+1)/2 \leq \text{rank}(R)$.

Proof. Using Lemma 6.4, proving a lower bound on the rank of R can be reduced to showing that the first $(p-1)/2$ columns of B are linearly independent. We use induction. Consider the sequence of matrices $B^{(k)}$ formed by the first k columns of B , where $2 \leq k \leq (p-1)/2$.

Base case: $B^{(2)}$, has rank 2 which is straightforward to verify.

Inductive case: we assume $B^{(k-1)}$ has rank $k - 1$. We have

$$B^{(k)} = \begin{bmatrix} 1 & 2 & \cdots & k \\ 2 & 4 & \cdots & 2k \\ 3 & 6 \bmod p & \cdots & (3k) \bmod p \\ & & \ddots & \end{bmatrix}. \quad (6.6)$$

We perform the following elementary operations:

1. $\text{Row}(p - 1) = (\text{Row}(p - 1) + \text{Row}(1)) / p$. Then we have $\text{Row}(p - 1) = [1, \dots, 1]$.
2. The first $(k - 1)$ columns have rank $k - 1$ by the inductive hypothesis. But the leading $(k - 1) \times (k - 1)$ submatrix is symmetric, so its row rank is also $k - 1$. Now reduce the first $k - 1$ rows to echelon form. The resulting matrix is

$$\begin{bmatrix} 1 & * & \cdots & \cdots & k \\ & 1 & \cdots & \cdots & 2k \\ & & \ddots & & \vdots \\ & & & 1 & \\ 1 & 1 & \cdots & 1 & 1 \end{bmatrix}. \quad (6.7)$$

3. If use the diagonal entries to eliminate the first $k - 1$ entries of the last row, we get:

$$\begin{bmatrix} 1 & * & \cdots & \cdots & k \\ & 1 & \cdots & \cdots & 2k \\ & & \ddots & & \vdots \\ & & & 1 & \\ 0 & \cdots & 0 & 1 & \end{bmatrix}. \quad (6.8)$$

The resulting matrix has column rank k which concludes the induction proof. \square

We are now ready to generalize Lemma 6.2 and prove Lemma 6.1.

Proof of Lemma 6.1. For the column vector $u \in \mathbb{Z}^{n \times 1}$, consider the matrix $\widehat{R} \in \mathbb{Z}^{n \times n} = uu^T \bmod p$, which is analogous to the matrix R of Lemma 6.2. The image of $u \bmod p$ has

entries from the interval $[0, p - 1]$. If $n > p$ then, by the pigeonhole principle, the vector $u \bmod p$ will contain duplicate (and zero) entries, which correspond to duplicate and zero rows in \widehat{R} . So up to row/column permutations, \widehat{R} contains R as a submatrix, and the extra rows/columns are duplicate and/or zero. Hence $\text{rank}(\widehat{R}) = \text{rank}(R)$. \square

6.2.2 A Note on Latin Squares

It is worth noting that Lemma 6.2 also implies a result on the ranks of Latin squares of certain orders and certain isotopy classes.

Recall that a *Latin square* of order n is an $n \times n$ matrix (or array) of n unique symbols arranged such that each symbol appears only once in each row and each column. Here are two examples of Latin squares on the symbols 1, 2, 3, 4:

$$\begin{bmatrix} 1 & 2 & 3 & 4 \\ 2 & 4 & 1 & 3 \\ 3 & 1 & 4 & 2 \\ 4 & 3 & 2 & 1 \end{bmatrix}, \begin{bmatrix} 4 & 1 & 2 & 3 \\ 3 & 2 & 4 & 1 \\ 2 & 3 & 1 & 4 \\ 1 & 4 & 3 & 2 \end{bmatrix}.$$

If one starts with a Latin square L and applies row and column permutations, then we obtain a Latin square L' . The two squares L, L' are said to be in the same *isotopy class*.

As before, let p be an odd prime, and let R be the $(p - 1) \times (p - 1)$ integer matrix whose (i, j) th entry is $ij \bmod p$. We show that R is a Latin square as follows. The matrix R is the Cayley multiplication table of the finite field $\mathbb{Z}/p\mathbb{Z}$ excluding the entries for the element 0. We have $ij \bmod p \neq i'j' \bmod p$ whenever $j \neq j'$, where $i, j, j' \in [1, p - 1]$. So every row and column of R has the residues $\{1, \dots, p - 1\}$ appearing only once, and R is a Latin square of order $p - 1$.

Although R has rank 1 over $\mathbb{Z}/p\mathbb{Z}$, Lemma 6.2 shows that R has a non-trivial rank over \mathbb{Z} . Also, all Latin squares generated by permuting the rows and columns of R will also have the same rank.

Corollary 6.1. *Let p and R be as above. Any Latin square in the isotopy class of R , taken as a $(p - 1) \times (p - 1)$ integer matrix, has rank $(p + 1)/2$.*

For example, the squares listed above are in the same isotopy class where the prime is 5. They both have rank $(5 + 1)/2 = 3$. However, not all Latin squares on $\{1, \dots, p - 1\}$ have rank $(p + 1)/2$. For example, The following Latin square has rank 4:

$$\begin{bmatrix} 2 & 3 & 4 & 1 \\ 1 & 4 & 3 & 2 \\ 3 & 1 & 2 & 4 \\ 4 & 2 & 1 & 3 \end{bmatrix}.$$

It is an interesting question to classify the isotopy classes of Latin squares on $\{1, \dots, p - 1\}$, and study the rank properties of each class.

Other permutation-invariant properties, such as the Smith normal form, might be of interest. Empirically, one finds that R has the following invariant factors:

$$1, \underbrace{p, \dots, p}_{\frac{p-1}{2}}, \underbrace{0, \dots, 0}_{\frac{p-3}{2}}.$$

The multiplicities of 1's and 0's trivially follow from the rank results presented here. Recall that if we add the i th column of R to the $(p - i)$ th column, we get the column $[p, p, \dots, p]$. Therefore the following vectors are in the right nullspace of R :

$$\begin{aligned} & [1, -1, 0, \dots, 0, -1, 1] \\ & [1, 0, -1, \dots, -1, 0, 1] \\ & [0, 1, -1, \dots, -1, 1, 0] \\ & \dots \end{aligned}$$

We can form $\binom{\frac{p-1}{2}}{2}$ such vectors up to sign changes. However, only $(p - 3)/2$ vectors are linearly independent by Lemma 6.2.

If we take any of the above vectors, and remove the negative signs and multiply it by $p^{i-1}R$, we do not get the zero vector. Instead, we get the vector $[2p^i, 2p^i, \dots, 2p^i]$ which is 0 modulo p^i . Similarly, if we multiply the vector $[1, 0, \dots, 0, 1]$ by $p^{i-1}R$ we get $[p^i, p^i, \dots, p^i] \equiv 0 \pmod{p^i}$. In total, we get $(p - 3)/2 + 1 = (p - 1)/2$ linearly independent nullspace vectors modulo p^i for all $i \geq 1$. We can deduce that the kernel of R over $\mathbb{Z}/p^i\mathbb{Z}$,

when put in Smith normal form, will have a diagonal block $p^{i-1}I$ of size $(p-1)/2 \times (p-1)/2$. Thus the Smith normal form of R will have a diagonal block pI of the same size, and we get the aforementioned multiplicities of R 's invariant factors.

However this (informal) argument does not exclude other primes from appearing in the integer Smith normal form of R . It would be interesting to establish that the only non-trivial invariant factors of R are powers of p .

6.2.3 Rank Theorem

Let A be an $n \times n$ matrix over \mathbb{Z} , $r = \text{rank}(A)$, $r_0 = \text{rank}_p(A)$, and assume $n > p^{r_0}$. Then

Lemma 6.6. $\text{rank}(A \text{ rem } p) \leq (p^{r_0} - 1)(p + 1)/(2(p - 1))$.

Proof. Let $A = USV$ be the Smith normal form of A , and $S = S_r + pS_q$ where $S_q = S \text{ quo } p$ and $S_r = S \text{ rem } p$. Then

$$A \text{ rem } p = USV \text{ rem } p = (US_rV + pUS_qV) \text{ rem } p = US_rV \text{ rem } p. \quad (6.9)$$

If $r_0 = \text{rank}_p(A)$ then $S_r = \text{diag}(\sigma_1, \dots, \sigma_{r_0}, 0, \dots, 0)$ where $\sigma_i \in [1, p-1]$ for all $1 \leq i \leq r_0$.

The j th column of $A \text{ rem } p$ is

$$A_{*,j} \text{ rem } p = \left(\sum_{\ell=1}^{r_0} \sigma_\ell v_{\ell,j} U_{*,\ell} \right) \text{ rem } p = \left(\sum_{\ell=1}^{r_0} c_{\ell,j} U_{*,\ell} \right) \text{ rem } p, \quad (6.10)$$

where $c_{\ell,j} \in [0, p-1]$. If we only consider the non-zero coefficients $c_{\ell,j}$, then the right-hand side of (6.10) is an i -term sum $(c_{\ell_1,j} U_{*,\ell_1} + \dots + c_{\ell_i,j} U_{*,\ell_i}) \text{ rem } p$, where $1 \leq i \leq r_0$ and $1 \leq \ell_1 < \ell_2 < \dots < \ell_i \leq r_0$. The coefficients $c_{\ell_k,j}$ are elements in $[1, p-1]$ which are units modulo p . In particular, we can factor $c_{\ell_1,j}$ from the sum, and re-write (6.10) as:

$$A_{*,j} \text{ rem } p = (c_{\ell_1,j} (U_{*,\ell_1} + \alpha_{\ell_2,j} U_{\ell_2,j} + \dots + \alpha_{\ell_i,j} U_{*,\ell_i})) \text{ rem } p, \quad (6.11)$$

where $\alpha_{\ell_k,j} \in [1, p-1]$ for all k .

Fix some i, j and some non-zero assignment of $\alpha_{\ell_2,j}, \dots, \alpha_{\ell_i,j}$ in (6.11) and let $\hat{u} = U_{*,\ell_1} + \alpha_{\ell_2,j} U_{\ell_2,j} \dots + \alpha_{\ell_i,j} U_{*,\ell_i}$. Then (6.11) becomes $A_{*,j} \text{ rem } p = (c_{\ell_1,j} \hat{u}) \text{ rem } p$. There are $p-1$ possible values for $c_{\ell_1,j}$ and hence the possible values of $A_{*,j} \text{ rem } p$ are:

$$\{\hat{u} \text{ rem } p, (2\hat{u}) \text{ rem } p, \dots, ((p-1)\hat{u}) \text{ rem } p\}. \quad (6.12)$$

We are interested in getting an upper bound on the rank of this set of vectors. First note that $(xy) \bmod p = (x \bmod p)(y \bmod p) \bmod p$. So $(i\hat{u}) \bmod p = (i(\hat{u} \bmod p)) \bmod p$ for $i \in [1, p-1]$. Hence the maximal rank one can achieve from (6.12) occurs when (up to permutation) $\hat{u} \bmod p = (0, 1, 2, \dots, p-1, \dots)$. The rest of the entries are duplicates from the same range $[0, p-1]$ by the pigeonhole principle. Now apply Lemma 6.1 to conclude that the vectors in (6.12) have rank at most $(p+1)/2$.

Thus for each i, j and non-zero assignment of $\alpha_{\ell_{2,j}}, \dots, \alpha_{\ell_{i,j}}$, there are at most $(p+1)/2$ linearly independent columns of $A \bmod p$. We now count the maximal possible number of distinct $A_{*,j}$'s. There are $\binom{r_0}{i}$ possible ways to select i different columns from the first r_0 columns of U . For each choice, there are $i-1$ coefficients: $\alpha_{\ell_{2,j}}, \dots, \alpha_{\ell_{i,j}}$, and $(p-1)^{i-1}$ possible ways to assign their non-zero values from $[1, p-1]$. Each choice gives a set of vectors as in (6.12) whose rank is at most $(p+1)/2$. Summing over all $i \in [1, r_0]$, the maximal possible rank from the span of columns in (6.10) is

$$\sum_{i=1}^{r_0} \binom{r_0}{i} (p-1)^{i-1} \frac{p+1}{2}. \quad (6.13)$$

We have

$$\begin{aligned} \sum_{i=1}^{r_0} \binom{r_0}{i} (p-1)^{i-1} &= \sum_{i=0}^{r_0} \binom{r_0}{i} (p-1)^{i-1} - \frac{1}{p-1} \\ &= \frac{1}{p-1} \sum_{i=0}^{r_0} \binom{r_0}{i} (p-1)^i - \frac{1}{p-1} \\ &= \frac{p^{r_0} - 1}{p-1} \end{aligned}$$

So

$$\sum_{i=1}^{r_0} \binom{r_0}{i} (p-1)^{i-1} \frac{p+1}{2} = \frac{p^{r_0} - 1}{p-1} \frac{p+1}{2}. \quad (6.14)$$

□

We are now ready to prove the main rank theorem.

Theorem 6.1. *Let A be an $n \times n$ matrix over \mathbb{Z} , $r = \text{rank}(A)$, $r_0 = \text{rank}_p(A)$, and assume $n > p^{r_0}$. Then*

$$(i) \text{rank}(A \bmod p) \leq (p^{r_0} - 1)(p+1)/(2(p-1));$$

$$(ii) \text{ rank}(A \text{ quo } p) \leq r + (p^{r_0} - 1)(p + 1)/(2(p - 1)).$$

Proof. Lemma 6.6 proves part (i). For part (ii), we have $A = (A \text{ rem } p) + p(A \text{ quo } p)$, or $p(A \text{ quo } p) = A - (A \text{ rem } p)$. For matrices $X = Y + Z$, rank is sub-additive and $\text{rank}(X) \leq \text{rank}(Y) + \text{rank}(Z)$. Scaling a matrix by p or -1 does not change its rank. So $\text{rank}(A \text{ quo } p) \leq \text{rank}(A) + \text{rank}(A \text{ rem } p) = r + \text{rank}(A \text{ rem } p)$. \square

6.3 p -Adic Matrices

Ranks in this section are over $\mathbb{Z}/p\mathbb{Z}$. For any prime p and any matrix $M \in \mathbb{Z}^{n \times n}$ with entries $|m_{i,j}| < \beta$, the p -adic expansion of M is $M = M^{[0]} + pM^{[1]} + \dots + p^s M^{[s]}$, where the entries of each matrix $M^{[i]}$ are between $[0, p - 1]$, and $s \leq \lceil \log_p \beta \rceil$. We call $M^{[i]}$ the i th p -adic matrix digit of M . We extend the superscript $[i]$ notation to vectors and integers in the obvious way. It should be noted that we do not use the p -adic metric. We use the term p -adic in the sense of p -adic expansion.

We will present results concerning the ranks of 2-adic matrix digits. For odd primes, we only present a conjecture. It is an open question to study the combinatorial structure of the column space of the p -adic matrix digits for odd primes.

6.3.1 Binary Code Matrices

Fix $p = 2$. The goal of this section is to show that for all $i \geq 1$, $\text{rank}_p(M^{[i]}) = \binom{r}{2^i}$ where $M = AA^T$ for some specially constructed A , which we call *binary code* matrix. We will generalize the construction of M in a subsequent section. For now, A is constructed as follows. Start with the $2^r \times r$ matrix whose i, j entry is the j th bit in the binary expansion of i . Then apply row permutations to A such that the first $\binom{r}{0}$ rows have exactly 0 non-zero entries, followed by $\binom{r}{1}$ rows which have exactly 1 non-zero entries, followed by $\binom{r}{2}$ rows which have exactly 2 non-zero entries and so on. See Figure 6.1 for an example where $r = 4$.

$$\begin{bmatrix}
0 & 0 & 0 & 0 \\
0 & 0 & 0 & 1 \\
0 & 0 & 1 & 0 \\
0 & 1 & 0 & 0 \\
1 & 0 & 0 & 0 \\
\hline
0 & 0 & 1 & 1 \\
0 & 1 & 0 & 1 \\
0 & 1 & 1 & 0 \\
1 & 0 & 0 & 1 \\
1 & 0 & 1 & 0 \\
1 & 1 & 0 & 0 \\
\hline
0 & 1 & 1 & 1 \\
1 & 0 & 1 & 1 \\
1 & 1 & 0 & 1 \\
1 & 1 & 1 & 0 \\
\hline
1 & 1 & 1 & 1
\end{bmatrix},
\begin{bmatrix}
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 1 & 0 & 0 & 0 & 1 & 1 & 0 & 1 & 0 & 0 & 1 \\
0 & 0 & 1 & 0 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 \\
0 & 0 & 0 & 1 & 0 & 0 & 1 & 1 & 0 & 0 & 1 & 1 \\
0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 1 & 1 & 1 & 1 \\
\hline
0 & 1 & 1 & 0 & 0 & 2 & 1 & 1 & 1 & 1 & 0 & 2 \\
0 & 1 & 0 & 1 & 0 & 1 & 2 & 1 & 1 & 0 & 1 & 2 \\
0 & 0 & 1 & 1 & 0 & 1 & 1 & 2 & 0 & 1 & 1 & 2 \\
0 & 1 & 0 & 0 & 1 & 1 & 1 & 0 & 2 & 1 & 1 & 2 \\
0 & 0 & 1 & 0 & 1 & 1 & 0 & 1 & 1 & 2 & 1 & 2 \\
0 & 0 & 0 & 1 & 1 & 0 & 1 & 1 & 1 & 1 & 2 & 2 \\
\hline
0 & 1 & 1 & 1 & 0 & 2 & 2 & 2 & 1 & 1 & 1 & 3 \\
0 & 1 & 1 & 0 & 1 & 2 & 1 & 1 & 2 & 2 & 1 & 3 \\
0 & 1 & 0 & 1 & 1 & 1 & 2 & 1 & 2 & 1 & 2 & 3 \\
0 & 0 & 1 & 1 & 1 & 1 & 1 & 2 & 1 & 2 & 2 & 3 \\
\hline
0 & 1 & 1 & 1 & 1 & 2 & 2 & 2 & 2 & 2 & 2 & 3
\end{bmatrix}.$$

Figure 6.1: An example of A (left) and $M = AA^T$ (right), where $r = 4$. The rows of A are partitioned by the number of non-zero entries in each row. The corresponding blocks in the symmetric matrix M are shown with borders. The column partitions of M are \mathbf{m}_0 , \mathbf{m}_1 , \mathbf{m}_2 , \mathbf{m}_3 , \mathbf{m}_4 . Finally, $\text{rank}_p(M^{[0]}) = \text{rank}_p(\mathbf{m}_1^{[0]}) = 4$, $\text{rank}_p(M^{[1]}) = \text{rank}_p(\mathbf{m}_2^{[1]}) = 6$, $\text{rank}_p(M^{[2]}) = \text{rank}_p(\mathbf{m}_4^{[2]}) = 1$.

The ℓ th column of M is given by:

$$M_{*,\ell} = a_{1,\ell}A_{*,1} + \dots + a_{r,\ell}A_{*,r} = \sum_{j \in J_\ell} A_{*,j}, \quad (6.15)$$

where $J_\ell \subseteq \{1, 2, \dots, r\}$ and the second equality holds because $a_{i,\ell} \in \{0, 1\}$. We call J_ℓ the *summing index set* of $M_{*,\ell}$. Let \mathbf{m}_k denote the $2^r \times \binom{r}{k}$ submatrix of M , which includes all columns of the form: $M_{*,\ell} = \sum_{j \in J_\ell} A_{*,j}$ where $J_\ell \subseteq \{1, 2, \dots, r\}$ and $|J_\ell| = k$. Then the

columns of M can be partitioned into:

$$M = \begin{bmatrix} \mathbf{m}_0 & \mathbf{m}_1 & \mathbf{m}_2 & \dots & \mathbf{m}_{2^i} & \mathbf{m}_{2^i+1} & \dots & \mathbf{m}_r \end{bmatrix}. \quad (6.16)$$

The next lemma shows that

$$M^{[i]} = \begin{bmatrix} \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{m}_{2^i}^{[i]} & \mathbf{m}_{2^i+1}^{[i]} & \dots & \mathbf{m}_r^{[i]} \end{bmatrix}. \quad (6.17)$$

Lemma 6.7. *If $k < 2^i$, then $\mathbf{m}_k^{[i]} = \mathbf{0}$ for all $i \geq 1$.*

Proof. Columns of \mathbf{m}_k are given by $\sum_{j \in J} A_{*,j}$ where $|J| = k$. The entries of A are either 0 or 1. So the largest entry in \mathbf{m}_k is $1 + \dots + 1 = k$. The result follows by appealing to the binary expansion of k . \square

We expect $\text{rank}_p(\mathbf{m}_{2^i}^{[i]}) \leq \binom{r}{2^i}$ since $\mathbf{m}_{2^i}^{[i]}$ is a matrix of dimension $2^r \times \binom{r}{2^i}$. The next lemma shows that the rank is, in fact, equal to this upper bound.

Lemma 6.8. *$\text{rank}_p(\mathbf{m}_{2^i}^{[i]}) = \binom{r}{2^i}$ for all $i \geq 1$.*

Proof. Let $c_1, \dots, c_{\binom{r}{2^i}}$ be the column indices of \mathbf{m}_{2^i} in M . Let $S(\mathbf{m}_{2^i})$ be the $\binom{r}{2^i} \times \binom{r}{2^i}$ submatrix of \mathbf{m}_{2^i} formed by the rows $c_1, \dots, c_{\binom{r}{2^i}}$, and $S(A)$ be the $\binom{r}{2^i} \times r$ submatrix of A formed by the rows $c_1, \dots, c_{\binom{r}{2^i}}$. Rows of $S(A)$ have exactly 2^i non-zero entries because of the construction of A . If we treat A and M as block matrices then $S(\mathbf{m}_{2^i}) = S(A)S(A)^T$ is the 2^i th diagonal block of M (See Figure 6.1).

The entries in row ρ of $S(\mathbf{m}_{2^i})$ are given by linear combinations of the entries in row ρ of $S(A)$. The summing index sets J_j , where $|J_j| = 2^i$, are exactly the locations of the non-zero entries of rows of $S(A)$, which are all *different* by construction. Hence there is only one entry in row ρ of $S(\mathbf{m}_{2^i})$ whose summing set matches the locations of the non-zero entries in row ρ of $S(A)$. The value of this entry is $1 + 1 + \dots + 1 = 2^i$. The other entries have values less than 2^i . The binary expansion of 2^i gives us that $S(\mathbf{m}_{2^i}^{[i]})$ is an identity (sub)matrix* of $\mathbf{m}_{2^i}^{[i]}$ whose size is $\binom{r}{2^i} \times \binom{r}{2^i}$. Therefore, $\mathbf{m}_{2^i}^{[i]}$ has rank $\binom{r}{2^i}$. \square

*This is true in the example of Figure 6.1 without any reordering, because we constructed the row blocks of A such that the binary expansion of i comes after the binary expansion of j whenever $i > j$. Without such ordering, the identity block assertion holds up to row and column permutations.

Next we will prove that $\text{rank}_p(M^{[i]}) = \text{rank}_p(\mathbf{m}_{2^i}^{[i]})$ by showing that all the columns of $\mathbf{m}_{2^i+1}^{[i]}, \mathbf{m}_{2^i+2}^{[i]}, \dots, \mathbf{m}_r^{[i]}$ are linearly *dependent* on those of $\mathbf{m}_{2^i}^{[i]}$. First, we need the following auxiliary statements.

Fact 6.1 (Kummer's Theorem). *The exact power of p dividing $\binom{a+b}{a}$ is equal to the number of carries when performing the addition of $(a+b)$ written in base p .*

A corollary of Kummer's theorem is that $\binom{a+b}{a}$ is odd (resp. even) if adding $(a+b)$ written in binary expansion generates no (resp. some) carries.

Lemma 6.9. $(2^i + k) \text{ quo } 2^i \equiv \binom{2^i+k}{2^i} \pmod{2}$.

Proof. We will show that $(2^i + k) \text{ quo } 2^i$ and $\binom{2^i+k}{2^i}$ have the same parity and hence equivalent modulo 2. Write $k = Q2^i + R$ for a quotient $Q \geq 0$ and a remainder $0 \leq R < 2^i$.

If Q is even, then the i th bit of k (i.e., the coefficient of 2^i in the binary expansion of k) is zero and hence no carries are generated when adding k and 2^i in base 2. So by Kummer's Theorem, $\binom{2^i+k}{2^i}$ is odd. If Q is odd, then the i th bit of k is 1 and the number of carries generated when adding $2^i + k$ in base 2 is at least 1. So by Kummer's theorem $\binom{2^i+k}{2^i}$ is even. Putting the two cases together implies that $\binom{2^i+k}{2^i} \equiv Q + 1 \pmod{2}$.

From $k = Q2^i + R$ we get $(2^i + k) \text{ quo } 2^i = (2^i + Q2^i + R) \text{ quo } 2^i = Q + 1$. Therefore $\binom{2^i+k}{2^i} \equiv (2^i + k) \text{ quo } 2^i \pmod{2}$. \square

Now we are ready to show that the ranks only depend on $\mathbf{m}_{2^i}^{[i]}$.

Lemma 6.10. *Consider any column m in \mathbf{m}_{2^i+z} , where $z \geq 1$. Then $m^{[i]}$ is a linear combination of columns of $\mathbf{m}_{2^i}^{[i]}$.*

Proof. Let J be the summing index set of m , where $|J| = 2^i + z$. Let \mathcal{I} be the set of all subsets of J of size 2^i , so $|\mathcal{I}| = \binom{2^i+z}{2^i}$. For every $I \in \mathcal{I}$, there is a unique corresponding column c_I in \mathbf{m}_{2^i} whose summing set is I . We will show that $m^{[i]}$ can be obtained by adding up c_I 's. In other words,

$$m^{[i]} \equiv \sum_{I \in \mathcal{I}} c_I^{[i]} \pmod{2}. \quad (6.18)$$

Let A_J denote the submatrix of A formed by the columns indexed by J . For any row ρ of A_J , let $2^i + k_\rho$ be the number of 1's in that row, where $-2^i \leq k_\rho \leq z$. First, if $k_\rho < 0$, then the corresponding sum of 1's at this row is less than 2^i . By Lemma 6.7, we have the corresponding entries in both $\mathbf{m}_{2^i}^{[i]}$ and $\mathbf{m}_{2^i+z}^{[i]}$ are zeros and (6.18) trivially holds. On the other hand, if $0 \leq k_\rho \leq z$, then the ρ th entry of the right-hand side of (6.18) is $1 + 1 + \dots + 1 \equiv \binom{2^i+k_\rho}{2^i} \pmod{2}$ since $|\mathcal{I}| = \binom{2^i+k_\rho}{2^i}$. (Recall that the number of non-zero entries in row ρ is $2^i + k_\rho$ rather than $2^i + z$.) The ρ th entry of the left-hand side of (6.18) is $(2^i + k_\rho) \text{ quo } 2^i$. The $(2^i + k_\rho)$ term corresponds to adding $(2^i + k_\rho)$ non-zero entries, and the $\text{quo } 2^i$ operation corresponds to the i th bit of the binary expansion of m . By Lemma 6.9, we have $(2^i + k_\rho) \text{ quo } 2^i \equiv \binom{2^i+k_\rho}{2^i} \pmod{2}$, and (6.18) holds. \square

6.3.2 Non-Symmetric Matrices

So far we have shown that $\text{rank}_p(M^{[i]}) = \text{rank}_p(\mathbf{m}_{2^i}^{[i]}) = \binom{r}{2^i}$, where $M = AA^T$ for some specially constructed A . We now put the results together into a more general form. Let $A \in \mathbb{Z}^{2^r \times r}$ be the binary code matrix as before.

Lemma 6.11. *Assume $U, S \in \mathbb{Z}^{n \times n}$, such that U is symmetric and has entries from $\{0, 1\}$, $\det U \not\equiv 0 \pmod{2}$, $S = \text{diag}(1, \dots, 1, 0, \dots, 0)$, $\text{rank}_p(S) = r$, and $n \geq 2^r$. If $M = USU \in \mathbb{Z}^{n \times n}$, then $\text{rank}_p(M^{[i]}) \leq \binom{r}{2^i}$ for all $i \geq 1$.*

Proof. Since $S = SS$, we have $M = USSU = LL^T$ where $L = US$. Apply row operations to L to annihilate the duplicate the rows. Since L has only 0, 1 entries, it must be that the number of unique non-zero rows at most 2^r . Reorder the rows of L such that row i in L is either identical to row i in A or zero (i.e., missing from L). Ensure that L has exactly 2^r rows by appending zero rows at the end, or removing zero rows. Let \bar{L} denote the resulting matrix. Note that rank of \bar{L} is equal to rank of L , and that \bar{L} has the same structure as A , where some rows might be replaced by zero.

Let $\bar{M} = \bar{L}\bar{L}^T = \begin{bmatrix} \bar{\mathbf{m}}_0 & \bar{\mathbf{m}}_1 & \dots & \bar{\mathbf{m}}_r \end{bmatrix}$. Now \bar{M} and AA^T have the same dimensions and structure. Every entry of \bar{M} is either identical to the corresponding entry in AA^T or zero. Essentially, \bar{M} is a copy of AA^T with the difference that some rows (and columns) might

have been entirely replaced by zero. This implies that $\bar{M}^{[i]}$ can be obtained by copying $(AA^T)^{[i]}$ and potentially replacing some rows and columns by zero. So $\text{rank}_p(\bar{M}^{[i]}) \leq \text{rank}_p((AA^T)^{[i]}) = \binom{r}{2^i}$. The lemma statement holds since we obtained \bar{M} from M using column and row operations, and augmenting and omitting zero rows and columns at the end of the matrix. \square

We can now generalize the result to non-symmetric matrices.

Theorem 6.2. *Assume $U, S, V \in \mathbb{Z}^{n \times n}$, such that U, V have entries from $\{0, 1\}$, $\det U, \det V$ are non-zero modulo 2, $S = \text{diag}(1, \dots, 1, 0, \dots, 0)$, $\text{rank}_p(S) = r$, and $n \geq 2^r$. If $M = USV \in \mathbb{Z}^{n \times n}$, then $\text{rank}_p(M^{[i]}) \leq \binom{r}{2^i}$ for all $i \geq 1$.*

Proof. Since $S = SS$, we have $M = USSV = LR$ where $L = US$ and $R = SV$. Apply row operations to L to annihilate the duplicate the rows from L . Since L have 0-1 entries, there are at most 2^r unique rows in L . Adjust the number of rows of L to 2^r by appending zero rows, or omitting the extra zero rows. Finally, reorder the rows of L such that L is a copy of A where zero or more rows being replaced by zero. Let \bar{L} denote the resulting matrix. Apply similar operations on the columns of R , and let \bar{R} denote the resulting matrix. Then $\bar{M} = \bar{L}\bar{R}$ is essentially a copy of AA^T where some rows are replaced by zero, and some columns are replaced by the zero vector.

Similarly, $\bar{M}^{[i]}$ can be obtained by copying $(AA^T)^{[i]}$ and potentially replacing some rows and columns by zeros. So $\text{rank}_p(\bar{M}^{[i]}) \leq \text{rank}_p((AA^T)^{[i]}) = \binom{r}{2^i}$. The result holds since we obtained \bar{M} from M using column and row operations, and augmenting and omitting zero rows and columns at the end of the matrix. \square

6.3.3 Odd Primes

For $p = 2$, the non-zero patterns of the binary code matrix A coincide with the summing indices in (6.15). This is not true for odd primes, where the linear combinations can have coefficients other than 0 and 1. It is an open question to devise a construction for odd primes (similar to the binary code matrices) which exposes the combinatorial structure of

the column space of $M = AA^T$. We present the following conjecture towards understanding the p -adic ranks for odd primes.

Conjecture 6.1. *Assume $p = 2k + 1$ is an odd prime, $U, S, V \in \mathbb{Z}^{n \times n}$ such that U, V have entries from $[0, p - 1]$, $\det U \det V \not\equiv 0 \pmod{p}$, S is a $0, 1$ diagonal matrix and $\text{rank}_p(S) = r$. Let $M = USV = M^{[0]} + M^{[1]}p + \dots$ where $M^{[i]} \in (\mathbb{Z}/p\mathbb{Z})^{n \times n}$. It is conjectured that*

$$\text{rank}_p(M^{[1]}) \leq \sum_{i=0}^k \binom{r + 2i}{2i + 1} + \binom{r + 2k - 1}{2k} - 2r \quad (6.19)$$

Furthermore, in the generic case where the entries of U, V are uniformly chosen uniformly at random from $[0, p - 1]$, and n is arbitrarily large, the ranks are equal to the stated bound.

This conjecture first appeared in [Elsheikh et al., 2012]. It shows that a product of matrices with “small” entries and “small” rank can still have very large rank, but not full, p -adic expansion. In other words, the “carries” from the product USV will impact many digits in the expanded product.

Chapter 7

Conclusion and Future Work

In this thesis we have studied the problem of computing the Smith normal form of sparse matrices over local rings and related problems. We have also extended the application of the Wedderburn rank reduction process to computing the Smith normal form, and nullspace sampling. On the other hand, we studied some of the related notions of matrix invariants and matrix properties. In particular, we gave a characterization of the invariants factors in terms of the spectrum. Finally we have started an interesting study of ranks of remainder and quotient matrices and ranks of base- p slices of matrices.

We believe that an important open problem is designing a quadratic-time algorithm to substitute for Wiedemann's method when working over local rings. Such an algorithm will lead to significant improvements in computation over local rings and over the non-localized rings as well. In this short conclusion, we will summarize the major problems we presented in this thesis.

Problem 1. *Given a sparse polynomial matrix $A \in \mathbb{F}[x]^{n \times n}$, and an irreducible polynomial $f \in \mathbb{F}[x]$ of degree d , find the local Smith normal form of A at f .*

In Chapter 2 we examined this problem for sparse matrices over $\mathbb{F}[x]/(f^e)$. We gave an algorithm to compute the invariant factors, which tries to minimize the fill-in and the expression swell in the intermediate computations. The approach we took is *linearization*. We transform the $n \times n$ polynomial matrix into a $den \times den$ matrix over the field \mathbb{F} . We

established the relationships between the multiplicities of the invariant factors and the rank of the corresponding matrix over the ground field. The linearization takes advantage of existing fast algorithms for computing ranks of sparse matrices over finite fields. If A has a linear number of entries, then the cost of evaluating the black-box for A is $\mu \in O(\text{den})$ operations in F . In this case, Algorithm 2.1 has complexity $O(n^2 d^2 e^3)$ operations in F . We expect an efficient algorithm in this setup to have complexity relative to $O(n^2 de)$. An interesting open question is to reduce the complexity of our algorithm by a factor of de or de^2 . In particular, Algorithm 2.1 computes Krylov subspace iterates for e different embeddings. One can hope to compute these iterates only once modulo f^e , and then reuse these iterates modulo powers of f . Such improvement would achieve a reduction in complexity by factor of at least e .

Problem 2. *Given a sparse integer matrix $A \in \mathbb{Z}^{n \times n}$, and a prime p , find the local Smith normal form of A at p .*

We examined this problem in Chapters 2, 3, and 4. The integer case is arguably harder than the polynomial case since carries can interfere with base- p expansion of the problem.

In Chapter 2 we proposed a hybrid sparse-dense algorithm which uses sparse nullspace sampling followed by dense elimination to compute the local invariant factors. The cost of Algorithm 2.2 is dominated by $O(kn\mu)$ operations in $\mathbb{Z}/p^e\mathbb{Z}$, where k is the number of non-trivial invariant factors, μ is the cost of evaluating the black-box of A modulo p . The algorithm is useful for the case when there are only few non-trivial invariant factors, i.e., when $\ell \ll n$. In some applications, the goal is to detect the first non-trivial invariant factor. Algorithm 2.4 solves this problem by lifting nullspace vectors. The cost is $O(en\mu)$ operations in $\mathbb{Z}/p^e\mathbb{Z}$.

In Chapter 4 we discussed a new approach to computing the Smith normal form. We applied an iterative elimination process using rank-1 and rank- k updates over $\mathbb{Z}/p^e\mathbb{Z}$. We showed that rank updates given by the Wedderburn rank reduction formula can modify the invariant factors in a controllable manner; it decreases the number of 1s and increases the 0s in the Smith normal form. We presented two algorithms. The iterative vector-based algorithm has complexity $O(n^3)$ operations over $\mathbb{Z}/p^e\mathbb{Z}$, and requires storage of $O(n^2)$

elements from $\mathbb{Z}/p^e\mathbb{Z}$. Clearly, this algorithm is not suitable for sparse matrices. On the other hand, the block version of this algorithm leverages sub-cubic matrix multiplication, and has an expected cost of $O(e^2 n^\omega)$ operations in $\mathbb{Z}/p^e\mathbb{Z}$. At the core of Algorithm 4.2 is a selection of two matrices X, Y such that $Y^T A X$ is invertible. It is an interesting open problem to use structured matrices for X, Y such that $(Y^T A X)^{-1}$ can be computed faster than n^ω and the resulting matrix has a linear or quasi-linear time black-box. Such a construction could potentially reduce the overall complexity below n^ω , and extend the applications of the algorithm to sparse matrices.

Problem 3. *Given a sparse integer matrix $A \in \mathbb{F}^{n \times n}$ of rank r , compute a random sample from the nullspace of A .*

Nullspace sampling arose as a related problem in some of our algorithms. In Chapter 3 we explored preconditioning matrices for sparse Smith normal form computation. In particular, we showed how to extend the application of well-known linear-time preconditioners to sparse Smith normal form. Additionally, we showed how to use the diagonal transpose preconditioner $A \mapsto D_1 A^T D_2 A$, and the Wiedemann sparse preconditioner to compute a random sample from the nullspace of A . The cost of Algorithm 3.1 is $O(r\mu + r^2)$ operations in \mathbb{F} . The algorithm is randomized Monte Carlo and its success probability requires the $|\mathbb{F}| \in O(n)$. For smaller primes, we can use field extensions with additional factor of $M(\log(n/\epsilon))$ in the cost. We can avoid constructing extension fields by using Wiedemann preconditioners and the cost becomes $O(r\mu + rn \log^2 n)$ operations in \mathbb{F} .

An open problem is to find linear-time preconditioners over small fields. This would eliminate the need for field extensions, and drive the complexity closer to $O(n^2)$.

On the other hand, in Chapter 4, we discovered an additional application of the rank- k reduction procedure to the problem of nullspace sampling. Algorithm 4.3 computes a random sample from the nullspace of A with expected cost of $O^\sim(n^\omega)$ operations in \mathbb{F} . As a future work, we outlined the potential speedup of this algorithm using Krylov matrices, in which the cost would be $O^\sim(n\mu)$ operations in \mathbb{F} , which can be sensitive to the sparsity of the input matrix.

Problem 4. *Let A be an $n \times n$ integer matrix, and let $\lambda_1, \dots, \lambda_n$ be its eigenvalues, and s_1, \dots, s_n be its invariant factors. What is the relationship between the eigenvalues and the invariant factors?*

In Chapter 5, we discussed this relationship at a given prime p . We presented a new characterization using the p -adic valuations as a measure of size. For our setup, we view the eigenvalues as p -adic algebraic integers in a finite-degree extension over \mathbb{Q}_p . We show that for most matrices there is a 1-1 correspondence between the p -adic valuations of the eigenvalues and the powers of p dividing the invariant factors. In particular, our results imply that this correspondence holds with high probability for random integer matrices whose entries are sampled uniformly from a large enough range. This result holds if p is large compared to n . This is mostly an artifact of our proofs which rely heavily on the Schwartz-Zippel lemma. Our numerical experiments suggest that the results are also true for smaller primes. It remains open to prove similar density estimates for small primes.

Problem 5. *Let A be an $n \times n$ integer matrix of rank r . Let p be a prime, and write $A = p(A \text{ quo } p) + (A \text{ rem } p)$ as the remainder and quotient expansions of A . Furthermore, let $A = A_0 + pA_1 + p^2A_2 + \dots$ be the element-wise p -adic expansion of A . What is the relationship between the rank of A and the ranks of the expansion matrices?*

This problem is somewhat of an independent interest. We first encountered this problem while attempting a p -adic approach to expanding the Smith normal form $A = USV$ into

$$A_0 + pA_1 + \dots = (U_0 + pU_1 + \dots)(S_0 + pS_1 + \dots)(V_0 + pV_1 + \dots).$$

In Chapter 6 we were able to show that $\text{rank}(A \text{ rem } p)$ can be arbitrarily large even if A has a small number of non-zero invariant factors. In particular, the upper bound is $(p^{r_0} - 1)(p + 1)/(2(p - 1))$ where r_0 is the number of invariant factors not divisible by p . Similarly, the rank of A_i when $p = 2$ can be as large as $\binom{r}{2^i}$. Our numerical experiments suggest that random matrices will attain this bound. Finally, we presented a conjecture for the p -adic ranks when $p > 2$.

References

- [Abaffy et al., 1984] Abaffy, J., Broyden, C., and Spedicato, E. (1984). A class of direct methods for linear systems. *Numerische Mathematik*, 45(3):361–376.
- [Abbott et al., 1999] Abbott, J., Bronstein, M., and Mulders, T. (1999). Fast deterministic computation of determinants of dense matrices. In *Proceedings of the 1999 International Symposium on Symbolic and Algebraic Computation*, ISSAC '99, pages 197–204, New York, NY, USA. ACM.
- [Alon and Yuster, 2010] Alon, N. and Yuster, R. (2010). Solving linear systems through nested dissection. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, pages 225–234.
- [Artin, 1991] Artin, M. (1991). *Algebra*. Featured Titles for Abstract Algebra Series. Prentice Hall.
- [Babson et al., 1999] Babson, E., Björner, A., Linusson, S., Shareshian, J., and Welker, V. (1999). Complexes of not i -connected graphs. *Topology*, 38(2):271–299.
- [Björner and Welker, 1999] Björner, A. and Welker, V. (1999). Complexes of directed graphs. *SIAM Journal on Discrete Mathematics*, 12(4):413–424.
- [Bostan et al., 2003] Bostan, A., Lecerf, G., and Schost, E. (2003). Tellegen’s principle into practice. In *Proceedings of the 2003 International Symposium on Symbolic and Algebraic Computation*, ISSAC '03, pages 37–44, New York, NY, USA. ACM.
- [Brauer, 1946] Brauer, A. (1946). Limits for the characteristic roots of a matrix. *Duke Math. J.*, 13(3):387–395.

- [Bunch and Hopcroft, 1974] Bunch, J. R. and Hopcroft, J. E. (1974). Triangular factorization and inversion by fast matrix multiplication. *Mathematics of Computation*, 28(125):231–236.
- [Carothers, 2000] Carothers, N. (2000). *Real Analysis*. Cambridge University Press.
- [Chandler et al., 2010] Chandler, D. B., Sin, P., and Xiang, Q. (2010). Incidence modules for symplectic spaces in characteristic two. *J. Algebra*, 323:3157–3181.
- [Chen et al., 2002] Chen, L., Eberly, W., Kaltofen, E., Saunders, B. D., Turner, W. J., and Villard, G. (2002). Efficient matrix preconditioners for black box linear algebra. *Linear Algebra and its Applications*, 343–344:119–146.
- [Chou and Collins, 1982] Chou, T. J. and Collins, G. E. (1982). Algorithms for the solution of systems of linear diophantine equations. *SIAM Journal on Computing*, 11(4):687–708.
- [Chu et al., 1995] Chu, M. T., Funderlic, R. E., and Golub, G. H. (1995). A rank-one reduction formula and its applications to matrix factorizations. *SIAM Rev.*, 37(4):512–530.
- [Cline and Funderlic, 1979] Cline, R. and Funderlic, R. (1979). The rank of a difference of matrices and associated generalized inverses. *Linear Algebra and its Applications*, 24:185–215.
- [Coppersmith, 1994] Coppersmith, D. (1994). Solving homogeneous linear equations over $\text{GF}(2)$ via block Wiedemann algorithm. *Mathematics of Computation*, 62(205):333–350.
- [Demillo and Lipton, 1978] Demillo, R. A. and Lipton, R. J. (1978). A probabilistic remark on algebraic program testing. *Information Processing Letters*, 7(4):193–195.
- [Dixon, 1982] Dixon, J. D. (1982). Exact solution of linear equations using p -adic expansions. *Numerische Mathematik*, 40:137–141.
- [Dumas et al., 2003] Dumas, J.-G., Heckenbach, F., Saunders, B. D., and Welker, V. (2003). Computing simplicial homology based on efficient Smith normal form algorithms. In *Algebra, Geometry, and Software Systems*, pages 177–206. Springer.

- [Dumas et al., 2000] Dumas, J.-G., Saunders, B. D., and Villard, G. (2000). Integer Smith form via the valence: Experience with large sparse matrices from homology. In *Proc. 2000 Internat. Symp. Symbolic Algebraic Comput. ISSAC'00*, pages 95–105. ACM Press.
- [Dumas et al., 2001] Dumas, J.-G., Saunders, B. D., and Villard, G. (2001). On efficient sparse integer matrix Smith normal form computations. *Journal of Symbolic Computation*, 32:71–99.
- [Dumas and Villard, 2002] Dumas, J.-G. and Villard, G. (2002). Computing the rank of large sparse matrices over finite fields. In *Computer Algebra in Scientific Computing (CASC)*, pages 47–61.
- [Eberly, 2000] Eberly, W. (2000). Black box frobenius decompositions over small fields. In *Proceedings of the 2000 International Symposium on Symbolic and Algebraic Computation*, ISSAC '00, pages 106–113, New York, NY, USA. ACM.
- [Eberly, 2003] Eberly, W. (2003). Early termination over small fields. In *Proceedings of the 2003 International Symposium on Symbolic and Algebraic Computation*, ISSAC '03, pages 80–87, New York, NY, USA. ACM.
- [Eberly, 2004] Eberly, W. (2004). Reliable Krylov-based algorithms for matrix null space and rank. In *Proceedings of the 2004 International Symposium on Symbolic and Algebraic Computation*, ISSAC '04, pages 127–134, New York, NY, USA. ACM.
- [Eberly et al., 2007] Eberly, W., Giesbrecht, M., Giorgi, P., Storjohann, A., and Villard, G. (2007). Faster inversion and other black box matrix computations using efficient block projections. In *Proc. International Symposium on Symbolic and Algebraic Computation (ISSAC'07)*, pages 143–150.
- [Eberly et al., 2000] Eberly, W., Giesbrecht, M., and Villard, G. (2000). On computing the determinant and Smith form of an integer matrix. In *Proc. 41st Annual IEEE Symposium on Foundations of Computer Science (FOCS'2000)*, pages 675–687.

- [Eberly and Kaltofen, 1997] Eberly, W. and Kaltofen, E. (1997). On randomized Lanczos algorithms. In *Proceedings of the 1997 International Symposium on Symbolic and Algebraic Computation*, ISSAC '97, pages 176–183, New York, NY, USA. ACM.
- [Elsheikh and Giesbrecht, 2015] Elsheikh, M. and Giesbrecht, M. (2015). Relating p-adic eigenvalues and the local smith normal form. *Linear Algebra and its Applications*, 481:330–349.
- [Elsheikh et al., 2012] Elsheikh, M., Giesbrecht, M., Novocin, A., and Saunders, B. D. (2012). Fast computation of Smith forms of sparse matrices over local rings. In *Proceedings of the 37th International Symposium on Symbolic and Algebraic Computation*, ISSAC '12, pages 146–153, New York, NY, USA. ACM.
- [Galántai, 2010] Galántai, A. (2010). The rank reduction procedure of Egerváry. *Central European Journal of Operations Research*, 18(1):5–24.
- [Geddes et al., 1992] Geddes, K., Czapor, S., and Labahn, G. (1992). *Algorithms for Computer Algebra*. Kluwer Academic Publishers.
- [George, 1973] George, A. (1973). Nested dissection of a regular finite element mesh. *SIAM Journal on Numerical Analysis*, 10(2):345–363.
- [Gerstein, 1977] Gerstein, L. J. (1977). A local approach to matrix equivalence. *Linear Algebra and its Applications*, 16(3):221 – 232.
- [Giesbrecht, 1995] Giesbrecht, M. (1995). Fast computation of the Smith normal form of an integer matrix. In *Proceedings of the 1995 International Symposium on Symbolic and Algebraic Computation*, ISSAC '95, pages 110–118, New York, NY, USA. ACM.
- [Giesbrecht, 1996] Giesbrecht, M. (1996). Probabilistic computation of the Smith normal form of a sparse integer matrix. In Cohen, H., editor, *Algorithmic Number Theory*, volume 1122 of *Lecture Notes in Computer Science*, pages 173–186. Springer Berlin Heidelberg.

- [Giesbrecht, 2001] Giesbrecht, M. (2001). Fast computation of the Smith form of a sparse integer matrix. *Computational Complexity*, 10(1):41–69.
- [Gohberg and Fel'dman, 1974] Gohberg, I. C. and Fel'dman, I. A. (1974). Convolution equations and projection methods for their solution. *Translations of Mathematical Monographs*, 41.
- [Gouvêa, 1997] Gouvêa, F. Q. (1997). *p-adic Numbers: an introduction*. Springer-Verlag Berlin, 2nd edition.
- [Hafner and McCurley, 1989] Hafner, J. L. and McCurley, K. S. (1989). A rigorous subexponential algorithm for computation of class groups. *Journal of the American Mathematical Society*, 2(4):837–850.
- [Householder, 1964] Householder, A. S. (1964). *The Theory of Matrices in Numerical Analysis*. Dover Books on Mathematics. Dover Publications.
- [Hovinen and Eberly, 2005] Hovinen, B. and Eberly, W. (2005). A reliable block Lanczos algorithm over small finite fields. In *Proceedings of the 2005 International Symposium on Symbolic and Algebraic Computation*, ISSAC '05, pages 177–184, New York, NY, USA. ACM.
- [Hu, 1969] Hu, T. C. (1969). *Integer programming and network flows*. Addison-Wesley, Reading, Mass.
- [Kailath, 1980] Kailath, T. (1980). *Linear Systems*. Prentice-Hall, Englewood Cliffs, NJ.
- [Kaltofen, 1995] Kaltofen, E. (1995). Analysis of Coppersmith's block Wiedemann algorithm for the parallel solution of sparse linear systems. *Mathematics of Computation*, 64(210):777–806.
- [Kaltofen et al., 1987] Kaltofen, E., Krishnamoorthy, M., and Saunders, B. D. (1987). Fast parallel computation of hermite and smith forms of polynomial matrices. *SIAM. J. on Algebraic and Discrete Methods*, 8:683–690.

- [Kaltofen et al., 2000] Kaltofen, E., Lee, W.-S., and Lobo, A. A. (2000). Early termination in Ben-Or/Tiwari sparse interpolation and a hybrid of Zippel’s algorithm. In *Proceedings of the 2000 International Symposium on Symbolic and Algebraic Computation, ISSAC ’00*, pages 192–201, New York, NY, USA. ACM.
- [Kaltofen and Saunders, 1991] Kaltofen, E. and Saunders, B. D. (1991). On Wiedemann’s method of solving sparse linear systems. In *Applied Algebra, Algebraic Algorithms and Error-Correcting Codes (AAECC ’91)*, volume 539 of *LNCS*, pages 29–38.
- [Kaltofen and Villard, 2005] Kaltofen, E. and Villard, G. (2005). On the complexity of computing determinants. *Computational Complexity*, 13(3–4):91–130.
- [Kaplansky, 1949] Kaplansky, I. (1949). Elementary divisors and modules. *Transactions of the American Mathematical Society*, 66(2):464–491.
- [Kasami et al., 1968] Kasami, T., Lin, S., and Peterson, W. (1968). New generalizations of the Reed-Muller codes—I: Primitive codes. *IEEE Transactions on Information Theory*, 14(2):189–199.
- [Koblitz, 1984] Koblitz, N. (1984). *p-adic Numbers, p-adic Analysis and Zeta-Functions*. Graduate Texts in Mathematics. Springer-Verlag, 2nd edition.
- [Labahn et al., 1990] Labahn, G., Choi, D. K., and Cabay, S. (1990). The inverses of block hankel and block toeplitz matrices. *SIAM Journal on Computing*, 19(1):98–123.
- [Labahn and Shalom, 1992] Labahn, G. and Shalom, T. (1992). Inversion of toeplitz matrices with only two standard equations. *Linear Algebra and its Applications*, 175:143–158.
- [Lambert, 1996] Lambert, R. (1996). *Computational Aspects of Discrete Logarithms*. PhD thesis, University of Waterloo, Waterloo, Ontario, Canada.
- [Le Gall, 2014] Le Gall, F. (2014). Powers of tensors and fast matrix multiplication. In *Proceedings of the 39th International Symposium on Symbolic and Algebraic Computation, ISSAC ’14*, pages 296–303, New York, NY, USA. ACM.

- [Lidl and Niederreiter, 1986] Lidl, R. and Niederreiter, H. (1986). *Introduction to finite fields and their applications*. Cambridge University Press, New York, NY, USA.
- [Lipton et al., 1979] Lipton, R. J., Rose, D. J., and Tarjan, R. E. (1979). Generalized nested dissection. *SIAM Journal on Numerical Analysis*, 16(2):346–358.
- [Lobo, 1995] Lobo, A. A. (1995). *Matrix-free linear system solving and applications to symbolic computation*. PhD thesis, Rensselaer Polytechnic Institute, Troy, NY, USA.
- [Lorenzini, 2008] Lorenzini, D. (2008). Smith normal form and Laplacians. *Journal of Combinatorial Theory, Series B*, 98(6):1271–1300.
- [Lübeck, 2002] Lübeck, F. (2002). On the computation of elementary divisors of integer matrices. *Journal of Symbolic Computation*, 33(1):57–65.
- [McDonald, 1974] McDonald, B. (1974). *Finite Rings with Identity*. Marcel Dekker, Inc., New York.
- [McMillan, 1952] McMillan, B. (1952). On systems of linear indeterminate equations and congruences. *Bell System Technical Journal*, 31:541–600.
- [Meyer, 2000] Meyer, C. (2000). *Matrix Analysis and Applied Linear Algebra*. Society for Industrial and Applied Mathematics.
- [Motwani and Raghavan, 1995] Motwani, R. and Raghavan, P. (1995). *Randomized Algorithms*. Cambridge University Press.
- [Newman, 1972] Newman, M. (1972). *Integral Matrices*. Academic Press, New York, NY, USA.
- [Newman and Thompson, 1991] Newman, M. and Thompson, R. C. (1991). Matrices over rings of algebraic integers. *Linear Algebra and its Applications*, 145:1–20.
- [Pan, 1988] Pan, V. (1988). Computing the determinant and the characteristic polynomial of a matrix via solving linear systems of equations. *Information Processing Letters*, 28(2):71–75.

- [Penfield Jr. et al., 1970] Penfield Jr., P., Spencer, R., and Duinker, S. (1970). *Tellegen's Theorem and Electrical Networks*. M.I.T. Press, Cambridge, MA, USA.
- [Raboky and Amiri, 2013a] Raboky, E. G. and Amiri, N. M. (2013a). Extended integer rank reduction formulas and smith normal form. *Linear and Multilinear Algebra*, 61(12):1641–1659.
- [Raboky and Amiri, 2013b] Raboky, E. G. and Amiri, N. M. (2013b). Extended rank reduction formulas containing Wedderburn and Abaffy-Broyden-Spedicato rank reducing processes. *Linear Algebra and its Applications*, 439(11):3318–3331.
- [Rushanan, 1995] Rushanan, J. J. (1995). Eigenvalues and the Smith normal form. *Linear Algebra and its Applications*, 216:177–184.
- [Schwartz, 1980] Schwartz, J. T. (1980). Fast probabilistic algorithms for verification of polynomial identities. *Journal of the ACM*, 27(4):701–717.
- [Shoup, 1994] Shoup, V. (1994). Fast construction of irreducible polynomials over finite fields. *Journal of Symbolic Computation*, 17(5):371–391.
- [Shoup, 1999] Shoup, V. (1999). Efficient computation of minimal polynomials in algebraic extensions of finite fields. In *Proceedings of the 1999 International Symposium on Symbolic and Algebraic Computation*, ISSAC '99, pages 53–58, New York, NY, USA. ACM.
- [Smith, 1861] Smith, H. J. S. (1861). On systems of linear indeterminate equations and congruences. *Philosophical Transactions of the Royal Society of London*, 151:293–326.
- [Stein et al., 2014] Stein, W. et al. (2014). *Sage Mathematics Software (Version 6.2.beta1)*. The Sage Development Team. <http://www.sagemath.org>.
- [Storjohann, 2000] Storjohann, A. (2000). *Algorithms for matrix canonical forms*. PhD thesis, Swiss Federal Institute of Technology – ETH, Zürich.
- [Storjohann, 2003] Storjohann, A. (2003). High-order lifting and integrality certification. *Journal of Symbolic Computation*, 36(3–4):613–648.

- [Storjohann and Labahn, 1997] Storjohann, A. and Labahn, G. (1997). A fast las vegas algorithm for computing the smith normal form of a polynomial matrix. *Linear Algebra and its Applications*, 253(1):155–173.
- [Strassen, 1969] Strassen, V. (1969). Gaussian elimination is not optimal. *Numer. Math.*, 13(4):354–356.
- [Tellegen, 1952] Tellegen, B. (1952). A general network theorem, with applications. *Philips Research Reports*, 7:259–269.
- [Villard, 1997] Villard, G. (1997). Further analysis of Coppersmith’s block Wiedemann algorithm for the solution of sparse linear systems (extended abstract). In *Proceedings of the 1997 International Symposium on Symbolic and Algebraic Computation, ISSAC ’97*, pages 32–39, New York, NY, USA. ACM.
- [Villard, 2000] Villard, G. (2000). Computing the Frobenius normal form of a sparse matrix. In Ganzha, V., Mayr, E., and Vorozhtsov, E., editors, *Computer Algebra in Scientific Computing*, pages 395–407. Springer Berlin Heidelberg.
- [von zur Gathen and Gerhard, 2003] von zur Gathen, J. and Gerhard, J. (2003). *Modern Computer Algebra*. Cambridge University Press, New York, NY, USA, 2nd edition.
- [Wallis et al., 1972] Wallis, W., Street, A. P., and Wallis, J. S. (1972). *Combinatorics: room squares, sum-free sets, Hadamard matrices*. Lecture notes in mathematics. Springer, Berlin.
- [Wedderburn, 1934] Wedderburn, J. H. M. (1934). *Lectures on Matrices*, volume XVII of *Lecture notes in mathematics. Colloquium Publications*. American Mathematical Society, New York.
- [Wiedemann, 1986] Wiedemann, D. (1986). Solving sparse linear equations over finite fields. *IEEE Transactions on Information Theory*, IT-32:54–62.
- [Wilkening and Yu, 2011] Wilkening, J. and Yu, J. (2011). A local construction of the Smith normal form of a matrix polynomial. *Journal of Symbolic Computation*, 46(1):1–22.

- [Yannakakis, 1981] Yannakakis, M. (1981). Computing the minimum fill-in is NP-complete. *SIAM Journal on Algebraic Discrete Methods*, 2(1):77–79.
- [Zhou et al., 2015] Zhou, W., Labahn, G., and Storjohann, A. (2015). A deterministic algorithm for inverting a polynomial matrix. *Journal of Complexity*, 31(2):162–173.
- [Zippel, 1979] Zippel, R. (1979). Probabilistic algorithms for sparse polynomials. In *Symbolic and Algebraic Computation*, volume 72 of *LNCS*, pages 216–226. Springer Berlin.