On the Probabilistic Analysis of Normal Form Computation of a Sparse Matrix

Bruce R. Donald*
David Renpan Chang**

TR 90-1180
December 1990

Department of Computer Science
Cornell University
Ithaca, NY 14853-7501

*Supported in part by the National Science Foundation under Grant No. IRI-8802390 and by a Presidential Young Investigator award, and in part by the Mathematical Sciences Institute.
**Supported by the National Science Foundation under the REU program.

On the Probabilistic Analysis of Normal Form Computation of a Sparse Matrix

Bruce Randall Donald* **David Renpan Chang†**
Computer Science Department, Cornell University

December 18, 1990

Abstract

An (s,t)-*sparse* matrix has s non-zero entries per column and t per row. (s,t)-sparse integer matrices arise in the computation of integral homology. In this paper a probabilistic analysis is given for diagonalizing an integer (s,t)-sparse matrix into normal formal. By *normal form* of a matrix, we mean the diagonalization of the matrix over the ring of integers. We prove that under high probability the expected running time is $O(n^2)$ where n is the size of the given (s,t)-sparse matrix, i.e. this expected running time can be achieved with probability very close to 1 when $s,t \ll n$.

1 Introduction

Donald in his paper [3] discusses the sparsity of the boundary matrices in homology-type computation of a triangulated geometric design(i.e. a finite dimensional simplicial complex). In this paper, we study the probabilistic complexity of (Smith) normal form computation of such a sparse matrix. In [5], a polynomial time $O(n^5)$ deterministic algorithm is given to compute the normal form of a matrix. However, in this paper, we will study the classical reduction algorithm (see, e.g., [?]) and show that this reduction algorithm runs fast(i.e. $O(n^2)$) probabilistically on a (s,t)-*sparse* matrix with non-zero entries uniformly distributed over the set $\mathbf{Z}_{[-\wp,\wp]}\backslash\{0\}$, where $\mathbf{Z}_{[-\wp,\wp]}$ is the set of integers in the interval $[-\wp,\wp]$, $\wp$ is a very small positive integer and n is the *size* of the matrix. In the case of integral homology computation, $\wp = 1$.

Definition 1 *We denote the* algebraic complexity *of a matrix* $\mathbf{A}$ *by* $\text{alg}(\mathbf{A})$. $\text{alg}(\mathbf{A}) = \max\{|a_{ij}| \mid a_{ij}$ *is an entry in* $\mathbf{A}\}$.

Definition 2 *Henceforth we will consider diagonalization of an integer matrix* $\mathbf{A}$ *over* $\mathbf{Z}$. *We will assume* $\mathbf{A}$ *has* n *rows and* m *columns, and without loss of generality we take* $n \geq m$. *We call* n *the* size *of the matrix.*

*Supported in part by the National Science Foundation under grant No. IRI-8802390 and by a Presidential Young Investigator award, and in part by the Mathematical Science Institute.

†Supported by the National Science Foundation under the REU program.

Definition 3 *An $n \times m$ matrix $\mathbf{M}$ is called (s,t)-*sparse *if each row (resp. column) of $\mathbf{M}$ has exactly t (resp. s) non-zero elements and $s \ll n$ and $t \ll m$. Furthermore, we require that the non-zero entries in $\mathbf{M}$ are uniformly distributed over the set $\mathbf{Z}_{[-\wp,\wp]} \setminus \{0\}$ where $\wp$ is a very small positive integer, and each entry of this matrix has equal chance of being non-zero. For convenience, we define $\alpha = \log_2 n$, $\beta = \log_2 m$ and assume $n \geq m$ wlog.*

1.1 Statements of Approach and Results

In this section, we give a brief overview of our approach in order to give the reader the general idea; a formal and careful exposition comes later in the paper. All claims are proven, but due to limited space we have relegated some proofs to the appendix, labeled A.

We proceed as follows. The approach Donald took [3] in analyzing the probabilistic complexity of normal form computation was to examine *pre-reduction* complexity. This operation reduces the computation on an $n \times m$ matrix to the computation on an $(n-1) \times (m-1)$ matrix.

Definition 4 [3] *When an arbitrary matrix is in the form of eq. (1), where $a_{1,1}$ divides each element of $\mathbf{B}$, we say the matrix is in* pre-reduced *form. We call the algorithmic process of bringing a matrix into pre-reduced form* pre-reduction. *We call the matrix $\mathbf{B}$ in eq. (1) the* remaining matrix *after a pre-reduction.*

$$\begin{pmatrix} a_{1,1} & 0 \cdots 0 \\ \begin{matrix} 0 \\ \vdots \\ 0 \end{matrix} & \boxed{\mathbf{B}} \end{pmatrix} \tag{1}$$

By a cursory examination of the reduction algorithm, it is evident that the change in algebraic complexity of an entry after a pre-reduction depends only on the algebraic complexity before the pre-reduction. Suppose we begin with an (s,t)-sparse $n \times m$ matrix $\mathbf{A}$. In this paper we first notice that our original (s,t)-sparse matrix gets denser after each pre-reduction but the algebraic complexity remains the same (probabilistically) for the first few pre-reductions. That is, initially, as the matrix is diagonalized, it remains "sparse enough" that a subsequent pre-reduction will increase its density, but not its algebraic complexity. However, the remaining matrix becomes increasingly dense, and this "*sparse enough*" property is eventually violated. At this point, we call the remaining matrix "*dense*". We show that the dense remaining matrix we obtain has uniformly distributed entries all of low algebraic complexity. From our point of view, the difference between "sparse enough" and "dense" is as follows:

Definition 5 "Sparse enough" *matrices have non-zero elements that are uniformly distributed and of low algebraic complexity. A pre-reduction of a* "sparse enough" *matrix will make it somewhat denser, but will not raise its expected algebraic complexity.*

Definition 6 *The entries of a* "dense" *matrix are uniformly distributed and of low algebraic complexity. A pre-reduction of a* "dense" *matrix is expected to raise its algebraic complexity.*

In this paper we show that the expected number of pre-reductions we can perform before obtaining a dense remaining matrix $\mathbf{B}$ is at least $n - \sqrt[4]{n}$, and hence the resulting dense remaining matrix $\mathbf{B}$ is of expected size at most $\sqrt[4]{n}$ (recall the definition of size, def. 2). Furthermore, the entries in this dense remaining matrix $\mathbf{B}$ of size at most $\sqrt[4]{n}$ are uniformly distributed over the integral interval $\mathbf{Z}_{[-\wp,\wp]}$ with $\wp$ a very small positive integer. From [3], we know that each pre-reduction of a "sparse enough" matrix of size n takes time $O(n)$, and hence we can obtain our dense remaining matrix $\mathbf{B}$ in time $O(n^2)$.

It then remains to analyze the complexity of diagonalizing the remaining dense matrix $\mathbf{B}$. First, we will prove that the algebraic complexity of $\mathbf{B}$ changes by a constant amount after each successive pre-reduction. We show that for the case of integral homology computation the expected value of this constant is in fact 1. Pre-reducing a general dense $n \times m$ matrix of uniform algebraic complexity $\wp$ can be done in time $O(\wp mn)$ (see [3]). Let $r = \sqrt[4]{n}$. Hence we can diagonalize an $r \times r$ dense remaining matrix $\mathbf{B}$ with initial algebraic complexity $\wp = 1$ in expected time

$$r^2 + 2(r-1)^2 + 3(r-2)^2 + \cdots \tag{2}$$

which is $O(r^4) = O(n)$. Hence $\mathbf{B}$ can be diagonalized in linear ($O(n)$) expected time. We conclude that the complexity of normal form computation for a (s,t)-sparse matrix $\mathbf{A}$ is $O(n^2)$ probabilistically where n is the size of $\mathbf{A}$, i.e. This expected running time can be achieved with probability very close to 1 when $s, t \ll n$.

In order to prove this result, we describe a number of tools. First we develop a technique for analyzing the combinatorics of diagonalization, by gathering successive pre-reduction steps together into "groups" called *groups of pre-reductions.* By *group pre-reduction* we mean a sequence of pre-reductions having the property that each pre-reduction in the sequence increases the number of non-zero entries in a row by the same expected amount. We then gather the group pre-reductions into sequences called "*phases*". By *a phase of group pre-reductions* we mean a sequence of at most $\sqrt{2\alpha}$ successive group pre-reductions, where $\alpha = \log_2 n$. Phases of group pre-reductions can be combinatorially cascaded in order to effect a complexity analysis. We use discrete random variables to model integral matrix entries, and thereby determine bounds on their expected growth and density. The probabilistic analysis is complicated by the destruction of uniformness and independence of the matrix entries after the dense matrix is obtained. However, we are able to show that the entries are nevertheless *conditionally independent* and *uniform* (on the *outer row* and *column*). This admits an inductive probabilistic argument (based on the recursive conditioning) that we use to derive our theorem on the constant algebraic complexity growth per pre-reduction in a dense matrix with low algebraic complexity.

In an appendix, labeled B, we generalize our probabilistic analysis to a non-uniform and dependent $n \times m$ (s,t)-sparse matrix, that is a (s,t)-sparse matrix whose entries have dependent and non-uniform probability distributions on being non-zero. We propose a pre-processing algorithm using *active randomization* to destroy the non-uniformness and dependence in order to obtain a uniform and independent (s,t)-sparse matrix, that is a (s,t)-sparse matrix whose entries have independent and uniform probability distributions on being non-zero. Then we can use the reduction algorithm to diagonalize this uniform and independent (s,t)-sparse matrix. The randomization corresponds to a random "change of basis", hence the Smith normal form will be the same. We show that the active randomization algorithm takes time $O(n+m)$. Hence, we prove that diagonalizing a non-uniform and dependent $n \times m$ (s,t)-sparse matrix can be done in expected time $O(n^2)$ with probability very close to 1 as n is large.

2 Normal Form Computation of a (s,t)-sparse Matrix

Consider a pre-reduction computation. For the given matrix $\mathbf{A} = (\ a_{i,j}\ )$, we can always find an element a with the smallest size and transform it to the $(1,1)$-position. Next, we can use elementary row operations to ensure that this element will divide all the entries in the first column. We can then subtract off multiples of rows to zero out the entries in the first column (except for $a_{1,1}$). This process halts after most $|a|$ row operations (as can be seen from the Euclidean division algorithm). The entries in the first row can be zeroed out similarly. We have

Proposition 1 [3] *Let ℓ be the smallest (in size) non-zero element of the initial matrix $\mathbf{A}$. Then pre-reducing the matrix can be done in time $O(\ell nm)$. For a sparse matrix, a pre-reduction can be done in time $O(\ell n)$.*

Without loss of generality, will assume $a_{1,1} = 1$, because this case can maximize the possible worst-case increase in algebraic complexity [3]. Now, let's consider the following step in a pre-reduction.

$$\begin{pmatrix} 1 & \dots & b & \dots \\ \dots & \dots & \dots & \dots \\ q & \dots & a & \dots \\ \dots & \dots & \dots & \dots \end{pmatrix} \mapsto \begin{pmatrix} 1 & \dots & b & \dots \\ \dots & \dots & \dots & \dots \\ 0 & \dots & a-bq & \dots \\ \dots & \dots & \dots & \dots \end{pmatrix} \mapsto \begin{pmatrix} 1 & \dots & 0 & \dots \\ \dots & \dots & \dots & \dots \\ 0 & \dots & a-bq & \dots \\ \dots & \dots & \dots & \dots \end{pmatrix} \tag{3}$$

Definition 7 *We call the step illustrated in eq. (3) a* basic step *in a pre-reduction.*

Definition 8 *We call the first row of the matrix $\mathbf{A}$ the* outer row, *and the row containing a and q* hit row. *The* outer *and* hit columns *are defined analogously.*

We are interested in the following two questions for a (s,t)-sparse matrix $\mathbf{A}$:

Question 1 *Will a (s,t)-sparse matrix get denser after a sequence of pre-reductions? If so, how fast?*

Question 2 *Will the* algebraic complexity $\mathrm{alg}(\mathbf{A})$ *grow after a sequence of pre-reductions? If so, how fast?*

2.1 Change of Sparsity

We notice that a row can get denser after performing a basic step if and only if some zero entries are converted to non-zero. In the following, we will define a random variable X to measure the density growth after one basic step.

2.1.1 Discrete Random Variable, Expected Value and Variance of a Basic Step

For a (s,t)-sparse matrix $\mathbf{A}$, the discrete random variable X is defined as a map from the set of indices of non-zero entries in the outer row (except the first entry on that row), to the set $\{0,1,...,t-1\}$(i.e. the set of possible number of zero entries being converted to non-zero in a basic step). We know that the probability of a entry being non-zero in a (s,t)-sparse matrix $\mathbf{A}$ is $P(a_{i,j} \neq 0) = 1-(1-\frac{s}{n})(1-\frac{t}{m}) = \frac{s}{n}+\frac{t}{m}-\frac{st}{mn}$. For convenience, we define p to be $\frac{s}{n}+\frac{t}{m}-\frac{st}{mn}$. Clearly, $p \to 0$ as $\frac{s}{n}, \frac{t}{m} \to 0$ asymptotically. So we have

Definition 9 *The probability distribution function f is defined as follows,*

$$f(z) = P(X = z) = \binom{t-1}{z} (1-p)^z (p)^{t-z-1} \tag{4}$$

f here is usually called the binomial distribution function.
We know that the expectation $E(X)$ and variance $V(X)$ can be computed easily as follows,

$$\begin{aligned} E(X) &= (t-1)(1-p) && (5) \\ &= (t-1)(1 - \frac{s}{n} + \frac{t}{m} - \frac{st}{mn}) \rightarrow t-1 \; as \; p \rightarrow 0 && (6) \end{aligned}$$

$$\begin{aligned} V(X) &= (t-1)p(1-p) && (7) \\ &< (t-1)p && (8) \\ &= \frac{t(t-1)}{n} + \frac{s(t-1)}{m} - \frac{st(t-1)}{mn} \rightarrow 0 \; as \; p \rightarrow 0 && (9) \end{aligned}$$

$E(X)$ and $V(X)$ tell us that, $t-1$ zero entries are converted to non-zero in a basic step of a pre-reduction with probability very close to 1 as $p \rightarrow 0$, i.e., the hit row now is expected to have $2t-2$ non-zero entries and no non-zero entries on the hit row are expected to be hit in a basic step. So, we can also conclude that, the algebraic complexity stays unchanged with probability very close to 1 as $p \rightarrow 0$. Since in a pre-reduction there are exactly s rows involved, so we have

Proposition 2 *After n/s pre-reductions, we obtain an $2^\alpha(1-\frac{1}{s}) \times (2^\beta - \frac{2^\alpha}{s})$ remaining matrix* **B** *of expected sparsity (2s-2, 2t-2) with probability very close to 1 as $\frac{s}{n}, \frac{t}{m} \rightarrow 0$ asymptotically where $n = 2^\alpha, m = 2^\beta, n > m$.*

Proof: See appendix A. □

2.1.2 Group Pre-reductions and Phases of Group Pre-reductions

Proposition 2 tells us that during the first n/s (expected number of) pre-reductions, the expected number of zero entries converted to non-zero in a hit row stays the same with probability very close to 1 as $\frac{s}{n}, \frac{t}{m} \rightarrow 0$ asymptotically. So,

Definition 10 *We group these pre-reductions together and call them the first* group pre-reduction. *In general,* group pre-reduction *means that we group a sequence of successive pre-reductions together when this sequence of pre-reductions has the property that the expected number of zero entries converted to non-zero in a hit row stays the same after each pre-reduction in this sequence of pre-reductions.*

We know that after the first group pre-reduction, we get a denser matrix, but the probability of any entry of this matrix being non-zero is very close to 0 as the ratios $(2s-2)/(2^\alpha(1-\frac{1}{s})), (2t-2)/(2^\beta - \frac{2^\alpha}{s}) \rightarrow 0$ asymptotically where $2t-2$ (resp. $2s-2$) is the expected number of non-zero entries in each row (resp. column) and $2^\alpha(1-\frac{1}{s})$ (resp. $2^\beta - \frac{2^\alpha}{s}$) is the number of rows (resp. columns) of the remaining matrix after the first group pre-reduction. In other words, the remaining matrix is still "sparse enough" in the sense that an additional pre-reduction will not change its expected algebraic complexity (we see this from the proof of proposition 1 above). Recall that (Def. 6) a *dense* matrix B is one where a pre-reduction of B results in an expected increase in algebraic complexity. In general, we want to know how many such group pre-reductions we can perform before we reach such a dense matrix. We observe that after one group pre-reduction, the number of non-zero entries in a row is nearly doubled and $P(a_{ij} \neq 0)$, where a_{ij} is any entry in the remaining matrix, is also doubled. But as long as $P(a_{ij} \neq 0) \to 0$ we can keep performing the next group pre-reduction.

From the proof of proposition 2 and the definition of group pre-reduction, we can derive that the expected number of pre-reductions in a group pre-reduction is n'/s' where n' is the size of the matrix before the group pre-reduction and s' is the number of non-zero entries in the outer column of the matrix before the group pre-reduction. We now derive the expected number of group pre-reductions we can perform before $P(a_{ij} \neq 0) \not\to 0$. Suppose the i^{th} group pre-reduction contains p_i pre-reductions. Then the total expected number of pre-reductions we can perform before obtaining a dense remaining matrix is $\sum_i p_i$. We bound this sum below.

For convenience in our analysis, we gather successive group pre-reductions into sequences called "phases".

Definition 11 *The* i^{th} phase of group pre-reductions *consists of a sequence of* n_i *group pre-reductions with* $n_i \leq \sqrt{2\alpha}$ *such that after this sequence of group pre-reductions we obtain a matrix of size* $2^{a\alpha}$ *with* $0 < a < 1$ *for some* a, *where* $n = 2^\alpha$ *is the size of the original matrix.*

Also, we will assume that after each group pre-reduction, the number of non-zero entries in each row is doubled. Now, we claim the following,

Theorem 1 *For a given* (s,t)*-sparse matrix of size* $n = 2^\alpha$*: The expected number of phases of group pre-reductions is* k *and the expected number of group pre-reductions we can perform before obtaining a dense remaining matrix is* $n_1 + n_2 + \cdots + n_k$ *where* $n_i \leq \sqrt{2\alpha}$ *is the number of group pre-reductions in the* i^{th} *phase of group pre-reductions for* $i = 1, 2, ..., k$ *and* $k < \frac{\sqrt{\alpha}}{8}$.

Corollary 1 *The remaining dense matrix is of size at most* $2^{\frac{\alpha}{4}} = \sqrt[4]{n}$.

In order to prove theorem 1, we will make several observations about the size (def. 2) of the remaining matrix after each group pre-reduction. Let us consider the following sequence of values, each of which represents the size of the remaining matrix after each successive group pre-reduction. We start out with a matrix of size $n = 2^\alpha$ for some α:

$$2^\alpha \quad \xrightarrow{after\ first\ group\ pre-reduction} \quad 2^\alpha(1 - \frac{1}{s}) \tag{10}$$

$$\xrightarrow{after\ second\ group\ pre-reduction} \quad 2^{\alpha-1}(1 - \frac{1}{s})(2 - \frac{1}{s}) \tag{11}$$

$$\xrightarrow{after\ third\ group\ pre-reduction} \quad 2^{\alpha-3}(1 - \frac{1}{s})(2 - \frac{1}{s})(4 - \frac{1}{s}) \tag{12}$$

$$\vdots \tag{13}$$

$$\xrightarrow{after\ i^{th}\ group\ pre-reduction} \quad 2^{\alpha-\gamma_i^{(1)}}(1 - \frac{1}{s})(2 - \frac{1}{s})(4 - \frac{1}{s})\cdots(2^i - \frac{1}{s}) \tag{14}$$

$$\vdots \tag{15}$$

In deriving the above sequence of sizes of the remaining matrix after each successive group pre-reduction, we obtained the following recurrence relation on $\gamma_i^{(1)}$:

$$\gamma_0^{(1)} = 0 \quad (16)$$

$$\gamma_i^{(1)} = \gamma_{i-1}^{(1)} + i \quad (17)$$

The $\gamma_i^{(1)}$ here is a combinatorial device to help us find the transition point from the first phase of group pre-reductions to the second phase of group pre-reductions. Precisely, when $\gamma_i^{(1)}$ reaches α, we obtain a sequence of i group pre-reductions, and we will show that $i \leq \sqrt{2\alpha}$. According to the definition of a phase of group pre-reductions, we will call this sequence of i group pre-reductions the *first phase* of group pre-reductions.

Proposition 3 *In the above sequence (10) - (15), if $\gamma_i^{(1)} = \alpha$, then α and i satisfy the relation $\alpha = \frac{i(i+1)}{2}$ and the size of the remaining matrix is $(1-\frac{1}{s})(2-\frac{1}{s})\cdots(2^i-\frac{1}{s}) = 2^{a_1\alpha}$ for some real $0 < a_1 < 1$.*

Proof: See appendix A. □

We call the above ((10) - (15)) sequence of group pre-reductions before $\gamma_i^{(1)}$ reaches α the first phase of group pre-reductions. For convenience, we denote n_1 to be i, i.e. n_1 is the number of group pre-reductions in the first phase of group pre-reductions. From proposition 3, we know that $n_1 \leq \sqrt{2\alpha}$. The remaining matrix obtained after the n_1^{th} group pre-reduction is of size $(1-\frac{1}{s})(2-\frac{1}{s})\cdots(2^{n_1}-\frac{1}{s})$, and the number of non-zero entries in a row (resp. column) of this remaining matrix is at most $2^{n_1}t$ (resp. $2^{n_1}s$) and the ratio

$$\begin{aligned} & \frac{2^{n_1}s}{(1-\frac{1}{s})(2-\frac{1}{s})\cdots(2^{n_1}-\frac{1}{s})} \\ < \; & \frac{1}{(1-\frac{1}{s})(2-\frac{1}{s})\cdots(2^{n_1-2}-\frac{1}{s})} \\ \rightarrow \; & 0 \textit{ as } \alpha \rightarrow \infty \textit{ and } n_1 \approx \sqrt{2\alpha} \textit{ by proposition } 3. \end{aligned}$$

For convenience, we define the following,

Definition 12 *We let n' (resp. m') denote the number of rows (resp. columns) of the remaining matrix after the first phase of group pre-reductions, and t' (resp. s') the number of non-zero entries in a row (resp. column).*

The total number of non-zero entries in the remaining matrix after the first phase of group pre-reductions is $n't' = m's'$. So, the ratio $\frac{t'}{m'} \rightarrow 0$ as $\alpha \rightarrow \infty$, i.e. the probability of any entry being non-zero in the remaining matrix after the first phase of group pre-reductions is very close to 0 as $\frac{s'}{n'}, \frac{t'}{m'} \rightarrow 0$ asymptotically. Hence, we can keep performing group pre-reductions.

Now we start out with the remaining matrix of size $2^{a_1\alpha}$ and obtain the following sequence of values, each of which represents the size of the remaining matrix after each successive group pre-reduction.

$$2^{a_1\alpha} \xrightarrow{after\ 1^{st}\ group\ pre-reduction} 2^{a_1\alpha-n_1}(2^{n_1}-\frac{1}{s}) \tag{18}$$

$$\xrightarrow{after\ 2^{nd}\ group\ pre-reduction} 2^{a_1\alpha-n_1-(n_1+1)}(2^{n_1}-\frac{1}{s})(2^{n_1+1}-\frac{1}{s}) \tag{19}$$

$$\xrightarrow{after\ 3^{rd}\ group\ pre-reduction} 2^{a_1\alpha-n_1-(n_1+1)-(n_1+2)}(2^{n_1}-\frac{1}{s})(2^{n_1+1}-\frac{1}{s})(2^{n_1+2}-\frac{1}{s}) \tag{20}$$

$$\vdots \tag{21}$$

$$\xrightarrow{after\ i^{th}\ group\ pre-reduction} 2^{a_1\alpha-\gamma_i^{(2)}}(2^{n_1}-\frac{1}{s})(2^{n_1+1}-\frac{1}{s})(2^{n_1+2}-\frac{1}{s})\cdots(2^{n_1+i-1}-\frac{1}{s}) \tag{22}$$

$$\vdots \tag{23}$$

In deriving the above sequence of sizes of the remaining matrix after each successive group pre-reduction, we obtained the following recurrence relation on $\gamma_i^{(2)}$:

$$\gamma_0^{(2)} = n_1 \tag{24}$$

$$\gamma_i^{(2)} = \gamma_{i-1}^{(2)} + n_1 + i - 1 \tag{25}$$

The $\gamma_i^{(2)}$ here is a combinatorial device to help us find the transition point from the second phase of group pre-reductions to the third phase of group pre-reductions. Precisely, when $\gamma_i^{(2)}$ reaches $a_1\alpha$, we obtain a sequence of i group pre-reductions, and we will show that $i \leq \sqrt{2\alpha}$. According to the definition of a phase of group pre-reductions, we will call this sequence of i group pre-reductions the *second phase* of group pre-reductions.

Proposition 4 *In the above sequence (18) - (23), if $\gamma_i^{(2)} = a_1\alpha$, then $a_1\alpha$ and i satisfy the relation $a_1\alpha = \frac{i(i-1)}{2} + n_1 i$ where n_1 is the number of group pre-reductions in the first phase of group pre-reductions and the size of the remaining matrix is $(2^{n_1}-\frac{1}{s})(2^{n_1+1}-\frac{1}{s})\cdots(2^{n_1+i-1}-\frac{1}{s}) = 2^{a_1a_2\alpha}$ for some real $0 < a_2 < 1$.*

Proof: See appendix A. □

We call the above ((18) - (23)) sequence of group pre-reductions before $\gamma_i^{(2)}$ reaches α the second phase of group pre-reductions. For convenience, we denote n_2 to be i, i.e. n_2 is the number of group pre-reductions in the second phase of group pre-reductions. From proposition 4, we know that $n_2 \leq \sqrt{2\alpha}$. The remaining matrix obtained after the n_2^{th} group pre-reduction is of size $(2^{n_1}-\frac{1}{s})(2^{n_1+1}-\frac{1}{s})\cdots(2^{n_1+n_2-1}-\frac{1}{s})$ and the number of non-zero entries in a row (resp. column) of this remaining matrix is at most $2^{n_1+n_2}t$ (resp. $2^{n_1+n_2}s$) and the ratio

$$\frac{2^{n_1+n_2}s}{(2^{n_1}-\frac{1}{s})(2^{n_1+1}-\frac{1}{s})\cdots(2^{n_1+n_2}-\frac{1}{s})}$$
$$< \frac{1}{(2^{n_1}-\frac{1}{s})(2^{n_1+1}-\frac{1}{s})\cdots(2^{n_1+n_2-2}-\frac{1}{s})}$$
$$\rightarrow 0 \ as\ a_1\alpha \rightarrow \infty\ and\ n_1 \approx \sqrt{2\alpha}\ by\ proposition\ 3.$$

For convenience, we define the following,

Definition 13 *we let n'' (resp. m'') denote the number of rows (resp. columns) of the remaining matrix after the second phase of group pre-reduction, and t'' (resp. s'') the number of non-zero entries in a row (resp. column).*

The total number of non-zero entries in the remaining matrix after the second phase of group pre-reductions is $n''t'' = m''s''$. So, the ratio $\frac{t''}{m''} \to 0$ as $\alpha \to \infty$, i.e. the probability of any entry being non-zero in the remaining matrix after the second phase of group pre-reductions is very close to 0 as $\frac{s''}{n''}, \frac{t''}{m''} \to 0$. Hence, we are ensured that we can perform at least $n_1 + n_2$ group pre-reductions.

In general, we can use this combinatorial device $\gamma_i^{(j-1)}$ to find the transition point from the $(j-1)^{th}$ phase of group pre-reductions to the j^{th} phase of group pre-reductions. After the j^{th} phase of group pre-reductions, we obtain a remaining matrix of size $2^{a_1 a_2 \cdots a_j \alpha}$ with $0 < a_i < 1$ for $i = 1, 2, ..., j$, and each row (resp. column) of this remaining matrix has $2^{n_1+n_2+\cdots+n_j}t$ (resp. $2^{n_1+n_2+\cdots+n_j}s$) non-zero entries. Since the remaining matrix becomes denser after each phase of group pre-reductions, the process of performing pre-reductions without changing the expected algebraic complexity of the remaining matrix has to stop after the k^{th} phase of group pre-reductions for some k, i.e. the remaining matrix we obtain after the k^{th} phase of group pre-reductions is dense (recall def. 6). Now, we are ready to prove theorem 1.

Proof of theorem 1:

We obtain a bound k on the number of expected phases of group pre-reductions. We know that after $n_1 + n_2 + \cdots + n_k$ of group pre-reductions, the size of the remaining matrix is $2^{a_1 a_2 \cdots a_k \alpha}$. Now, suppose $k > \frac{\sqrt{\alpha}}{8}$. We derive a contradiction. Since $\alpha \to \infty$, $\frac{\sqrt{\alpha}}{8} \to \infty$. Hence, $\prod_{i=1}^{k} a_i \to 0$ as $k \to \infty$ because $0 < a_i < 1, i = 1, 2, ..., k$. Then $2^{n_1} \geq 2^{a_1 a_2 \cdots a_k \alpha}$ because $n_1 \approx \sqrt{2\alpha}$ and lemma 1 below. Now, the number of non-zero entries after k phases is $s_k = 2^{n_1+n_2+\cdots+n_k}$, and the size of the remaining matrix is $N_k = 2^{a_1 a_2 \cdots a_k \alpha}$. Clearly, $s_k > N_k$ as $k \to \infty$, which is a contradiction. Thus the assumption that that $k > \frac{\sqrt{\alpha}}{8}$ is false. Hence, $k \leq \frac{\sqrt{\alpha}}{8}$. Therefore, we can perform at most $\frac{\sqrt{\alpha}}{8}$ phases of group pre-reductions, i.e. k is at most $\frac{\sqrt{\alpha}}{8}$. By propositions 3 & 4, we see that $n_i < \sqrt{2\alpha}, i = 1, 2, ..., k$, therefore $n_1 + n_2 + \cdots + n_k < \frac{\alpha}{4}$. Hence, the remaining dense matrix is of size at most $2^{\frac{\alpha}{4}} = \sqrt[4]{n}$. □

Lemma 1 *Let $n_1 \approx \sqrt{2\alpha}$ and $\prod_{i=1}^{k} a_i \to 0$ as $k \to \infty$ where $0 < a_i < 1$ for $i = 1, 2, ...$ Then $2^{n_1} \geq 2^{a_1 a_2 \cdots a_k \alpha}$ as $k \to \infty$.*

Proof: See appendix A. □

Theorem 1 essentially tells us that as long as the size of the original matrix (recall def. 2) $n = 2^\alpha$ is large, the expected number of phases of group pre-reductions we can perform is at most $\frac{\sqrt{\alpha}}{8}$ and the expected size of the remaining dense matrix is at most $2^{\frac{\alpha}{4}} = \sqrt[4]{n}$. In the next section, we consider diagonalizing this remaining matrix. We will show that the algebraic complexity is increased by a constant after each pre-reduction on a dense matrix with low algebraic complexity. In the case of integral homology group computation, the algebraic complexity is initially 1, and we show that the expected algebraic complexity in this case is actually increased by 1 after each pre-reduction. Recall proposition 1, that each pre-reduction on a dense matrix takes time $O(\ell mn)$ where ℓ the smallest (in size) non-zero entry of the dense matrix and n (resp. m) is the number of rows (resp. columns) of this dense matrix. So, the total expected running time for a dense matrix of size $r = \sqrt[4]{n}$ with initial algebraic complexity 1 is given by eq. (2) which is $O(n)$. Hence, the dense remaining matrix obtained from pre-reducing a sparse enough initial matrix of size n can be diagonalized in expected linear time. This expected running time is achieved with probability very close to 1 when n is large.

2.2 Change of Algebraic Complexity

From last section, we know that after $k \leq \frac{\sqrt{\alpha}}{8}$ phases of group pre-reductions, i.e., after $\sum_{i=1}^{k} n_i$ group pre-reductions with $n_i \leq \sqrt{2\alpha}$. We are left with a dense matrix of entries uniformly distributed over $\mathbf{Z}_{[-p,p]}$. Pre-reducing a dense matrix will change the algebraic complexity since the non-zero entries will be hit. We want to find a fast probabilistic bound on the growth of algebraic complexity. In the following, we will prove that, the algebraic complexity is increased by a expected constant per pre-reduction with probability very close to 1 when the size of the original matrix n is large.

2.2.1 Pre-reducing a Dense Matrix With Low Algebraic Complexity

In this section, we only consider a dense matrix $\mathbf{A}$ having entries uniformly distributed over $\mathbf{Z}_{[-1,1]}$. We would like to know whether after a pre-reduction, the remaining matrix has independent entries that are uniformly distributed.

Proposition 5 *After the first pre-reduction of the dense matrix* $\mathbf{A}$*, the uniformness and independence of entries in* $\mathbf{B}$ *are destroyed where* $\mathbf{B}$ *is the remaining matrix obtained after first pre-reduction.*

Proof: See appendix A. □

Proposition 5 tells us that if we only consider the probability distribution for the algebraic complexity of entries in the remaining matrix $\mathbf{B}$, it will be very complicated to derive the probability distribution of algebraic complexity growth after in turn pre-reducing this matrix $\mathbf{B}$, since we don't have independence and uniformness on the entries in matrix $\mathbf{B}$. However, we notice that the change of algebraic complexity of entries in matrix $\mathbf{B}$ depends solely on the outer row and column of matrix $\mathbf{A}$, and the entries in matrix $\mathbf{A}$ are independent and uniform. So, in the next section, we will introduce *conditional independence* and *uniformness* of entries in the remaining matrix (conditioned on the outer row and column). This essentially enables us to derive a theorem on constant growth of algebraic complexity after each pre-reduction by an inductive probabilistic argument.

2.2.2 Conditional Independence and Uniformness After Pre-reducing a Dense matrix

Definition 14 *Events A and B are called* conditionally independent on event C *if* $P(AB|C) = P(A|C)P(B|C)$.

Definition 15 *Events $A_1, A_2, ..., A_n$ are called* conditionally uniform on event C *if* $P(A_1|C) = P(A_2|C) = \cdots = P(A_n|C)$.

Since we have that entries in the outer row and column of the matrix $\mathbf{A}$ are uniformly distributed and independent, we claim the following

Proposition 6 *The entries in the remaining matrix* $\mathbf{B}$ *obtained after pre-reducing the dense matrix* $\mathbf{A}$ *are conditionally uniform and independent on the outer row and column in* $\mathbf{A}$.

Proof: See appendix A. □

In general, the uniformness and independence of entries in $\mathbf{A}_i$, which is the remaining matrix obtained after i^{th} pre-reduction, is recursively conditioned on the outer row and column in $\mathbf{A}_{i-1}$ which is the remaining matrix obtained after $(i-1)^{th}$ pre-reduction. There are two cases. In the first, the dense remaining matrix is "small". In the second, it is "large". We must show that in both cases, it can be quickly diagonalized. To do this we must show that the algebraic complexity grows slowly.

Remark 1 *As we see from the sparsity analysis in [3](lemma C.5) and the proof of proposition 5 above, during the first few (i.e. O(1)) pre-reductions of a dense matrix of low algebraic complexity, the algebraic complexity is increased by a expected constant after each successive pre-reduction. So, for a constant size (i.e.* $O(1)$*) dense matrix, we can diagonalize this matrix in constant time* $O(1)$. *On the other hand, when the size of a dense matrix of low algebraic complexity is large, Theorem 2 below ensures us that the expected algebraic complexity is still only increased by a constant amount after each successive pre-reduction.*

Hence we conclude this section with the following theorem.

Theorem 2 *The algebraic complexity of a dense matrix increases by 1 after each pre-reduction with probability very close to* 1 *when the size of the original dense matrix is asymptotically large.*

Proof: See appendix A. □

3 Conclusions

In this paper we gave a probabilistic analysis of diagonalizing a (s,t)-sparse matrix (recall def. 3). From [3], we know that each pre-reduction of a "sparse enough" (recall def. 5) matrix of size n takes time $O(n)$. Let n'(resp. m') be the number of rows (resp. columns) and t'(resp. s') the number of non-zero entries in a row (resp. column) in the remaining matrix after one pre-reduction. This remaining matrix is still "sparse enough" as long as the probability of any entry being non-zero in this remaining matrix is very close to 0 as $\frac{s'}{n'}, \frac{t'}{m'} \to 0$ asymptotically i.e. a subsequent pre-reduction of this matrix can increase the number of non-zero entries in a row (resp. column) by some expected amount, but can not increase its expected algebraic complexity. We then introduced an combinatorial tool called group pre-reduction (recall def. 10), that is, we group together all the pre-reductions that increase the number of non-zero entries in a row (resp. column) by the *same* expected amount. The "sparse enough" property of a matrix ensures us that we can keep performing group pre-reductions on a "sparse enough" matrix. This process stops when the remaining matrix turns out to be dense (recall def. 6) after some number of group pre-reductions. We showed that the expected number of pre-reductions we can perform before we obtain a dense matrix is at least $n - \sqrt[4]{n}$ where n is the size of the original matrix, and the expected size of the remaining dense matrix is at most $\sqrt[4]{n}$. Since pre-reducing a "sparse enough" takes time $O(n)$, therefore it takes time $O(n^2)$ to perform $n - \sqrt[4]{n}$ pre-reductions before we obtain a dense matrix. Pre-reducing a dense matrix will raise its expected algebraic complexity. Moreover, pre-reducing a dense matrix with independent entries uniformly distributed over some integral interval will destroy the independence and uniformness of entries in the remaining dense matrix. This complicates our probabilistic analysis of pre-reducing the remaining dense matrix. In order to overcome this difficulty, we introduced conditional independence and uniformness of entries in the remaining dense matrix. That is, the entries in the remaning dense matrix are conditionally independent and uniformly distributed (conditioned on the outer row and column). We made use of this conditional independence and uniformness and gave an inductive probabilistic argument (based on the recursive conditioning) that if we start out with a dense matrix of low algebraic complexity, then the expected algebraic complexity grows by a constant amount with probability very close to 1 when the size of this dense matrix is large asymptotically. In the case of integral homology computation, the expected value of this constant is actually 1. Recall proposition 1, which says pre-reducing a dense matrix can be done in time $O(\ell mn)$ where n (resp. m) is the number of rows (resp. columns) and ℓ is the smallest (in size) non-zero entry in the matrix. Now, let us consider the dense remaning matrix of size $\sqrt[4]{n}$ we obtained after performing $n - \sqrt[4]{n}$ pre-reductions on a "sparse enough" matrix. Let $r = \sqrt[4]{n}$. The total time for diagonalizing this dense remaining matrix is $r^2 + 2(r-1)^2 + 3(r-2)^2 + \cdots$ which is $O(r^4) = O(n)$. Therefore, diagonalizing a (s,t)-sparse matrix takes expected time $O(n^2)$.

Acknowledgements

We are very grateful to Chee Yap for his comments and suggestions on active randomization, which were very helpful to us. We would also like to thank Peter Kahn, Henry Kesten, Dexter Kozen and Leslie Trotter for their very useful comments, suggestions and technical conversations in general.

References

[1] Chung, K. L., *Elementary Probability Theory with Stochastic Processes,* Springer-Verlag, (1979).

[2] Domich, P. D., *Residual Methods for Computing Hermite and Smith Normal Forms* Ph.D. Thesis, Cornell University, (1985).

[3] Donald, B. R., *On the Complexity of Computing the Homology Type of a Triangulation*, Computer Science Dept. Cornell University, Submitted to ACM STOC, (1991).

[4] Feller, W., *An Introduction to Probability Theory and its Applications,* John Wiley & Sons, (1968).

[5] Iliopoulos, C. S., *Worst-Case Complexity Bounds on Algorithms for Computing the Canonical Structure of finite Abelian Groups and the Hermite and Smith Normal Forms of an Integer Matrix*, SIAM journal on computing **18** (1989) , 658-669.

[6] Jacobson, N., *Basic Algebra I*, Freeman, (1985).

[7] Munkres, J. R., Elements of Algebraic Topology, Benjamin/Cummings, (1984).

[8] Tarasov, L., *The World Is Built On Probability,* MirPublishers Moscow, (1988).

APPENDICES

A Proofs

We now prove the lemmas and claims that for reasons of space could not be included in the body of this paper.

Proposition 2 *After n/s pre-reductions, we obtain an $2^{\alpha}(1-\frac{1}{s}) \times (2^{\beta}-\frac{2^{\alpha}}{s})$ remaining matrix* **B** *of expected sparsity (2s-2, 2t-2) with probability very close to 1 as $\frac{s}{n}, \frac{t}{m} \rightarrow 0$ asymptotically where $n = 2^{\alpha}, m = 2^{\beta}, n > m$.*

Proof:
Let us consider the remaining matrix **B** after the first pre-reduction. The matrix looks like the one in eq. (1). We know that after the first pre-reduction, some rows of **B** have $2t-2$ non-zero entries and some columns of **B** have $2s-2$ non-zero entries. In order that the outer row (resp. column) of **B** has $2t-2$ (resp. $2s-2$) non-zero entries, $a_{1,2}$ and $a_{2,1}$ have to be non-zero in the original matrix **A**. But we know that $P(a_{1,2} \neq 0) \rightarrow 0$ and $P(a_{2,1} \neq 0) \rightarrow 0$ as $\frac{s}{n}, \frac{t}{m} \rightarrow 0$. So, the outer row (resp. column) of **B** will still have t (resp. s) non-zero entries with probability very close to 1 as $\frac{s}{n}, \frac{t}{m} \rightarrow 0$. Then, the second pre-reduction will convert $t-1$ (expected number of) zero entries to non-zero in each of its hit rows. By the same analysis, we know that we can perform an expected number n/s of pre-reductions such that during each pre-reduction only $t-1$ (expected number of) zero entries are converted to non-zero in a hit row with probability very close to 1 as $\frac{s}{n}, \frac{t}{m} \rightarrow 0$. After an expected number n/s of pre-reductions, each row will have at most $2t-2$ expected non-zero entries and each column will have at most $2s-2$ expected non-zero entries. □

Recall the sequence of values (10) - (15) in the body of the paper above, representing the size of the remaining matrix after each successive group pre-reduction. We call such a sequence a *size sequence.* The following proposition about the size sequence (10) - (15) employs the recurrence relation in equations (16) - (17), above.

Proposition 3 *In the above size sequence (10) - (15), if $\gamma_i^{(1)} = \alpha$, then α and i satisfy the relation $\alpha = \frac{i(i+1)}{2}$ and the size of the remaining matrix is $(1-\frac{1}{s})(2-\frac{1}{s})\cdots(2^i-\frac{1}{s}) = 2^{a_1\alpha}$ for some real $0 < a_1 < 1$.*

Proof:
From recurrence relation given by the equations (16) - (17) above, we obtain

$$\begin{aligned}\gamma_i^{(1)} &= \gamma_{i-1}^{(1)} + i \\ &= 1 + 2 + \cdots + i \\ &= \frac{i(i+1)}{2}\end{aligned}$$

So, after i^{th} group pre-reduction, the remaining matrix has $2^{a_1\alpha}$ number of rows with $0 < a_1 < 1$ for some real a_1 and

$$(1-\frac{1}{s})(2-\frac{1}{s})\cdots(2^i-\frac{1}{s}) = 2^{a_1\alpha} \rightarrow \infty \text{ as } \alpha \rightarrow \infty$$

and the number of non-zero entries in a row (resp. column) is $2^i t$ (resp. $2^i s$) $\square$

The following proposition about the size sequence (18) - (23) employs the recurrence relation in equations (24) - (25), above.

Proposition 4 *In the above size sequence (18) - (23), if $\gamma_i^{(2)} = a_1\alpha$, then $a_1\alpha$ and i satisfy the relation $a_1\alpha = \frac{i(i-1)}{2} + n_1 i$ where n_1 is the number of group pre-reductions in the first phase of group pre-reductions and the size of the remaining matrix is $(2^{n_1} - \frac{1}{s})(2^{n_1+1} - \frac{1}{s})\cdots(2^{n_1+i-1} - \frac{1}{s}) = 2^{a_1 a_2 \alpha}$ for some real $0 < a_2 < 1$.*

Proof:

From the recurrence relation given by the equations (24) - (25) above, we get

$$\begin{aligned}
\gamma_i^{(2)} &= \gamma_{i-1}^{(2)} + n_1 + i - 1 \\
&= n_1 + (n_1+1) + (n_1+2) + \cdots + (n_1 + i - 1) \\
&= \frac{i(i-1)}{2} + n_1 i
\end{aligned}$$

So, after i^{th} group pre-reduction, the remaining matrix has $2^{a_1 a_2 \alpha}$ number of rows with $0 < a_2 < 1$ for some real a_1 and

$$(2^{n_1} - \frac{1}{s})(2^{n_1+1} - \frac{1}{s})\cdots(2^{n_1+i-1} - \frac{1}{s}) \;=\; 2^{a_1 a_2 \alpha} \to \infty \text{ as } \alpha \to \infty.$$

and the number of non-zero entries in a row (resp. column) is $2^{n_1+i}t$ (resp. $2^{n_1+i}s$). $\square$

Lemma 1 *Let $n_1 \approx \sqrt{2\alpha}$ and $\prod_{i=1}^{k} a_i \to 0$ as $k \to \infty$ where $0 < a_i < 1$ for $i = 1, 2, \ldots$. Then $2^{n_1} \geq 2^{a_1 a_2 \cdots a_k \alpha}$ as $k \to \infty$.*

Proof:

Since we have

$$\begin{aligned}
\log_\alpha 2 &> 0 \\
\log_\alpha \prod_{i=1}^{k} a_i &\to -\infty \text{ as } k \to \infty
\end{aligned}$$

Therefore,

$$\begin{aligned}
\log_\alpha \sqrt{2\alpha} &= \frac{1}{2}(\log_\alpha 2 + 1) \\
&> 1 + \log_\alpha \prod_{i=1}^{k} a_i \text{ as } k \to \infty \\
&= \log_\alpha a_1 a_2 \cdots a_k \alpha
\end{aligned}$$

Hence, $\sqrt{2\alpha} > a_1 a_2 \cdots a_k \alpha$ i.e, $2^{n_1} \geq 2^{a_1 a_2 \cdots a_k \alpha}$ as $k \to \infty$. $\square$

Proposition 5 *After the first pre-reduction of the dense matrix* $\mathbf{A}$*, the uniformness and independence of entries in* $\mathbf{B}$ *are destroyed where* $\mathbf{B}$ *is the remaining matrix obtained after first pre-reduction.*

Proof:

Let us consider eq. (3). We have over all 27 ways to choose the triple (a, b, q) from $\mathbf{Z}_{[-1,1]} \times \mathbf{Z}_{[-1,1]} \times \mathbf{Z}_{[-1,1]}$. Among these 27 triples, 9 triples will result in $|a - bq| = 0$, 14 triples will result in $|a - bq| = 1$ and 4 will result in $|a - bq| = 2$ after a basic step in a pre-reduction. We conclude that, the uniformness of entries in $\mathbf{B}$ is not preserved after a pre-reduction.

Let us now consider the following basic step in a pre-reduction:

$$\begin{pmatrix} 1 & \dots & X & \dots & Y & \dots \\ \dots & \dots & \dots & \dots & \dots & \dots \\ Z & \dots & U & \dots & V & \dots \\ \dots & \dots & \dots & \dots & \dots & \dots \end{pmatrix} \mapsto \begin{pmatrix} 1 & \dots & 0 & \dots & 0 & \dots \\ \dots & \dots & \dots & \dots & \dots & \dots \\ 0 & \dots & U - XZ & \dots & V - YZ & \dots \\ \dots & \dots & \dots & \dots & \dots & \dots \end{pmatrix} \tag{26}$$

where X, Y, Z, U and V are discrete random variables taking values in $\mathbf{Z}_{[-1,1]}$. In order to determine whether $U - XZ$ and $V - YZ$ are independent, we will check whether $E(U - XZ)E(V - YZ)$ is equal to $E((U - XZ)(V - YZ))$. We have

$$\begin{aligned} & E(U - XZ)E(V - YZ) \\ = \; & (E(U) - E(X)E(Z))(E(V) - E(Y)E(Z)) \\ = \; & E(U)E(V) - E(U)E(Y)E(Z) - E(V)E(X)E(Z) + E(X)E(Y)E(Z)^2 \\ & E((U - XZ)(V - YZ)) \\ = \; & E(UV - UYZ - VXZ + XYZ^2) \\ = \; & E(U)E(V) - E(U)E(Y)E(Z) - E(V)E(X)E(Z) + E(X)E(Y)E(Z^2) \end{aligned}$$

But we know that $E(Z^2) \neq E(Z)^2$, so $E(U - XZ)E(V - YZ) \neq E((U - XZ)(V - YZ))$, i.e. $U - XZ$ and $V - YZ$ are not independent. $\Box$

Proposition 6 *The entries in the remaining matrix* $\mathbf{B}$ *obtained after pre-reducing the dense matrix* $\mathbf{A}$ *are conditionally uniform and independent on the outer row and column in* $\mathbf{A}$.

Proof:

Let us first prove the conditional uniformness of entries in $\mathbf{B}$. We consider the following basic step in a pre-reduction.

$$\begin{pmatrix} 1 & \dots & Y & \dots \\ \dots & \dots & \dots & \dots \\ X & \dots & Z & \dots \\ \dots & \dots & \dots & \dots \end{pmatrix} \mapsto \begin{pmatrix} 1 & \dots & 0 & \dots \\ \dots & \dots & \dots & \dots \\ 0 & \dots & Z - XY & \dots \\ \dots & \dots & \dots & \dots \end{pmatrix} \tag{27}$$

where X, Y and Z are discrete random variables taking values in $\mathbf{Z}_{[-1,1]}$. By conditioned on X and Y, we mean that we fix X and Y to be some constants c_1 and c_2 respectively. We know that X is uniformly distributed before the basic step is performed. Now subtracting a constant c_1c_2 from X simply means shifting the uniform interval by c_1c_2. So, $Z - c_1c_2$ is uniformly distributed. Hence, $Z - XY$ is conditionally uniform on the outer row and column.

Now let us derive the conditional independence of entries in $\mathbf{B}$. We have three cases to consider. They are 1) entries in the same row, 2) entries in the same column, and 3) entries in different rows and columns.

Case 1) independence of entries in the same row:

We consider the following basic step in a pre-reduction.

$$\begin{pmatrix} 1 & \dots & X & \dots & Y & \dots \\ \dots & \dots & \dots & \dots & \dots & \dots \\ Z & \dots & X' & \dots & Y' & \dots \\ \dots & \dots & \dots & \dots & \dots & \dots \end{pmatrix} \mapsto \begin{pmatrix} 1 & \dots & 0 & \dots & 0 & \dots \\ \dots & \dots & \dots & \dots & \dots & \dots \\ 0 & \dots & X'-XZ & \dots & Y'-YZ & \dots \\ \dots & \dots & \dots & \dots & \dots & \dots \end{pmatrix} \tag{28}$$

where X, X', Y, Y' and Z are discrete random variables taking values in $\mathbf{Z}_{[-1,1]}$. Again, by conditioning on the outer row and column, we fix these discrete random variables X, Y and Z to be some constants c_1, c_2 and c_3 respectively. Since X' and Y' are independent, so $X' - c_1c_3$ and $Y' - c_2c_3$ are independent. Hence, $X' - XZ$ and $Y' - YZ$ are independent conditioned on the outer row and column.

Case 2) independence of entries in the same column:

This case is symmetric to case 1), so the analysis is similar to that of case 1).

Case 3) independence of entries in different rows and columns:

We again consider the following basic step in a pre-reduction.

$$\begin{pmatrix} 1 & \dots & X & \dots & Y & \dots \\ \dots & \dots & \dots & \dots & \dots & \dots \\ Z & \dots & X' & \dots & \dots & \dots \\ \dots & \dots & \dots & \dots & \dots & \dots \\ Z' & \dots & \dots & \dots & Y' & \dots \\ \dots & \dots & \dots & \dots & \dots & \dots \end{pmatrix} \mapsto \begin{pmatrix} 1 & \dots & 0 & \dots & 0 & \dots \\ \dots & \dots & \dots & \dots & \dots & \dots \\ 0 & \dots & X'-XZ & \dots & \dots & \dots \\ \dots & \dots & \dots & \dots & \dots & \dots \\ 0 & \dots & \dots & \dots & Y'-YZ' & \dots \\ \dots & \dots & \dots & \dots & \dots & \dots \end{pmatrix} \tag{29}$$

where X, X', Y, Y', Z and Z' are discrete random variables taking values in $\mathbf{Z}_{[-1,1]}$. Again, by conditioning on the outer row and column, we fix these discrete random variables X, Y, Z and Z' to be some constants c_1, c_2, c_3, and c_4 respectively. Since X' and Y' were independent, so $X' - c_1c_3$ and $Y' - c_2c_4$ are independent. Hence, $X' - XZ$ and $Y' - YZ'$ are independent conditioned on the outer row and column. □

Theorem 2 *The algebraic complexity of a dense matrix increases by 1 after each pre-reduction with probability very close to* 1 *when the size of the original dense matrix is asymptotically large.*

Proof:

We proceed by induction on the number of pre-reductions.

Base case: entries $a_{i,j}$ in $\mathbf{A}_1$ are independent and uniformly distributed over $\mathbf{Z}_{[-2,2]}$ conditioning on the outer row and column of matrix $\mathbf{A}$ by proposition 5.

Inductive hypothesis: $\text{alg}(\mathbf{A}_i) = \text{alg}(\mathbf{A}_{i-1}) + 1$ with high probability.

Inductive step: We perform a pre-reduction on $\mathbf{A}_i$. Recall the basic step as illustrated in eq. (3), if $q, a, b \in \mathbf{Z}_{[-i+1,i-1]}$, then by inductive hypothesis, with very high probability $a - bq \in \mathbf{Z}_{[-i,i]}$. By enlarging $\mathbf{Z}_{[-i+1,i-1]}$ to $\mathbf{Z}_{[-i,i]}$, we want to know how many ways there are to choose q, a, b so that $|a - bq| > i + 2$. We have the following three cases to consider:

Case 1) We fix q to be i, then we have $4i^2$ ways to choose a and b so that $|a - bq| > i + 2$. Similarly for fixing q to be $-i$.

Case 2) We fix b to be i, then we have $4i^2$ ways to choose a and b so that $|a - bq| > i + 2$. Similarly for fixing b to be $-i$.

Case 3) We fix a to be i, then we have $4(i-1)^2$ ways to choose q and b so that $|a - bq| > i + 2$. Similarly for fixing a to be $-i$.

So, totally we have $16i^2 + 8(i-1)^2$ ways to choose a, b and q so that $|a - bq| > i + 2$. But there exist $8i^3$ ways of choosing a, b, q over all. since $\frac{16i^2+8(i-1)^2}{8i^3} \approx \frac{1}{i} \to 0$ *as* $i \to \infty$, and entries of $\mathbf{A}_i$ are independent and uniformly distributed over $\mathbf{Z}_{[-i+1,i-1]}$ conditioned on the outer row and column of $\mathbf{A}_{i-1}$, so the probability of having a, b and q so that $|a - bq| > i + 2$ is very close to 0 as $i \to \infty$.

Hence, we have $\text{alg}(\mathbf{A}_{i+1}) = \text{alg}(\mathbf{A}_i) + 1$ with probability very close to 1 when the size of the original dense matrix is large. $\square$

B Active Randomization of a (s, t)-sparse Matrix

B.1 Introduction

The discussion in Donald [3] concerns "random" sparse simplicial complexes. Our analysis has concerned "random" (s, t)-sparse integer matrices. Unfortunately, these two categories are not in one-to-one correspondence because not every random (s, t)-sparse integer matrix corresponds to a legal boundary matrix of a triangulation. We now show how using *active randomization* we can get around this problem. Specifically, while the boundary matrices arising in homology-type computation of a triangulated geometric design are sparse as discussed in Donald [3], they can have non-uniform distribution and dependence on the probability of their entries being non-zero. Namely, two q-simplices in a triangulation can intersect at some common face, hence constraining some entries in the boundary matrix to be zero. This constraint arises from the simple fact that two simplices of the same dimension must have at least one different vertex. Hence the boundary matrices have non-uniform distribution and dependence on the probability of their entries being non-zero (see Donald [3] for discussion on the boundary matrices). However, in our probabilistic analysis of normal form computation of a (s, t)-sparse matrix, we assumed that a given (s, t)-sparse matrix has uniform distribution and independent probabilities on its entries being non-zero. Now, we want to relax this assumption. We show that our results go through for a given (s, t)-sparse matrix with non-uniform distribution and dependence on the probability of its entries being non-zero. For convenience, we will adopt the following definitions.

Definition 16 *A (s, t)-sparse matrix is called* non-uniform *and* dependent *if it has non-uniform distribution and dependence on the probability of its entries being non-zero. It is called* uniform *and* independent *if it has uniform distribution and independent probabilities on its entries being non-zero.*

Definition 17 *A permutation σ on n digits is called* random *if σ is uniform among all $n!$ permutations on n digits, i.e., for $j \in \mathbf{Z}_{[1,n]}$, the probability $P(\sigma(i) = j) = \frac{1}{n}$ and for $a_1, a_2, ..., a_{i-1}$ all distinct and different from j, $P(\sigma(i) = j \mid \sigma(1) = a_1, \sigma(2) = a_2, ..., \sigma(i-1) = a_{i-1}) = \frac{1}{n-i+1}$ with $1 \le i \le n$.*

In order to cope with non-uniform and dependent boundary matrices in our probabilistic analysis of normal form computation, we propose to *actively randomize* a given non-uniform and dependent (s, t)-sparse matrix. By *active randomization*, we mean the following:

Definition 18 *Consider a $n \times m$ (s,t)-sparse matrix A, that is non-uniform and dependent. We define a new $n \times m$ matrix A' as follows. We generate a random permutation σ on n digits, and initialize the $\sigma(i)^{th}$ row in matrix A' to be equal to the i^{th} row of matrix A for all i with $1 \leq i \leq n$. This is called a* random row permutation.

Definition 19 *A* random column permutation *is defined analogously.* Active randomization *is a sequence of random row and column permutations.*

We will derive the number β of random row and column permutations we need to perform during active randomization in order to obtain a uniform and independent $n \times m$ (s,t)-sparse matrix. We show that β is two, i.e., one random row permutation and one random column permutation. Because of the (s,t)-sparseness of the matrix, each random permutation takes linear time (i.e., $O(n)$). So, even a non-uniform and dependent (s,t)-sparse matrix can be diagonalized into normal form in expected time $O(n^2)$ with very high probability, i.e., this probability is very close to 1 as n is large.

B.2 Active Randomization of a Non-uniform and Dependent (s,t)-sparse Matrix

We first describe a pre-processing algorithm employing active randomization to convert a $n \times m$ non-uniform and dependent (s,t)-sparse matrix into a uniform and independent one with two random row and column permutations. In other words, we perform one random row permutation and one random column permutation. Then we will show that these two operations suffice to perform to actively randomize a non-uniform and dependent (s,t)-sparse matrix so that a uniform and independent (s,t)-sparse matrix can be obtained.

Lemma 2 *A random permutation on n digits can be generated in linear time (i.e. $O(n)$).*

Proof:
We start out with a permutation σ such that $\sigma(i) = i$ for $1 \leq i \leq n$. We loop $n-1$ times. At the i^{th} iteration, we pick a random number j from $\mathbf{Z}_{[i,n]}$ and interchange $\sigma(i)$ with $\sigma(j)$. At the end, we obtain a new permutation σ. For any $k \in \mathbf{Z}_{[1,n]}$, the probability $P(\sigma(i) = k) = \frac{1}{n}$ and for $a_1, a_2, ..., a_{i-1} \in \mathbf{Z}_{[1,n]}$ all distinct and different from k, $P(\sigma(i) = k \mid \sigma(1) = a_1, \sigma(2) = a_2, ..., \sigma(i-1) = a_{i-1}) = \frac{1}{n-i+1}$, so the newly obtained σ is a random premutation. Clearly, the above process of generating a random permutation takes linear time, i.e., $O(n)$. □

B.2.1 The Algorithm for Active Randomization

We proceed as follows. Let A be the given non-uniform and dependent $n \times m$ (s,t)-sparse matrix and A', A'' be dummy $n \times m$ matrices

Algorithm 1 *(Active Randomization)*

Step 1) We generate a random permutation σ on n digits. We loop n times. At the i^{th} iteration, we replace the $\sigma(i)^{th}$ row of matrix A' with the i^{th} row of matrix A.

Step 2) We generate a random permutation τ on m digits. We loop m times. At the i^{th} iteration, we replace the $\tau(i)^{th}$ column of A'' with the i^{th} column of matrix A'.

When diagonalizing a given non-uniform and dependent (s,t)-sparse matrix, we first apply the active randomization algorithm as a pre-processing step to obtain a uniform and independent (s,t)-sparse matrix. Then we apply the reduction algorithm to obtain the normal form of the original non-uniform and dependent (s,t)-sparse matrix.

B.2.2 Probabilistic Analysis of Active Randomization

Clearly, the normal form of a matrix is unchanged under any permutations of rows and columns since such permutations are part of the elementary row and column operations in the reduction algorithm. So, the normal form of a matrix is preserved under random permutations of rows and columns. More importantly, we are interested in the following question:

Question 3 *Can we perform active randomization on a non-uniform and dependent (s,t)-sparse matrix such that a finite number of random row and column permutations results in a uniform and independent matrix? If so, what is the number of random row and column permutations we need to perform?*

For the purpose of this paper, we can wlog assume that the non-zero entries in the non-uniform and dependent (s,t)-sparse matrix are taken from the set $\{-1,1\}$. We can easily generalize to non-uniform and dependent (s,t)-sparse matrices with small algebraic complexity.

In order to answer the above question, let us first introduce some standard tools from probability theory, namely, discrete random variables. Let X be a discrete random variable defined on $\mathbf{Z}_{[1,n]} \times \mathbf{Z}_{[1,m]}$ and taking values from the set $\mathbf{Z}_{[-1,1]}$. Let X_i (resp. X_j) be discrete random variable defined on $\mathbf{Z}_{[1,n]}$ (resp. $\mathbf{Z}_{[1,m]}$) and taking values from $\mathbf{Z}_{[1,n]}$ (resp. $\mathbf{Z}_{[1,m]}$). Clearly, the composite of X with X_i and X_j is also a discrete random variable. In addition, we require that the discrete random variables X have the following probability distribution,

$$P(X=0) = 1-\frac{s}{n}-\frac{t}{m}+\frac{st}{nm} \tag{30}$$

$$P(X\neq 0) = \frac{s}{n}+\frac{t}{m}-\frac{st}{nm} \tag{31}$$

$$P(X=-1) = P(X=1), \tag{32}$$

X_i and X_j have uniform probability distribution, and X, X_i, X_j are independent discrete random variables. For a given non-uniform and dependent matrix B, we want to know whether we can actively randomize A to obtain a matrix A' so that

$$P(A'_{i,j}\neq 0) = P(X(i,j)\neq 0) \tag{33}$$

and the probabilities of its entries being non-zero are independent. Recall the probability distribution of entries being non-zero for a uniform and independent (s,t)-sparse matrix in section 2.1.1. By requiring the above probability distribution (i.e., eq. (30), (31) & (32)) and independence for the discrete random variables X, X_i and X_j, we see that the composite of X with X_i and X_j, i.e., $X(X_i,X_j)$, in fact describes the independent probability distribution of an entry in a uniform and independent (s,t)-sparse matrix. So, if the probabilities of entries in the matrix (after active randomization) satisfy eq. (33), we obtain a uniform and independent (s,t)-sparse matrix. Also we will see later that X_i and X_j capture a random position to which an entry in the original non-uniform and dependent $n \times m$ (s,t)-sparse matrix is moved after one random row permutation and one random column permutation.

We know that after applying the active randomization algorithm, each entry in the original matrix A is moved to somewhere. The following lemma will tell us that each entry of A is moved to a random position.

Lemma 3 *Let A be a non-uniform and dependent $n \times m$ (s,t)-sparse matrix and A' the resulting matrix. Then $A'_{X_i(k),X_j(l)} = A_{k,l}$ where X_i and X_j are discrete random variables defined above.*

Proof:
We wlog assume $A'_{u,v} = A_{k,l}$. Let $\diamond$ sit at the (k,l)-position and $\triangle$ at the (u,v)-position. we have the following picture:

$$\begin{pmatrix} \cdots & \cdots & \cdots & \cdots & \cdots \\ \cdots & \triangle & \cdots & \cdots & \cdots \\ \cdots & \cdots & \cdots & \cdots & \cdots \\ \cdots & \cdots & \cdots & \diamond & \cdots \\ \cdots & \cdots & \cdots & \cdots & \cdots \end{pmatrix} \tag{34}$$

Setting X_i (resp. X_j) to some value u (resp. v) with probability $\frac{1}{n}$ (resp. $\frac{1}{m}$) corresponds to randomly permuting to a row (resp. column) which happens to be the u^{th} row (resp. v^{th} column) and moving $\diamond$ to the (u,v)-position. So, pair (X_i, X_j) captures a random position to which an entry in the original non-uniform and dependent (s,t)-sparse matrix is moved. □

Theorem 3 *After actively randomizing a given non-uniform and dependent $n \times m$ (s,t)-sparse matrix by performing one random row permutation and one random column permutation, we obtain an (s,t)-sparse matrix with uniform distribution on the probability of its entries being non-zero.*

Proof:
Let A be a given non-uniform and dependent (s,t)-sparse matrix. We want to show that after one random row permutation and one random column permutation, we can obtain a resulting matrix A' with

$$P(A'_{k,l} \neq 0) = P(X(k,l) \neq 0) \tag{35}$$

where $1 \leq k \leq n, 1 \leq l \leq m$. In order to achieve the equality in eq. (35), we certainly need to assign

$$A'_{X_i(k),X_j(l)} := A_{k,l}. \tag{36}$$

By lemma 3, eq. (36) is guaranteed to be achieved. □

Theorem 4 *After actively randomizing a given non-uniform and dependent $n \times m$ (s,t)-sparse matrix by performing one random row permutation and one random column permutation, we obtain a (s,t)-sparse matrix with independent probabilities of its entries being non-zero.*

Proof:
Let A be our given non-uniform and dependent (s,t)-sparse matrix. We want to show that after one random row permutation and one random column permutation, we can obtain a resulting matrix

A' such that $P(A'_{k,l} \neq 0)$ and $P(A'_{u,v} \neq 0)$ are independent with $1 \leq k, u \leq n, 1 \leq l, v \leq m$ and $k \neq u$ or $l \neq v$. In order to achieve these independent probabilities, we certainly need to have

$$A'_{k,l} = X(k,l) \tag{37}$$
$$A'_{u,v} = X(u,v). \tag{38}$$

In order to achieve eq. (38) & (38), we need to assign

$$A'_{X_i(k),X_j(l)} := A_{k,l} \tag{39}$$
$$A'_{X_i(u),X_j(v)} := A_{u,v}. \tag{40}$$

By lemma 3, eq. (39) & (40) are guaranteed to be achieved. □

Theorems 3 & 4 prove the correctness of our active randomization algorithm. We have already noticed that uniformness and independence are stable in the sense that performing any additional active randomization on a uniform and independent (s,t)-sparse matrix will not destroy its uniformness and independence. Therefore, we conclude that performing active randomization by using one random row permutation and one random column permutation on a $n \times m$ non-uniform and dependent (s,t)-sparse matrix will result in a uniform and independent (s,t)-sparse matrix.