Math. Systems Theory **28**, 215–228 (1995)

Mathematical Systems Theory

Communication Complexity of Matrix Computation over Finite Fields*

J. I. Chu[1] and G. Schnitger[2]

[1]Computer Science Department, New Mexico State University, Las Cruces, NM 88003, USA

[2]Computer Science Department, Pennsylvania State University, University Park, PA 16802, USA

Abstract. We investigate the communication complexity of *singularity testing* in a finite field, where the problem is to determine whether a given square matrix M is singular. We show that, for $n \times n$ matrices whose entries are elements of a finite field of size p, the communication complexity of this problem is $\Theta(n^2 \log p)$. Our results imply tight bounds for several other problems like *determining the rank* and *computing the determinant.*

1. Introduction

Parallel computing saves time by spreading the work load among a number of active elements. However, this may create the necessity of communication among those active elements. Yao [12], [13] introduced the deterministic model of *communication complexity* to measure this additional computational resource. In this model there are two agents performing the desired computation cooperatively. The agents have their own memory and can exchange messages with each other. When the computation starts, each agent receives one-half of the input bits according to a fixed partition rule π. Then they execute a fixed protocol P which leads the two agents to perform local computations and communicate messages according to their own information and previously obtained messages. At the end of the computation, each agent knows the values of the output variables for which it is responsible. The number of bits exchanged represents the communication requirement for this particular computation.

* This research was supported in part by NSF Grant CCR-8805978 and AFOSR Grant 87-0-400.

Let f be the function to be computed by a parallel system and let N be the number of input bits. Call any partition rule which divides these N bits into two even portions an even partition. Every even partition, π, together with a correct protocol, P, represents a possible design of a system computing f. Let $C(\pi, P)$ be the maximum number of bits to be exchanged for some input instance of size N under partition π and protocol P. The communication complexity of f under partition π, $C(\pi)$, is the minimum of $C(\pi, P)$ over all correct protocols P for π. Then the communication complexity of f is defined to be the minimum of $C(\pi)$ over all even partitions π.

Thompson's [9] area–time tradeoffs can be best expressed by communication complexity. Suppose f has N input bits and a VLSI chip computing f has area A and worst-case computation time T. Thompson shows the tradeoff $AT^2 = \Omega(I^2)$. Subsequently, the area bound $A = \Omega(I)$ was proven [2], [10], [13]. Combining these two formulas, we get $AT^{2\alpha} = \Omega(I^{1+\alpha})$, $0 \leq \alpha \leq 1$.

We investigate the communication complexity of testing whether a matrix over a finite field is singular.

Theorem 1.1. *Let F_p be a finite field with p elements. Let M be a square matrix of dimension n, where each entry consists of $\lceil \log p \rceil$ bits encoding an element of F_p. (All logarithms in this paper are base 2.) The communication complexity of "deciding whether M is nonsingular" is $\Theta(n^2 \log p)$.*

Chu and Schnitger [4] showed that when entries of M are integers of up to k bit long, the corresponding communication complexity is $\Theta(kn^2)$. The proof for the integer case relies on the ability of encoding vectors of very large components by matrices. That is not applicable in a finite field as any value can be represented by $\log p$ bits.

The commonly used *transitivity* approach of Vuillemin [10] does not seem to work for our problem, either. Vuillemin's approach is successful for many functions or languages that are powerful enough to express the *identity* problem (given two strings x and y, are x and y identical?). However, there does not seem to be a large enough identity problem embedded in the singularity testing problem.

We utilize the following property of vector spaces over a finite field to bound the size of monochromatic submatrices in the truth matrix of our problem (see Section 2). Let F_p be a finite field of p elements. Assume S_1 and S_2 are two sets of dimension m subspaces in F_p^{2m} such that, for any $A \in S_1$ and any $B \in S_2$, $A \cap B$ contains only the zero vector. Then $\# S_1 \cdot \# S_2 \leq p^{m^2+2m}$ (Lemma 3.7).

We prove Theorem 1.1 in Section 3. Theorem 1.1 also establishes the communication complexity of the following problems.

Corollary 1.2. *Let M be defined as in Theorem* 1.1. *The communication complexity of the following problems is $\Theta(n^2 \log p)$.*

(a) *Computing the determinant of M.*
(b) *Computing the rank of M.*

Roundoff errors are inevitable when computing in the ring of real numbers using digital computers. In applications where Fourier transform is used to compute a convolution efficiently, an exact result is often required. However, the fast Fourier transform (FFT) implementation of cyclic convolution introduces significant amounts of roundoff errors [6]. These errors can be avoided by computing over a suitable finite field. Agarwal and Burrus [1] give efficient implementation of finite transforms over rings of integers modulo Fermat numbers. They report a much faster performance over the best FFT implementation. Convolution over a finite field also speeds up multiplications of long integers [7]. Another important potential of computing in finite fields is the use of residue number systems in arithmetic computations [8]. VLSI implementation computing complex inner products using residue arithmetic has been proposed [11].

Let X be a finite set of vectors from a vector space U and let L be the set of subspaces $\{V: V \text{ is spanned by some subset of } X\}$. The *Vector Space Span* problem is defined as follows: Given two elements of L, V_1 and V_2, decide whether their union spans U.

Lovász and Saks [5] determined the *fixed-partition* communication complexity of the vector space span problem to be $\Theta(\log(\#L))$, where one agent reads V_1 and the other reads V_2. Theorem 1.1 also establishes the *unrestricted* communication complexity of this problem when X is chosen to be the set of vectors over some finite field.

Corollary 1.3. *When $X = \boldsymbol{F}_p^n$, the communication complexity of the vector space span problem is $\Theta(n^2 \log p)$.*

We present a sketch of our proof techniques in Section 2. The proof of Theorem 1.1 is given in Section 3.

2. Proof Techniques

Let f be some function to be computed using the communication complexity model. Let us fix the size of the input to be N bits and the input partition to be π. Then the computation can be viewed as computing a function of two arguments, in which the first argument consists of the $N/2$ input bits given to the first agent and the second argument consists of the remaining bits.

In case the output consists of one bit, we can characterize the computation of a two-argument boolean function by a *Truth Matrix*. Each possible instance of the first argument takes up one row of the truth matrix and each possible instance of the second argument takes up one column. The matrix entry at the intersection of a particular pair of row and column contains the output bit corresponding to the arguments assigned to this row and this column.

A submatrix of the truth matrix is a *monochromatic submatrix* if it contains only a single value. More specifically, if this value is 1, the monochromatic submatrix is called 1-chromatic; otherwise it is called 0-chromatic. Yao [12] shows that, given the input partition π (and hence the truth matrix) of some boolean-valued

function f, the communication complexity of f under partition π is at least $\log d(f) - 2$, where $d(f)$ is the minimal number of disjoint monochromatic submatrices that partition the truth matrix of f.

Suppose when the input dimension, n, of the singularity testing is odd, the communication complexity is $o(n^2 \log p)$. Then the $o(n^2 \log p)$ bound also holds for the case when n is even. This is because, for any even n, the agents can pretend that there are one additional row and one additional column in the input matrix. All the entries on the $(n + 1)$st row and the $(n + 1)$st column are 0 except that the $[n + 1, n + 1]$ entry is nonzero. Therefore, we only consider even input dimensions when proving Theorem 1.1.

For the proof of Theorem 1.1, we first tackle the case of a particular input partition as defined below; then we show that arbitrary partitions do not change the communication complexity asymptotically.

Definition 2.1. Assume the input is a $2m \times 2m$ matrix. Let π_0 denote the following input partition: the first agent receives all the bits encoding the entries in the first m columns; the second agent receives the other half.

Given π_0, we define an entry of the truth matrix to be "one" if the corresponding input matrix is nonsingular. We show that there is a large submatrix, T, of the truth matrix such that:

(1) T contains a large number ($>p^{n^2/2}$) of "one" entries.
(2) Every 1-chromatic submatrix of T is of relatively small size ($\leq p^{n^2/4+n}$).

This allows us to apply Yao's lower bound method.

The first property of T as given above is easy to prove. We use induction to establish the second property. The inductive step involves proving a $\Theta(m \log p)$ bound on the communicaton complexity of "the inner product over F_p." In this new problem each agent reads a vector in F_p^m and tries to decide whether the two input vectors have a nonzero inner product ("one" entries in its truth matrix represent nonzero products). In this proof we use techniques of Chor and Goldreich [3]. In particular, we utilize that the truth matrix of the inner product function can be made orthogonal if we replace the value 0 by $1 - p$.

3. Singularity Testing over Finite Fields

In this section we prove Theorem 1.1 Let F_p be a finite field with p elements. Our input, an $n \times n$ matrix over F_p, can be represented by $n^2 \log p$ bits. First, we establish the $\Theta(n^2 \log p)$ bound on the communication complexity of this problem under the assumption that the input is partitioned according to π_0 (see Definition 2.1). Note that the upper bound is trivial.

We restrict our attention to a submatrix T of the truth matrix. Interpret each column of the input matrix as a vector in F_p^n. Then each row or column in the truth matrix corresponds to a subspace spanned by $n/2$ vectors. We select for T exactly one row and one column that correspond to each dimension $n/2$ subspaces of F_p^n.

The entry $[i, j]$ in T is set to 1 whenever the two subspaces corresponding to the ith row and and jth column intersect only at the zero vector, i.e., whenever the input matrix is nonsingular.

Lemma 3.1. *The number of "one" entries in T is at least $p^{n^2/2}$.*

Proof. First we count the number of "one" entries in a row of T. Assume that a row is labeled by a vector space W of dimension $n/2$. Consider the number of distinct bases which span subspaces of dimension $n/2$ and intersect W only at the zero vector. Let us construct such bases by picking $n/2$ base vectors one by one.

If $V \subset \boldsymbol{F}_p^n$ is a subspace of dimension d, there are exactly $p^n - p^d$ vectors outside of V. We have $\#(\boldsymbol{F}_p^n \backslash W) = p^n - p^{n/2}$ choices for the first base vector and $p^n - p^{n/2+1}$ choices for the second base vector after the first one is chosen. In general, once the first i base vectors are chosen we have $p^n - p^{n/2+i}$ choices for the $(i+1)$st base vector. So the total number of bases is $\prod_{i=n/2}^{n-1} (p^n - p^i)$.

The same counting argument shows that there are $\prod_{j=0}^{n/2-1}(p^{n/2} - p^j)$ distinct bases for each subspace of dimension $n/2$. Therefore, the number of "one" entries in any row of T is equal to

$$\prod_{j=0}^{n/2-1} \frac{p^n - p^{n/2+j}}{p^{n/2} - p^j} = \prod_{j=0}^{n/2-1} p^{n/2} = p^{n^2/4}.$$

Since T is symmetric, there are $p^{n^2/4}$ "one" entries in each column of T. This means T has at least $p^{n^2/4}$ rows. Therefore, the number of "one" entries in T is at least $p^{n^2/2}$. □

Let $IPZ(m, p)$ denote the problem of deciding whether the inner product of two vectors in $\boldsymbol{F}_p^m$ is nonzero. Let π_1 be the input partition of $IPZ(m, p)$ such that each agent reads one of the input vectors. We now establish an intermediate result which states that $IPZ(m, p)$ has communication complexity $\Theta(m \log p)$ under partition π_1. This generalizes a result of Chor and Goldreich [3] for the inner product over $\boldsymbol{F}_2$.

Let Ψ be the set of dimension $m-1$ subspaces of $\boldsymbol{F}_p^m$. Ψ contains $\#\Psi = (p^m - 1)/(p-1)$ elements because a hyperplane is specified by one linear equation which is unique up to the $p-1$ multiples of its coefficients. We set up an auxiliary matrix Z as follows. Create one row for each space in Ψ and one column for each vector in $\boldsymbol{F}_p^m$. For the entry at row Q and column $\boldsymbol{v}$, set its value to $(1-p)$ if $\boldsymbol{v} \in Q$; otherwise, set it to 1.

A vector $\boldsymbol{v}$ is orthogonal to a subspace Q if and only if the inner product $\boldsymbol{v} \cdot \boldsymbol{t}$ is 0 for all $\boldsymbol{t} \in Q$. Each row of Z also corresponds to the p vectors orthogonal to Q. Another way to interpret Z is that each row or column of Z corresponds to a vector; an entry is $(1-p)$ if and only if the two vectors have inner product 0. Hence, Z is also a submatrix of the truth matrix of $IPZ(m, p)$ under π_1 if we replace the value $(1-p)$ by 0.

Lemma 3.2. *The rows of Z are pairwise orthogonal.*

Proof. We use $\boldsymbol{r}_i$ to denote the ith row vector of Z. Pick any two different row vectors of Z, say $\boldsymbol{r}_i$ *and* $\boldsymbol{r}_j$. Let Q_i and Q_j be the elements of Ψ which correspond to these two row vectors, respectively. Then $\boldsymbol{r}_i \cdot \boldsymbol{r}_j = \sum_{k=1}^{p^m} (Z[i, k] \cdot Z[j, k])$. Let $\boldsymbol{v}$ be the vector corresponding to the kth column of Z. Then

$$Z[i,k] \cdot Z[j,k] = \begin{cases} (1-p)^2, & \text{if } \boldsymbol{v} \in (Q_i \cap Q_j), \\ 1, & \text{if } \boldsymbol{v} \notin (Q_i \cup Q_j), \\ 1-p, & \text{otherwise.} \end{cases}$$

Note that $\#(Q_i \cap Q_j) = p^{m-2}$ and $\#(Q_i \cup Q_j) = 2p^{m-1} - p^{m-2}$. We obtain $\boldsymbol{r}_i \cdot \boldsymbol{r}_j = (p-1)^2 p^{m-2} + (p^m - 2p^{m-1} + p^{m-2}) - (p-1)(2p^{m-1} - 2p^{m-2}) = p^{m-2}(2(p-1)^2 - 2(p-1)^2) = 0$. □

Lemma 3.3. *The size of any 1-chromatic submatrix of Z cannot exceed* $(p-1)p^m$.

Proof. Pick any a rows and b columns of Z. Since we do not demand any ordering of rows and columns, we assume the first a rows and the first b columns are chosen in order to simplify notation. Define $x = \sum_{i=1}^{a} \sum_{j=1}^{b} Z[i, j]$. If the submatrix defined by these a rows and b columns is 1-chromatic, then x^2 must equal $(a\ b)^2$. It suffices to prove that $x^2 \leq ab(p-1)p^m$.

Observe from the proof of Lemma 3.2 that, for every i, $\boldsymbol{r}_i \cdot \boldsymbol{r}_i = (p-1)^2 p^{m-1} + (p^m - p^{m-1}) = (p-1)p^{m-1}(p-1+1) = (p-1)p^m$. Recall Cauchy's inequality: $(\boldsymbol{v} \cdot \boldsymbol{u})^2 \leq \|\boldsymbol{v}\|^2 \cdot \|\boldsymbol{u}\|^2$.

$$\begin{aligned} x^2 &= \left(\sum_{j=1}^{b} \left(1 \cdot \sum_{i=1}^{a} Z[i,j] \right) \right)^2 && \text{interpret the outer summation as inner product and apply Cauchy's inequality} \\ &\leq b \sum_{j=1}^{b} \left(\sum_{i=1}^{a} Z[i,j] \right)^2 \\ &\leq b \sum_{j=1}^{p^m} \left(\sum_{i=1}^{a} Z[i,j] \right)^2 \\ &= b \left(\sum_{i=1}^{a} \boldsymbol{r}_i \cdot \sum_{i=1}^{a} \boldsymbol{r}_i \right) && \text{apply orthogonality} \\ &= b \sum_{i=1}^{a} (\boldsymbol{r}_i \cdot \boldsymbol{r}_i) \\ &= ab(p-1)p^m. \end{aligned}$$

Theorem 3.4. *The communication complexity of IPZ(m, p) under partition* π_1 *is* $\Theta(m \log p)$.

Proof. Since Z has $\#\Psi = (p^m - 1)/(p - 1)$ rows and each rows has $p^m - p^{m-1}$ "one" entries, there are $(p^m - 1)p^{m-1}$ "one" entries in Z. Applying Yao's lower-bound method, we obtain that the communication complexity of $IPZ(m, p)$ under π_1 is at least $\Omega(\log((p^m - 1)/p(p - 1)) = \Omega(m \log p)$. On the other hand, the input read by each agent has $m \log p$ bits, which implies a trivial upper bound. □

We restate Lemma 3.3 in a form needed in subsequent proofs.

Lemma 3.5. *Let $Q_1, \ldots, Q_t$ be elements of Ψ. Define D to be the complement of $Q_1 \cup \cdots \cup Q_t$ with respect to F_p^n. Then $t \times (\#D) \leq (p - 1)p^m$.*

Proof. The r rows of Z corresponding to Q_i, $1 \leq i \leq t$, and the $\#D$ columns corresponding to vectors in D form a 1-chromatic submatrix of Z. □

Let Y_1 be any subset of F_p^n which contains $\mathbf{0}$ (the zero vector). Let Y_2 denote the set $(F_p^n \backslash Y_1) \cup \{\mathbf{0}\}$. Observe that any choice of Y_1 specifies a 1-chromatic submatrix, S, of T—the rows (columns) of S correspond to the dimension $n/2$ subspaces contained in Y_1 (Y_2). Conversely, every 1-chromatic submatrix induces canonically a set Y_1.

Defintion 3.6. Let i be any integer in the range $[0, n/2]$. Set $m := i + n/2$. Let Y_1 and Y_2 be any two subsets of F_p^m such that $\mathbf{0} \in Y_1$ and $Y_2 = (F_p^m \backslash Y_1) \cup \{\mathbf{0}\}$. We call any dimension $n/2$ subspace a Y_1-space if it is contained in Y_1. Also, any dimension i subspace contained in Y_2 is called a Y_2-space. Let $s(Y_1)$ and $t(Y_2)$ be the number of Y_1-spaces and Y_2-spaces, respectively. Define $E(i)$ as the maximum, over all choices of Y_1, of the product $s(Y_1) \cdot t(Y_2)$.

We now investigate the fixed partition case of the singularity testing problem. Note that when $i = n/2$, every Y_2-space corresponds to a column in some 1-chromatic submatrix of T, and every Y_1-space corresponds to a row. Hence, $E(n/2)$ is the maximum size of any 1-chromatic submatrix.

Lemma 3.7. *The maximum size of any 1-chromatic submatrix of T is at most $p^{n^2/4+n}$.*

Proof. We claim that, for all $0 \leq i \leq n/2$, $E(i) \leq p^{i \cdot n/2+2i}$. Hence, $E(n/2) \leq p^{n^2/4+n}$. This claim is proved by induction on i.

Basis: $i = 0$.

We have either $Y_1 = F_p^{n/2}$ or there are no Y_1-spaces. In the former case $s(Y_1) = 1$ and $t(Y_2) = 1$. Therefore, $E(0) \leq 1$.

Induction Step: Assume that $E(i - 1) \leq p^{(i-1)n/2+2i-2}$ for some $i \leq n/2$. We prove that $E(i) \leq p^{i\,n/2+2i}$.

Let m, Y_1, and Y_2 be defined as in Definition 3.6 such that $E(i)$ is maximized. Let Ψ be the set of dimension $m - 1$ subspaces of F_p^m. Recall that a Y_1-space is a dimension $n/2$ subspace contained in Y_1 and Y_2-space is a dimension i subspace contained in Y_2.

Let Q be the element of Ψ which contains the largest number of Y_1 spaces. Let this number be $r > 0$. Consider the set of dimension $i - 1$ spaces contained in $Y_2 \cap Q$. Assume there are c of them. Then $r \cdot c \leq p^{(i-1)n/2+2i-2}$ by the induction hypothesis. Also, each Y_2-space intersects Q at exactly one of these c spaces. To see this, let R_1 be any Y_1-space in Q and let K be the intersection of any Y_2-space and Q. Then K must be a vector space. Since R_1 and K only intersect at the zero vector, the dimension of K is no more than $i - 1$. Since R_1 and any Y_2-space span $\boldsymbol{F}_p^m$, the dimension of K is no less than $i - 1$.

Let K be an arbitrary dimension $i - 1$ vector space contained in $Y_2 \cap Q$. In order to bound $t(Y_2)$, the total number of Y_2-spaces, we first bound the number of Y_2-spaces containing K.

We can generate a Y_2-space containing K by adding one more base vector, $\boldsymbol{b}$, to the basis of K. Let R_1 be any Y_1-space. Then $\boldsymbol{b}$ cannot belong to the dimension $m - 1$ supspace $R_1 + K$. Assume that there are l different spaces $R_i + K$ (for Y_1-spaces $R_1, \ldots, R_l$). We call l the *stiffness* of K. Then, by Lemma 3.5, there are at most $(p - 1)p^m/l$ choices for $\boldsymbol{b}$.

Note that each Y_2-space thus generated from K contains $(p - 1)p^{i-1}$ possible choices of $\boldsymbol{b}$. This is because a Y_2-space contains p^i vectors and each of them, except those p^{i-1} vectors in K, may be used as $\boldsymbol{b}$. Therefore, the total number of Y_2-spaces containing K is at most $(p - 1)p^m/(l(p - 1)p^{i-1}) = p^{n/2+1}/l$.

Assume that K_0 minimizes the stiffness among all dimensions $i - 1$ subspaces contained in $Y_2 \cap Q$ and assume that its stiffness equals l_0. then $t(Y_2) \leq cp^{n/2+1}/l_0$.

Let R_j, $1 \leq j \leq l_0$, be the Y_1-spaces defining the stiffness of K_0. Then every Y_1-space R is contained in some dimension $m - 1$ space $R_j + K_0$. Since r is the largest number of Y_1-spaces any dimension $m - 1$ subspace may contain, we have $s(Y_1) \leq r \cdot l_0$.

Combining the results in the above two paragraphs, we have $s(Y_1) \cdot t(Y_2) \leq r \cdot c \cdot p^{n/2+1} \leq p^{(i-1)n/2+2i-2} \cdot p^{n/2+1} = p^{i \cdot n/2+2i-1}$. Therefore, $E(i) \leq p^{i \cdot n/2+2i}$. □

The $\Theta(n^2 \log p)$ lower bound for the input partition π_0 now follows directly from Lemmas 3.1 and 3.7.

We now prove our results for arbitrary partitions that evenly divide the input between the two agents. Since permuting the columns of a matrix does not change its rank, we only need to consider the *proper* partitions defined below.

Definition 3.8. An even input partition is a proper partition if and only if, for any pair of integers i and j such that $1 \leq i < j \leq n$, at least as many input bits from column i are assigned to the first agent as from column j.

Let π be any proper partition. Then π only differs from π_0 on the assignment of up to $(n^2 \log p)/2$ input bits. If the values of these input bits are fixed, then the difference between π and π_0 is eliminated.

Definition 3.9. Given any proper partition π, we use Δ to denote the set of input bit positions which are assigned differently under π_0 and π. A function $f: \Delta \to \{0, 1\}$ is called a fixing function for π.

No matter what fixing function we use to restrict the computation, the resulting truth matrix is a submatrix of the truth matrix defined by π_0. To analyze the communication complexity of this restricted computation, we select a submatrix T' in the same way we selected T. Accordingly, T' is a submatrix of T. So, the maximum size of 1-chromatic matrices in T' is still bounded according to Lemma 3.7. Therefore, it suffices to show that, for every proper partition, a fixing function exists such that the resulting matrix T' contains at least $p^{n^2/3}$ "one" entries. This will establish Theorem 1.1.

We first prove that such fixing functions exist for some of the proper partitions.

Defintion 3.10. A proper partition is called a *nice* partition if and only if it partitions the input bits in each column of the input matrix evenly between the agents.

In other words, if π is a nice partition, then Δ contains exactly $\frac{1}{2}n \log p$ bit positions from each column of the input matrix. We use Δ_1 to denote those bit positions of Δ in the first $n/2$ columns and Δ_2 to denote the other half of Δ. Accordingly, for any fixing function f of π, f_1 (f_2) is the restriction of f relative to Δ_1 (Δ_2).

Definition 3.11. Let X be any dimension $n/2$ subspace of $\boldsymbol{F}_p^n$, let π be a nice partition, and let f be a fixing function for π. We say X is *encodable* under f_1 if and only if there is an instance of the input matrix whose bits in Δ_1 are fixed according to f_1 and whose first $n/2$ columns span X. Each of such instances is an *encoding* of X under f_1.

The notion "encodable under f_2" is defined in a similar way. If X is encodable under f_1 and Y is encodable under f_2, then we say the pair (X, Y) is encodable under f.

By Lemma 3.1, there are more than $p^{n^2/2}$ pairs of dimensions $n/2$ subspaces which introduce "one" entries in T. Our strategy is to show that for any nice partition there is a collection of $p^{2\varepsilon n^2}$ fixing functions such that almost all such pairs of subspaces are encodable under some function in this collection. Then, by the pigeonhole principle, there is one fixing function which yields $p^{(1/2-2\varepsilon)n^2}$ "one" entries in T'.

Assume that we are given a nice partition π. Our first step is to show that most dimension $n/2$ subspaces have not too many encodings under any fixing function and hence are encodable under a large number of fixing functions.

Definition 3.12. Set $\varepsilon := \frac{1}{16}$. A dimension $n/2$ subspace is *concise* if and only if it does not have more than $p^{2\varepsilon n^{7/4}}$ encodings under any fixing function f_1.

Using the first $n/2$ columns of the input matrix to encode any dimension $n/2$ subspace X, we have $\prod_{i=0}^{n/2-1}(p^{n/2} - p^i) \geq p^{n^2/4} \cdot e^{-2/(p-i)}$ different ways to do so. (This inequality can be proved by using some suitable inequalities appearing in the proof of Lemma 3.14.)

Proposition 3.13. *If X is concise, it must be encodable under more than $p^{n^2/4-2\varepsilon n^{7/4}} \cdot e^{-2/(p-1)}$ fixing functions.*

Lemma 3.14. *Set $\delta := p^{-\varepsilon^2 n^{3/2}}$. Assume $p \geq 8$ and $n \geq 128$. All the dimension $n/2$ subspaces in $\boldsymbol{F}_p^n$ except for a δ fraction of them, are concise.*

Proof. After the bit proposition in Δ_1 are fixed according to any given f_1, each column in the left half of the input matrix can only encode $p^{n/2}$ vectors. Let V_i be the set of vectors encodable by the ith column, $1 \leq i \leq n/2$. To encode a dimension $n/2$ subspace X, we must find one base vector of X in every V_i. So the number of encodings of X under f_1 is not larger than $\prod_{i=1}^{n/2} \#(X \cap V_i)$.

If X is not concise, then there must be an f_1 and an i such that $\#(X \cap V_i) > p^{4\varepsilon n^{3/4}}$. Given f_1 and i, let U be the set $\{X\colon \#(X \cap V_i) > p^{4\varepsilon n^{3/4}}\}$. Let *Total* represent the total number of dimension $n/2$ subspaces in $\boldsymbol{F}_p^n$. So

$$Total = \prod_{j=0}^{n/2-1} \frac{p^n - p^j}{p^{n/2} - p^j}.$$

Each subspace in U must have some basis containing $4\varepsilon n^{3/4}$ base vectors from V_i. Let us generate a basis for the subspaces in U by first choosing any $4\varepsilon n^{3/4}$ linearly independent vectors in V_i and then adding $n/2 - 4n^{3/4}$ base vectors from $\boldsymbol{F}_p^n$ one by one. The number of different basis representations is at most

$$\binom{p^{n/2}}{4\varepsilon n^{3/4}} \prod_{j=4\varepsilon n^{3/4}}^{n/2-1} (p^n - p^j). \tag{3.1}$$

However, each subspace in U contains at least

$$\prod_{j=0}^{4\varepsilon n^{3/4}-1} \frac{p^{4\varepsilon n^{3/4}} - p^j}{4\varepsilon n^{3/4} - j} \prod_{j=4\varepsilon n^{3/4}}^{n/2-1} (p^{n/2} - p^j) \tag{3.2}$$

sets of base vectors counted in (3.1). So, $\#U$ is at most the quantity in (3.1) divided by the quantity in (3.2), which is

$$\prod_{j=0}^{4\varepsilon n^{3/4}-1} \frac{p^{n/2} - j}{p^{4\varepsilon n^{3/4}} - p^j} \prod_{j=4\varepsilon n^{3/4}}^{n/2-1} \frac{p^n - p^j}{p^{n/2} - p^j}.$$

So, $\#U/Total$ is at most

$$\prod_{j=0}^{4\varepsilon n^{3/4}-1} \frac{p^{n/2} - j}{p^{4\varepsilon n^{3/4}} - p^j} \frac{p^{n/2} - p^j}{p^n - p^j}. \tag{3.3}$$

Note that

$$\prod_{j=0}^{m-1} (p^m - p^j) = p^{m^2} \prod_{j=0}^{m-1} (1 - p^{j-m}) = p^{m^2} \prod_{j=1}^{m} (1 - p^{-j}).$$

Also,

$$e^{-2p^{-j}} = 1 - 2p^{-j} + 2p^{-2j} - 4p^{-3j}/3 + \cdots < 1 - p^{-j}.$$

So,

$$\prod_{j=0}^{4\varepsilon n^{3/4}-1} (p^{4\varepsilon n^{3/4}} - p^j)^{-1} = p^{-16\varepsilon^2 n^{3/2}} \cdot \prod_{j=1}^{4\varepsilon n^{3/4}} (1 - p^{-j})^{-1}$$

$$< p^{-16\varepsilon^2 n^{3/2}} \cdot \prod_{j=1}^{4\varepsilon n^{3/4}} e^{2p^{-j}} < p^{-16\varepsilon^2 n^{3/2}} \cdot e^{2/(p-1)}.$$

Since $(p^{n/2} - p^j)/(p^n - p^j) \leq p^{-n/2}$, we have

$$(3.3) < p^{-16\varepsilon^2 n^{3/2}} e^{2/(p-1)} \cdot \prod_{j=0}^{4\varepsilon n^{3/4}-1} (1 - j \cdot p^{n/2}) < p^{16\varepsilon^2 n^{3/2}} e^{2/(p-1)}.$$

Because there are only $p^{n/2}$ different ways to assign values to the bit positions in the intersection of Δ_1 and the ith column of the input matrix, for any given i, the set

$$\{S : \text{there exists } f_1 \text{ such that } \#(S \cap V_i) > p^{4\varepsilon n^{3/4}}\}$$

has size less than $(p^{n/2}\, p^{-16\varepsilon^2 n^{3/2}} e^{2/(p-1)}) \times \textit{Total}$. Therefore the total number of nonconcise subspaces is less than a $\delta' = (n/2)p^{n/2}p^{-16\varepsilon^2 n^{3/2}} e^{2/(p-1)}$ fraction of *Total*. When $p \geq 8$ and $n \geq 128$, we have $\delta' < \delta$. □

Next, we prove that almost all concise subspaces are encodable under a small collection of fixing functions.

Lemma 3.15. *Given any Δ, there is a set of $p^{\varepsilon n^2}$ fixing functions, F_1, restricted to Δ_1 such that all but a p^{-p^n} fraction of concise subspaces are encodable under some $f_1 \in F_1$.*

Proof. Given a group of m concise subspaces, we count the number of pairs (f_1, X) such that f_1 is a fixing function restricted to Δ_1, and X is a concise subspace in this group encodable under f_1. By Proposition 3.13, there are at least $m \cdot p^{n^2/4 - 2\varepsilon n^{7/4}} \cdot e^{-2/(p-1)}$ such pairs. Since there are $p^{n^2/4}$ different restricted fixing functions, one of them must be included in at least $m \cdot p^{-2\varepsilon n^{7/4}} \cdot e^{-2/(p-1)}$ pairs. This means we can always find a restricted fixing function f_1 such that a $p^{-2\varepsilon n^{7/4}} \cdot e^{-2/(p-1)}$ fraction of subspaces in this group is encodable under f_1.

Starting with all the concise subspaces, we can find one restricted fixing function f_1 such that a $p^{-2\varepsilon n^{7/4}} \cdot e^{-2/(p-1)}$ fraction of the concise subspaces are encodable under f_1. Then we can find another fixing function f_1' such that a $p^{-2\varepsilon n^{7/4}} \cdot e^{-2/(p-1)}$ fraction of all the remaining concise subspaces are encodable under f_1'.

We can repeat this process so that after the ith step, a $1 - (1 - p^{-2\varepsilon n^{7/4}} \cdot e^{-2/(p-1)})^i$ fraction of all the concise subspaces are encodable under some of the i restricted fixing functions we have found. Carrying out this process $p^{\varepsilon n^2}$ steps gives us F_1.

The fraction of concise subspaces which are not encodable under any restricted fixing function in F_1 is at most $(1 - p^{-2\varepsilon n^{7/4}} \cdot e^{-2/(p-1)})^{p^{\varepsilon n^2}}$. This quantity is smaller than p^{-p^n}. □

Lemmas 3.14 and 3.15 establish that there is a set of $p^{\varepsilon n^2}$ restricted fixing functions F_1 such that almost all dimension $n/2$ subspaces in $\boldsymbol{F}_p^n$ are encodable under some $f_1 \in F_1$. The exceptions are less than a $p^{-\varepsilon^2 n^{3/2}} + p^{-p^n}$ fraction of all such subspaces.

We repeat this procedure for the last $n/2$ columns of the input matrix. This means we have $p^{\varepsilon n^2}$ restricted fixing functions for Δ_2 such that every dimension $n/2$ subspace in $\boldsymbol{F}_p^n$ is encodable under some of these functions, with the exception of less than a $p^{-\varepsilon^2 n^{3/2}} + p^{-p^n}$ fraction of all such subspaces.

Therefore given a nice partition π', we have a set of $p^{2\varepsilon n^2}$ fixing functions such that almost all pairs of dimension $n/2$ subspaces (X, Y) are encodable under some of these functions. Since $p^{-\varepsilon^2 n^{3/2}} + p^{-p^n} \leq 2p^{-\varepsilon^2 n^{3/2}}$, the number of exceptions is less than $4p^{-\varepsilon^2 n^{3/2}} \times (Total)^2$, where

$$Total = \prod_{j=0}^{n/2-1} \frac{p^n - p^j}{p^{n/2} - p^j}$$

is the total number of dimension $n/2$ subspace in $\boldsymbol{F}_p^n$. (Note that $Total < p^{n^2/4} e^{2/(p-1)}$, see the inequality before Proposition 3.13.)

Each of the subspace pairs corresponds to an entry in the truth matrix T. Also, there are at least $p^{n^2/2}$ "one" entries in T by Lemma 3.1. So, there must be one fixing function which defines a submatrix, T', of T that contains at least $p^{(1/2-2\varepsilon)n^2} - p^{(1/2-2\varepsilon)n^2 - \varepsilon^2 n^{3/2}} e^{O(1)}$ "one" entries. This result, together with Lemma 3.7, establishes the following lemma by an application of Yao's lower-bound method.

Lemma 3.16. *Under any nice partition π', the singularity test over the finite field $\boldsymbol{F}_p$ requires at least $(n^2 \log p)/8 - n \log p - O(1)$ bits of communication.*

Now we need to show that $\Theta(n^2 \log p)$ bits of communication are required for any proper partition π.

Definition 3.17. (Continuing from Definition 3.9). For $1 \leq i \leq n$, let δ_i be the set of input bit positions in Δ which belong to the ith column of the input matrix.

In case $\#\delta_i \leq (n \log p)/2$ for all i, we can proceed as in the nice partition case by picking $(n \log p)/2$ bit positions to be fixed in every column, which include all positions in δ_i, $1 \leq i \leq n$. The $\Theta(n^2 \log p)$ bound follows directly.

Now assume that $\#\delta_i > (n \log p)/2$ for some i. Then either $\#\delta_{n/2} > (n \log p)/2$ or $\#\delta_{n/2+1} > (n \log p)/2$, exclusively (Definition 3.8). We need to address only one of these cases, say $\#\delta_{n/2} > (n \log p)/2$.

Let $m \leq n/2$ be the index such that $\#\delta_m > (n \log p)/2$ and $\#\delta_{m-1} \leq (n \log p)/2$. For each $m \leq i \leq n/2$, define d_i to be $(\#\delta_i) - (n \log p)/2$. Pick arbitrary d_i bit positions from each δ_i, $m \leq i \leq n/2$, and call them *characteristic* positions of π. Let $D = \sum_{i=m}^{n/2} d_i$ denote the total number of these characteristic positions.

The proper partition π can be reduced to a nice partition π' by reassigning all the characteristics positions to the first agent and compensate the second agent with positions in column 1 through column $(m-1)$ which are originally assigned to the first agent and which are to be fixed. Therefore, the amount of communication required under π can be less than that required under π' by at most D bits.

Lemma 3.18. *The number of characteristic positions, D of any proper partition is less than* $(\frac{1}{16})n^2 \log p$.

Proof. By Definition 3.8, $d_i \leq d_{i+1}$ for $m \leq i < n/2$. Let $l = n/2 - m + 1$. Then $D \leq l \cdot d_{n/2}$. The second agent receives at least $\#\delta_{n/2}$ positions from each column in the right half of the input matrix. The total number of bit positions assigned to the second agent is at least $\sum_{i=m}^{n/2} \delta_i + (n/2)(\#\delta_{n/2})$. This value cannot exceed $(n^2 \log p)/2$. So we have the following inequality:

$$\frac{nl \log p}{2} + D + \frac{n^2 \log p}{4} + \frac{n}{2} d_{n/2} \leq \frac{n^2 \log p}{2}. \tag{3.4}$$

Solve for $d_{n/2}$ in (3.4) and we get

$$d_{n/2} \leq \frac{n \log p}{2} - l \log p - \frac{2D}{n}. \tag{3.5}$$

Since $D \leq l \cdot d_{n/2}$, multiplying both sides of the inequality (4.5) with l gives us

$$D \leq \frac{l(n-2l)}{(2l+n)2n} n^2 \log p. \tag{3.6}$$

Assume $l = cn$. Solve (3.6) for the maximum of D. We get $D = (\frac{3}{4} - \sqrt{2}/2)n^2 \log p$, when $c = (-1 + \sqrt{2})/2$. This maximum value is less than $(\frac{1}{16})n^2 \log p$. □

Therefore, by Lemmas 3.16 and 3.18, the communication complexity of the singularity test over the finite field $\boldsymbol{F}_p$ is at least $(n^2 \log p)/16 - n \log p - O(1)$, which is in $\Theta(n^2 \log p)$.

References

[1] R. C. Agarwal, C. S. Burrus, Fast convolution using Fermat number transforms with application to digital filtering, *IEEE Trans. Acoust. Speech Signal Process.*, **22** (1974), 87–97.

[2] R. P. Brent, H. T. Kung, The chip complexity of binary arithmetic, *J. Assoc. Comput. Mech.*, **28** (1981), 521–534.

[3] B. Chor, O. Goldreich, Unbiased bits from sources of weak randomness and probabilistic communication complexity, *Proc. 26th Annual IEEE Symp. on Foundations of Computer Science*, 1985, pp. 429–442.

[4] J. I. Chu, G. Schnitger, The communication complexity of several problems in matrix computation, *J. Complexity*, **7** (1991), 395–407.
[5] L. Lovász, M. Saks, Lattices, möbius functions and communication complexity, *Proc. 29th Annual IEEE Symp. on Foundations of Computer Science*, 1988, pp. 81-90.
[6] A. V. Oppenheim, C. Weinstein, Effects of finite register length in digital filtering and the fast Fourier transform, *Proc. IEEE*, **60** (1972), 957–976.
[7] J. M. Pollard, The fast Fourier transform in a finite field, *Math. Comp.*, **25** (1971), 365–374.
[8] N. S. Szabo, R. I. Tanaka, *Residue Arithmetic and Its Applications to Computer Technology*, McGraw-Hill, New York, 1967.
[9] C. D. Thompson, Area time complexity for VLSI, *Proc. 11th Annual ACM Symp. on Theory of Computing*, 1979, pp. 81–88.
[10] J. Vuillemin, A combinatorial limit to the computing power of VLSI circuits, *IEEE Trans. Comput.*, **32** (1983), 294–300.
[11] N. M. Wigley, G. A. Jullien, On modulus replication for residue arithmetic computations of complex inner products, *IEEE Trans. Comput.*, **39** (1990), 1065–1076.
[12] A. C. Yao, Some complexity questions related to distributive computing, *Proc. 11th Annual ACM Symp. on Theory of Computing*, 1979, pp. 209–213.
[13] A. C. Yao, The entropic limitations of VLSI computation, *Proc. 13th Annual ACM Symp. on Theory of Computing*, 1981, pp. 308–311.

Received October 24, 1991, and in final form October 26, 1992.