The Communication Complexity of Approximating Matrix Rank

Alexander A. Sherstov
Computer Science Department
University of California, Los Angeles
Los Angeles, CA, USA
Email: sherstov@cs.ucla.edu

Andrey A. Storozhenko
Computer Science Department
University of California, Los Angeles
Los Angeles, CA, USA
Email: storozhenko@cs.ucla.edu

***Abstract*—We fully determine the communication complexity of approximating matrix rank, over any finite field $\mathbb{F}$. We study the most general version of this problem, where $0 \leqslant r < R \leqslant n$ are given integers and Alice and Bob need to determine whether their respective matrices $A, B \in \mathbb{F}^{n\times n}$ satisfy $\mathrm{rk}(A+B) = r$ versus $\mathrm{rk}(A+B) = R$. We show that this problem has communication cost $\Omega(r^2 \log |\mathbb{F}|)$, which is optimal. Our lower bound holds even for quantum protocols and even for error probability $\frac{1}{2}(1 - |\mathbb{F}|^{-r/3})$, which too is optimal because this problem has a two-bit classical protocol with error $\frac{1}{2}(1 - |\mathbb{F}|^{-\Theta(r)})$. Prior to our work, lower bounds were known only for *constant-error* protocols and only for *consecutive* integers r and R, with no implication for the approximation of matrix rank. We also settle an analogous question for subspaces, where Alice has a subspace S, Bob has a subspace T, and they need to approximate the dimension of the subspace $S + T$ generated by S and T (equivalently, approximate the dimension of $S \cap T$). As an application, we obtain an $\Omega(n^2 \log |\mathbb{F}|)/k$ memory lower bound for any streaming algorithm with k passes that approximates the rank of an input matrix $M \in \mathbb{F}^{n\times n}$ within a factor of $\sqrt{2} - \delta$, for any $\delta > 0$. Our result is an exponential improvement in k over previous work.*

***Index Terms*—Approximation of matrix rank, communication complexity, quantum computation, subspace intersection problem, subspace sum problem**

I. Introduction

The exact and approximate computation of matrix rank is a fundamental problem in theoretical computer science, studied for its intrinsic importance as well as its connections to other algorithmic and complexity-theoretic questions. In particular, a large body of research has focused on the communication complexity of the matrix rank problem in Yao's two-party model [1], [2], with both classical and quantum communication. In this problem, the two parties Alice and Bob receive matrices $A, B \in \mathbb{F}^{n\times n}$, respectively, over a finite field $\mathbb{F}$ and are tasked with determining the rank of $A + B$ using minimal communication. The first result in this line of research was obtained three decades ago by Chu and Schnitger [3], who proved a lower bound of $\Omega(kn^2)$ for the deterministic communication complexity of computing the rank of $A + B$ when the matrix entries are k-bit integers. Several years later, Chu and Schnitger [4] further showed that this communication problem has deterministic complexity $\Omega(n^2 \log p)$ when the matrix entries are in $\mathbb{F}_p$, the finite field with p elements. The first result on the *randomized* communication complexity of the matrix rank problem was obtained by Sun and Wang [5], who proved that determining whether $A + B$ is singular requires $\Omega(n^2 \log p)$ bits of communication for matrices A, B over the finite field $\mathbb{F}_p$ for prime p. In a follow-up paper, Li, Sun, Wang, and Woodruff [6] showed that this $\Omega(n^2 \log p)$ lower bound holds even for a promise version of the matrix rank problem, where the matrix $A + B$ is guaranteed to have rank either $n - 1$ or n. The lower bounds of [5], [6] further apply to quantum communication.

Despite these exciting developments, no progress has been made on lower bounds for *approximating* matrix rank. Our main contribution is the complete resolution of the approximate matrix rank problem. In what follows, we state our results for matrix rank and several other approximation problems, and present applications of our work to streaming complexity.

A. Matrix rank problem

We study the problem of approximating matrix rank in its most general form. Specifically, let $\mathbb{F}$ be any finite field. For integer parameters n, m, R, r such that $\min\{n, m\} \geqslant R > r \geqslant 0$, we consider the promise communication problem defined on pairs of matrices $A, B \in \mathbb{F}^{n\times m}$ by

$$\mathrm{RANK}_{r,R}^{\mathbb{F},n,m}(A, B) = \begin{cases} -1 & \text{if } \mathrm{rk}(A+B) = r, \\ 1 & \text{if } \mathrm{rk}(A+B) = R, \\ * & \text{otherwise,} \end{cases}$$

where the asterisk indicates that the communication protocol may exhibit arbitrary behavior when $\mathrm{rk}(A + B) \notin \{r, R\}$. In words, the problem amounts to distinguishing input pairs with $\mathrm{rk}(A + B) = r$ from those with $\mathrm{rk}(A + B) = R$. The corresponding *total* communication problem is given by

$$\mathrm{RANK}_{r}^{\mathbb{F},n,m}(A, B) = \begin{cases} -1 & \text{if } \mathrm{rk}(A+B) \leqslant r, \\ 1 & \text{otherwise.} \end{cases}$$

Clearly, the total problem $\mathrm{RANK}_{r}^{\mathbb{F},n,m}$ is more challenging than the promise problem $\mathrm{RANK}_{r,R}^{\mathbb{F},n,m}$. Prior to our work, the strongest known result was the $\Omega(n^2 \log p)$ lower bound of [6] on the bounded-error quantum communication complexity of $\mathrm{RANK}_{n-1,n}^{\mathbb{F}_p,n,n}$ for fields $\mathbb{F}_p$ of prime order. Unfortunately, this lower bound has no implications for the approximation of

Work supported by NSF grant CCF-2220232.

matrix rank because the ratio $(n-1)/n$ rapidly tends to 1. We resolve this question in full in the following theorem.

Theorem I.1 (Lower bound for rank problem)**.** *There is an absolute constant $c > 0$ such that for all finite fields $\mathbb{F}$ and all integers n, m, R, r with $\min\{n, m\} \geqslant R > r \geqslant 0$,*

$$Q^*_{\frac{1}{2} - \frac{1}{4|\mathbb{F}|^{r/3}}}(\mathrm{RANK}^{\mathbb{F},n,m}_{r,R}) \geqslant c(1 + r^2 \log|\mathbb{F}|).$$

In particular,

$$Q^*_{1/4}(\mathrm{RANK}^{\mathbb{F},n,m}_{r,R}) \geqslant c(1 + r^2 \log|\mathbb{F}|).$$

In the statement above, Q^*_ε denotes ε-error quantum communication complexity with arbitrary prior entanglement, which is the most powerful model of probabilistic computation. Clearly, all our lower bounds apply to the randomized (classical) model as well. Two other remarks are in order. Even in the special case of $r = n-1$ and $R = n$, our result is a significant improvement on previous work because our theorem is proved in the *large-error regime,* with the error probability exponentially close to $1/2$. This should be contrasted with the communication lower bounds of [5], [6], which were proved for error probability $1/3$. Moreover, Theorem I.1 is the first result of its kind because it allows for an arbitrary gap between r and R. In particular, Theorem I.1 shows for the first time that approximating the matrix rank to any constant factor requires $\Omega(n^2 \log|\mathbb{F}|)$ bits of communication, even for protocols that succeed with exponentially small probability (take $R = n$ and $r = cn$ for a small constant $c > 0$).

Theorem I.1 is optimal in a very strong sense. Specifically, we have the following matching upper bound, which we prove by adapting Clarkson and Woodruff's streaming algorithm for matrix rank [7]. In the statement below, R_ε denotes randomized (classical) communication complexity with error ε.

Theorem I.2 (Upper bound for rank problem)**.** *There is an absolute constant $c > 0$ such that for all finite fields $\mathbb{F}$ and all integers n, m, r with $\min\{n, m\} > r \geqslant 0$,*

$$R_{1/3}(\mathrm{RANK}^{\mathbb{F},n,m}_{r}) \leqslant c(1 + r^2 \log|\mathbb{F}|),$$
$$R_{\frac{1}{2} - \frac{1}{32|\mathbb{F}|^r}}(\mathrm{RANK}^{\mathbb{F},n,m}_{r}) \leqslant 2.$$

This result shows that the lower bound of Theorem I.1 is tight not only for quantum protocols solving the partial problem $\mathrm{RANK}^{\mathbb{F},n,m}_{r,R}$ but even for *classical*, *bounded-error* protocols solving the *total* problem $\mathrm{RANK}^{\mathbb{F},n,m}_{r}$. Moreover, Theorem I.2 shows that the error regime for which we prove our lower bound in Theorem I.1 is also optimal, in that the total rank problem has a classical protocol with cost only 2 bits and error probability $\frac{1}{2} - |\mathbb{F}|^{-\Theta(r)}$.

B. Streaming complexity

The streaming complexity of matrix rank has received extensive attention in the literature [5]–[11]. In this model, an algorithm with limited memory is presented with a matrix M of order n over a given field, in row-major order. The objective is to compute or approximate the rank of M, using either a single pass or multiple passes over M. Via standard reductions, our Theorem I.1 implies an essentially optimal lower bound on the streaming complexity of approximating matrix rank. Unlike previous work, our result remains valid even for polynomially many passes and even for correctness probability exponentially close to $1/2$. Stated in its most general form, our result is as follows.

Theorem I.3. *Let n, r, R be nonnegative integers with $n/2 \leqslant r < R \leqslant n$, and let $\mathbb{F}$ be a finite field. Define $f\colon \mathbb{F}^{n\times n} \to \{-1, 1, *\}$ by*

$$f(M) = \begin{cases} -1 & \text{if } \operatorname{rk} M = r, \\ 1 & \text{if } \operatorname{rk} M = R, \\ * & \text{otherwise.} \end{cases}$$

Let $\mathscr{A}$ be any randomized streaming algorithm for f with error probability $\frac{1}{2} - \frac{1}{4}|\mathbb{F}|^{-(r-\lceil n/2\rceil)/3}$ that uses s bits of memory and k passes. Then

$$sk = \Omega\left(\left(r - \left\lceil \frac{n}{2} \right\rceil\right)^2 \log|\mathbb{F}|\right).$$

By way of notation, recall that f in the above statement is a *partial* function, and the algorithm is allowed to exhibit arbitrary behavior on matrices M where $f(M) = *$.

Corollary I.4. *Fix arbitrary constants $\varepsilon \in (0, 2]$ and $\delta \in (1/2, 1)$. Let $\mathbb{F}$ be a finite field. Then no randomized $o(n^\varepsilon)$-pass streaming algorithm with $n^{2-\varepsilon}\log|\mathbb{F}|$ bits of memory, on input a matrix $M \in \mathbb{F}^{n\times n}$, can distinguish between the cases $\operatorname{rk} M = n$ and $\operatorname{rk} M = \lfloor \delta n \rfloor$ with probability of correctness greater than $\frac{1}{2}(1 + |\mathbb{F}|^{-\Theta(n)})$.*

Proof. Take $R = n$ and $r = \lfloor \delta n \rfloor$ in Theorem I.3, for any n larger than a certain constant. □

The memory lower bound in Corollary I.4 is essentially optimal since the rank of a matrix $M \in \mathbb{F}^{n\times n}$ can be computed exactly by a trivial, single-pass algorithm with memory $O(n^2 \log|\mathbb{F}|)$. Prior to our work, the strongest streaming lower bound for approximating matrix rank was due to Chen et al. [11]. For all constants $\varepsilon > 0$ and $\delta > 0$, they proved that no $o(\sqrt{\log n})$-pass algorithm with space $n^{2-\varepsilon}$ can distinguish between the cases $\operatorname{rk} M = n$ and $\operatorname{rk} M \leqslant \delta n$ with probability $2/3$, where M is an input matrix of order n over a finite field of size $\omega(n)$. Our Corollary I.4 is an exponential improvement on [11] in the number of passes. Moreover, Corollary I.4 is valid for all finite fields regardless of size, and holds even when the correctness probability is exponentially close to $1/2$.

C. Determinant problem

Recall that a square matrix over a field $\mathbb{F}$ has full rank if and only if its determinant is nonzero. As a result, the problem of computing the determinant has received considerable attention in previous work on matrix rank, e.g., [4]–[6]. We are interested in the most general form of the determinant problem, where Alice and Bob receive as input matrices $A, B \in \mathbb{F}^{n\times n}$, respectively, and need to determine whether the determinant of $A + B$ equals a or b. The problem parameters a and b are

distinct field elements that are fixed in advance. Formally, the determinant problem is the partial communication problem on matrix pairs (A, B) given by

$$\mathrm{DET}_{a,b}^{\mathbb{F},n}(A,B) = \begin{cases} -1 & \text{if } \det(A+B) = a, \\ 1 & \text{if } \det(A+B) = b, \\ * & \text{otherwise.} \end{cases}$$

Prior to our work, the strongest result on the determinant problem was due to Sun and Wang [5], who proved a tight lower bound of $\Omega(n^2 \log |\mathbb{F}|)$ for the randomized and quantum communication complexity of $\mathrm{DET}_{a,b}^{\mathbb{F},n}$ for nonzero a, b over any finite field $\mathbb{F}$ of prime order. They conjectured the same lower bound for the case of arbitrary a, b. To see why the case of nonzero a, b is rather special, observe that the number of matrices with determinant a is always the same as the number of matrices with determinant b, with natural bijections between these two sets; but this is no longer true if one of a, b is zero. This asymmetry suggests that the determinant problem requires a substantially different approach when one of a, b is zero. In this work, we develop sufficiently strong techniques to solve the determinant problem in full, thereby settling Sun and Wang's conjecture in the affirmative.

Theorem I.5. *There is an absolute constant $c > 0$ such that for every finite field $\mathbb{F}$, every pair of distinct elements $a, b \in \mathbb{F}$, and all integers $n \geqslant 2$,*

$$Q^*_{\frac{1}{2} - \frac{1}{4|\mathbb{F}|^{(n-1)/3}}}(\mathrm{DET}_{a,b}^{\mathbb{F},n}) \geqslant cn^2 \log |\mathbb{F}|.$$

The communication lower bound of Theorem I.5 is best possible, up to the multiplicative constant c. It matches the trivial, deterministic protocol where Alice sends her input matrix A to Bob using $n^2 \lceil \log |\mathbb{F}| \rceil$ bits, at which point Bob computes $\det(A + B)$ and announces the output of the protocol. Furthermore, the error regime in Theorem I.5 is also essentially optimal because, for example, the problem $\mathrm{DET}_{0,b}^{\mathbb{F},n}$ has a randomized protocol with only 2 bits of communication and error probability $\frac{1}{2} - \Theta(|\mathbb{F}|^{n-1})$, by taking $r = n - 1$ and $R = m = n$ in Theorem I.2. Lastly, we note that the requirement that $n \geqslant 2$ in Theorem I.5 is also necessary because the determinant problem for 1×1 matrices reduces to the equality problem with domain $\mathbb{F} \times \mathbb{F}$ and therefore has randomized communication complexity $O(1)$.

We prove Theorem I.5 for all a, b from first principles, without relying on the work of Sun and Wang [5]. In the case of nonzero a, b, we give a new proof that is quite short and uses only basic Fourier analysis, unlike the rather technical proof of [5]. To settle the complementary case where one of a, b is zero, we prove a stronger result of independent interest. Here, we introduce a natural problem that we call $\mathrm{RANKDET}_{r,a}^{\mathbb{F},n}$, which combines features of the matrix rank and determinant problems. It is parameterized by a nonzero field element $a \in \mathbb{F}$ and a nonnegative integer $r < n$, and Alice and Bob's objective is to distinguish input pairs (A, B) with $\mathrm{rk}(A + B) = r$ from those with $\det(A + B) = a$. We prove the following.

Theorem I.6. *There is an absolute constant $c > 0$ such that for every finite field $\mathbb{F}$, every field element $a \in \mathbb{F} \setminus \{0\}$, and all integers $n > r \geqslant 0$,*

$$Q^*_{\frac{1}{2} - \frac{1}{4|\mathbb{F}|^{r/3}}}(\mathrm{RANKDET}_{r,a}^{\mathbb{F},n}) \geqslant c(1 + r^2 \log |\mathbb{F}|).$$

Taking $r = n - 1$ in this result settles Theorem I.5 when one of a, b is zero, as desired. Theorem I.6 is optimal in a strong sense: even the *total* problem $\mathrm{RANK}_r^{\mathbb{F},n,n}$, which subsumes $\mathrm{RANKDET}_{r,a}^{\mathbb{F},n}$, has bounded-error classical communication complexity $O(1 + r^2 \log |\mathbb{F}|)$ by Theorem I.2. Theorem I.6 for the $\mathrm{RANKDET}_{r,a}^{\mathbb{F},n}$ problem significantly strengthens our main result, Theorem I.1, for the matrix rank problem $\mathrm{RANK}_{r,n}^{\mathbb{F},n,n}$ (in the former problem, Alice and Bob distinguish rank r from determinant $a \neq 0$; in the latter problem, they distinguish rank r from rank n).

D. Subspace sum and intersection problems

There are two natural ways to recast the computation of matrix rank as a communication problem. One approach, discussed in detail above, is to assign matrices A and B to Alice and Bob, respectively, and require them to compute the rank of $A + B$. Alternatively, one can require Alice and Bob to compute the rank of the matrix $\begin{bmatrix} A & B \end{bmatrix}$. This alternative approach is best described in the language of linear subspaces: letting S and T stand for the column space of A and B, respectively, the rank of $\begin{bmatrix} A & B \end{bmatrix}$ is precisely the dimension of the linear subspace $S + T$ generated by S and T. Here, we may assume that the dimensions of S and T are known in advance because this information can be communicated at negligible cost.

In this way, one arrives at the *subspace sum problem* over a finite field $\mathbb{F}$, where Alice receives as input an m-dimensional linear subspace $S \subseteq \mathbb{F}^n$ and Bob receives an ℓ-dimensional linear subspace $T \subseteq \mathbb{F}^n$. The integers m and ℓ are part of the problem specification and are fixed in advance. In the promise version of the subspace sum problem, the objective is to distinguish subspace pairs with $\dim(S + T) = d_1$ from those with $\dim(S + T) = d_2$, for distinct integers d_1, d_2 fixed in advance. This corresponds to the partial function given by

$$\mathrm{SUM}_{d_1,d_2}^{\mathbb{F},n,m,\ell}(S,T) = \begin{cases} -1 & \text{if } \dim(S+T) = d_1, \\ 1 & \text{if } \dim(S+T) = d_2, \\ * & \text{otherwise.} \end{cases}$$

The corresponding total communication problem is that of determining whether $S + T$ has dimension at most d, for an integer d fixed in advance:

$$\mathrm{SUM}_{d}^{\mathbb{F},n,m,\ell}(S,T) = \begin{cases} -1 & \text{if } \dim(S+T) \leqslant d, \\ 1 & \text{otherwise.} \end{cases}$$

The total problem is more challenging than the promise problem in that $\mathrm{SUM}_{d_1,d_2}^{\mathbb{F},n,m,\ell}$ is a restriction of $\mathrm{SUM}_{d_1}^{\mathbb{F},n,m,\ell}$, for any integers $d_1 < d_2$. As noted by many authors, from the standpoint of communication complexity, computing the dimension of the subspace sum $S + T$ is equivalent to computing the dimension of the subspace intersection $S \cap T$.

This equivalence follows from the identity $\dim(S+T) = \dim(S) + \dim(T) - \dim(S \cap T)$.

Despite the syntactic similarity between the matrix sum $A + B$ and the corresponding subspace sum $S + T$, the subspace sum problem appears to be significantly more subtle and technical. Previous work has focused on a special case that we call *subspace disjointness* (determining whether Alice and Bob's subspaces have trivial intersection, $\{0\}$) and the dual problem that we call *vector space span* (determining if the sum of Alice and Bob's subspaces is the entire vector space). These two problems were studied in [4], [12], with an optimal lower bound of $\Omega(n^2 \log p)$ on their deterministic communication complexity over a field with p elements. Sun and Wang [5] showed that the $\Omega(n^2 \log p)$ lower bound for subspace disjointness remains valid even for randomized and quantum communication. In follow-up work, Li, Sun, Wang, and Woodruff [6] proved an $\Omega(n^2 \log p)$ quantum lower bound for a promise version of subspace disjointness, where Alice and Bob's inputs are $n/2$-dimensional subspaces that either have trivial intersection or intersect in a one-dimensional subspace. The authors of [13] considered an asymmetric problem where Alice receives an n-bit vector, Bob receives a subspace, and their objective is to determine whether Alice's vector is contained in Bob's subspace. They showed that in any randomized one-way protocol that solves this problem, either Alice sends $\Omega(n)$ bits, or Bob sends $\Omega(n^2)$ bits.

In summary, all previous lower bounds for two-way communication complexity have focused on subspace disjointness or vector space span. The general problem, where Alice and Bob need to distinguish between the cases $\dim(S+T) = d_1$ and $\dim(S+T) = d_2$, is substantially harder and has remained unsolved. The difficulty is that previous results [5], [6] are based on a reduction from the matrix rank problem to subspace disjointness, and this straightforward strategy does not produce optimal results for the subspace sum problem with arbitrary parameters. In this paper, we approach the subspace sum problem from first principles and solve it completely. Our solution settles both the promise version of subspace sum and the corresponding total version. For clarity, we first state our result in the regime of constant error.

Theorem I.7. *Let $\mathbb{F}$ be a finite field with $q = |\mathbb{F}|$ elements, and let n, m, ℓ, d, D be nonnegative integers with $\max\{m, \ell\} \leqslant d < D \leqslant \min\{m + \ell, n\}$. If $m = \ell = d$, then*

$$R_{1/3}(\mathrm{SUM}_d^{\mathbb{F},n,m,\ell}) = O(1).$$

If m, ℓ, d are not all equal, then

$$Q^*_{1/3}(\mathrm{SUM}_{d,D}^{\mathbb{F},n,m,\ell}) = \Theta((d-m+1)(d-\ell+1)\log q),$$
$$R_{1/3}(\mathrm{SUM}_d^{\mathbb{F},n,m,\ell}) = \Theta((d-m+1)(d-\ell+1)\log q).$$

Several remarks are in order. Recall that in $\mathbb{F}^n$, the sum of an m-dimensional subspace and an ℓ-dimensional subspace has dimension between $\max\{m, \ell\}$ and $\min\{m+\ell, n\}$. This justifies the above requirement that $d, D \in [\max\{m, \ell\}, \min\{m + \ell, n\}]$. Theorem I.7 shows that the promise version of the subspace sum problem has the same communication complexity as the total version, up to a constant factor. Moreover, the theorem shows that this communication complexity is the same, up to a constant factor, for quantum and classical communication protocols. Both the lower and upper bounds in Theorem I.7 require substantial effort. Lastly, the degenerate case $d = m = \ell$ of the subspace sum problem is easily seen to be equivalent to the equality problem, which explains the $O(1)$ bound in the theorem statement.

In addition to the constant-error regime of Theorem I.7, we are able to determine the communication complexity of subspace sum for essentially all settings of the error parameter, as follows.

Theorem I.8. *Let $\mathbb{F}$ be a finite field with $q = |\mathbb{F}|$ elements, and let n, m, ℓ, d, D be nonnegative integers with $\max\{m, \ell\} \leqslant d < D \leqslant \min\{m + \ell, n\}$. If $m = \ell = d$, then*

$$R_{1/3}(\mathrm{SUM}_d^{\mathbb{F},n,m,\ell}) = O(1).$$

If m, ℓ, d are not all equal, then for all $\gamma \in [\frac{1}{3} q^{-(2d-m-\ell)/5}, \frac{1}{3}]$,

$$Q^*_{\frac{1-\gamma}{2}}(\mathrm{SUM}_{d,D}^{\mathbb{F},n,m,\ell}) = \Theta((\log_q \lceil q^{d-m}\gamma \rceil + 1)(\log_q \lceil q^{d-\ell}\gamma \rceil + 1)\log q),$$
$$R_{\frac{1-\gamma}{2}}(\mathrm{SUM}_d^{\mathbb{F},n,m,\ell}) = \Theta((\log_q \lceil q^{d-m}\gamma \rceil + 1)(\log_q \lceil q^{d-\ell}\gamma \rceil + 1)\log q)$$

and moreover

$$R_{\frac{1}{2} - \frac{1}{16q^{2d-m-\ell+16}}}(\mathrm{SUM}_d^{\mathbb{F},n,m,\ell}) \leqslant 2. \tag{1}$$

Theorem I.8 determines the communication complexity of subspace sum for every error probability in $[\frac{1}{3}, \frac{1}{2} - \Theta(|\mathbb{F}|^{-(2d-m-\ell)/5})]$. This is essentially the complete range of interest because by (1), the communication cost drops to 2 bits when the error probability is set to $\frac{1}{2} - |\mathbb{F}|^{-(2d-m-\ell)-\Theta(1)}$. Analogous to the constant-error regime, Theorem I.8 shows that the communication complexity of subspace sum for any error in $[\frac{1}{3}, \frac{1}{2} - \Theta(|\mathbb{F}|^{-(2d-m-\ell)/5})]$ is the same, up to a constant factor, for both the partial and total versions of the problem, and for both quantum and classical communication. Theorems I.7 and I.8 reveal a rather subtle dependence of the communication complexity on the problem parameters d, m, ℓ, particularly as one additionally varies the error parameter. This explains why we were not able to obtain these theorems via a reduction from the matrix rank problem, as was done in previous work [5], [6] in the special case of subspace disjointness.

In view of the aforementioned identity $\dim(S+T) = \dim(S) + \dim(T) - \dim(S \cap T)$, our results for subspace sum can be equivalently stated in terms of subspace intersection. Formally, the *subspace intersection problem* requires Alice and Bob to distinguish subspace pairs (S, T) with $\dim(S \cap T) = d_1$ from those with $\dim(S \cap T) = d_2$, where S is an m-dimensional subspace given as input to Alice, T is an ℓ-dimensional subspace given to Bob, and d_1, d_2 are distinct

integers fixed in advance. This corresponds to the partial function

$$\mathrm{INTERSECT}^{\mathbb{F},n,m,\ell}_{d_1,d_2}(S,T) = \begin{cases} -1 & \text{if } \dim(S\cap T) = d_1, \\ 1 & \text{if } \dim(S\cap T) = d_2, \\ * & \text{otherwise.} \end{cases}$$

The total version of the subspace intersection problem is given by

$$\mathrm{INTERSECT}^{\mathbb{F},n,m,\ell}_{d}(S,T) = \begin{cases} -1 & \text{if } \dim(S\cap T) \geqslant d, \\ 1 & \text{otherwise,} \end{cases}$$

where d is a problem parameter fixed in advance. Theorem I.8 fully settles the complexity of the subspace intersection problem, as follows.

Theorem I.9. *Let $\mathbb{F}$ be a finite field with $q = |\mathbb{F}|$ elements, and let n, m, ℓ, r, R be nonnegative integers with $\max\{0, m+\ell-n\} \leqslant r < R \leqslant \min\{m,\ell\}$. If $m = \ell = R$, then*

$$R_{1/3}(\mathrm{INTERSECT}^{\mathbb{F},n,m,\ell}_{R}) = O(1).$$

If m, ℓ, R are not all equal, then for all $\gamma \in [\frac{1}{3}q^{-(m+\ell-2R)/5}, \frac{1}{3}]$,

$$\begin{aligned} &Q^*_{\frac{1-\gamma}{2}}(\mathrm{INTERSECT}^{\mathbb{F},n,m,\ell}_{r,R}) \\ &\qquad = \Theta((\log_q\lceil q^{m-R}\gamma\rceil + 1)(\log_q\lceil q^{\ell-R}\gamma\rceil + 1)\log q), \\ &R_{\frac{1-\gamma}{2}}(\mathrm{INTERSECT}^{\mathbb{F},n,m,\ell}_{R}) \\ &\qquad = \Theta((\log_q\lceil q^{m-R}\gamma\rceil + 1)(\log_q\lceil q^{\ell-R}\gamma\rceil + 1)\log q) \end{aligned}$$

and moreover

$$R_{\frac{1}{2}-\frac{1}{16q^{m+\ell-2R+16}}}(\mathrm{INTERSECT}^{\mathbb{F},n,m,\ell}_{R}) \leqslant 2.$$

A moment's reflection (see Proposition II.25) shows that in $\mathbb{F}^n$, the intersection of an m-dimensional subspace and an ℓ-dimensional subspace is a subspace of dimension between $\max\{0, m+\ell-n\}$ and $\min\{m,\ell\}$, hence the requirement that $r, R \in [\max\{0, m+\ell-n\}, \min\{m,\ell\}]$. Remarks analogous to those for subspace sum apply to Theorem I.9 as well. Specifically, Theorem I.9 determines the ε-error communication complexity of subspace intersection for all $\varepsilon \in [\frac{1}{3}, \frac{1}{2} - \Theta(|\mathbb{F}|^{-(m+\ell-2R)/5})]$, which is essentially the complete range of interest because the communication cost drops to 2 bits when the error probability is set to $\frac{1}{2} - |\mathbb{F}|^{-(m+\ell-2R)-\Theta(1)}$. Also, Theorem I.9 shows that in this range of interest, the ε-error communication complexity of subspace intersection is the same (up to a constant factor) for both the partial and total versions of the problem, and for both quantum and classical communication.

E. Previous approaches

A powerful tool for proving lower bounds on randomized and quantum communication complexity is the *approximate trace norm* [2], [14]–[17]. In more detail, let $F\colon X \times Y \to \{-1, 1\}$ be a given communication problem, and let $M = [F(x,y)]_{x,y}$ be its characteristic matrix. The *δ-approximate trace norm of M*, denoted $\|M\|_{\Sigma,\delta}$, is the minimum trace norm of a real matrix $\widetilde{M}$ that approximates M entrywise within δ. The approximate trace norm bound states that

$$Q^*_\varepsilon(F) \geqslant \frac{1}{2}\log\left(\frac{\|M\|_{\Sigma,2\varepsilon}}{3\sqrt{|X||Y|}}\right) \tag{2}$$

for all $\varepsilon \geqslant 0$, making it possible to prove communication lower bounds by analyzing the approximate trace norm of M. To bound the approximate trace norm from below, it is useful to appeal to its dual formulation as a maximization problem, whereby

$$\|M\|_{\Sigma,2\varepsilon} \geqslant \frac{\langle M, \Phi\rangle - 2\varepsilon\|\Phi\|_1}{\|\Phi\|} \tag{3}$$

for every nonzero real matrix Φ. As a result, proving a communication lower bound reduces to constructing a matrix Φ whose spectral norm and ℓ_1 norm are small relative to the inner product of Φ with the communication matrix M. The matrix Φ is often referred to as a *dual matrix* or a *witness*. The lower bound (2) remains valid for partial functions $F\colon X \times Y \to \{-1, 1, *\}$ and their characteristic matrices M, in which case the dual characterization of the approximate trace norm is given by

$$\begin{aligned} \|M\|_{\Sigma,2\varepsilon} \geqslant \frac{1}{\|\Phi\|}\Bigg(&\sum_{\operatorname{dom} F} M_{x,y}\Phi_{x,y} \\ &-2\varepsilon\|\Phi\|_1 - \sum_{\overline{\operatorname{dom} F}} |\Phi_{x,y}|\Bigg) \end{aligned} \tag{4}$$

for all $\Phi \neq 0$. In this equation, $\operatorname{dom} F = \{(x,y) : F(x,y) \neq *\}$ denotes the domain of the partial function F. Comparing this dual characterization with the original one (3) for total functions, we notice that the inner product is now restricted to the domain of F, and there is an additional penalty term for any weight placed by Φ outside the domain of F. For more background on the use of duality in proving communication lower bounds, we refer the reader to the surveys [18], [19].

Main idea in [5], [6]*:* Constructing a good witness Φ can be very challenging. Sun and Wang [5] studied the *nonsingularity problem* over fields $\mathbb{F}_p$ of prime order p, where Alice and Bob's inputs are matrices $A, B \in \mathbb{F}_p^{n\times n}$, respectively, and they are required to output -1 if $A+B$ is nonsingular and 1 otherwise. Let M be the characteristic matrix of this communication problem. To analyze the approximate trace norm of M, the authors of [5] use the witness $\Phi = [(-1)^n\hat{g}(A+B)]_{A,B}$, where $\hat{g}$ is the Fourier transform of the function $g\colon \mathbb{F}_p^{n\times n} \to \{0,1\}$ given by $g(X) = 1$ if and only if X is nonsingular. The calculations in [5] reveal the following, where $C \geqslant 6$ is an absolute constant:

(i) $\|\Phi\| = 1$;
(ii) $\langle M, \Phi\rangle = 2p^{n^2-n}\prod_{i=1}^n(p^i - 1)$;
(iii) $\|\Phi\|_1 \leqslant Cp^{n^2-n}\prod_{i=1}^n(p^i - 1)$.

Using this witness Φ in (3) with a sufficiently small error parameter ε, Sun and Wang obtain $\|M\|_{\Sigma,2\varepsilon} = \Omega(p^{n^2}p^{n(n-1)/2})$, which in view of (2) gives an $\Omega(n^2\log p)$ lower bound on the

bounded-error communication complexity of the nonsingularity problem.

In follow-up work, Li, Sun, Wang, and Woodruff [6] studied the partial communication problem $F = \mathrm{RANK}_{n-1,n}^{\mathbb{F}_p,n,n}$. Let M' denote its characteristic matrix. The authors of [6] used the same witness Φ as Sun and Wang [5] and obtained the following refinements:

(i) $\|\Phi\| = 1$;
(ii) $\sum_{\mathrm{dom}\, F} M'_{A,B}\Phi_{A,B} = p^{n^2-n}(1+\frac{p-p^{-n+1}}{p-1})\prod_{i=1}^{n}(p^i-1)$;
(iii) $\|\Phi\|_1 = p^{n^2-n}\prod_{i=0}^{n-1}(1+p^{-i})\cdot\prod_{i=1}^{n}(p^i-1)$;
(iv) $\sum_{\overline{\mathrm{dom}\, F}}|\Phi_{A,B}| \leqslant \|\Phi\|_1 - \sum_{\mathrm{dom}\, F} M'_{A,B}\Phi_{A,B}$.

Making these substitutions in (4) and setting ε to a sufficiently small constant, the authors [6] obtain $\|M'\|_{\Sigma,2\varepsilon} = \Omega(p^{n^2}p^{n(n-1)/2})$, which along with (2) results in an $\Omega(n^2\log p)$ lower bound on the quantum communication complexity of $F = \mathrm{RANK}_{n-1,n}^{\mathbb{F}_p,n,n}$. We note that we have described the work of [5], [6] in the framework that we adopt in our paper, which differs somewhat from the original presentation in [5], [6]. These differences do not affect any of the ideas or bounds in question.

Unfortunately, the above analyses rely heavily on ε being set to a small constant. This is because $\|\Phi\|_1$ is too large compared to the inner product $\langle M, \Phi\rangle$ and the correlation $\sum_{\mathrm{dom}\, F} M'_{A,B}\Phi_{A,B}$, which makes setting ε close to $1/2$ impossible. Since the authors of [6] determined $\|\Phi\|_1$ and $\sum_{\mathrm{dom}\, F} M'_{A,B}\Phi_{A,B}$ exactly, with equality, there is no room for improved analysis and no possibility of setting ε close to $1/2$ with this choice of witness Φ. This rules out the use of Φ for proving Theorem I.1 even in the special case of $\mathrm{RANK}_{n-1,n}^{\mathbb{F},n,n}$. When it comes to the more general problem $\mathrm{RANK}_{r,n}^{\mathbb{F},n,n}$ with $r \leqslant n-2$, the witness Φ does not produce any meaningful results at all, even for small constant ε. The fundamental obstacle is that the ℓ_1 norm of Φ is concentrated on matrix pairs (A, B) for which $A + B$ has rank n or $n-1$, whereas the domain of $\mathrm{RANK}_{r,n}^{\mathbb{F},n,n}$ with $r \leqslant n-2$ consists of matrix pairs whose sum has rank n or r. This makes the contribution of $\sum_{\mathrm{rk}(A+B)=n-1}|\Phi_{A,B}|$ to the summation $\sum_{\overline{\mathrm{dom}\, F}}|\Phi_{A,B}|$ too large, and renders the resulting lower bound worthless. Our attempts at simple modifications to Φ were not successful.

F. Our approach

Our techniques depart substantially from the previous work in [5], [6]. Instead of attempting to guess a good witness Φ and analyzing its metric and analytic properties, we determine how exactly these properties depend on the choice of a witness. In this way, we are able to construct essentially optimal witnesses for the matrix rank, determinant, subspace sum, and subspace intersection problems. We first discuss the matrix rank problem, over an arbitrary finite field $\mathbb{F}$. In this overview, we focus on the canonical case $F = \mathrm{RANK}_{k,n}^{\mathbb{F},n,n}$, where Alice and Bob receive square matrices $A, B \in \mathbb{F}^{n\times n}$, respectively, and need to distinguish between the cases $\mathrm{rk}(A+B) = k$ and $\mathrm{rk}(A+B) = n$. This special case captures the matrix rank problem in its full generality via straightforward reductions.

Reducing the degrees of freedom: We will call a witness Φ *symmetric* if each entry $\Phi_{A,B}$ is fully determined by the rank of $A+B$. In searching for a good witness for the matrix rank problem, we will only consider symmetric witnesses Φ. This restriction is without loss of generality: since $F(A, B)$ depends only on the rank of $A+B$, it is not hard to verify that any witness for F can be "symmetrized" without harming the corresponding value of the approximate trace norm bound, (4). The resulting witness matrix Φ has only $n+1$ degrees of freedom, corresponding to every possible value of the rank of $A+B$.

Let $i \in \{0, 1, \ldots, n\}$ be given. Consider the matrix whose rows and columns are indexed by elements of $\mathbb{F}^{n\times n}$, and whose (A, B) entry is defined to be 1 if $\mathrm{rk}(A+B) = i$ and zero otherwise. Normalize this matrix to have ℓ_1 norm 1, and call the resulting matrix E_i. Then any symmetric witness matrix is a linear combination of $E_0, E_1, \ldots, E_n$. With this in mind, for any real function $\varphi\colon \{0, 1, \ldots, n\} \to \mathbb{R}$, we define

$$E_\varphi = \varphi(0)E_0 + \varphi(1)E_1 + \cdots + \varphi(n)E_n.$$

Taking $\Phi = E_\varphi$ in the approximate trace norm bound (4) and simplifying, we arrive at the following bound for the characteristic matrix M of F:

$$\|M\|_{\Sigma,2\varepsilon} \geqslant \frac{1}{\|E_\varphi\|}\left(\varphi(n) - \varphi(k) - 2\varepsilon\|\varphi\|_1 - \sum_{i\notin\{k,n\}}|\varphi(i)|\right). \tag{5}$$

Our challenge now is to understand how φ affects the spectral norm of E_φ.

By analyzing the singular values of E_φ, we prove that

$$\|E_\varphi\| = q^{-n^2}\max_{s=0,1,\ldots,n}\left|\sum_{t=0}^{n}\varphi(t)\Gamma_n(s,t)\right|, \tag{6}$$

where q is the order of the finite field $\mathbb{F}$, and Γ_n is an auxiliary function. In more detail, we define

$$\Gamma_n(s,t) = \mathop{\mathbf{E}}_{\substack{\mathrm{rk}\, A=s\\ \mathrm{rk}\, B=t}}\omega^{\langle A,B\rangle},$$

where ω is a primitive root of unity of order equal to the characteristic of $\mathbb{F}$, with the operation $x \mapsto \omega^x$ for field elements $x \in \mathbb{F}$ deferred to Section II-D. An exact expression for $\Gamma_n(n, t)$ can be obtained from the analysis of the Fourier spectrum of the nonsingularity function in [5]. Understanding $\Gamma_n(s, t)$ for general s, t, however, is rather nontrivial. To this end, we derive the representation

$$\Gamma_n(s,t) = \sum_{r=0}^{n}P_n(s,t,r)\Gamma_n(n,r),$$

where $P_n(s, t, r)$ is the probability that the upper-left $s \times t$ quadrant of a uniformly random nonsingular matrix of order n has rank r. By explicitly calculating the probabilities

$P_n(s,t,r)$ and combining them with the closed-form expression for $\Gamma_n(n,r)$, we obtain the upper bound $|\Gamma_n(s,t)| \leqslant cq^{-st/2}$ for an absolute constant c. In addition to this *analytic* property, we establish the following *algebraic* result: for n, s fixed, $\Gamma_n(s,t)$ as a function of $t \in \{0,1,\ldots,n\}$ is a polynomial in q^{-t} of degree at most s. These two properties play a central role in our analysis. In what follows, we will refer to a polynomial in q^{-t} as a *hyperpolynomial in* t.

Univariate object for the rank problem: Since (5) is invariant under multiplication of φ by a positive factor, we will normalize φ such that $\varphi(n) = 1$. To achieve a large value on the right-hand side of (5), we will construct a function φ that is negative at k, has ℓ_1 norm concentrated on $\{k, n\}$, and results in E_φ having a small spectral norm. In view of (6), the spectral norm requirement amounts to a bound on $\max_s |\sum_{t=0}^n \varphi(t)\Gamma_n(s,t)|$. Quantitatively speaking, to obtain an asymptotically optimal lower bound for the matrix rank problem, we need φ to satisfy the following constraints:

(i) $\varphi(n) = 1$;
(ii) $\varphi(k) < 0$;
(iii) $\sum_{i \notin \{k,n\}} |\varphi(i)| = q^{-\Omega(k)}$;
(iv) $|\sum_{t=0}^n \varphi(t)\Gamma_n(s,t)| = q^{-\Omega(k^2)}$ for every integer $s \in \{0,1,\ldots,n\}$.

The last requirement states that φ needs to be almost orthogonal to each $\Gamma_n(s,t)$, viewed as a function of t with fixed s. Recall from our earlier discussion that for s and n fixed, $\Gamma_n(s,t)$ is a hyperpolynomial of low degree, namely, a polynomial in q^{-t} of degree at most s. To achieve orthogonality to hyperpolynomials of low degree, we leverage the *Cauchy binomial theorem* [20, eqn. (1.87)], which implies that

$$\sum_{t=0}^{n} \binom{n}{t}_q (-1)^t q^{\binom{t}{2}} g(q^{-t}) = 0 \tag{7}$$

for every polynomial g of degree less than n. In particular, defining $\varphi(t) = \binom{n}{t}_q (-1)^t q^{\binom{t}{2}}$ for $t = 0,1,\ldots,n$ ensures that φ is exactly orthogonal to each hyperpolynomial $\Gamma_n(s,t)$ for $s < n$. Unfortunately, this choice of φ does not satisfy our constraint on the distribution of the ℓ_1 norm because most of it would be concentrated on the values $\varphi(t)$ at points $t \approx n$. To overcome this difficulty, we apply a hyperpolynomial of low degree to achieve the desired distribution of the ℓ_1 norm. Specifically, we set

$$\varphi(t) = \binom{n}{t}_q (-1)^{t-n} q^{\binom{t}{2} - \binom{n}{2}} \zeta(q^{-t})$$

for a carefully constructed polynomial ζ; the factor $(-1)^{-n} q^{-\binom{n}{2}}$ in this formula serves to normalize φ and ensure the proper signs. As we increase the degree of ζ, we improve the distribution of the ℓ_1 norm of φ at the expense of a weaker orthogonality guarantee, for now φ is orthogonal only to hyperpolynomials of degree less than $n - \deg\zeta$. With an appropriate choice of ζ, we are able to ensure all four desiderata (i)–(iv) for the univariate function φ. The most technical part of the analysis is the upper bound in (iv). For s small, our construction guarantees (iv) as a consequence of the Cauchy binomial theorem, with $\sum_{t=0}^n \varphi(t)\Gamma_n(s,t) = 0$. For s large, we use the pointwise bounds for φ and Γ_n and show that $\sum_{t=0}^n |\varphi(t)|\,|\Gamma_n(s,t)| = q^{-\Omega(k^2)}$.

By combining equations (5) and (6) with the properties (i)–(iv) of the univariate function φ, we derive the following bound on the approximate trace norm: $\|M\|_{\Sigma,2\varepsilon} \geqslant (1 - 2\varepsilon - q^{-\Omega(k)})q^{n^2} q^{\Omega(k^2)}$. Applying the approximate trace norm method (4), we obtain the sought lower bound of $\Omega(k^2 \log q)$ on the quantum communication complexity of F for error $\varepsilon = \frac{1}{2} - q^{-\Theta(k)}$. To achieve the error probability as stated in Theorem I.1, we derive bounds for φ with explicit constants, which we did not discuss in this proof sketch.

The determinant problem: To solve the determinant problem $\mathrm{DET}_{a,b}^{\mathbb{F},n}$ for all field elements a, b, we combine our approach to the matrix rank problem presented above with additional Fourier-theoretic ideas. Recall that we tackle the determinant problem from first principles, without relying on the partial solution for nonzero a, b due to Sun and Wang [5]. With this in mind, we will first discuss the case of nonzero a, b. Consider the function $g_{a,b} \colon \mathbb{F}^{n\times n} \to \{-1, 1, 0\}$ given by

$$g_{a,b}(X) = \begin{cases} -1 & \text{if } \det X = a, \\ 1 & \text{if } \det X = b, \\ 0 & \text{otherwise.} \end{cases}$$

A simple argument reveals that the Fourier coefficients of $g_{a,b}$ corresponding to singular matrices are zero, whereas those corresponding to nonsingular matrices M depend only on $\det(M)$. By applying Parseval's identity, we obtain a strong upper bound on the absolute value of every Fourier coefficient of $g_{a,b}$:

$$\|\widehat{g_{a,b}}\|_\infty \leqslant \frac{1}{\sqrt{|\mathrm{SL}(\mathbb{F},n)|}},$$

where $\mathrm{SL}(\mathbb{F},n)$ denotes the special linear group of order-n matrices over $\mathbb{F}$. Consider now the matrix $\Phi_{a,b}$ whose rows and columns are indexed by elements of $\mathbb{F}^{n\times n}$ and whose entries are given by $\Phi_{a,b}(A,B) = g_{a,b}(A+B)$. The spectral norm of $\Phi_{a,b}$ is governed by the Fourier coefficients of $g_{a,b}$, with

$$\|\Phi_{a,b}\| = q^{n^2} \|\widehat{g_{a,b}}\|_\infty \leqslant \frac{q^{n^2}}{\sqrt{|\mathrm{SL}(\mathbb{F},n)|}}.$$

Observe that $\Phi_{a,b}$ is precisely the characteristic matrix of $\mathrm{DET}_{a,b}^{\mathbb{F},n}$ with the $*$ entries replaced by zeroes. Using $\Phi_{a,b}$ as a witness in the approximate trace norm method, we immediately obtain Theorem I.5 for nonzero a, b.

Consider now the complementary case when one of a, b is zero, say, $a \neq 0$ and $b = 0$. Here, we study the rank versus determinant problem $\mathrm{RANKDET}_{k,a}^{\mathbb{F},n}$, which in this case is a subproblem of the determinant problem. Its parameters are an integer $k \in \{0,1,\ldots,n-1\}$ and a nonzero field element $a \in \mathbb{F}$. Recall that in this problem, Alice and Bob are given matrices $A, B \in \mathbb{F}^{n\times n}$, respectively, and are called upon to distinguish between the cases $\mathrm{rk}(A+B) = k$ and $\det(A+B) = a$. To construct a witness for $\mathrm{RANKDET}_{k,a}^{\mathbb{F},n}$, we combine our solutions to the matrix rank problem and

the determinant problem for nonzero field elements. In more detail, consider the witness Φ for the problem $\mathrm{RANK}_{k,n}^{\mathbb{F},n,n}$ that we sketched above. Recall that $\Phi_{A,B}$ depends only on the rank of $A+B$, and moreover the ℓ_1 norm of Φ is concentrated on matrix pairs (A,B) with $\mathrm{rk}(A+B) \in \{k,n\}$. To turn Φ into a witness for $\mathrm{RANKDET}_{k,a}^{\mathbb{F},n}$, we form a *linear combination* of Φ with the matrices $\Phi_{a,b}$ for all $b \in \mathbb{F} \setminus \{0,a\}$, constructed in the previous paragraph for the determinant problem with nonzero field elements. The coefficients in this linear combination are chosen so as to transfer the ℓ_1 weight placed by Φ on matrix pairs with $\det(A+B) \notin \{0,a\}$ to the matrix pairs with $\det(A+B) = a$, without affecting any other entries of Φ. The resulting dual witness has low spectral norm (being a combination of matrices with low spectral norm) and has its ℓ_1 norm concentrated on matrix pairs (A,B) for which $A+B$ has rank k or determinant a, ensuring strong correlation with the partial function $\mathrm{RANKDET}_{k,a}^{\mathbb{F},n}$. By applying the approximate trace norm method, we obtain the claimed communication lower bounds for $\mathrm{RANKDET}_{k,a}^{\mathbb{F},n}$.

Subspace sum and intersection: We now present the main ideas in our solution to the subspace sum and subspace intersection problems. Since these problems are equivalent, we will discuss the intersection problem alone. As before, we work with an arbitrary finite field $\mathbb{F}$, whose order we denote by q. Also by way of notation, recall that m and ℓ stand for the dimensions of Alice's subspace S and Bob's subspace T, respectively. For simplicity, we will assume in this overview that the dimension n of the ambient vector space satisfies $n \geqslant m+\ell$, which ensures that $\dim(S \cap T)$ takes on every possible value in $\{0,1,2,\ldots,\min\{m,\ell\}\}$ as one varies the subspaces S,T. We will focus on the canonical case of the subspace intersection problem where Alice and Bob need to distinguish subspace pairs with $\dim(S \cap T) = 0$ from those with $\dim(S \cap T) = R$, for an integer R with $0 < R \leqslant \min\{m,\ell\}$. In what follows, we let $F = \mathrm{INTERSECT}_{0,R}^{\mathbb{F},n,m,\ell}$ stand for this communication problem of interest. The general case of the subspace intersection problem, which we will not discuss in this overview, reduces to this canonical case.

As before, the challenge is to construct a dual matrix Φ that witnesses a strong lower bound on the approximate trace norm of the characteristic matrix M of F. Note that the rows of Φ are indexed by m-dimensional subspaces, and the columns are indexed by ℓ-dimensional subspaces. Analogous to the matrix rank problem, we start with the methodological observation that the symmetry of F greatly reduces the number of degrees of freedom in Φ. Specifically, $F(S,T)$ by definition depends only on $\dim(S \cap T)$. A moment's thought now shows that any dual matrix Φ for the subspace intersection problem can be "symmetrized" such that its (S,T) entry depends only on $\dim(S \cap T)$, and this symmetrization can only improve the resulting lower bound on the approximate trace norm in (4).

For $r = 0,1,\ldots,\min\{m,\ell\}$, let $J_r^{n,m,\ell}$ stand for the matrix whose rows are indexed by m-dimensional subspaces of $\mathbb{F}^n$, whose columns are indexed by ℓ-dimensional subspaces of $\mathbb{F}^n$, and whose (S,T) entry is 1 if $\dim(S \cap T) = r$ and zero otherwise. Put another way, $J_r^{n,m,\ell}$ is the characteristic matrix of subspace pairs whose intersection has dimension r. For an arbitrary function $\psi\colon \{0,1,\ldots,\min\{m,\ell\}\} \to \mathbb{R}$, we define

$$J_\psi^{n,m,\ell} = \sum_{r=0}^{\min\{m,\ell\}} \psi(r) J_r^{n,m,\ell}.$$

We refer to this family of matrices, whose (S,T) entry depends only on $\dim(S \cap T)$, as *subspace matrices*. It will also be helpful to have notation for normalized versions of these matrices, as follows:

$$\overline{J}_r^{n,m,\ell} = \frac{1}{\|J_r^{n,m,\ell}\|_1} J_r^{n,m,\ell},$$
$$\overline{J}_\psi^{n,m,\ell} = \sum_{r=0}^{\min\{m,\ell\}} \frac{\psi(r)}{\|J_r^{n,m,\ell}\|_1} J_r^{n,m,\ell}.$$

In this notation, we are looking to construct a dual witness of the form $\Phi = \overline{J}_\psi^{n,m,\ell}$ for some function ψ. This matrix has $\min\{m,\ell\}+1$ degrees of freedom, corresponding to every possible value that $\dim(S \cap T)$ can take. Setting $\Phi = \overline{J}_\psi^{n,m,\ell}$ in the approximate trace norm bound (4) and simplifying, one obtains the following bound for the characteristic matrix M of F:

$$\|M\|_{\Sigma,2\varepsilon} \geqslant \frac{1}{\|\overline{J}_\psi^{n,m,\ell}\|} \left(-\psi(0) + \psi(R) - 2\varepsilon\|\psi\|_1 - \sum_{i \notin \{0,R\}} |\psi(i)| \right). \tag{8}$$

At first glance, this equation looks similar to the corresponding equation (5) for the matrix rank problem. However, there is a major difference: the spectral norm of E_φ is now replaced with the spectral norm of $\overline{J}_\psi^{n,m,\ell}$, and there is no reason to expect that these quantities depend on their corresponding univariate objects φ and ψ in a similar way. Indeed, our spectral analysis of $\overline{J}_\psi^{n,m,\ell}$ is quite different and significantly more technical than that of E_φ.

Analyzing the spectrum of subspace matrices: Symmetric subspace matrices $J_\psi^{n,m,m}$ are classical objects whose eigenvectors and eigenvalues have been studied in numerous works, e.g., [21]–[24]. However, these previous analyses do not seem to apply to the general, asymmetric case of interest to us, namely, that of subspace matrices $J_\psi^{n,m,\ell}$ for arbitrary m,ℓ. One way to reduce the analysis of the spectral norm of $J_\psi^{n,m,\ell}$ to the symmetric case is to express the product $J_\psi^{n,m,\ell}(J_\psi^{n,m,\ell})^{\mathsf{T}} = J_\psi^{n,m,\ell} J_\psi^{n,\ell,m}$ as the sum of symmetric subspace matrices and then apply known results for the symmetric case. Unfortunately, multiplying these subspace matrices leads to expressions so unwieldy and complicated that this is clearly not the method of choice.

Instead, our analysis is inspired by a result of Knuth [25] on what he called *combinatorial matrices*. Specifically, Knuth investigated the eigenvalues of symmetric matrices of order $\binom{n}{t}$ whose rows and columns are indexed by t-element subsets of $\{1,2,\ldots,n\}$ and whose (A,B) entry depends only on

$|A \cap B|$. To determine the eigenvectors of a combinatorial matrix, Knuth studied certain homogeneous linear systems with variables indexed by subsets of a fixed cardinality s, and the equations themselves corresponding to sets of cardinality $s-1$. He showed that any solution to such a system for $s \in \{1, 2, \ldots, t\}$ is an eigenvector for every combinatorial matrix of order $\binom{n}{t}$. Knuth also proved that for any given s, the space of solutions has a basis supported on the variables indexed by what he called *basic sets*. These sets have a simple combinatorial description, which the author of [25] used to prove that the eigenvectors arising from the homogeneous systems for $s = 1, 2, \ldots, t$, together with the all-ones vector, form an exhaustive description of the eigenvectors of each combinatorial matrix. Once the eigenvectors are determined, one readily calculates their associated eigenvalues and in particular the spectral norm.

With some effort, we are able to adapt Knuth's ideas to the context of subspaces. Along the way, we encounter several obstacles. To begin with, counting problems that are straightforward for sets become challenging for subspaces, and some intuitive combinatorial principles no longer work. For example, the inclusion-exclusion formula $\dim(S+T) = \dim(S) + \dim(T) - \dim(S \cap T)$ has no analogue for three or more subspaces. Another obstacle is that Knuth's notion of a basic set does not seem to have a meaningful analogue for subspaces. For this reason, we reformulate Knuth's ideas in a purely linear-algebraic way and sidestep much of the combinatorial machinery in [25]. The final hurdle is extending Knuth's analysis to the asymmetric case. Ultimately, we are able to determine the spectral norm of every subspace matrix $J_\psi^{n,m,\ell}$ and in particular its normalized version $\overline{J}_\psi^{n,m,\ell}$. We prove that

$$\|\overline{J}_\psi^{n,m,\ell}\| = \max_{s=0,1,\ldots,\min\{m,\ell\}} \left| \sum_{r=0}^{\min\{m,\ell\}} \psi(r) \overline{\Lambda}_r^{n,m,\ell}(s) \right|^{1/2} \times \left| \sum_{r=0}^{\min\{m,\ell\}} \psi(r) \overline{\Lambda}_r^{n,\ell,m}(s) \right|^{1/2}, \quad (9)$$

where $\overline{\Lambda}_r^{n,m,\ell}$ and $\overline{\Lambda}_r^{n,\ell,m}$ are functions with algebraic and analytic properties analogous to those of the Γ_n function in our solution to the matrix rank problem. Specifically, we have:

(i) for n, m, ℓ, s fixed, $\overline{\Lambda}_r^{n,m,\ell}(s)$ as a function of $r \in \{0, 1, \ldots, \min\{m, \ell\}\}$ is a polynomial in q^r of degree at most s;

(ii) $|\overline{\Lambda}_r^{n,m,\ell}(s)| \leqslant 8\binom{n}{m}_q^{-1} q^{-s(m-r)/2}$ for every integer $r = 0, 1, \ldots, \min\{m, \ell\}$.

By swapping the roles of m and ℓ, one obtains analogous properties for $\overline{\Lambda}_r^{n,\ell,m}(s)$.

This spectral result gives us fine-grained control over the spectrum of $J_\psi^{n,m,\ell}$ via the univariate function ψ. Our construction of ψ is based on the Cauchy binomial theorem and is conceptually similar to our univariate function φ in the matrix rank problem. In particular, we use the algebraic property (i) to bound the product in (9) for small s, and the analytic property (ii) to bound it for large s. We further ensure that the ℓ_1 norm of ψ is highly concentrated on $\{0, R\}$, with $\psi(0) < 0$ and $\psi(R) > 0$. This results in a strong lower bound in (8), which in turn leads to an optimal lower bound on the communication complexity of F by virtue of the approximate trace norm method.

II. Preliminaries

A. General notation

We view Boolean functions as mappings $X \to \{-1, 1\}$, where X is a nonempty finite set and the range elements $-1, 1$ correspond to "true" and "false," respectively. A *partial* Boolean function is a mapping $f\colon X \to \{-1, 1, *\}$, whose *domain* is defined as $\operatorname{dom} f = \{x \in X : f(x) \neq *\}$. Recall that for an arbitrary function $f\colon X \to Y$, the restriction of f to a subset $X' \subseteq X$ is defined to be the mapping $f|_{X'}\colon X' \to Y$ given by $(f|_{X'})(x) = f(x)$.

We adopt the shorthand $[n] = \{1, 2, \ldots, n\}$. We use the letters p and q throughout this manuscript to refer to a prime number and a prime power, respectively. As usual, $\mathbb{F}_q$ stands for the Galois field $\mathrm{GF}(q)$, the q-element field which is unique up to isomorphism. For a given set X, the *Kronecker delta* $\delta_{x,y}$ is defined for $x, y \in X$ by

$$\delta_{x,y} = \begin{cases} 1 & \text{if } x = y, \\ 0 & \text{otherwise.} \end{cases}$$

For a function $f\colon X \to \mathbb{C}$, we use the familiar norms $\|f\|_1 = \sum_{x \in X} |f(x)|$ and $\|f\|_\infty = \max_{x \in X} |f(x)|$. Similarly, for a real or complex matrix M, one defines $\|M\|_1 = \sum |M_{i,j}|$ and $\|M\|_\infty = \max |M_{i,j}|$. The norms $\|v\|_1$ and $\|v\|_\infty$ for a real or complex vector v are defined analogously. The Euclidean norm is given by $\|v\|_2 = \sqrt{\sum |v_i|^2}$. We denote the base-$q$ logarithm of x by $\log_q x$. In the special case of the binary logarithm, we write simply $\log x$ rather than $\log_2 x$.

B. Linear-algebraic preliminaries

Let $\mathbb{F}$ be a given field. We denote the set of $n \times m$ matrices over $\mathbb{F}$ by $\mathbb{F}^{n \times m}$. We use the standard notation $\operatorname{rk} A$, $\ker A$, and A^{T} for the rank, null space, and transpose of the matrix A. As usual, the determinant of $A \in \mathbb{F}^{n \times n}$ is denoted $\det A$. The trace of a matrix $A \in \mathbb{F}^{n \times n}$ is denoted $\operatorname{tr} A$ and defined as the sum of the diagonal elements of A. The commutativity of the trace operator is often helpful: $\operatorname{tr}(AB) = \operatorname{tr}(BA)$ for square matrices A, B. We let $\operatorname{diag}(a_1, a_2, \ldots, a_n)$ denote the diagonal matrix of order n with diagonal entries $a_1, a_2, \ldots, a_n$. Recall that I_n normally denotes the identity matrix of order n, whereas I denotes the identity matrix whose order is to be inferred from the context. We generalize the meaning of I_n somewhat by defining

$$I_n = \operatorname{diag}(\underbrace{1, 1, \ldots 1}_{n}, 0, \ldots 0),$$

where the order of the matrix (and hence the number of zeroes on the diagonal) will be clear from the context. We let J and

$\mathbf{1}$ denote the all-ones matrix and all-ones vector, respectively, whose dimensions will be clear from the context.

Fact II.1. *For square matrices A, B of order n over a given field $\mathbb{F}$,*

$$\operatorname{rk} AB \geqslant \operatorname{rk} A + \operatorname{rk} B - n.$$

Proof: Recall that the dimension of $\ker AB$ is at most the sum of the dimensions of $\ker A$ and $\ker B$. By the rank-nullity theorem, this is equivalent to the claimed inequality. ∎

For $\mathbb{F}$ a finite field or the field of real numbers, the inner product operation on vectors and matrices is defined as usual by $\langle x, y\rangle = \sum x_i y_i$ and $\langle A, B\rangle = \sum A_{i,j} B_{i,j}$. For $\mathbb{F} = \mathbb{C}$, the modified definitions $\langle x, y\rangle = \sum x_i \overline{y_i}$ and $\langle A, B\rangle = \sum A_{i,j} \overline{B_{i,j}}$ are used instead. For complex-valued functions $f, g\colon X \to \mathbb{C}$, we write $\langle f, g\rangle = \sum_{x\in X} f(x)\overline{g(x)}$. Again for $\mathbb{F} = \mathbb{C}$, the conjugate transpose of a matrix $A = [A_{i,j}]_{i,j}$ is denoted by $A^* = [\overline{A_{j,i}}]_{i,j}$, and a matrix $A \in \mathbb{C}^{n\times n}$ is called *unitary* if $A^*A = AA^* = I$. The following useful fact relates the inner product and trace operators.

Fact II.2. *Let A, B, C, D be matrices of order n over $\mathbb{R}$ or a finite field. Then:*

(i) $\langle A, B\rangle = \operatorname{tr}(AB^{\mathsf{T}}) = \operatorname{tr}(A^{\mathsf{T}}B)$,
(ii) $\langle A, C_1 B C_2\rangle = \langle C_1^{\mathsf{T}} A C_2^{\mathsf{T}}, B\rangle$.

Proof: Item (i) is immediate from the definition of matrix multiplication, whereas (ii) follows from (i) and the commutativity of the trace operator: $\langle A, C_1BC_2\rangle = \operatorname{tr}(AC_2^{\mathsf{T}}B^{\mathsf{T}}C_1^{\mathsf{T}}) = \operatorname{tr}(C_1^{\mathsf{T}}AC_2^{\mathsf{T}}B^{\mathsf{T}}) = \langle C_1^{\mathsf{T}}AC_2^{\mathsf{T}}, B\rangle$. ∎

For any field $\mathbb{F}$, we let $e_1, e_2, \ldots, e_n$ denote as usual the vectors of the standard basis for $\mathbb{F}^n$. For any subset $S \subseteq \mathbb{F}^n$, recall that its span over $\mathbb{F}$ is denoted $\operatorname{span} S$. For a linear subspace S, the symbols $\dim S$ and $S^\perp$ refer as usual to the dimension of S and the orthogonal complement of S, respectively. For a linear transformation M, we let $M(S) = \{Mx : x \in S\}$ denote the image of S under M. Recall that the *sum* of linear subspaces S and T is defined as $S + T = \{x + y : x \in S, y \in T\}$ and is the smallest subspace that contains both S and T. In expressions involving subspaces, we adopt the convention that the union $\cup$ and intersection $\cap$ operators have higher precedence than the subspace sum operator $+$. For a vector space V and an integer k, we adopt the notation $\mathscr{S}(V, k)$ for the set of all subspaces of V of dimension k. For arbitrary subspaces S, T in a finite-dimensional vector space, the following identity is well-known, and we use it extensively in our proofs without further mention:

$$\dim(S + T) = \dim(S) + \dim(T) - \dim(S \cap T). \tag{10}$$

This equation is one of the few instances when subspaces behave in ways analogous to sets. Such instances are rare. For example, unlike sets, general subspaces S, T, U need not satisfy $S \cap (T + U) = S \cap T + S \cap U$. The equality requires additional hypotheses, as recorded below.

Fact II.3. *For any linear subspaces S, S', T with $S' \subseteq S$,*

$$S \cap (S' + T) = S' + S \cap T.$$

Proof: It is clear that $S' + S \cap T$ is a subspace of both S and $S' + T$. It remains to prove the opposite inclusion, $S \cap (S' + T) \subseteq S' + S \cap T$. For this, consider an arbitrary vector $u+v \in S$ with $u \in S'$ and $v \in T$. Then $v \in S+u = S$. As a result, $v \in S \cap T$ and therefore $u + v \in S' + S \cap T$ as claimed. ∎

We continue with a fact that relates the dimension of $S \cap T$ to that of $S^\perp \cap T^\perp$.

Fact II.4. *Let $S, T \subseteq \mathbb{F}^n$ be subspaces over a given field $\mathbb{F}$. Then*

$$(S + T)^\perp = S^\perp \cap T^\perp, \tag{11}$$
$$(S \cap T)^\perp = S^\perp + T^\perp. \tag{12}$$

Moreover,

$$\dim(S \cap T) = \dim(S) + \dim(T) + \dim(S^\perp \cap T^\perp) - n. \tag{13}$$

Proof: To begin with,

$$\begin{aligned} S^\perp \cap T^\perp &= \{x : \langle x, y\rangle = 0 \text{ for all } y \in S\} \\ &\qquad \cap \{x : \langle x, y\rangle = 0 \text{ for all } y \in T\} \\ &= \{x : \langle x, y\rangle = 0 \text{ for all } y \in S \cup T\} \\ &= \{x : \langle x, y\rangle = 0 \text{ for all } y \in S + T\} \\ &= (S + T)^\perp, \end{aligned}$$

where the third step uses the linearity of inner product. This settles (11). Applying (11) to the orthogonal complements of S and T results in $(S^\perp + T^\perp)^\perp = S \cap T$, which upon orthogonal complementation of both sides yields (12). Equation (13) is also a straightforward consequence of (11), as follows:

$$\begin{aligned} \dim(S^\perp \cap T^\perp) &= \dim((S + T)^\perp) \\ &= n - \dim(S + T) \\ &= n - \dim(S) - \dim(T) + \dim(S \cap T). \end{aligned}$$

This completes the proof. ∎

It is well-known that for a symmetric real matrix, any pair of eigenvectors corresponding to distinct eigenvalues are orthogonal. For completeness, we state this simple fact with a proof below.

Fact II.5. *Let M be a symmetric real matrix. Let u, v be eigenvectors of M corresponding to different eigenvalues. Then $\langle u, v\rangle = 0$.*

Proof: Suppose that $Mu = \alpha u$ and $Mv = \beta v$, where $\alpha \neq \beta$. Then $(\alpha - \beta)\langle u, v\rangle = \langle \alpha u, v\rangle - \langle u, \beta v\rangle = \langle Mu, v\rangle - \langle u, Mv\rangle = 0$, where the last step uses $M = M^{\mathsf{T}}$. This forces $\langle u, v\rangle = 0$, as claimed. ∎

C. Matrix norms

Associated with every matrix $A \in \mathbb{C}^{n\times m}$ are $\min\{n, m\}$ nonnegative reals that are called the *singular values of A*, denoted $\sigma_1(A) \geqslant \sigma_2(A) \geqslant \cdots \geqslant \sigma_{\min\{n,m\}}(A)$. Every matrix $A \in \mathbb{C}^{n\times m}$ has a *singular value decomposition* $A = U\Sigma V^*$, where U and V are unitary matrices of order n and m,

respectively, and Σ is a rectangular diagonal matrix whose diagonal entries are $\sigma_1(A), \sigma_2(A), \ldots, \sigma_{\min\{n,m\}}(A)$. In the case of real matrices A, the matrices U and V in the singular value decomposition can be taken to be real. An alternative characterization of the singular values is given by

Fact II.6. *Let $A \in \mathbb{C}^{n\times m}$ be given, with $n \leqslant m$. Then the singular values of A are precisely the square roots of the eigenvalues of AA^*, counting multiplicities.*

The spectral norm, trace norm, and Frobenius norm of A are defined in terms of the singular values as follows:

$$\|A\| = \sigma_1(A), \tag{14}$$

$$\|A\|_\Sigma = \sum \sigma_i(A), \tag{15}$$

$$\|A\|_{\mathrm{F}} = \sqrt{\sum \sigma_i(A)^2}. \tag{16}$$

Equivalently,

$$\|A\| = \max_{x:\|x\|_2=1} \|Ax\|_2, \tag{17}$$

$$\|A\|_{\mathrm{F}} = \sqrt{\sum |A_{ij}|^2}. \tag{18}$$

These equations agree with (14) and (16) because the Euclidean norm on vectors is invariant under unitary transformations.

Fact II.7. *For any matrices $A, B \in \mathbb{C}^{n\times m}$,*

$$|\langle A, B\rangle| \leqslant \|A\| \, \|B\|_\Sigma.$$

Fact II.7 follows directly from (17) and the singular value decomposition of B. We now recall a relationship between the trace norm and Frobenius norm; see, e.g., [17, Prop. 2.4].

Fact II.8. *For all matrices A and B of compatible dimensions,*

$$\|AB\|_\Sigma \leqslant \|A\|_{\mathrm{F}} \, \|B\|_{\mathrm{F}}.$$

Recall that a *sign matrix* is a real matrix with entries in $\{-1, 1\}$. A *partial* sign matrix, then, is a matrix with entries in $\{-1, 1, *\}$. We define the *domain* of a partial sign matrix F by $\operatorname{dom} F = \{(i, j) : F_{ij} \neq *\}$. The *$\varepsilon$-approximate trace norm of F*, denoted $\|F\|_{\Sigma,\varepsilon}$, is the least trace norm of a real matrix $\tilde{F}$ that satisfies

$$|F_{ij} - \tilde{F}_{ij}| \leqslant \varepsilon \qquad \text{if } F_{ij} \in \{-1, 1\}, \tag{19}$$

$$|\tilde{F}_{ij}| \leqslant 1 + \varepsilon \qquad \text{if } F_{ij} = *. \tag{20}$$

The following lower bound on the approximate trace norm is well known [17], [19], [26]. For reader's convenience, we include a proof.

Proposition II.9. *For any partial sign matrix F and $\varepsilon \geqslant 0$,*

$$\|F\|_{\Sigma,\varepsilon} \geqslant \sup_{\Phi\neq 0} \frac{1}{\|\Phi\|} \left(\sum_{(i,j)\in \operatorname{dom} F} F_{ij}\Phi_{ij} - \varepsilon\|\Phi\|_1 - \sum_{(i,j)\notin \operatorname{dom} F} |\Phi_{ij}| \right).$$

Proof: Let $\tilde{F}$ be a real matrix that approximates F in the sense of (19) and (20). Then for any $\Phi \neq 0$,

$$\begin{aligned}
\langle \tilde{F}, \Phi\rangle &= \sum_{\operatorname{dom} F} F_{ij}\Phi_{ij} + \sum_{\operatorname{dom} F} (\tilde{F}_{ij} - F_{ij})\Phi_{ij} + \sum_{\overline{\operatorname{dom} F}} \tilde{F}_{ij}\Phi_{ij} \\
&\geqslant \sum_{\operatorname{dom} F} F_{ij}\Phi_{ij} - \sum_{\operatorname{dom} F} |\tilde{F}_{ij} - F_{ij}|\,|\Phi_{ij}| \\
&\qquad\qquad - \sum_{\overline{\operatorname{dom} F}} |\tilde{F}_{ij}|\,|\Phi_{ij}| \\
&\geqslant \sum_{\operatorname{dom} F} F_{ij}\Phi_{ij} - \sum_{\operatorname{dom} F} \varepsilon|\Phi_{ij}| - \sum_{\overline{\operatorname{dom} F}} (1+\varepsilon)|\Phi_{ij}| \\
&= \sum_{\operatorname{dom} F} F_{ij}\Phi_{ij} - \varepsilon\|\Phi\|_1 - \sum_{\overline{\operatorname{dom} F}} |\Phi_{ij}|.
\end{aligned}$$

On the other hand, Fact II.7 shows that $\langle \tilde{F}, \Phi\rangle \leqslant \|\tilde{F}\|_\Sigma \, \|\Phi\|$. Combining these two bounds for $\langle \tilde{F}, \Phi\rangle$ gives

$$\|\tilde{F}\|_\Sigma \geqslant \frac{1}{\|\Phi\|} \left(\sum_{\operatorname{dom} F} F_{ij}\Phi_{ij} - \varepsilon\|\Phi\|_1 - \sum_{\overline{\operatorname{dom} F}} |\Phi_{ij}| \right).$$

Taking the supremum over $\Phi \neq 0$ completes the proof. ∎

D. Fourier transform

Consider a prime power $q = p^k$, with p a prime and k a positive integer. Recall that the additive group of $\mathbb{F}_q$ is isomorphic to the Abelian group $\mathbb{Z}_p^k$. Fix any such isomorphism ψ. Let $\omega = e^{2\pi i/p}$, a primitive p-th root of unity. For $x \in \mathbb{F}_q$, define $\omega^x = \omega^{x_1}\omega^{x_2}\cdots\omega^{x_k}$, where $(x_1, x_2, \ldots, x_k)$ is the image of x under ψ. Then for all $x, y \in \mathbb{F}_q$,

$$\omega^{x+y} = \omega^x\omega^y, \tag{21}$$

$$\omega^{-x} = \overline{\omega^x}. \tag{22}$$

One further calculates $\sum_{x\in\mathbb{F}_q} \omega^x = \prod_{i=1}^k (1 + \omega + \omega^2 + \cdots + \omega^{p-1}) = 0$, which in turn generalizes to

$$\sum_{x\in\mathbb{F}_q} \omega^{ax} = 0, \qquad a \in \mathbb{F}_q \setminus \{0\} \tag{23}$$

since $x \mapsto ax$ is a permutation on $\mathbb{F}_q$.

Let n be a positive integer. For $A \in \mathbb{F}_q^{n\times n}$, define a corresponding character $\chi_A : \mathbb{F}_q^{n\times n} \to \mathbb{C}$ by

$$\chi_A(X) = \omega^{\langle A,X\rangle}.$$

It follows from (21) that

$$\chi_A(X + Y) = \chi_A(X)\chi_A(Y), \tag{24}$$

making χ_A a homomorphism of the additive group $\mathbb{F}_q^{n\times n}$ into the multiplicative group of $\mathbb{C}$. Using (21) and (22), one obtains $\langle \chi_A, \chi_B\rangle = \sum_X \omega^{\langle A,X\rangle}\overline{\omega^{\langle B,X\rangle}} = \sum_X \omega^{\langle A,X\rangle - \langle B,X\rangle} = \sum_X \omega^{\langle A-B,X\rangle}$, which along with (23) leads to

$$\langle \chi_A, \chi_B\rangle = \begin{cases} q^{n^2} & \text{if } A = B, \\ 0 & \text{otherwise.} \end{cases} \tag{25}$$

Hence, the characters χ_A for $A \in \mathbb{F}_q^{n\times n}$ form an orthogonal basis for the complex vector space of functions $\mathbb{F}_q^{n\times n} \to \mathbb{C}$.

In particular, every function $f\colon \mathbb{F}_q^{n\times n} \to \mathbb{C}$ has a unique representation as a linear combination of the characters:

$$f(X) = \sum_{A\in\mathbb{F}_q^{n\times n}} \widehat{f}(A)\chi_A(X). \tag{26}$$

The numbers $\widehat{f}(A)$ are called the *Fourier coefficients of* f. They are given by

$$\widehat{f}(A) = q^{-n^2}\langle f, \chi_A\rangle = \mathop{\mathbf{E}}_{X\in\mathbb{F}_q^{n\times n}} f(X)\omega^{-\langle A,X\rangle}. \tag{27}$$

where the first step is justified by (25), and the second step uses (22). An immediate consequence of (25) and (26) is that $\langle f, f\rangle = q^{n^2}\sum_A |\widehat{f}(A)|^2$. This result is known as *Parseval's identity,* and it is typically written in the form

$$\mathop{\mathbf{E}}_{X\in\mathbb{F}_q^{n\times n}}[|f(X)|^2] = \sum_{A\in\mathbb{F}_q^{n\times n}} |\widehat{f}(A)|^2. \tag{28}$$

With $\widehat{f}$ viewed as a complex-valued function on $\mathbb{F}_q^{n\times n}$, the linear transformation that sends $f \mapsto \widehat{f}$ is called the *Fourier transform*. Its matrix representation is easy to describe. Specifically, define

$$H_n = q^{-n^2/2}[\omega^{\langle A,B\rangle}]_{A,B},$$

where the row and column indices range over all matrices in $\mathbb{F}_q^{n\times n}$. Analogous to (25), one shows that H_n is unitary:

$$H_nH_n^* = H_n^*H_n = I. \tag{29}$$

Then the Fourier transform $f \mapsto \widehat{f}$, given by (27), corresponds to the linear transformation $q^{-n^2/2}H_n^*$. Analogously, the inverse transformation $\widehat{f} \mapsto f$ of (26) corresponds to $q^{n^2/2}H_n$.

The following well-known fact relates the singular values of a matrix $[\varphi(A+B)]_{A,B}$ to the Fourier spectrum of the outer function φ. We include a proof adapted from [6] and generalized to the case of $\mathbb{F}_q$.

Fact II.10 (adapted from Li et al., Lemma 20)**.** *Let* $\varphi : \mathbb{F}_q^{n\times n} \to \mathbb{C}$ *be given. Define*

$$\Phi = [\varphi(X+Y)]_{X,Y\in\mathbb{F}_q^{n\times n}}.$$

Then

$$\Phi = H_nDH_n,$$

where D *is the diagonal matrix given by* $D_{A,A} = q^{n^2}\widehat{\varphi}(A)$. *In particular, the singular values of* Φ *are* $q^{n^2}|\widehat{\varphi}(A)|$ *for* $A \in \mathbb{F}_q^{n\times n}$.

Proof: Using the homomorphic property (24) of the characters,

$$\begin{aligned}\varphi(X+Y) &= \sum_{A\in\mathbb{F}_q^{n\times n}} \widehat{\varphi}(A)\chi_A(X+Y)\\ &= \sum_{A\in\mathbb{F}_q^{n\times n}} \widehat{\varphi}(A)\chi_A(X)\chi_A(Y).\end{aligned}$$

Restated in matrix form, this equation becomes $\Phi = [\chi_A(X)]_{X,A}\,\mathrm{diag}(\ldots,\widehat{\varphi}(A),\ldots)\,[\chi_A(Y)]_{A,Y} = H_nDH_n$, as desired. ∎

E. Gaussian binomial coefficients

Gaussian binomial coefficients, also known as *q-binomial coefficients*, are defined by

$$\binom{n}{m}_q = \frac{(q^n-1)(q^n-q)\cdots(q^n-q^{m-1})}{(q^m-1)(q^m-q)\cdots(q^m-q^{m-1})} \tag{30}$$

$$= \frac{(q^n-1)(q^{n-1}-1)\cdots(q^{n-m+1}-1)}{(q^m-1)(q^{m-1}-1)\cdots(q-1)} \tag{31}$$

for all nonnegative integers n, m and real numbers $q > 1$. Observe that $\binom{n}{0}_q = 1$ since the above product is empty for $m = 0$. Note further that $\binom{n}{m}_q = 0$ whenever $m > n$. One recovers standard binomial coefficients from this definition via

$$\lim_{q\searrow 1}\binom{n}{m}_q = \binom{n}{m}.$$

As a matter of convenience, one generalizes Gaussian binomial coefficients to arbitrary integers n, m by defining

$$\binom{n}{m}_q = 0 \qquad \text{if } \min\{n,m\} < 0.$$

With this convention, one has the familiar identity

$$\binom{n}{m}_q = \binom{n}{n-m}_q, \qquad n, m \in \mathbb{Z}. \tag{32}$$

Gaussian binomial coefficients play an important role in enumerative combinatorics. In particular, we recall the following classical fact.

Fact II.11. *Fix a prime power* q *and integers* $n \geqslant m \geqslant 0$. *Then the number of* m*-dimensional subspaces of* $\mathbb{F}_q^n$ *is exactly* $\binom{n}{m}_q$.

Proof: This result is clearly true for $m = 0$. For $m \geqslant 1$, there are $(q^n-1)(q^n-q)\cdots(q^n-q^{m-1})$ ordered bases $(v_1, v_2, \ldots, v_m)$ of vectors in $\mathbb{F}_q^n$. Each such basis defines an m-dimensional subspace. Conversely, every m-dimensional subspace has exactly $(q^m-1)(q^m-q)\cdots(q^m-q^{m-1})$ ordered bases. Thus, the number of m-dimensional subspaces is (30), as claimed. ∎

The following monotonicity property of q-binomial coefficients is well-known. We provide a proof for convenience.

Fact II.12. *Let* $n \geqslant m \geqslant 0$ *be given integers. Then for all integers* $\ell \in [m, n-m]$ *and reals* $q > 1$,

$$\binom{n}{m}_q \leqslant \binom{n}{\ell}_q. \tag{33}$$

Proof: The defining equation (31) gives

$$\binom{n}{\ell}_q = \binom{n}{m}_q \cdot \prod_{i=m+1}^{\ell} \frac{q^{n-i+1}-1}{q^i-1}.$$

If $\ell \leqslant n/2$, then every fraction in the above product is greater than 1. As a result, (33) holds in this case. In the complementary case $\ell > n/2$, we have $n-\ell \in [m, n/2]$ and therefore

$$\binom{n}{m}_q \leqslant \binom{n}{n-\ell}_q$$

by the first part of the proof. Since $\binom{n}{n-\ell}_q = \binom{n}{\ell}_q$, we again arrive at (33). ∎

We will use the next proposition to accurately estimate Gaussian binomial coefficients.

Proposition II.13. *For any set I of positive integers, and any real number $x \geqslant 2$,*

$$\frac{1}{4} \leqslant \prod_{i \in I} \left(1 - \frac{1}{x^i}\right) \leqslant 1.$$

Proof: The upper bound is trivial. For the lower bound, we may clearly assume that $I = \{1, 2, 3 \ldots\}$. A simple inductive argument shows that $(1-a_1)\cdots(1-a_n) \geqslant 1-a_1-\cdots-a_n$ for any $a_1, \ldots, a_n \in (0,1)$. It follows that

$$\prod_{i=2}^{\infty} \left(1 - \frac{1}{x^i}\right) \geqslant 1 - \frac{1}{x^2} - \frac{1}{x^3} - \ldots = 1 - \frac{1}{x(x-1)}$$

and therefore

$$\prod_{i=1}^{\infty} \left(1 - \frac{1}{x^i}\right) \geqslant \left(1 - \frac{1}{x}\right)\left(1 - \frac{1}{x(x-1)}\right) \geqslant \frac{1}{4},$$

where the last step uses $x \geqslant 2$. ∎

Corollary II.14. *For any integers $n \geqslant m \geqslant 0$ and any real number $q \geqslant 2$,*

$$q^{m(n-m)} \leqslant \binom{n}{m}_q \leqslant 4q^{m(n-m)}.$$

Proof: The lower bound follows directly from the fact that $(q^n - q^i)/(q^m - q^i) \geqslant q^n/q^m$ for $n \geqslant m$. The upper bound can be verified as follows:

$$\begin{aligned}\binom{n}{m}_q &= \frac{(q^n-1)(q^n-q)\ldots(q^n-q^{m-1})}{(q^m-1)(q^m-q)\ldots(q^m-q^{m-1})} \\ &\leqslant \frac{q^{nm}}{q^{m^2}\prod_{i=1}^{m}(1-q^{-i})} \\ &\leqslant 4q^{m(n-m)},\end{aligned}$$

where the last step applies Proposition II.13. ∎

We now recall a classical result known as the *Cauchy binomial theorem*, see, e.g., [20, eqn. (1.87)].

Fact II.15. *For any integer $n \geqslant 1$ and real number $q > 1$, the following identity holds in $\mathbb{R}[t]$:*

$$(1+t)(1+qt)\ldots(1+q^{n-1}t) = \sum_{i=0}^{n} q^{\binom{i}{2}} \binom{n}{i}_q t^i. \quad (34)$$

Corollary II.16. *For any integer $n \geqslant 1$ and real number $q > 1$, and any real polynomial g of degree less than n,*

$$\sum_{i=0}^{n} (-1)^i q^{\binom{i}{2}} \binom{n}{i}_q g(q^{-i}) = 0. \quad (35)$$

Proof: For $d = 0, 1, \ldots, n-1$, take $t = -1/q^d$ in (34) to obtain

$$\sum_{i=0}^{n} (-1)^i q^{\binom{i}{2}} \binom{n}{i}_q (q^{-i})^d = 0. \quad (36)$$

This establishes (35) when g is a *monomial* of degree less than n. The general case follows by linearity: multiply (36) by the degree-d coefficient in g and sum over d. ∎

F. Counting and generating matrices of given rank

For a field $\mathbb{F}$, we let $\mathscr{M}_r^{\mathbb{F},n,m}$ denote the set of matrices in $\mathbb{F}^{n\times m}$ of rank r. Since we mostly use $\mathbb{F} = \mathbb{F}_q$ in this work, we will usually omit the reference to the field and write simply $\mathscr{M}_r^{n,m}$. As a matter of convenience, we adopt the convention that for any $n \geqslant 0$ there is exactly one "matrix" of size $0 \times n$ and exactly one "matrix" of size $n \times 0$, both of rank 0. The role of these empty matrices is to ensure that

$$|\mathscr{M}_0^{0,n}| = |\mathscr{M}_0^{n,0}| = 1, \qquad n \geqslant 0,$$

which simplifies the statement of several lemmas in this paper. Analogously, we define

$$\mathscr{M}_r^{n,m} = \varnothing \qquad \text{if } \min\{n, m, r\} < 0. \quad (37)$$

For nonsingular matrices of order $n \geqslant 1$, we adopt the shorthand $\mathscr{M}_n = \mathscr{M}_n^{n,n}$.

Proposition II.17. *Let n, m, r be nonnegative integers with $r \leqslant \min\{n, m\}$. Then*

$$|\mathscr{M}_r^{n,m}| = \binom{n}{r}_q (q^m - 1)(q^m - q)\ldots(q^m - q^{r-1}). \quad (38)$$

Proof: If $r = 0$, then the right-hand side of (38) evaluates to 1. This is consistent with our convention that $|\mathscr{M}_0^{n,m}| = 1$ for all $n, m \geqslant 0$.

We now consider the complementary case $r \geqslant 1$, which forces n and m to be positive. Fix an arbitrary r-dimensional subspace $S \subseteq \mathbb{F}_q^n$ and consider the subset $\mathscr{M}_S \subseteq \mathscr{M}_r^{n,m}$ of matrices whose column space is S. Fix an $n \times r$ matrix A with column space S. Since the columns of A are linearly independent, every matrix in $\mathscr{M}_S$ has a unique representation of the form AB for some $B \in \mathscr{M}_r^{r,m}$. Conversely, any product AB with $B \in \mathscr{M}_r^{r,m}$ is a matrix in $\mathscr{M}_S$. Therefore,

$$|\mathscr{M}_S| = |\mathscr{M}_r^{r,m}|. \quad (39)$$

Recall that $\mathscr{M}_r^{n,m}$ is the disjoint union of $\mathscr{M}_S$ over r-dimensional subspaces $S \subseteq \mathbb{F}_q^n$, and there are precisely $\binom{n}{r}_q$ such subspaces (Fact II.11). With this in mind, (39) leads to

$$|\mathscr{M}_r^{n,m}| = \binom{n}{r}_q |\mathscr{M}_r^{r,m}|. \quad (40)$$

Finally, the number of $r \times m$ matrices of rank r is precisely the number of bases $(v_1, v_2, \ldots, v_r)$ of row vectors in $\mathbb{F}_q^m$, whence $|\mathscr{M}_r^{r,m}| = (q^m-1)(q^m-q)\cdots(q^m-q^{r-1})$. Making this substitution in (40) completes the proof. ∎

Using Proposition II.13 and Corollary II.14 to estimate the right-hand side of (38), we obtain:

Corollary II.18. *Let m, n, r be nonnegative integers with $r \leqslant \min\{n, m\}$. Then*

$$\frac{1}{4} q^{r(n+m-r)} \leqslant |\mathscr{M}_r^{n,m}| \leqslant 4q^{r(n+m-r)}.$$

The following fact is well-known; cf. [6].

Proposition II.19. *Let $n \geqslant 1$ be a given integer. Let X, Y be random matrices distributed independently and uniformly on $\mathscr{M}_n$. Then:*

(i) *for any fixed $A \in \mathscr{M}_n$, the matrices XA and AX are distributed uniformly on $\mathscr{M}_n$;*

(ii) *for any $r \in \{0, 1, \ldots, n\}$ and fixed $A \in \mathscr{M}_r^{n,n}$, the matrix XAY is distributed uniformly on $\mathscr{M}_r^{n,n}$.*

Proof: (i) For any $B \in \mathscr{M}_n$, we have $\mathbf{P}[XA = B] = \mathbf{P}[X = BA^{-1}] = 1/|\mathscr{M}_n|$. Therefore, XA is distributed uniformly on $\mathscr{M}_n$. The argument for AX is analogous.

(ii) Fix $B \in \mathscr{M}_r^{n,n}$ arbitrarily. Then B can be obtained from A by a series of elementary row and column operations, so that $B = M_1 A M_2$ for nonsingular M_1, M_2. As a result,

$$\begin{aligned}\mathbf{P}[XAY = B] &= \mathbf{P}[M_1^{-1} XAY M_2^{-1} = A] \\ &= \mathbf{P}[XAY M_2^{-1} = A] \\ &= \mathbf{P}[XAY = A],\end{aligned}$$

where the last two steps are valid by part (i). To summarize, XAY takes on every value in $\mathscr{M}_r^{n,n}$ with the same probability. Since $XAY \in \mathscr{M}_r^{n,n}$, the proof is complete. ∎

G. Random projections

Given a collection of subspaces $S_1, S_2, \ldots, S_m$ in a vector space, we use random projections to reduce the dimension of the ambient space while preserving algebraic relationships among the S_i. This is done by choosing a uniformly random matrix X and replacing $S_1, S_2, \ldots, S_m$ with the subspaces $X(S_1), X(S_2), \ldots, X(S_m)$, respectively. The following lemma provides quantitative details.

Lemma II.20. *Let n and d be positive integers, $\mathbb{F}$ a finite field with $|\mathbb{F}| = q$ elements, and $S \subseteq \mathbb{F}^n$ a subspace. Then for every integer $t \leqslant \min\{\dim(S), d\}$,*

$$\mathop{\mathbf{P}}_{X \in \mathbb{F}^{d \times n}}[\dim(X(S)) \leqslant t] \leqslant 4q^{-(\dim(S)-t)(d-t)}. \tag{41}$$

In particular, for every integer $T \leqslant \min\{\dim(S), d\}$,

$$\mathop{\mathbf{E}}_{X \in \mathbb{F}^{d \times n}} q^{T - \min\{T, \dim(X(S))\}} \leqslant 1 + 8q^{-(\dim(S)-T+1)(d-T+1)+1}. \tag{42}$$

Proof: Equations (41) and (42) hold trivially for negative t and T, respectively. As a result, we may assume that $t \geqslant 0$ and $T \geqslant 0$. Abbreviate $k = \dim(S)$. Fix a basis $v_1, v_2, \ldots, v_k$ for S and extend it to a basis $v_1, v_2, \ldots, v_n$ for $\mathbb{F}^n$. Let $A \in \mathbb{F}^{n \times n}$ be the unique matrix such that $Av_i = e_i$ for each $i = 1, 2, \ldots, n$. In particular, $A(S) = \operatorname{span}\{e_1, e_2, \ldots, e_k\}$. Now, let $X \in \mathbb{F}^{d \times n}$ be uniformly random. Then the rows of XA are independent random variables, each a uniformly random linear combination of the rows of A. Since A is nonsingular of order n, it follows that the rows of XA are independent random vectors in $\mathbb{F}^n$. Put another way, $XA \in \mathbb{F}^{d \times n}$ has the same distribution as X. As a result,

$$\begin{aligned}&\mathbf{P}[\dim(X(S)) \leqslant t] \\ &= \mathbf{P}[\dim(XA(S)) \leqslant t] \\ &= \mathbf{P}[\dim(X(A(S))) \leqslant t] \\ &= \mathbf{P}[\dim(\operatorname{span}\{Xe_1, Xe_2, \ldots, Xe_k\}) \leqslant t] \\ &= \mathbf{P}[\exists B \in \mathscr{S}(\mathbb{F}^d, t) \text{ such that } Xe_1, Xe_2, \ldots, Xe_k \in B] \\ &\leqslant \sum_{B \in \mathscr{S}(\mathbb{F}^d, t)} \mathbf{P}[Xe_1, Xe_2, \ldots, Xe_k \in B], \end{aligned} \tag{43}$$

where the third step uses $A(S) = \operatorname{span}\{e_1, e_2, \ldots, e_k\}$, and the last step applies a union bound. Now

$$\begin{aligned}\mathbf{P}[\dim(X(S)) \leqslant t] &\leqslant \sum_{\mathscr{S}(\mathbb{F}^d, t)} \left(\frac{q^t}{q^d}\right)^k \\ &= \binom{d}{t}_q q^{-k(d-t)} \\ &\leqslant 4q^{t(d-t)} q^{-k(d-t)} \\ &= 4q^{-(k-t)(d-t)},\end{aligned}$$

where the first step is justified by (43) and the fact that $Xe_1, Xe_2, \ldots, Xe_k$ are independent and uniformly random vectors in $\mathbb{F}^d$; the second step applies Fact II.11; and the third step uses Corollary II.14. This settles (41). Now (42) can be verified as follows:

$$\begin{aligned}&\mathbf{E}\, q^{T - \min\{T, \dim(X(S))\}} \\ &\leqslant 1 + \sum_{t=0}^{T-1} q^{T-t}\, \mathbf{P}[\dim(X(S)) = t] \\ &\leqslant 1 + \sum_{t=0}^{T-1} q^{T-t} \cdot 4q^{-(k-t)(d-t)} \\ &= 1 + \sum_{t=1}^{T} q^t \cdot 4q^{-(k-T+t)(d-T+t)} \\ &= 1 + \sum_{t=1}^{T} q^t \cdot 4q^{-(k-T+1)(d-T+1) - (t-1)(d+k+t-2T+1)} \\ &\leqslant 1 + \sum_{t=1}^{\infty} q^t \cdot 4q^{-(k-T+1)(d-T+1) - (t^2-1)} \\ &\leqslant 1 + 4q^{-(k-T+1)(d-T+1)+1} \cdot \frac{q}{q-1} \\ &\leqslant 1 + 8q^{-(k-T+1)(d-T+1)+1},\end{aligned}$$

where the third step is a change of variable, the next-to-last step bounds the series by a geometric series, and the last step is valid due to $q \geqslant 2$. ∎

The previous lemma gives an analogous results for matrices:

Lemma II.21. *Let n, m, d be positive integers, $\mathbb{F}$ a finite field with $|\mathbb{F}| = q$ elements, and $M \in \mathbb{F}^{n \times m}$ a given matrix. Then for every integer $t \leqslant \min\{\operatorname{rk} M, d\}$:*

(i) *$\mathbf{P}[\operatorname{rk}(XM) \leqslant t] \leqslant 4q^{-(\operatorname{rk}(M)-t)(d-t)}$ for a uniformly random matrix $X \in \mathbb{F}^{d \times n}$;*

(ii) *$\mathbf{P}[\operatorname{rk}(MY) \leqslant t] \leqslant 4q^{-(\operatorname{rk}(M)-t)(d-t)}$ for a uniformly random matrix $Y \in \mathbb{F}^{m \times d}$.*

Proof: Let S be the column span of M. Then $\operatorname{rk}(XM) = \dim(X(S))$, and (i) follows from Lemma II.20. For (ii), rewrite the probability of interest as $\mathbf{P}[\operatorname{rk}(Y^{\mathsf{T}}M^{\mathsf{T}}) \leqslant t]$ and apply (i). ∎

H. Communication complexity

An excellent reference on communication complexity is the monograph by Kushilevitz and Nisan [27]. In this overview, we will limit ourselves to key definitions and notation. The *public-coin randomized model*, due to Yao [1], features two players Alice and Bob and a (possibly partial) Boolean function $F\colon X \times Y \to \{-1, 1, *\}$ for finite sets X and Y. Alice is given as input an element $x \in X$, Bob is given $y \in Y$, and their objective is to evaluate $F(x, y)$. To this end, Alice and Bob communicate by sending bits according to a protocol agreed upon in advance. Moreover, they have an unlimited supply of shared random bits which they can use when deciding what message to send at any given point in the protocol. An *ε-error protocol* for F is one which, on every input $(x, y) \in \operatorname{dom} F$, produces the correct answer $F(x, y)$ with probability at least $1 - \varepsilon$. The protocol's behavior on inputs outside $\operatorname{dom} F$ can be arbitrary. The *cost* of a protocol is the total bit length of the messages exchanged by Alice and Bob in the worst-case execution of the protocol. The *ε-error randomized communication complexity of F*, denoted $R_\varepsilon(F)$, is the least cost of an ε-error randomized protocol for F. The standard setting of the error parameter is $\varepsilon = 1/3$, which can be replaced by any other constant in $(0, 1/2)$ with only a constant-factor change in communication cost.

A far-reaching generalization of the randomized model is Yao's *quantum model* [2], where Alice and Bob exchange *quantum* messages. As before, their objective is to evaluate a fixed function $F\colon X \times Y \to \{-1, 1, *\}$ on any given input pair (x, y), where Alice receives as input x and Bob receives y. We allow arbitrary *prior entanglement* at the start of the communication, which is the quantum analogue of shared randomness. A measurement at the end of the protocol produces a one-bit answer, which is interpreted as the protocol output. An *ε-error protocol for F* is required to output, on every input $(x, y) \in \operatorname{dom} F$, the correct value $F(x, y)$ with probability at least $1 - \varepsilon$. As before, the protocol can exhibit arbitrary behavior on inputs outside $\operatorname{dom} F$. The *cost* of a quantum protocol is the total number of quantum bits exchanged in the worst-case execution. The *ε-error quantum communication complexity of F*, denoted $Q^*_\varepsilon(F)$, is the least cost of an ε-error quantum protocol for F. The asterisk in $Q^*_\varepsilon(F)$ indicates that the parties can share arbitrary prior entanglement. As before, the standard setting of the error parameter is $\varepsilon = 1/3$. For a detailed formal description of the quantum model, we refer the reader to [15], [17], [28]. For any protocol Π, quantum or otherwise, we write $\operatorname{cost}(\Pi)$ for the communication cost of Π.

The following theorem, due to Linial and Shraibman [16, Lem. 10], states that the matrix of the acceptance probabilities of a quantum protocol has an efficient factorization with respect to the Frobenius norm. Closely analogous statements were established earlier by Yao [2], Kremer [14], and Razborov [15].

Theorem II.22. *Let X, Y be finite sets. Let P be a quantum protocol (with or without prior entanglement) with cost C qubits and input sets X and Y. Then*

$$\Big[\mathbf{P}[P(x, y) = 1]\Big]_{x \in X, y \in Y} = AB$$

for some real matrices A, B with $\|A\|_{\mathrm{F}} \leqslant 2^C\sqrt{|X|}$ and $\|B\|_{\mathrm{F}} \leqslant 2^C\sqrt{|Y|}$.

Theorem II.22 provides a transition from quantum protocols to matrix factorization, which is by now a standard technique that has been used by various authors in various contexts. Among other things, Theorem II.22 gives the following *approximate trace norm method* for quantum lower bounds; see, e.g., [15, Thm. 5.5]. For the reader's convenience, we state and prove this result in the generality that we require.

Theorem II.23 (Approximate trace norm method)**.** *Let $F\colon X \times Y \to \{-1, 1, *\}$ be a (possibly partial) communication problem. Then*

$$4^{Q^*_\varepsilon(F)} \geqslant \frac{\|M\|_{\Sigma, 2\varepsilon}}{3\sqrt{|X|\,|Y|}},$$

where $M = [F(x, y)]_{x \in X, y \in Y}$ is the characteristic matrix of F.

Proof: Let P be a quantum protocol with prior entanglement that computes F with error ε and cost C. Put

$$\Pi = \Big[\mathbf{P}[P(x, y) = 1]\Big]_{x \in X,\, y \in Y}.$$

Then the matrix $\widetilde{M} = 2\Pi - J$ satisfies $|\widetilde{M}_{x,y}| \leqslant 1$ for all $(x, y) \in X \times Y$ and $|M_{x,y} - \widetilde{M}_{x,y}| \leqslant 2\varepsilon$ for all $(x, y) \in \operatorname{dom} M$. In particular,

$$\|M\|_{\Sigma, 2\varepsilon} \leqslant \|\widetilde{M}\|_\Sigma. \tag{44}$$

On the other hand, Theorem II.22 guarantees the existence of matrices A and B with $AB = \Pi$ and $\|A\|_{\mathrm{F}}\,\|B\|_{\mathrm{F}} \leqslant 4^C\sqrt{|X|\,|Y|}$. Therefore,

$$\begin{aligned}
\|\widetilde{M}\|_\Sigma &= \|2AB - J\|_\Sigma \\
&\leqslant 2\|AB\|_\Sigma + \|J\|_\Sigma \\
&\leqslant 2\|A\|_{\mathrm{F}}\|B\|_{\mathrm{F}} + \|J\|_\Sigma \\
&\leqslant 2 \cdot 4^C\sqrt{|X|\,|Y|} + \|J\|_\Sigma \\
&= 2 \cdot 4^C\sqrt{|X|\,|Y|} + \sqrt{|X|\,|Y|},
\end{aligned} \tag{45}$$

where the third step uses Fact II.8. Equations (44) and (45) give $\|M\|_{\Sigma, 2\varepsilon} \leqslant (2 \cdot 4^C + 1)\sqrt{|X|\,|Y|}$, which implies the claimed lower bound on 4^C. ∎

A *distinguisher* for a communication problem $F\colon X \times Y \to \{-1, 1, *\}$ is a communication protocol Π for which the expected output on every input in $F^{-1}(-1)$ is less than the expected output on every input in $F^{-1}(1)$. We will use the following proposition to convert any distinguisher for F into a communication protocol that computes F.

Proposition II.24. *Let $F\colon X\times Y\to\{-1,1,*\}$ be a (possibly partial) communication problem. Suppose that Π is a cost-c randomized protocol with output ±1 such that*

$$\mathbf{E}[\Pi(x,y)]\leqslant\alpha \qquad \text{for all } (x,y)\in F^{-1}(-1), \tag{46}$$

$$\mathbf{E}[\Pi(x,y)]\geqslant\beta \qquad \text{for all } (x,y)\in F^{-1}(1), \tag{47}$$

where α,β are reals with $-1\leqslant\alpha\leqslant\beta\leqslant1$. Then

$$R_{\frac{1}{2}-\frac{1}{8}(\beta-\alpha)}(F)\leqslant c.$$

Proof: For a real number t, define $\widetilde{\operatorname{sgn}}\,t$ to be 1 if $t\geqslant0$ and -1 if $t<0$. Set $p=|\alpha+\beta|/(2+|\alpha+\beta|)$ and consider the following randomized protocol Π' with input $(x,y)\in X\times Y$: with probability p, Alice and Bob output $-\widetilde{\operatorname{sgn}}(\alpha+\beta)$ without any communication; with the complementary probability $1-p$, they execute the original protocol Π on (x,y) and output its answer. Clearly, Π' has the same cost as Π. On every $(x,y)\in F^{-1}(-1)$,

$$\begin{aligned}\mathbf{E}[\Pi'(x,y)]&\leqslant-p\,\widetilde{\operatorname{sgn}}(\alpha+\beta)+(1-p)\alpha\\&=\frac{-(\alpha+\beta)+2\alpha}{2+|\alpha+\beta|}\\&=\frac{\alpha-\beta}{2+|\alpha+\beta|}\\&\leqslant-\frac{\beta-\alpha}{4},\end{aligned}$$

where the first step uses (46), and the last step uses $-1\leqslant\alpha\leqslant\beta\leqslant1$. Analogously, on every $(x,y)\in F^{-1}(1)$,

$$\begin{aligned}\mathbf{E}[\Pi'(x,y)]&\geqslant-p\,\widetilde{\operatorname{sgn}}(\alpha+\beta)+(1-p)\beta\\&=\frac{-(\alpha+\beta)+2\beta}{2+|\alpha+\beta|}\\&=\frac{\beta-\alpha}{2+|\alpha+\beta|}\\&\geqslant\frac{\beta-\alpha}{4},\end{aligned}$$

where the first step uses (47). We have shown that $\mathbf{E}[\Pi'(x,y)F(x,y)]\geqslant(\beta-\alpha)/4$ on the domain of F, which is another way of saying that Π' computes F with error at most $\frac{1}{2}-\frac{1}{8}(\beta-\alpha)$. ∎

I. Communication problems defined

Let $\mathbb{F}$ be a given field. For nonnegative integers n,m,r with $r\leqslant\min\{n,m\}$, the *rank problem* is the communication problem in which Alice and Bob are given matrices $A,B\in\mathbb{F}^{n\times m}$, respectively, and their objective is to determine whether $\operatorname{rk}(A+B)\leqslant r$. Formally, this problem corresponds to the Boolean function $\operatorname{RANK}_r^{\mathbb{F},n,m}\colon\mathbb{F}^{n\times m}\times\mathbb{F}^{n\times m}\to\{-1,1\}$ given by

$$\operatorname{RANK}_r^{\mathbb{F},n,m}(A+B)=-1 \qquad\Leftrightarrow\qquad \operatorname{rk}(A+B)\leqslant r.$$

We also study the corresponding partial problem $\operatorname{RANK}_{r,R}^{\mathbb{F},n,m}$ for nonnegative integers n,m,r,R with $r<R\leqslant\min\{n,m\}$, defined on $\mathbb{F}^{n\times m}\times\mathbb{F}^{n\times m}$ by

$$\operatorname{RANK}_{r,R}^{\mathbb{F},n,m}(A,B)=\begin{cases}-1 & \text{if } \operatorname{rk}(A+B)=r,\\ 1 & \text{if } \operatorname{rk}(A+B)=R,\\ * & \text{otherwise.}\end{cases}$$

For a positive integer n and a pair of distinct field elements $a,b\in\mathbb{F}$, the *determinant problem* $\operatorname{DET}_{a,b}^{\mathbb{F},n}\colon\mathbb{F}^{n\times n}\times\mathbb{F}^{n\times n}\to\{-1,1,*\}$ is given by

$$\operatorname{DET}_{a,b}^{\mathbb{F},n}(A,B)=\begin{cases}-1 & \text{if } \det(A+B)=a,\\ 1 & \text{if } \det(A+B)=b,\\ * & \text{otherwise.}\end{cases}$$

The *rank versus determinant problem* is a hybrid inspired by the previous two problems. Specifically, for a number $r\in\{0,1,\dots,n-1\}$ and a nonzero field element $a\in\mathbb{F}\setminus\{0\}$, we define $\operatorname{RANKDET}_{r,a}^{\mathbb{F},n}\colon\mathbb{F}^{n\times n}\times\mathbb{F}^{n\times n}\to\{-1,1,*\}$ by

$$\operatorname{RANKDET}_{r,a}^{\mathbb{F},n}(A,B)=\begin{cases}-1 & \text{if } \operatorname{rk}(A+B)=r,\\ 1 & \text{if } \det(A+B)=a,\\ * & \text{otherwise.}\end{cases}$$

Note that $\operatorname{RANKDET}_{r,a}^{\mathbb{F},n}$ is a *subproblem* of both $\operatorname{RANK}_{r,n}^{\mathbb{F},n,n}$ and $\operatorname{DET}_{0,a}^{\mathbb{F},n}$, in the sense that the domain of $\operatorname{RANKDET}_{r,a}^{\mathbb{F},n}$ is a subset of the domain of each of these other two problems and it agrees on its domain with those problems.

Consider now the setting where Alice is given an m-dimensional subspace $S\subseteq\mathbb{F}^n$ and Bob is given an ℓ-dimensional subspace $T\subseteq\mathbb{F}^n$, for some nonnegative integers n,m,ℓ with $\max\{m,\ell\}\leqslant n$. In the *subspace intersection problem* with parameter d, Alice and Bob need to determine whether $S\cap T$ has dimension at least d. In the *subspace sum problem*, they need to determine whether $S+T$ has dimension at most d. Formally, these problems correspond to the Boolean functions $\operatorname{INTERSECT}_d^{\mathbb{F},n,m,\ell}$ and $\operatorname{SUM}_d^{\mathbb{F},n,m,\ell}$ that are defined on $\mathscr{S}(\mathbb{F}^n,m)\times\mathscr{S}(\mathbb{F}^n,\ell)$ by

$$\begin{aligned}\operatorname{INTERSECT}_d^{\mathbb{F},n,m,\ell}(S,T)=-1 \quad&\Leftrightarrow\quad \dim(S\cap T)\geqslant d,\\ \operatorname{SUM}_d^{\mathbb{F},n,m,\ell}(S,T)=-1 \quad&\Leftrightarrow\quad \dim(S+T)\leqslant d.\end{aligned}$$

Their partial counterparts $\operatorname{INTERSECT}_{d_1,d_2}^{\mathbb{F},n,m,\ell}$ and $\operatorname{SUM}_{d_1,d_2}^{\mathbb{F},n,m,\ell}$, for any pair of distinct integers d_1,d_2, are defined on $\mathscr{S}(\mathbb{F}^n,m)\times\mathscr{S}(\mathbb{F}^n,\ell)$ by

$$\operatorname{INTERSECT}_{d_1,d_2}^{\mathbb{F},n,m,\ell}(S,T)=\begin{cases}-1 & \text{if } \dim(S\cap T)=d_1,\\ 1 & \text{if } \dim(S\cap T)=d_2,\\ * & \text{otherwise,}\end{cases}$$

$$\operatorname{SUM}_{d_1,d_2}^{\mathbb{F},n,m,\ell}(S,T)=\begin{cases}-1 & \text{if } \dim(S+T)=d_1,\\ 1 & \text{if } \dim(S+T)=d_2,\\ * & \text{otherwise.}\end{cases}$$

These partial functions are well-defined for any d_1,d_2 with $d_1\neq d_2$. Their communication complexity, however, is zero unless both d_1 and d_2 have meaningful values for the problem in question. Specifically, one must have $d_1,d_2\in[\max\{m,\ell\},\min\{m+\ell,n\}]$ for the subspace sum problem and $d_1,d_2\in[\max\{0,m+\ell-n\},\min\{m,\ell\}]$ for the subspace intersection problem. We record this simple fact as a proposition below.

Proposition II.25. *Let $\mathbb{F}$ be a field. Let n, m, ℓ be nonnegative integers with $\max\{m,\ell\} \leqslant n$. Then:*

(i) *there exist $S \in \mathscr{S}(\mathbb{F}^n, m)$ and $T \in \mathscr{S}(\mathbb{F}^n, \ell)$ with $\dim(S+T) = d$ if and only if d is an integer with $\max\{m,\ell\} \leqslant d \leqslant \min\{m+\ell, n\}$;*

(ii) *there exist $S \in \mathscr{S}(\mathbb{F}^n, m)$ and $T \in \mathscr{S}(\mathbb{F}^n, \ell)$ with $\dim(S \cap T) = d$ if and only if d is an integer with $\max\{0, m+\ell-n\} \leqslant d \leqslant \min\{m,\ell\}$.*

Proof: (i) For any subspaces $S, T \subseteq \mathbb{F}^n$, we have the trivial bounds $\max\{\dim(S), \dim(T)\} \leqslant \dim(S+T) \leqslant \min\{\dim(S)+\dim(T), n\}$. This proves the "only if" part of (i). For the converse, let d be any integer with $\max\{m,\ell\} \leqslant d \leqslant \min\{m+\ell, n\}$. Then the sets $A = \{1, 2, \ldots, m\}$ and $B = \{d-\ell+1, \ldots, d-1, d\}$ satisfy $A, B \subseteq \{1, 2, \ldots, n\}$ (because $\ell \leqslant d \leqslant n$) and $A \cup B = \{1, 2, \ldots, d\}$ (because $m \leqslant d \leqslant m+\ell$). As a result, $\operatorname{span}\{e_1, e_2, \ldots, e_m\}$ and $\operatorname{span}\{e_{d-\ell+1}, \ldots, e_{d-1}, e_d\}$ are a pair of subspaces in $\mathbb{F}^n$ of dimension m and ℓ, respectively, whose sum has dimension d.

(ii) Recall that $\dim(S \cap T) = \dim(S) + \dim(T) - \dim(S+T)$ for any subspaces S, T. As a result,

$$\begin{aligned}
&\{\dim(S \cap T) : S \in \mathscr{S}(\mathbb{F}^n, m), T \in \mathscr{S}(\mathbb{F}^n, \ell)\} \\
&= \{m + \ell - \dim(S+T) : S \in \mathscr{S}(\mathbb{F}^n, m), T \in \mathscr{S}(\mathbb{F}^n, \ell)\} \\
&= \{m + \ell - d : d \in \mathbb{Z},\ \max\{m,\ell\} \leqslant d \leqslant \min\{m+\ell, n\}\} \\
&= \{\max\{0, m+\ell-n\}, \ldots, \min\{m,\ell\} - 1, \min\{m,\ell\}\},
\end{aligned}$$

where the second step uses (i). ∎

Let $F\colon X \times Y \to \{-1, 1, *\}$ and $F'\colon X' \times Y' \to \{-1, 1, *\}$ be (possibly partial) communication problems. A *communication-free reduction from F to F'* is a pair of mappings $\alpha\colon X \to X'$ and $\beta\colon Y \to Y'$ such that $F(x,y) = F'(\alpha(x), \beta(y))$ for all $(x,y) \in \operatorname{dom} F$. We indicate the existence of a communication-free reduction from F to F' by writing $F' \succeq F$. In this case, it is clear that the communication complexity of F' in any given model is bounded from below by the communication complexity of F in the same model.

Proposition II.26. *Let n, m, ℓ, r, R be integers with $0 \leqslant r < R \leqslant \min\{m,\ell\}$ and $\max\{m,\ell\} \leqslant n$. Then*

$$\mathrm{INTERSECT}^{\mathbb{F},n,m,\ell}_{r,R} \succeq \mathrm{INTERSECT}^{\mathbb{F},n-r,m-r,\ell-r}_{0,R-r}.$$

Proof: Consider the injective linear map $\varphi\colon \mathbb{F}^{n-r} \to \mathbb{F}^n$ that takes any vector and extends it with r zero components to obtain a vector in $\mathbb{F}^n$. Given arbitrary subspaces $S, T \subseteq \mathbb{F}^{n-r}$ of dimension $m-r$ and $\ell-r$, respectively, define $S' = \operatorname{span}(\varphi(S) \cup \{e_{n-r+1}, \ldots, e_{n-1}, e_n\})$ and $T' = \operatorname{span}(\varphi(T) \cup \{e_{n-r+1}, \ldots, e_{n-1}, e_n\})$. Then clearly

$$\begin{aligned}
&\dim(S' \cap T') \\
&\qquad = \dim(S') + \dim(T') - \dim(S'+T') \\
&\qquad = \dim(S) + r + \dim(T) + r - \dim(S+T) - r \\
&\qquad = \dim(S) + \dim(T) - \dim(S+T) + r \\
&\qquad = \dim(S \cap T) + r,
\end{aligned}$$

whence the reduction $\mathrm{INTERSECT}^{\mathbb{F},n-r,m-r,\ell-r}_{0,R-r}(S,T) = \mathrm{INTERSECT}^{\mathbb{F},n,m,\ell}_{r,R}(S',T')$. ∎

III. The matrix rank problem

In this section, we prove a tight lower bound on the randomized and quantum communication complexity of the rank problem. As discussed in the introduction, we obtain this lower bound by constructing a dual matrix Φ with certain properties, namely, low spectral norm, low ℓ_1 norm, and high correlation with the characteristic matrix of the rank problem. We start in Section III-A by analyzing the probabilities P_n that arise in the recurrence relation for the Γ_n function. The latter plays an important role in our proof and is studied in Section III-B. Section III-C constructs a univariate dual object φ defined on $\{0, 1, \ldots, n\}$ and studies its analytic and metric properties. We build on φ to construct a dual matrix E_φ in Section III-D, and discuss how the properties of φ give rise to analogous properties of E_φ. Sections III-E and III-F establish lower bounds for the approximate trace norm of the characteristic matrix and the communication complexity of the rank problem, with $\Phi = E_\varphi$ used as the dual witness. We prove a matching communication upper bound in Section III-G. Section III-H concludes our study of the rank problem with an application to streaming complexity.

Throughout this section, the underlying field is $\mathbb{F}_q$ for an arbitrary prime power q. The root of unity ω and the notation ω^x for $x \in \mathbb{F}_q$ are as defined in Section II-D.

A. The P_n function

The P_n function, defined next, conveys useful information about random nonsingular matrices of order n over a given field.

Definition III.1. Let $n \geqslant 1$ be a given integer. For nonnegative integers $s, t, r \in \{0, 1, \ldots, n\}$, define $P_n(s,t,r)$ to be the probability that the upper-left $s \times t$ quadrant of a uniformly random nonsingular matrix in $\mathbb{F}_q^{n \times n}$ has rank r:

$$P_n(s,t,r) = \mathop{\mathbf{P}}_{X \in \mathscr{M}_n}[\operatorname{rk}(I_s X I_t) = r]. \tag{48}$$

To derive a closed-form expression for P_n, we essentially need to count the number of ways to complete a given $s \times t$ matrix of rank r to a nonsingular matrix of order n. We break this counting task into two steps, where the first step is to count the number of completions of an $s \times t$ matrix of rank r to an $s \times n$ matrix of rank s.

Lemma III.2. *Let s, t, r, m be nonnegative integers with $r \leqslant \min\{s,t\}$. Let $A \in \mathscr{M}_r^{s,t}$ be given. Then the number of matrices $B \in \mathbb{F}_q^{s \times m}$ for which $\operatorname{rk}\begin{bmatrix} A & B \end{bmatrix} = s$ is*

$$q^{rm} \left| \mathscr{M}_{s-r}^{s-r,m} \right|.$$

Proof: If $r = 0$, then $\operatorname{rk}\begin{bmatrix} A & B \end{bmatrix} = \operatorname{rk} B$. As a result, $\operatorname{rk}\begin{bmatrix} A & B \end{bmatrix} = s$ if and only if $B \in \mathscr{M}_s^{s,m}$. Therefore, the lemma holds in this case. In what follows, we consider $r \geqslant 1$, which forces s and t to be positive integers.

Define the matrices A' and A'' to be the top r rows of A and the bottom $s-r$ rows of A, respectively. We first consider

the possibility when A'' is zero or empty. Here, the column span of A' is necessarily all of $\mathbb{F}_q^r$. Given an $s \times m$ matrix B, partition it into B' and B'' conformably with the partition of A. Then

$$\begin{aligned}\operatorname{rk}\begin{bmatrix}A & B\end{bmatrix} &= \operatorname{rk}\begin{bmatrix}A' & B'\\ 0 & B''\end{bmatrix}\\ &= \operatorname{rk}\begin{bmatrix}A' & 0\\ 0 & B''\end{bmatrix}\\ &= \operatorname{rk}(A') + \operatorname{rk}(B'')\\ &= r + \operatorname{rk}(B'').\end{aligned}$$

Thus, $\begin{bmatrix}A & B\end{bmatrix}$ has rank s if and only if $\operatorname{rk}(B'') = s - r$. This means that there are $|\mathscr{M}_{s-r}^{s-r,m}|$ ways to choose B'', and independently q^{rm} ways to choose B', such that $\operatorname{rk}\begin{bmatrix}A & B\end{bmatrix} = s$.

It remains to examine the case of a general matrix A of rank $r \geqslant 1$. Let V be an invertible matrix such that the bottom $s-r$ rows of VA are zero. Let $\mathscr{M}$ be the set of $s \times m$ matrices M for which $\operatorname{rk}\begin{bmatrix}VA & M\end{bmatrix} = s$. Then $\operatorname{rk}\begin{bmatrix}A & B\end{bmatrix} = s$ if and only if $VB \in \mathscr{M}$. In particular, the number of matrices B for which $\operatorname{rk}\begin{bmatrix}A & B\end{bmatrix} = s$ is $|\mathscr{M}|$. Since $|\mathscr{M}| = q^{rm}\,|\mathscr{M}_{s-r}^{s-r,m}|$ by the previous paragraph, we are done. ∎

We now derive an exact expression for P_n and establish its relevant algebraic and analytic properties.

Lemma III.3. *Let $n \geqslant 1$ be a given integer. Then for all $s,t,r \in \{0,1,\dots,n\}$:*

(i) $P_n(s,t,r) = 0$ if $r > \min\{s,t\}$ or $r < s+t-n$;
(ii) $P_n(s,t,r) = q^{r(n-t)}|\mathscr{M}_r^{s,t}|\,|\mathscr{M}_{s-r}^{s-r,n-t}|/((q^n-1)(q^n-q)\cdots(q^n-q^{s-1}))$;
(iii) for any fixed values of n,s,r, the quantity $P_n(s,t,r)$ as a function of $t \in \{0,1,\dots,n\}$ is a polynomial in q^{-t} of degree at most s;
(iv) $P_n(s,t,r) \leqslant 16q^{-(s-r)(t-r)}$.

Proof: (i) Since the quadrant of interest is an $s \times t$ matrix, the first inequality is trivial. For the second inequality, observe that the matrix $I_s X I_t$ in the defining equation (48) satisfies $\operatorname{rk}(I_s X I_t) \geqslant \operatorname{rk} I_s + \operatorname{rk}(X I_t) - n = s + t - n$ by Fact II.1.

(ii) If $r > \min\{s,t\}$, then the left-hand side and right-hand side of (ii) both vanish due to (i) and the definition of $\mathscr{M}_r^{s,t}$. We now treat the case $r \leqslant \min\{s,t\}$. Letting $\mathscr{M}$ stand for the set of nonsingular matrices of order n whose upper-left $s \times t$ quadrant has rank r, we have

$$P_n(s,t,r) = \frac{|\mathscr{M}|}{|\mathscr{M}_n|}. \tag{49}$$

A matrix in $\mathscr{M}$ can be chosen by the following three-step process: choose a matrix in $\mathscr{M}_r^{s,t}$ for the upper-left quadrant; extend the quadrant to a matrix in $\mathscr{M}_s^{s,n}$, which by Lemma III.2 can be done in $q^{r(n-t)}\,|\mathscr{M}_{s-r}^{s-r,n-t}|$ ways; and finally add $n-s$ rows to obtain an invertible matrix, which can be done in $(q^n-q^s)(q^n-q^{s+1})\cdots(q^n-q^{n-1})$ ways. Altogether, we obtain

$$\begin{aligned}|\mathscr{M}| = |\mathscr{M}_r^{s,t}| \cdot q^{r(n-t)}\,|\mathscr{M}_{s-r}^{s-r,n-t}|\\ \times (q^n-q^s)(q^n-q^{s+1})\cdots(q^n-q^{n-1}),\end{aligned}$$

whereas Proposition II.17 gives

$$|\mathscr{M}_n| = (q^n-1)(q^n-q)\cdots(q^n-q^{n-1}).$$

Making these substitutions in (49) completes the proof.

(iii) We claim that for all $s,t,r \in \{0,1,\dots,n\}$,

$$\begin{aligned}P_n(s,t,r) = q^{r(n-t)}\binom{s}{r}_q (q^t-1)(q^t-q)\cdots(q^t-q^{r-1})\\ \times \frac{(q^{n-t}-1)(q^{n-t}-q)\cdots(q^{n-t}-q^{s-r-1})}{(q^n-1)(q^n-q)\cdots(q^n-q^{s-1})}.\end{aligned} \tag{50}$$

Indeed, in the case when $r > \min\{s,t\}$ or $r < s+t-n$, the right-hand side vanishes and therefore the equality holds due to (i). In the complementary case, Proposition II.17 gives closed-form expressions for $|\mathscr{M}_r^{s,t}|$ and $|\mathscr{M}_{s-r}^{s-r,n-t}|$ which, when substituted in (ii), result in (50). This settles (50) for all $s,t,r \in \{0,1,\dots,n\}$.

Rewrite (50) to obtain

$$\begin{aligned}P_n(s,t,r) = q^{rn}\binom{s}{r}_q (1-q^{-t})(1-q^{-t+1})\cdots(1-q^{-t+r-1})\\ \times \frac{(q^{n-t}-1)(q^{n-t}-q)\cdots(q^{n-t}-q^{s-r-1})}{(q^n-1)(q^n-q)\cdots(q^n-q^{s-1})}.\end{aligned} \tag{51}$$

Now, fix n,s,r arbitrarily. If $r \leqslant s$, then (51) makes it clear that $P_n(s,t,r)$ is a polynomial in q^{-t} of degree at most $r + (s-r) = s$. If $r > s$, then $P_n(s,t,r)$ is identically zero and thus trivially a polynomial in q^{-t} of degree at most s.

(iv) For $r > s$, we have $P_n(s,t,r) = 0$ by (i) and therefore the claimed upper bound holds trivially. In the complementary case, simplify (50) to obtain

$$\begin{aligned}&P_n(s,t,r)\\ &\quad\leqslant q^{r(n-t)}\binom{s}{r}_q q^{tr} \cdot \frac{q^{(n-t)(s-r)}}{(q^n-1)(q^n-q)\cdots(q^n-q^{s-1})}\\ &\quad\leqslant q^{r(n-t)}\binom{s}{r}_q q^{tr} \cdot 4q^{(n-t)(s-r)}q^{-ns}\\ &\quad\leqslant q^{r(n-t)} \cdot 4q^{r(s-r)}q^{tr} \cdot 4q^{(n-t)(s-r)}q^{-ns}\\ &\quad= 16q^{-(s-r)(t-r)},\end{aligned}$$

where the second and third steps apply Proposition II.13 and Corollary II.14, respectively. ∎

B. The Γ_n function

A basic building block in our construction is the characteristic function of matrices in $\mathbb{F}_q^{n\times n}$ of a given rank. Its Fourier spectrum is best understood in terms of what we call the Γ_n function.

Definition III.4. Let $n \geqslant 1$ be a given integer. For $s,t \in \{0,1,\dots,n\}$, define

$$\Gamma_n(s,t) = \mathop{\mathbf{E}}_{\substack{\operatorname{rk} A = s\\ \operatorname{rk} B = t}} \omega^{\langle A,B\rangle},$$

where the expectation is taken with respect to the uniform distribution on $\mathscr{M}_s^{n,n} \times \mathscr{M}_t^{n,n}$.

Sun and Wang [5] studied the Fourier spectrum of the non-singularity function on $\mathbb{F}_q^{n\times n}$, defined to be 1 on nonsingular matrices and 0 otherwise. In our notation, they established the following result.

Lemma III.5. *For any integers $n \geqslant 1$ and $r \in \{0,1,\ldots,n\}$,*

$$\Gamma_n(n,r) = \frac{(-1)^r q^{\binom{r}{2}}}{(q^n-1)(q^n-q)\cdots(q^n-q^{r-1})}.$$

The proof of Sun and Wang [5] is stated for fields $\mathbb{F}_p$ with prime p, but their analysis readily extends to fields of cardinality a prime power. In the full version of this paper [29], we prove Lemma III.5 from scratch in our desired generality, using a simpler proof than that of [5].

Our next lemma collects crucial properties of $\Gamma_n(s,t)$ for general values of s,t.

Lemma III.6. *Let $n \geqslant 1$ be a given integer. Then for all $s,t \in \{0,1,\ldots,n\}$:*

(i) $|\Gamma_n(s,t)| \leqslant 1$;
(ii) $\Gamma_n(s,t) = \Gamma_n(t,s)$;
(iii) $\Gamma_n(s,t) = \sum_{r=0}^{n} P_n(s,t,r)\Gamma_n(n,r)$;
(iv) for n,s fixed, $\Gamma_n(s,t)$ as a function of $t \in \{0,1,\ldots,n\}$ is a polynomial in q^{-t} of degree at most s;
(v) $|\Gamma_n(s,t)| \leqslant 128q^{-st/2}$.

Proof: (i) Using $|\omega| = 1$ and the triangle inequality,

$$|\Gamma_n(s,t)| = \left|\mathop{\mathbf{E}}_{A,B} \omega^{\langle A,B\rangle}\right| \leqslant \mathop{\mathbf{E}}_{A,B}\left|\omega^{\langle A,B\rangle}\right| = 1.$$

(ii) The symmetry of Γ_n follows from the independence of A and B in the defining equation for Γ_n, and the symmetry of the inner product over $\mathbb{F}_q$.

(iii) We have:

$$\begin{aligned}
&\Gamma_n(s,t)\\
&\quad= \mathop{\mathbf{E}}_{\substack{A\in\mathscr{M}_s^{n,n}\\ B\in\mathscr{M}_t^{n,n}}} \omega^{\langle A,B\rangle}\\
&\quad= \mathop{\mathbf{E}}_{X,Y,Z_1,Z_2,W\in\mathscr{M}_n} \omega^{\langle XI_sY,Z_1Z_2I_tW\rangle}\\
&\quad= \mathop{\mathbf{E}}_{X,Y,Z_1,Z_2,W\in\mathscr{M}_n} \omega^{\langle XI_sYW^{\mathsf{T}}I_tZ_2^{\mathsf{T}},Z_1\rangle}\\
&\quad= \mathop{\mathbf{E}}_{X,U,Z_1,Z_2\in\mathscr{M}_n} \omega^{\langle X(I_sUI_t)Z_2^{\mathsf{T}},Z_1\rangle}\\
&\quad= \sum_{r=0}^{n} \mathop{\mathbf{P}}_{U\in\mathscr{M}_n}[\mathrm{rk}(I_sUI_t)=r]\\
&\qquad\times \mathop{\mathbf{E}}_{X,U,Z_1,Z_2\in\mathscr{M}_n}\left[\omega^{\langle X(I_sUI_t)Z_2^{\mathsf{T}},Z_1\rangle} \;\middle|\; \mathrm{rk}(I_sUI_t)=r\right]\\
&\quad= \sum_{r=0}^{n} \mathop{\mathbf{P}}_{U\in\mathscr{M}_n}[\mathrm{rk}(I_sUI_t)=r] \mathop{\mathbf{E}}_{\substack{B\in\mathscr{M}_r^{n,n}\\ Z_1\in\mathscr{M}_n}} \omega^{\langle B,Z_1\rangle}\\
&\quad= \sum_{r=0}^{n} P_n(s,t,r)\Gamma_n(n,r),
\end{aligned}$$

where the first step restates the definition of Γ_n, the second step uses Proposition II.19, the third step applies Fact II.2(ii), the fourth and sixth steps again use Proposition II.19, and the last step is immediate from the definitions of P_n and Γ_n.

(iv) Immediate from (iii) and Lemma III.3(iii).

(v) We have:

$$\begin{aligned}
|\Gamma_n(s,t)| &= \left|\sum_{r=0}^{n} P_n(s,t,r)\Gamma_n(n,r)\right|\\
&\leqslant \sum_{r=0}^{n} P_n(s,t,r)|\Gamma_n(n,r)|\\
&= \sum_{r=\max\{0,s+t-n\}}^{n} P_n(s,t,r)\,|\Gamma_n(n,r)|\\
&\leqslant \sum_{r=\max\{0,s+t-n\}}^{n} 16q^{-(s-r)(t-r)}\\
&\qquad\times \frac{q^{\binom{r}{2}}}{(q^n-1)(q^n-q)\cdots(q^n-q^{r-1})}\\
&\leqslant \sum_{r=\max\{0,s+t-n\}}^{n} 64q^{-(s-r)(t-r)+\binom{r}{2}-nr}\\
&\leqslant 128q^{-st/2},
\end{aligned}$$

where the first step appeals to (iii), the third step is valid by Lemma III.3(i), the fourth step uses Lemma III.3(iv) and Lemma III.5, the fifth step applies Proposition II.13, and the last step which completes the proof is justified by the following claim. ∎

Claim III.7. *For any integers $n \geqslant 1$ and $s,t \in \{0,1,\ldots,n\}$,*

$$\sum_{r=\max\{0,s+t-n\}}^{\infty} q^{-(s-r)(t-r)+\binom{r}{2}-nr} \leqslant 2q^{-st/2}. \tag{52}$$

Proof: The exponent of q on the left-hand side of (52) is given by the function

$$A(r) = -(s-r)(t-r) + \binom{r}{2} - nr \tag{53}$$

$$= -st - \frac{1}{2}\left(r+n-s-t+\frac{1}{2}\right)^2 + \frac{1}{2}\left(n-s-t+\frac{1}{2}\right)^2. \tag{54}$$

The first equality shows that $A(r)$ is always an integer, whereas the second shows that $A(r)$ is a strictly decreasing function in the variable $r \in [\max\{0,s+t-n\},\infty)$. These two facts lead to

$$A(\max\{0,s+t-n\}+i) \leqslant A(\max\{0,s+t-n\}) - i, \quad i = 0,1,2,\ldots. \tag{55}$$

We will now prove that

$$A(\max\{0,s+t-n\}) \leqslant -\frac{st}{2}. \tag{56}$$

There are two cases to consider. If $s+t \leqslant n$, then $A(\max\{0,s+t-n\}) = A(0) = -st$ and therefore (56) holds.

The complementary case $s+t \geqslant n+1$ is more challenging. Here, we have

$$A(\max\{0, s+t-n\}) = A(s+t-n) \leqslant -st + \frac{1}{2}\left(n-s-t+\frac{1}{2}\right)^2,$$

where the second step uses (54). Thus, the proof of (56) will be complete once we show that

$$\left(n-s-t+\frac{1}{2}\right)^2 - st \leqslant 0. \tag{57}$$

To prove (57), suppose that of all pairs $(s,t) \in \{0,1,\dots,n\}^2$ with $s+t \geqslant n+1$, the left-hand side of (57) is maximized at a pair (s^*,t^*). By symmetry, we may assume that $s^* \leqslant t^*$. If we had $t^* \leqslant n-1$, then it would follow that $s^* \geqslant 2$ (due to the requirement that $s^*+t^* \geqslant n+1$); as a result, the left-hand side of (57) would be larger for the pair $(s,t) = (s^*-1, t^*+1)$ than it is for the pair $(s,t) = (s^*,t^*)$, an impossibility. Therefore, $t^* = n$. In addition, we have $s^* \geqslant 1$ (due to the requirement that $s^*+t^* \geqslant n+1$). Evaluating the right-hand side of (57) at this pair $(s^*,t^*) = (s^*,n)$, we obtain $(s^*-\frac{1}{2})^2 - s^*n$, which is clearly negative due to $s^* \in \{1,2,\dots,n\}$. This completes the proof of (57) and therefore that of (56).

Now,

$$\begin{aligned}
&\sum_{r=\max\{0,s+t-n\}}^{\infty} q^{-(s-r)(t-r)+\binom{r}{2}-nr} \\
&\qquad = \sum_{r=\max\{0,s+t-n\}}^{\infty} q^{A(r)} \\
&\qquad = \sum_{i=0}^{\infty} q^{A(\max\{0,s+t-n\}+i)} \\
&\qquad \leqslant q^{A(\max\{0,s+t-n\})} \sum_{i=0}^{\infty} q^{-i} \\
&\qquad \leqslant q^{-st/2} \cdot \frac{q}{q-1},
\end{aligned}$$

where the first step uses the definition of $A(r)$, the third step applies (55), and the final step appeals to (56) and a geometric series. Since $q \geqslant 2$, this completes the proof of (52). ∎

C. Univariate dual object

Our construction of the univariate dual object is based on the Cauchy binomial theorem along with a certain "correcting" polynomial ζ. The next lemma presents ζ as parametrized by two numbers ℓ and m and gives its basic properties.

Lemma III.8. *Let n,k,ℓ,m be nonnegative integers such that $\ell+m \leqslant k < n$. Define a univariate polynomial ζ by*

$$\zeta(t) = \prod_{i=0}^{\ell-1} \frac{t-q^{-i}}{q^{-n}-q^{-i}} \cdot \prod_{i=k-m}^{k-1} \frac{t-q^{-i}}{q^{-n}-q^{-i}} \times \prod_{i=k+1}^{n-1} \frac{t-q^{-i}}{q^{-n}-q^{-i}}. \tag{58}$$

Then:

(i) $\zeta(q^{-n}) = 1$;
(ii) $\operatorname{sgn}\zeta(q^{-k}) = (-1)^{n-k-1}$;
(iii) $\zeta(q^{-r}) = 0$ for $r \in \{0,1,\dots,n\} \setminus (\{\ell,\ell+1,\dots,k-m-1\} \cup \{k,n\})$;
(iv) $\deg\zeta = n+\ell+m-k-1$;
(v) $|\zeta(q^{-r})| \leqslant 4q^{-r(n-k+m-1)+\binom{n}{2}-k-\binom{k-m}{2}}$ for $r \in \{\ell,\ell+1,\dots,k-m-1\}$.

Proof: Items (i), (iii), and (iv) are immediate from the defining equation for ζ. Item (ii) holds because for $t = q^{-k}$, the first and second products in (58) contain only positive factors, whereas the third product contains exactly $n-k-1$ factors all of which are negative. For (v),

$$\begin{aligned}
&|\zeta(q^{-r})| \\
&= \left|\prod_{i=0}^{\ell-1} \frac{q^{-r}-q^{-i}}{q^{-n}-q^{-i}} \cdot \prod_{i=k-m}^{k-1} \frac{q^{-r}-q^{-i}}{q^{-n}-q^{-i}} \cdot \prod_{i=k+1}^{n-1} \frac{q^{-r}-q^{-i}}{q^{-n}-q^{-i}}\right| \\
&= \prod_{i=0}^{\ell-1} \frac{1-q^{i-r}}{1-q^{-(n-i)}} \cdot \prod_{i=k-m}^{k-1} \frac{q^{i-r}-1}{1-q^{-(n-i)}} \cdot \prod_{i=k+1}^{n-1} \frac{q^{i-r}-1}{1-q^{-(n-i)}} \\
&\leqslant \prod_{i=0}^{\ell-1} \frac{1}{1-q^{-(n-i)}} \cdot \prod_{i=k-m}^{k-1} \frac{q^{i-r}}{1-q^{-(n-i)}} \cdot \prod_{i=k+1}^{n-1} \frac{q^{i-r}}{1-q^{-(n-i)}}.
\end{aligned}$$

The product of the numerators in the last expression is $q^{-r(n-k+m-1)+\binom{n}{2}-k-\binom{k-m}{2}}$, whereas the product of the denominators is at least $1/4$ by Proposition II.13. ∎

With ζ in hand, we are now in a position to construct the promised univariate dual object φ. The properties of φ established in the lemma below will give rise to analogous properties in the dual matrix E_φ.

Lemma III.9. *Let n,k,ℓ,m be nonnegative integers such that $\ell+m \leqslant k < n$. Then there is a function $\varphi\colon \{0,1,\dots,n\} \to \mathbb{R}$ such that:*

(i) $\varphi(n) = 1$;
(ii) $\varphi(k) < 0$;
(iii) $\varphi(r) = 0$ for $r \in \{0,1,\dots,n\} \setminus (\{\ell,\ell+1,\dots,k-m-1\} \cup \{k,n\})$;
(iv) $\sum_{r=0}^{n} \varphi(r)\xi(q^{-r}) = 0$ for every polynomial ξ of degree at most $k-\ell-m$;
(v) $\sum_{r\in\{0,\dots,n\}\setminus\{k,n\}} |\varphi(r)| \leqslant 32q^{-m-1}$.

Proof: Define

$$\varphi(r) = \binom{n}{r}_q (-1)^{r-n} q^{\binom{r}{2}-\binom{n}{2}} \zeta(q^{-r}),$$

where ζ is the univariate polynomial from Lemma III.8. Then items (i)–(iii) are immediate from the corresponding items (i)–(iii) of Lemma III.8.

For (iv), fix a univariate polynomial ξ of degree at most $k-\ell-m$. In view of Lemma III.8(iv), the product of ζ and ξ has degree less than n. As a result, the Cauchy binomial theorem (Corollary II.16) implies that

$$\sum_{r=0}^{n} \varphi(r)\xi(q^{-r})$$

$$= (-1)^{-n} q^{-\binom{n}{2}} \sum_{r=0}^{n} \binom{n}{r}_q (-1)^r q^{\binom{r}{2}} \zeta(q^{-r}) \xi(q^{-r})$$
$$= 0.$$

For (v), fix any $r \in \{\ell, \ell+1, \dots, k-m-1\}$. Then

$$\begin{aligned} |\varphi(r)| &= \binom{n}{r}_q q^{\binom{r}{2}-\binom{n}{2}} |\zeta(q^{-r})| \\ &\leqslant 4q^{r(n-r)} \cdot q^{\binom{r}{2}-\binom{n}{2}} \cdot 4q^{-r(n-k+m-1)+\binom{n}{2}-k-\binom{k-m}{2}} \\ &= 16q^{-\binom{k-m-r+1}{2}-m}, \end{aligned} \tag{59}$$

where in the second step we bound the q-binomial coefficient via Corollary II.14 and $|\zeta(q^{-r})|$ via Lemma III.8(v). Now

$$\begin{aligned} \sum_{r \in \{0,\dots n\} \setminus \{k,n\}} |\varphi(r)| &= \sum_{r=\ell}^{k-m-1} |\varphi(r)| \\ &\leqslant \sum_{r=\ell}^{k-m-1} 16q^{-\binom{k-m-r+1}{2}-m} \\ &\leqslant \sum_{i=2}^{\infty} 16q^{-\binom{i}{2}-m} \\ &\leqslant \frac{16q^{-m-1}}{1-\frac{1}{q}}, \end{aligned}$$

where the first step is valid by (iii), the second step uses (59), and the fourth step uses a geometric series with $\binom{i}{2} \geqslant i-1$ for $i \geqslant 2$. Since $q \geqslant 2$, this completes the proof of (v). ∎

D. From univariate dual objects to dual matrices

En route to the main result of this section, we now show how to convert a univariate dual object φ, such as the one constructed in Lemma III.9, into a dual matrix E_φ.

Definition III.10. Let $n \geqslant 1$ be a given integer. For $r = 0, 1, \dots, n$, define E_r to be the matrix with rows and columns indexed by matrices in $\mathbb{F}_q^{n \times n}$, and entries given by

$$(E_r)_{A,B} = \begin{cases} q^{-n^2} |\mathscr{M}_r^{n,n}|^{-1} & \text{if } \operatorname{rk}(A+B) = r, \\ 0 & \text{otherwise.} \end{cases}$$

For a function $\varphi\colon \{0, 1, \dots, n\} \to \mathbb{R}$, define

$$E_\varphi = \sum_{r=0}^{n} \varphi(r) E_r.$$

As one would expect, the metric and analytic properties of E_φ are closely related to those of φ.

Lemma III.11 (Metric properties of E_φ)**.** *Let $n \geqslant 1$ be an integer and $\varphi\colon \{0, 1, \dots, n\} \to \mathbb{R}$ a given function. Then*

$$\sum_{A,B:\operatorname{rk}(A+B)=r} (E_\varphi)_{A,B} = \varphi(r), \qquad r = 0, 1, \dots, n, \tag{60}$$

$$\sum_{A,B:\operatorname{rk}(A+B)=r} |(E_\varphi)_{A,B}| = |\varphi(r)|, \quad r = 0, 1, \dots, n. \tag{61}$$

In particular,

$$\|E_\varphi\|_1 = \|\varphi\|_1.$$

Proof: Recall that for any fixed matrix $A \in \mathbb{F}_q^{n \times n}$, the mapping $B \mapsto A + B$ is a permutation on $\mathbb{F}_q^{n \times n}$. As a result, for any fixed matrix A, there are exactly $|\mathscr{M}_r^{n,n}|$ matrices B such that $\operatorname{rk}(A+B) = r$. Altogether, there are $q^{n^2} |\mathscr{M}_r^{n,n}|$ matrix pairs (A, B) with $\operatorname{rk}(A+B) = r$. With this in mind, Definition III.10 implies the following for each r:

$$\sum_{\operatorname{rk}(A+B)=r} (E_r)_{A,B} = \sum_{\operatorname{rk}(A+B)=r} |(E_r)_{A,B}| = 1. \tag{62}$$

Now for each r,

$$\begin{aligned} \sum_{\operatorname{rk}(A+B)=r} (E_\varphi)_{A,B} &= \sum_{\operatorname{rk}(A+B)=r} \sum_{i=0}^{n} \varphi(i) (E_i)_{A,B} \\ &= \sum_{\operatorname{rk}(A+B)=r} \varphi(r) (E_r)_{A,B} \\ &= \varphi(r), \end{aligned}$$

where the second step uses $(E_i)_{A,B} = 0$ for $i \neq r$, and the final step applies (62). Analogously,

$$\begin{aligned} \sum_{\operatorname{rk}(A+B)=r} |(E_\varphi)_{A,B}| &= \sum_{\operatorname{rk}(A+B)=r} \left| \sum_{i=0}^{n} \varphi(i) (E_i)_{A,B} \right| \\ &= \sum_{\operatorname{rk}(A+B)=r} |\varphi(r)| \, |(E_r)_{A,B}| \\ &= |\varphi(r)|. \end{aligned}$$

This establishes (60) and (61). Summing (61) over r gives $\|E_\varphi\|_1 = \|\varphi\|_1$. ∎

To discuss the spectrum of E_φ, we first describe the Fourier spectrum of the characteristic function of matrices of a given rank. This is where the significance of the Γ_n function becomes evident.

Lemma III.12. *Let $n \geqslant 1$ be a given integer. For $r \in \{0, 1, \dots, n\}$, define $f_r\colon \mathbb{F}_q^{n \times n} \to \{0, 1\}$ by $f_r(X) = 1$ if and only if $\operatorname{rk} X = r$. Then for all $M \in \mathbb{F}_q^{n \times n}$,*

$$\widehat{f}_r(M) = \frac{|\mathscr{M}_r^{n,n}|}{q^{n^2}} \cdot \Gamma_n(\operatorname{rk} M, r).$$

Proof: We have

$$\begin{aligned} \widehat{f}_r(M) &= \mathop{\mathbf{E}}_{X \in \mathbb{F}_q^{n \times n}} \omega^{-\langle M, X \rangle} f_r(X) \\ &= q^{-n^2} \sum_{X \in \mathscr{M}_r^{n,n}} \omega^{-\langle M, X \rangle} \\ &= q^{-n^2} |\mathscr{M}_r^{n,n}| \mathop{\mathbf{E}}_{X \in \mathscr{M}_r^{n,n}} \omega^{-\langle M, X \rangle} \\ &= q^{-n^2} |\mathscr{M}_r^{n,n}| \mathop{\mathbf{E}}_{\substack{X \in \mathscr{M}_r^{n,n} \\ U, V \in \mathscr{M}_n}} \omega^{-\langle M, UXV \rangle} \\ &= q^{-n^2} |\mathscr{M}_r^{n,n}| \mathop{\mathbf{E}}_{\substack{X \in \mathscr{M}_r^{n,n} \\ U, V \in \mathscr{M}_n}} \omega^{\langle -U^{\mathsf{T}} M V^{\mathsf{T}}, X \rangle} \\ &= q^{-n^2} |\mathscr{M}_r^{n,n}| \mathop{\mathbf{E}}_{\substack{X \in \mathscr{M}_r^{n,n} \\ Y \in \mathscr{M}_{\operatorname{rk} M}^{n,n}}} \omega^{\langle Y, X \rangle} \\ &= q^{-n^2} |\mathscr{M}_r^{n,n}| \, \Gamma_n(\operatorname{rk} M, r), \end{aligned}$$

where the second step uses the definition of f_r, the fourth step is valid by Proposition II.19, the fifth step invokes Fact II.2(ii), the sixth step uses Proposition II.19 once more, and the last step applies the definition of Γ_n. ∎

We are now ready to describe the spectrum of E_φ in terms of φ and the Γ_n function.

Lemma III.13 (Singular values of E_φ)**.** *Let $n \geqslant 1$ be an integer and $\varphi\colon \{0,1,\dots,n\} \to \mathbb{R}$ a given function. Then the singular values of E_φ are*

$$q^{-n^2}\left|\sum_{t=0}^{n} \varphi(t)\Gamma_n(s,t)\right|, \qquad s = 0,1,\dots,n,$$

with corresponding multiplicities $|\mathscr{M}_s^{n,n}|$ for $s = 0,1,\dots,n$.

Proof: For $t = 0,1,\dots,n$, define f_t as in Lemma III.12. In this notation,

$$\begin{aligned} E_\varphi &= \sum_{t=0}^{n} \varphi(t) E_t \\ &= \sum_{t=0}^{n} \varphi(t) \left[\frac{1}{q^{n^2}|\mathscr{M}_t^{n,n}|} \cdot f_t(A+B)\right]_{A,B} \\ &= [f(A+B)]_{A,B}, \end{aligned}$$

where

$$f = \sum_{t=0}^{n} \frac{\varphi(t)}{q^{n^2}|\mathscr{M}_t^{n,n}|} \cdot f_t.$$

By Fact II.10, the singular values of E_φ are $q^{n^2}|\widehat{f}(M)|$ for $M \in \mathbb{F}_q^{n\times n}$. Calculating,

$$\begin{aligned} q^{n^2}|\widehat{f}(M)| &= q^{n^2}\left|\sum_{t=0}^{n} \frac{\varphi(t)}{q^{n^2}|\mathscr{M}_t^{n,n}|} \cdot \widehat{f_t}(M)\right| \\ &= q^{-n^2}\left|\sum_{t=0}^{n} \varphi(t)\Gamma_n(\operatorname{rk} M, t)\right|, \end{aligned}$$

where the first step uses the linearity of the Fourier transform, and the second step applies Lemma III.12. Grouping these singular values according to $\operatorname{rk} M$ shows that the spectrum of E_φ is as claimed. ∎

E. Approximate trace norm of the rank problem

Using the machinery developed in previous sections, we now prove a lower bound on the approximate trace norm of the characteristic matrix of the rank problem. Combined with the approximate trace norm method, this will allow us to obtain our communication lower bounds for the rank problem.

Theorem III.14. *Let $n > k \geqslant 0$ be given integers. Let F be the matrix with rows and columns indexed by elements of $\mathbb{F}_q^{n\times n}$, and entries given by*

$$F_{A,B} = \begin{cases} 1 & \text{if } \operatorname{rk}(A+B) = n, \\ -1 & \text{if } \operatorname{rk}(A+B) = k, \\ * & \text{otherwise.} \end{cases}$$

Then for all reals $\delta \geqslant 0$ and all nonnegative integers ℓ, m with $\ell + m \leqslant k$,

$$\|F\|_{\Sigma,\delta} \geqslant \frac{1}{150}\left(1 - \delta - \frac{64}{q^{m+1}}\right) q^{\ell(k-\ell-m+1)/2}\, q^{n^2}, \quad (63)$$

$$\|F\|_{\Sigma,\delta} \geqslant \frac{1-\delta}{150} \cdot q^{k/2}\, q^{n^2}. \quad (64)$$

Proof: Let $\varphi\colon \{0,1,\dots,n\} \to \mathbb{R}$ be the function constructed in Lemma III.9. Then

$$\begin{aligned} \sum_{\operatorname{dom} F} F_{A,B}(E_\varphi)_{A,B} - \delta\|E_\varphi\|_1 - \sum_{\overline{\operatorname{dom} F}} |(E_\varphi)_{A,B}| \\ &= \sum_{\operatorname{rk}(A+B)=n} (E_\varphi)_{A,B} - \sum_{\operatorname{rk}(A+B)=k} (E_\varphi)_{A,B} \\ &\quad - \delta\|E_\varphi\|_1 - \sum_{\operatorname{rk}(A+B)\notin\{n,k\}} |(E_\varphi)_{A,B}| \\ &= \varphi(n) - \varphi(k) - \delta\|\varphi\|_1 - \sum_{r\notin\{n,k\}} |\varphi(r)| \\ &= |\varphi(n)| + |\varphi(k)| - \delta\|\varphi\|_1 - \sum_{r\notin\{n,k\}} |\varphi(r)| \\ &= (1-\delta)\|\varphi\|_1 - 2\sum_{r\notin\{n,k\}} |\varphi(r)| \\ &\geqslant \left(1 - \delta - 2\sum_{r\notin\{n,k\}} |\varphi(r)|\right)\|\varphi\|_1, \end{aligned} \quad (65)$$

where the second step uses Lemma III.11, the third step is valid by Lemma III.9(i)–(ii), and the last step is justified by Lemma III.9(i).

We now analyze the spectral norm of E_φ. Recall from Lemma III.6(iv) that for any fixed values of n and s, the quantity $\Gamma_n(s,t)$ as a function of $t \in \{0,1,\dots,n\}$ is a polynomial in q^{-t} of degree at most s. In this light, Lemma III.9(iv) implies that

$$\max_{s\in\{0,1,\dots,k-\ell-m\}} \left|\sum_{t=0}^{n} \varphi(t)\Gamma_n(s,t)\right| = 0. \quad (66)$$

Continuing,

$$\begin{aligned} &\max_{s\in\{k-\ell-m+1,\dots,n-1,n\}} \left|\sum_{t=0}^{n} \varphi(t)\Gamma_n(s,t)\right| \\ &= \max_{s\in\{k-\ell-m+1,\dots,n-1,n\}} \left|\sum_{t=\ell}^{n} \varphi(t)\Gamma_n(s,t)\right| \\ &\leqslant \max_{s\in\{k-\ell-m+1,\dots,n-1,n\}} \left\{\|\varphi\|_1 \max_{t\in\{\ell,\ell+1,\dots,n\}} |\Gamma_n(s,t)|\right\} \\ &\leqslant \max_{s\in\{k-\ell-m+1,\dots,n-1,n\}} \left\{\|\varphi\|_1 \max_{t\in\{\ell,\ell+1,\dots,n\}} 128q^{-st/2}\right\} \\ &= 128\|\varphi\|_1 q^{-\ell(k-\ell-m+1)/2}, \end{aligned} \quad (67)$$

where the first step uses Lemma III.9(iii), and the third step applies the bound of Lemma III.6(v). By (66), (67), and Lemma III.13,

$$\|E_\varphi\| \leqslant 128\|\varphi\|_1\, q^{-\ell(k-\ell-m+1)/2}\, q^{-n^2}. \quad (68)$$

Proposition II.9 with $\Phi = E_\varphi$ implies, in view of (65) and (68), that

$$\|F\|_{\Sigma,\delta} \geqslant \frac{1}{128} \cdot \left(1 - \delta - 2 \sum_{r \notin \{n,k\}} |\varphi(r)|\right) \times q^{\ell(k-\ell-m+1)/2}\, q^{n^2}. \quad (69)$$

Since $\sum_{r \notin \{n,k\}} |\varphi(r)| \leqslant 32q^{-m-1}$ by Lemma III.9(v), this settles (63). The alternate lower bound (64) follows from (69) by taking $\ell = k$ and $m = 0$ and noting that $\sum_{r \notin \{n,k\}} |\varphi(r)| = 0$ in this case (by Lemma III.9(iii)). ∎

F. Communication lower bounds

We will now use our newly obtained lower bound on the approximate trace norm to prove the main result of this section, a tight lower bound on the communication complexity of the rank problem. We will first examine the canonical case of distinguishing rank-k matrices in $\mathbb{F}^{n \times n}$ from full-rank matrices.

Theorem III.15. *There is an absolute constant $c > 0$ such that for all finite fields $\mathbb{F}$ and all integers $n > k \geqslant 0$,*

$$Q^*_{\frac{1}{2} - \frac{1}{4|\mathbb{F}|^{k/3}}}(\mathrm{RANK}^{\mathbb{F},n,n}_{k,n}) \geqslant c(1 + k^2 \log |\mathbb{F}|). \quad (70)$$

Proof: Abbreviate $q = |\mathbb{F}|$ and $\varepsilon = \frac{1}{2} - \frac{1}{4q^{k/3}}$. Since $\mathrm{RANK}^{\mathbb{F},n,n}_{k,n}$ is a nonconstant function, we have the trivial lower bound

$$Q^*_\varepsilon(\mathrm{RANK}^{\mathbb{F},n,n}_{k,n}) \geqslant 1. \quad (71)$$

Let F be the characteristic matrix of this communication problem. We first examine the case $k \leqslant 50$. Here, taking $\delta = 2\varepsilon$ in equation (64) of Theorem III.14 shows that $\|F\|_{\Sigma,2\varepsilon} \geqslant q^{k/6} q^{n^2}/300 \geqslant q^{k^2/300} q^{n^2}/300$, where the last step uses $k \leqslant 50$. It follows from Theorem II.23 that

$$Q^*_\varepsilon(\mathrm{RANK}^{\mathbb{F},n,n}_{k,n}) \geqslant \frac{1}{2} \log \frac{q^{k^2/300}}{3 \cdot 300} \geqslant \frac{1}{600}\, k^2 \log q - 5.$$

Taking a weighted arithmetic average of this lower bound and (71) settles (70).

Consider now the complementary case $k > 50$. Taking $\delta = 2\varepsilon$, $\ell = \lceil k/3 \rceil$, and $m = \lfloor k/2 \rfloor$ in equation (63) of Theorem III.14 gives

$$\begin{aligned}
\|F\|_{\Sigma,2\varepsilon} &\geqslant \frac{1}{150} \left(\frac{1}{2q^{k/3}} - \frac{64}{q^{\lfloor k/2 \rfloor + 1}}\right) \times q^{\lceil k/3 \rceil (k - \lceil k/3 \rceil - \lfloor k/2 \rfloor + 1)/2} q^{n^2} \\
&\geqslant \frac{1}{300} \left(1 - \frac{128}{q^{k/6}}\right) q^{-k/3} q^{\lceil k/3 \rceil k/12} q^{n^2} \\
&\geqslant \frac{1}{600}\, q^{-k/3} q^{\lceil k/3 \rceil k/12} q^{n^2} \\
&\geqslant \frac{1}{600}\, q^{k^2/48} q^{n^2},
\end{aligned}$$

where the last two steps use $k > 50$. As a result, Theorem II.23 guarantees that

$$Q^*_\varepsilon(\mathrm{RANK}^{\mathbb{F},n,n}_{k,n}) \geqslant \frac{1}{2} \log \frac{q^{k^2/48}}{3 \cdot 600} \geqslant \frac{1}{96}\, k^2 \log q - 6.$$

Taking a weighted arithmetic average of this lower bound and (71) settles (70). ∎

We now establish our main lower bound for the rank problem in its full generality.

Theorem (restatement of Theorem I.1)**.** *There is an absolute constant $c > 0$ such that for all finite fields $\mathbb{F}$ and all integers n, m, R, r with $\min\{n, m\} \geqslant R > r \geqslant 0$,*

$$Q^*_{\frac{1}{2} - \frac{1}{4|\mathbb{F}|^{r/3}}}(\mathrm{RANK}^{\mathbb{F},n,m}_{r,R}) \geqslant c(1 + r^2 \log |\mathbb{F}|). \quad (72)$$

In particular,

$$Q^*_{1/4}(\mathrm{RANK}^{\mathbb{F},n,m}_{r,R}) \geqslant c(1 + r^2 \log |\mathbb{F}|). \quad (73)$$

Proof: There is a communication-free reduction from $\mathrm{RANK}^{\mathbb{F},R,R}_{r,R}$ to $\mathrm{RANK}^{\mathbb{F},n,m}_{r,R}$, where Alice and Bob pad their input matrices $A, B \in \mathbb{F}^{R \times R}$ with zeroes to obtain matrices $A', B' \in \mathbb{F}^{n \times m}$ with $\mathrm{rk}(A + B) = \mathrm{rk}(A' + B')$. Therefore, $Q^*_\varepsilon(\mathrm{RANK}^{\mathbb{F},n,m}_{r,R}) \geqslant Q^*_\varepsilon(\mathrm{RANK}^{\mathbb{F},R,R}_{r,R})$ for all ε. Now Theorem III.15 implies (72), which in turn implies (73). ∎

G. Communication upper bounds

To finalize our study of the rank problem, we will prove a matching upper bound on its communication complexity. We emphasize that our upper bound is achieved by a randomized (classical) protocol, whereas our lower bound is valid even for quantum communication.

Theorem III.16. *Let n, m, R be nonnegative integers with $\min\{n, m\} \geqslant R \geqslant 0$. Let $\mathbb{F}$ be a finite field with $|\mathbb{F}| = q$ elements. Then for all $\varepsilon \in (0, 1)$, there is a two-party randomized communication protocol which:*

- *takes as input a pair of matrices $A, B \in \mathbb{F}^{n \times m}$ for Alice and Bob, respectively;*
- *computes $\min\{\mathrm{rk}(A + B), R\}$ with probability of error at most ε; and*
- *has communication cost $O((R + \lceil \log_q(1/\varepsilon) \rceil)^2 \log q)$.*

Proof: We may assume that $n, m \geqslant 1$ since the theorem is trivial otherwise. Set $\Delta = \lceil \log_q(8/\varepsilon) \rceil$. On input $A, B \in \mathbb{F}^{n \times m}$, the protocol is as follows. Alice and Bob use their shared randomness to pick a pair of independent and uniformly random matrices $X \in \mathbb{F}^{(R+\Delta) \times n}$ and $Y \in \mathbb{F}^{m \times (R+\Delta)}$. Then Alice sends the matrix $XAY \in \mathbb{F}^{(R+\Delta) \times (R+\Delta)}$ to Bob, who announces $\min\{\mathrm{rk}(X(A + B)Y), R\}$ as the output. The communication cost is $O((R + \Delta)^2 \log q)$ as claimed, due to $X(A + B)Y = XAY + XBY$. It is also clear that this protocol always outputs a *lower* bound on the correct value $\min\{\mathrm{rk}(A + B), R\}$, due to $\mathrm{rk}(X(A + B)Y) \leqslant \mathrm{rk}(A + B)$ for all X, Y. It remains to show that

$$\mathbf{P}[\mathrm{rk}(X(A + B)Y) \geqslant \min\{\mathrm{rk}(A + B), R\}] \geqslant 1 - \varepsilon. \quad (74)$$

Abbreviate $C = A + B$. Conditioned on X, we have $\mathrm{rk}(XCY) \geqslant \min\{\mathrm{rk}(XC), R\}$ with probability at least $1 - 4q^{-\Delta - 1} \geqslant 1 - \varepsilon/2$ (apply Lemma II.21(ii) with $M = XC$ and $t = \min\{\mathrm{rk}(XC), R\} - 1$). Similarly, $\mathrm{rk}(XC) \geqslant \min\{\mathrm{rk}\, C, R\}$ with probability at least $1 - \varepsilon/2$ (apply Lemma II.21(i) with $M = C$ and $t = \min\{\mathrm{rk}\, C, R\} - 1$). The union

bound now gives $\mathbf{P}[\mathrm{rk}(XCY) \geqslant \min\{\mathrm{rk}\, C, R\}] \geqslant 1-\varepsilon$, settling (74). ∎

Corollary III.17. *Let n, m, r be integers with $\min\{n,m\} > r \geqslant 0$. Let $\mathbb{F}$ be a finite field with $|\mathbb{F}| = q$ elements. Then for all $\varepsilon \in (0, 1/2)$,*

$$R_\varepsilon(\mathrm{RANK}_r^{\mathbb{F},n,m}) = \begin{cases} O(\log(1/\varepsilon)) & \text{if } r = 0, \\ O((r + \lceil \log_q(1/\varepsilon)\rceil)^2 \log q) & \text{otherwise.} \end{cases} \tag{75}$$

Proof: Observe that $\mathrm{RANK}_0^{\mathbb{F},n,m}(A,B) = -1$ if and only if $A = -B$. Thus, $\mathrm{RANK}_0^{\mathbb{F},n,m}$ is equivalent to the equality problem with domain $\mathbb{F}^{n\times m} \times \mathbb{F}^{n\times m}$. It is well known [27] that the ε-error randomized communication complexity of equality is $O(\log(1/\varepsilon))$. Thus, (75) holds for $r = 0$.

For $r \geqslant 1$, we have $\mathrm{RANK}_r^{\mathbb{F},n,m}(A,B) = -1$ if and only if $\min\{\mathrm{rk}(A+B), r+1\} \leqslant r$. To compute $\min\{\mathrm{rk}(A+B), r+1\}$ on input A, B to error ε, Alice and Bob can use the randomized protocol of Theorem III.16 with $R = r+1$, with communication cost $O((r + \lceil \log_q(1/\varepsilon)\rceil)^2 \log q)$. ∎

We now prove an alternate communication upper bound, showing that even a two-bit protocol can solve the rank problem with nontrivial advantage.

Theorem III.18. *Let n, m, r be integers with $\min\{n,m\} > r \geqslant 0$. Let $\mathbb{F}$ be a finite field with $|\mathbb{F}| = q$ elements. Then*

$$R_{\frac{1}{2} - \frac{1}{32q^r}}(\mathrm{RANK}_r^{\mathbb{F},n,m}) \leqslant 2. \tag{76}$$

Proof: Consider the following auxiliary protocol Π. On input $A, B \in \mathbb{F}^{n\times m}$, Alice and Bob use their shared randomness to pick a pair of independent and uniformly random vectors $x \in \mathbb{F}^n$ and $y \in \mathbb{F}^m$, as well as a uniformly random function $H\colon \mathbb{F} \to \{-1, 1\}$. They exchange $H(x^\mathsf{T} A y)$ and $H(-x^\mathsf{T} B y)$ using 2 bits of communication and output $-H(x^\mathsf{T} A y) H(-x^\mathsf{T} B y)$.

We now analyze the expected output of $\Pi(A,B)$ on a given matrix pair A, B. To begin with,

$$\mathbf{E}[\Pi(A,B) \mid x, y] = \begin{cases} -1 & \text{if } x^\mathsf{T}(A+B)y = 0, \\ 0 & \text{otherwise.} \end{cases} \tag{77}$$

Indeed, if $x^\mathsf{T}(A+B)y = 0$ then $x^\mathsf{T} A y = -x^\mathsf{T} B y$ and therefore Π outputs -1. If, on the other hand, $x^\mathsf{T}(A+B)y \neq 0$ then $x^\mathsf{T} A y \neq -x^\mathsf{T} B y$, which means that $H(x^\mathsf{T} A y)$ and $H(-x^\mathsf{T} B y)$ are independent and their product has expected value 0. Equation (77) implies that $\mathbf{E}\,\Pi(A,B) = -\mathbf{P}[x^\mathsf{T}(A+B)y = 0]$, which can be expanded as

$$\begin{aligned} \mathbf{E}\,\Pi(A,B) = &-\mathbf{P}[x^\mathsf{T}(A+B) = 0] \\ &-\mathbf{P}[x^\mathsf{T}(A+B) \neq 0]\,\mathbf{P}[x^\mathsf{T}(A+B)y = 0 \mid x^\mathsf{T}(A+B) \neq 0]. \end{aligned}$$

The event $x^\mathsf{T}(A+B) = 0$ is equivalent to x being in the orthogonal complement of the column span of $A+B$, which happens with probability $q^{n-\mathrm{rk}(A+B)}/q^n = q^{-\mathrm{rk}(A+B)}$. Conditioned on $x^\mathsf{T}(A+B) \neq 0$, the field element $x^\mathsf{T}(A+B)y$ is uniformly random and in particular is 0 with probability $1/q$. As a result,

$$\begin{aligned} \mathbf{E}\,\Pi(A,B) &= -\frac{1}{q^{\mathrm{rk}(A+B)}} - \left(1 - \frac{1}{q^{\mathrm{rk}(A+B)}}\right) \cdot \frac{1}{q} \\ &= -\frac{1}{q} - \frac{q-1}{q^{\mathrm{rk}(A+B)+1}}. \end{aligned}$$

Therefore, the expected value of $\Pi(A,B)$ is at most $-1/q - (q-1)/q^{r+1}$ when $\mathrm{rk}(A+B) \leqslant r$ and is at least $-1/q - (q-1)/q^{r+2}$ when $\mathrm{rk}(A+B) > r$. Proposition II.24 now shows that $\mathrm{RANK}_r^{\mathbb{F},n,m}$ has a communication protocol with the same cost as Π and error at most $\frac{1}{2} - \frac{1}{8}(q-1)^2/q^{r+2}$. This settles (76) since $q \geqslant 2$. ∎

Corollary III.17 (with $\varepsilon = 1/3$) and Theorem III.18 settle Theorem I.2 from the introduction.

H. Streaming complexity

Fix a finite field $\mathbb{F}$ and a (possibly partial) function $f\colon \mathbb{F}^{n\times n} \to \{-1, 1, *\}$. A streaming algorithm for f receives as input a matrix $M \in \mathbb{F}^{n\times n}$ in row-major order. We say that $\mathscr{A}$ *computes f with error ε* if for every input in the domain of f, the output of $\mathscr{A}$ agrees with f with probability at least $1-\varepsilon$. We will now use a well-known reduction [6] to transform our communication lower bound for the matrix rank problem into a lower bound on its streaming complexity.

Theorem (restatement of Theorem I.3)**.** *Let n, r, R be non-negative integers with $n/2 \leqslant r < R \leqslant n$, and let $\mathbb{F}$ be a finite field. Define $f\colon \mathbb{F}^{n\times n} \to \{-1, 1, *\}$ by*

$$f(M) = \begin{cases} -1 & \text{if } \mathrm{rk}\, M = r, \\ 1 & \text{if } \mathrm{rk}\, M = R, \\ * & \text{otherwise.} \end{cases}$$

Let $\mathscr{A}$ be any randomized streaming algorithm for f with error probability $\frac{1}{2} - \frac{1}{4}|\mathbb{F}|^{-(r-\lceil n/2\rceil)/3}$ that uses s bits of memory and k passes. Then

$$sk = \Omega\left(\left(r - \left\lceil \frac{n}{2} \right\rceil\right)^2 \log |\mathbb{F}|\right). \tag{78}$$

Proof: Abbreviate $m = \lfloor n/2 \rfloor$ and $F = \mathrm{RANK}^{\mathbb{F},m,m}_{r-\lceil n/2\rceil, R-\lceil n/2\rceil}$. We will use a reduction from communication to streaming due to Li, Sun, Wang, and Woodruff [6, Thm. 29]. Specifically, let $A, B \in \mathbb{F}^{m\times m}$ be Alice and Bob's inputs, respectively, for F. Define

$$M = \begin{bmatrix} A & -I_m & 0 \\ B & I_m & 0 \\ 0 & 0 & I_{n-2m} \end{bmatrix},$$

where I_m and I_{n-2m} stand for the identity matrices of order m and $n-2m$, respectively (in particular, I_{n-2m} is empty for even n). We have

$$\begin{aligned} \mathrm{rk}\, M &= \mathrm{rk} \begin{bmatrix} A+B & 0 & 0 \\ B & I_m & 0 \\ 0 & 0 & I_{n-2m} \end{bmatrix} \\ &= \mathrm{rk}(A+B) + n - m \end{aligned}$$

$$= \operatorname{rk}(A+B) + \left\lceil \frac{n}{2} \right\rceil.$$

As a result, for all matrix pairs (A, B) with $\operatorname{rk}(A+B) \in \{r - \lceil n/2 \rceil, R - \lceil n/2 \rceil\}$, we have $F(A, B) = f(M)$. This makes it possible for Alice and Bob to compute F by simulating $\mathscr{A}$ on M. Alice starts the simulation by running $\mathscr{A}$ on the first m rows of M, which depend only on her input A. She then sends Bob the contents of $\mathscr{A}$'s memory, and Bob runs $\mathscr{A}$ on the remaining $n - m$ rows of M. This completes the first pass. Next, Bob sends Alice the contents of $\mathscr{A}$'s memory, and they continue as before until they simulate all k passes. At the end of the k-th pass, Bob announces the output of $\mathscr{A}$ as the protocol output. The error probability of the described protocol is the same as that of $\mathscr{A}$, and the communication cost is $s(2k-1)+1$ bits. Therefore,

$$R_{\frac{1}{2} - \frac{1}{4}|\mathbb{F}|^{-(r - \lceil n/2 \rceil)/3}}(F) \leqslant s(2k-1) + 1.$$

Since the left-hand side is at least $\Omega((r - \lceil n/2 \rceil)^2 \log |\mathbb{F}| + 1)$ by Theorem I.1, the claimed trade-off (78) follows. ∎

IV. The determinant problem

In this section, we establish our lower bound on the communication complexity of the determinant problem. We begin in Section IV-A with technical results on characteristic functions of matrices with a given determinant value. In Section IV-B, we give our own proof of the lower bound for distinguishing two nonzero values of the determinant, which is simpler and more elementary than the proof in [5]. In Section IV-C, we prove an optimal lower bound for the general case of distinguishing two arbitrary values of the determinant, solving an open problem from [5]. Throughout this section, we use a generic finite field $\mathbb{F}$ with q elements, where q is an arbitrary prime power. The root of unity ω and the notation ω^x for $x \in \mathbb{F}$ are as defined in Section II-D.

A. Auxiliary results

Fix a finite field $\mathbb{F}$ and a positive integer n. Recall that the determinant function on $\mathbb{F}^{n\times n}$ is multiplicative, with $\det(AB) = \det(A)\det(B)$. As a result, the set of matrices in $\mathbb{F}^{n\times n}$ with nonzero determinants form a group under matrix multiplication, called the *general linear group* and denoted by $\mathrm{GL}(\mathbb{F}, n)$. Analogously, the matrices in $\mathbb{F}^{n\times n}$ with determinant 1 also form a group, called the *special linear group* and denoted by $\mathrm{SL}(\mathbb{F}, n)$. The multiplicativity of the determinant further implies that $\mathrm{SL}(\mathbb{F}, n)$ is a normal subgroup of $\mathrm{GL}(\mathbb{F}, n)$, with quotient isomorphic to the multiplicative group of the field: $\mathrm{GL}(\mathbb{F}, n)/\mathrm{SL}(\mathbb{F}, n) \cong \mathbb{F}^\times$. For any given field element $u \neq 0$, the set of matrices with determinant u form a coset of $\mathrm{SL}(\mathbb{F}, n)$ in $\mathrm{GL}(\mathbb{F}, n)$. In particular,

$$|\{X \in \mathbb{F}^{n\times n} : \det X = u\}| = |\mathrm{SL}(\mathbb{F}, n)| = \frac{|\mathscr{M}_n|}{|\mathbb{F}| - 1}, \qquad u \in \mathbb{F} \setminus \{0\}. \quad (79)$$

Recall that for each $Y \in \mathbb{F}^{n\times n}$, the mapping $X \mapsto X + Y$ is a permutation on $\mathbb{F}^{n\times n}$. As a result, the previous equation implies that

$$|\{(X, Y) \in \mathbb{F}^{n\times n} \times \mathbb{F}^{n\times n} : \det(X+Y) = u\}| = |\mathbb{F}|^{n^2} |\mathrm{SL}(\mathbb{F}, n)|, \qquad u \in \mathbb{F} \setminus \{0\}. \quad (80)$$

To understand the spectral norm of the determinant problem, we now introduce a relevant function on $\mathbb{F}^{n\times n}$ and discuss its Fourier coefficients.

Lemma IV.1. *Let n be a positive integer, $\mathbb{F}$ a finite field. For a pair of distinct elements $u, v \in \mathbb{F} \setminus \{0\}$, define $g_{u,v} \colon \mathbb{F}^{n\times n} \to \{-1, 1, 0\}$ by*

$$g_{u,v}(X) = \begin{cases} -1 & \text{if } \det X = u, \\ 1 & \text{if } \det X = v, \\ 0 & \text{otherwise.} \end{cases}$$

Then:

(i) $\widehat{g_{u,v}}(A) = 0$ *for every singular matrix* A;
(ii) $\widehat{g_{u,v}}(A) = \widehat{g_{u,v}}(B)$ *whenever* $\det A = \det B$;
(iii) $\|\widehat{g_{u,v}}\|_\infty \leqslant 1/\sqrt{|\mathrm{SL}(\mathbb{F}, n)|}$.

Proof: (i) In view of (79), we have

$$\begin{aligned} \widehat{g_{u,v}}(A) &= \mathop{\mathbf{E}}_{X \in \mathbb{F}^{n\times n}} g_{u,v}(X) \omega^{-\langle A, X\rangle} \\ &= |\mathbb{F}|^{-n^2} \cdot \frac{|\mathscr{M}_n|}{|\mathbb{F}| - 1} \left(\mathop{\mathbf{E}}_{X : \det X = v} \omega^{-\langle A, X\rangle} - \mathop{\mathbf{E}}_{X : \det X = u} \omega^{-\langle A, X\rangle} \right). \end{aligned}$$

It remains to show that the expectations in the last expression are equal. Since A is singular, there exist nonsingular matrices P and Q such that $A = PI_sQ$ for $s = \operatorname{rk} A < n$. Consider the order-$n$ diagonal matrix $Z = \operatorname{diag}(1, 1, \ldots, 1, u^{-1}v)$. Using $I_s = I_sZ$, we obtain $A = PI_sZQ = PI_sQQ^{-1}ZQ = AQ^{-1}ZQ$. As a result,

$$\begin{aligned} \mathop{\mathbf{E}}_{X : \det X = u} \omega^{-\langle A, X\rangle} &= \mathop{\mathbf{E}}_{X : \det X = u} \omega^{-\langle AQ^{-1}ZQ, X\rangle} \\ &= \mathop{\mathbf{E}}_{X : \det X = u} \omega^{-\langle A, X(Q^{-1}ZQ)^{\mathsf{T}}\rangle} \\ &= \mathop{\mathbf{E}}_{Y : \det Y = v} \omega^{-\langle A, Y\rangle}, \end{aligned}$$

where the second step uses Fact II.2(ii), and the last step is valid because the mapping $X \mapsto X(Q^{-1}ZQ)^{\mathsf{T}}$ is a bijection from the set of matrices with determinant u onto the set of matrices with determinant $u \cdot \det((Q^{-1}ZQ)^{\mathsf{T}}) = v$.

(ii) For singular A and B, the claim is immediate from (i). In the complementary case,

$$\begin{aligned} \widehat{g_{u,v}}(A) &= \mathop{\mathbf{E}}_{X \in \mathbb{F}^{n\times n}} g_{u,v}(X) \omega^{-\langle A, X\rangle} \\ &= \mathop{\mathbf{E}}_{X \in \mathbb{F}^{n\times n}} g_{u,v}((BA^{-1})^{\mathsf{T}}X) \omega^{-\langle A, (BA^{-1})^{\mathsf{T}}X\rangle} \\ &= \mathop{\mathbf{E}}_{X \in \mathbb{F}^{n\times n}} g_{u,v}(X) \omega^{-\langle A, (BA^{-1})^{\mathsf{T}}X\rangle} \\ &= \mathop{\mathbf{E}}_{X \in \mathbb{F}^{n\times n}} g_{u,v}(X) \omega^{-\langle B, X\rangle} \\ &= \widehat{g_{u,v}}(B), \end{aligned}$$

where the second step is valid because $(BA^{-1})^{\mathsf{T}}$ is invertible and hence $X \mapsto (BA^{-1})^{\mathsf{T}}X$ is a permutation on

$\mathbb{F}^{n\times n}$; the third step is justified by $\det((BA^{-1})^{\mathsf{T}}X) = \det(B)\det(X)/\det(A) = \det X$; and the fourth step is an application of Fact II.2(ii).

(iii) Let M be a matrix with $|\widehat{g_{u,v}}(M)| = \|\widehat{g_{u,v}}\|_\infty$. By (i), we know that $\det M \neq 0$. Now

$$\begin{aligned}
1 &\geqslant \mathop{\mathbf{E}}_{X\in\mathbb{F}^{n\times n}}[|g_{u,v}(X)|^2] \\
&= \sum_{A\in\mathbb{F}^{n\times n}} |\widehat{g_{u,v}}(A)|^2 \\
&\geqslant \sum_{A:\det A=\det M} |\widehat{g_{u,v}}(A)|^2 \\
&= |\{A : \det A = \det M\}|\,|\widehat{g_{u,v}}(M)|^2 \\
&= |\mathrm{SL}(\mathbb{F},n)|\,\|\widehat{g_{u,v}}\|_\infty^2,
\end{aligned}$$

where the second step applies Parseval's inequality (28), the fourth step is justified by (ii), and the fifth step uses $\det M \neq 0$ along with (79). ∎

B. Determinant problem for nonzero field elements

As an application of the previous lemma, we now prove that the characteristic matrix of the determinant problem $\mathrm{DET}_{a,b}^{\mathbb{F},n}$ for any two nonzero field elements a, b has small spectral norm.

Lemma IV.2. *Let $\mathbb{F}$ be a finite field with $|\mathbb{F}| = q$ elements. For each $u \in \mathbb{F}\setminus\{0\}$, define G_u to be the matrix with rows and columns indexed by elements of $\mathbb{F}^{n\times n}$, and entries given by*

$$(G_u)_{X,Y} = \begin{cases} q^{-n^2}|\mathrm{SL}(\mathbb{F},n)|^{-1} & \text{if } \det(X+Y) = u, \\ 0 & \text{otherwise.} \end{cases} \tag{81}$$

Then for all $u, v \in \mathbb{F}\setminus\{0\}$,

$$\|G_u\|_1 = 1, \tag{82}$$

$$\|G_v - G_u\| \leqslant |\mathrm{SL}(\mathbb{F},n)|^{-3/2} \leqslant 8q^{-3(n^2-1)/2}. \tag{83}$$

Proof: Equation (82) follows from (80). For (83), there are two cases to consider. If $u = v$, then $G_v - G_u = 0$ and thus $\|G_v - G_u\| = 0$. If $u \neq v$, write $G_v - G_u = [q^{-n^2}|\mathrm{SL}(\mathbb{F},n)|^{-1}g_{u,v}(X+Y)]_{X,Y}$ with $g_{u,v}$ as defined in Lemma IV.1. Then

$$\|G_v - G_u\| = \frac{\|\widehat{g_{u,v}}\|_\infty}{|\mathrm{SL}(\mathbb{F},n)|} \leqslant \frac{1}{|\mathrm{SL}(\mathbb{F},n)|^{3/2}}, \tag{84}$$

where the first step applies Fact II.10, and the second step uses Lemma IV.1(iii). It remains to simplify the bound of (84):

$$\begin{aligned}
\frac{1}{|\mathrm{SL}(\mathbb{F},n)|^{3/2}} &= \left(\frac{|\mathscr{M}_n|}{q-1}\right)^{-3/2} \\
&= \left(q^{n-1}\prod_{i=0}^{n-2}(q^n - q^i)\right)^{-3/2} \\
&\leqslant 8q^{-3(n^2-1)/2},
\end{aligned}$$

where the first step uses (79), the second step applies Proposition II.17, and the last step uses Proposition II.13. ∎

Lemma IV.2 was originally obtained by Sun and Wang [5] using a different and rather technical proof. By contrast, the proof presented above is short and uses only basic Fourier analysis. With this newly obtained bound on the spectral norm of the characteristic matrix of $\mathrm{DET}_{a,b}^{\mathbb{F},n}$ for nonzero a, b, we can use the approximate trace norm method to obtain a tight communication lower bound for this special case of the determinant problem.

Theorem IV.3. *Let $\mathbb{F}$ be a finite field, and n a positive integer. Then for every pair of distinct elements $a, b \in \mathbb{F}\setminus\{0\}$ and every $\gamma \in (0,1)$,*

$$Q^*_{(1-\gamma)/2}(\mathrm{DET}_{a,b}^{\mathbb{F},n}) \geqslant \frac{1}{4}(n^2-3)\log|\mathbb{F}| - \frac{1}{2}\log\frac{12}{\gamma}. \tag{85}$$

Proof: Let F be the characteristic matrix of $\mathrm{DET}_{a,b}^{\mathbb{F},n}$. For $u \in \mathbb{F}\setminus\{0\}$, define G_u as in Lemma IV.2. Since G_a and G_b are supported on disjoint sets of entries, (82) leads to

$$\|G_b - G_a\|_1 = \|G_b\|_1 + \|G_a\|_1 = 2. \tag{86}$$

Taking $\Phi = G_b - G_a$ in Proposition II.9, we obtain

$$\begin{aligned}
\|F\|_{\Sigma,1-\gamma} &\geqslant \frac{1}{\|G_b - G_a\|}\left(\sum_{\mathrm{dom}\,F} F_{A,B}(G_b - G_a)_{A,B}\right. \\
&\qquad \left. -(1-\gamma)\|G_b - G_a\|_1 - \sum_{\overline{\mathrm{dom}\,F}} |(G_b - G_a)_{A,B}|\right) \\
&= \frac{1}{\|G_b - G_a\|}\left(\sum_{\mathrm{dom}\,F} |(G_b - G_a)_{A,B}|\right. \\
&\qquad \left. - (1-\gamma)\|G_b - G_a\|_1\right) \\
&= \frac{\gamma\|G_b - G_a\|_1}{\|G_b - G_a\|} \\
&\geqslant \frac{1}{4}\gamma|\mathbb{F}|^{3(n^2-1)/2},
\end{aligned} \tag{87}$$

where the second and third steps are valid because $G_b - G_a$ by definition coincides in sign with F on $\mathrm{dom}\,F$ and vanishes on $\overline{\mathrm{dom}\,F}$; and the last step uses (83) and (86). Now (85) follows from (87) in view of Theorem II.23. ∎

We remind the reader that Theorem IV.3 was obtained with different techniques by Sun and Wang [5], who settled the determinant problem $\mathrm{DET}_{a,b}^{\mathbb{F},n}$ for nonzero a, b and left open the complementary case when one of a, b is zero.

C. Determinant problem for arbitrary field elements

Recall that the *rank versus determinant problem,* $\mathrm{RANKDET}_{k,a}^{\mathbb{F},n}$, is a hybrid problem that naturally generalizes the matrix rank problem $\mathrm{RANK}_{k,n}^{\mathbb{F},n,n}$ and the determinant problem $\mathrm{DET}_{0,a}^{\mathbb{F},n}$. Specifically, the rank versus determinant problem requires Alice and Bob to distinguish matrix pairs with $\mathrm{rk}(A+B) = k$ from those with $\det(A+B) = a$, where a is a nonzero field element, k is an integer with $k < n$, and A, B are Alice and Bob's respective inputs. We

will now construct a dual matrix for $\mathrm{RANKDET}_{k,a}^{\mathbb{F},n}$ and thereby obtain a lower bound on its approximate trace norm. As a dual matrix, we will use a linear combination of the dual matrices from our analyses of the rank and determinant problems.

Theorem IV.4. *Let $n > k \geqslant 1$ be given integers. Let $\mathbb{F}$ be a finite field with $|\mathbb{F}| = q$ elements, and let $a \in \mathbb{F} \setminus \{0\}$. Let F be the characteristic matrix of $\mathrm{RANKDET}_{k,a}^{\mathbb{F},n}$. Then for all reals $\delta \geqslant 0$ and all nonnegative integers ℓ, m with $\ell + m \leqslant k$,*

$$\|F\|_{\Sigma,\delta} \geqslant \frac{1}{150}\left(1 - \delta - \frac{64}{q^{m+1}}\right) q^{\ell(k-\ell-m+1)/2}\, q^{n^2}, \quad (88)$$

$$\|F\|_{\Sigma,\delta} \geqslant \frac{1-\delta}{150} \cdot q^{k/2}\, q^{n^2}. \quad (89)$$

Proof: This proof combines our ideas in Theorems III.14 and IV.3, and our dual matrix here will be a linear combination of the dual matrices used in those theorems.

Fix nonnegative integers ℓ, m with $\ell + m \leqslant k$, and let $\varphi\colon \{0, 1, \ldots, n\} \to \mathbb{R}$ be the corresponding function constructed in Lemma III.9. This univariate function gives rise to a matrix E_φ, described in Definition III.10. To restate equation (68) from our proof of Theorem III.14,

$$\|E_\varphi\| \leqslant 128\|\varphi\|_1\, q^{-\ell(k-\ell-m+1)/2}\, q^{-n^2}. \quad (90)$$

For $u \in \mathbb{F} \setminus \{0\}$, define G_u as in Lemma IV.2. As our dual matrix, we will use

$$\Phi = E_\varphi + \sum_{b \in \mathbb{F}\setminus\{0,a\}} \frac{\varphi(n)}{q-1}\,(G_a - G_b). \quad (91)$$

Claim IV.5. *For every matrix pair (A, B),*

$$\Phi_{A,B} = \begin{cases} (E_\varphi)_{A,B} & \text{if } \det(A+B) = 0, \\ \varphi(n) q^{-n^2} |\mathrm{SL}(\mathbb{F}, n)|^{-1} & \text{if } \det(A+B) = a, \\ 0 & \text{otherwise.} \end{cases}$$

Proof: If $\det(A+B) = 0$, then by definition $(G_u)_{A,B} = 0$ for every nonzero field element u. As a result, (91) gives $\Phi_{A,B} = (E_\varphi)_{A,B}$ in this case.

In what follows, we treat the complementary case when $\det(A+B) \neq 0$. For all such matrix pairs,

$$\begin{aligned}(E_\varphi)_{A,B} &= \sum_{i=0}^{n} \varphi(i)(E_i)_{A,B} \\ &= \varphi(n)(E_n)_{A,B} \\ &= \frac{\varphi(n)}{q^{n^2}|\mathscr{M}_n|} \\ &= \frac{\varphi(n)}{q^{n^2}(q-1)|\mathrm{SL}(\mathbb{F}, n)|},\end{aligned}$$

where the first three steps are immediate from Definition III.10, and the last step uses (79). In particular,

$$\begin{aligned}\Phi_{A,B} = \frac{\varphi(n)}{q^{n^2}(q-1)|\mathrm{SL}(\mathbb{F}, n)|} \\ + \sum_{b \in \mathbb{F}\setminus\{0,a\}} \frac{\varphi(n)}{q-1}\,((G_a)_{A,B} - (G_b)_{A,B}). \quad (92)\end{aligned}$$

If $\det(A+B) = a$, then by definition $(G_a)_{A,B} = q^{-n^2}|\mathrm{SL}(\mathbb{F}, n)|^{-1}$ and $(G_b)_{A,B} = 0$ for all $b \in \mathbb{F} \setminus \{0, a\}$, so that (92) gives

$$\begin{aligned}\Phi_{A,B} &= \frac{\varphi(n)}{q^{n^2}(q-1)|\mathrm{SL}(\mathbb{F}, n)|} + \sum_{b \in \mathbb{F}\setminus\{0,a\}} \frac{\varphi(n)}{(q-1)q^{n^2}|\mathrm{SL}(\mathbb{F}, n)|} \\ &= \frac{\varphi(n)}{q^{n^2}|\mathrm{SL}(\mathbb{F}, n)|}.\end{aligned}$$

If, on the other hand, $\det(A+B) = c$ for some $c \in \mathbb{F} \setminus \{0, a\}$, then by definition $(G_a)_{A,B} = 0$ and likewise $(G_b)_{A,B} = 0$ for every $b \neq c$, so that (92) simplifies to

$$\Phi_{A,B} = \frac{\varphi(n)}{q^{n^2}(q-1)|\mathrm{SL}(\mathbb{F}, n)|} - \frac{\varphi(n)}{(q-1)}(G_c)_{A,B} = 0.$$

This completes the proof of the claim. ∎

We proceed to establish key analytic and metric properties of Φ. To begin with,

$$\begin{aligned}\|\Phi\| &\leqslant \|E_\varphi\| + \sum_{b \in \mathbb{F}\setminus\{0,a\}} \frac{|\varphi(n)|}{q-1}\,\|G_a - G_b\| \\ &\leqslant \|E_\varphi\| + \sum_{b \in \mathbb{F}\setminus\{0,a\}} \frac{\|\varphi\|_1}{q-1}\,\|G_a - G_b\| \\ &\leqslant 128\|\varphi\|_1\, q^{-\ell(k-\ell-m+1)/2}\, q^{-n^2} \\ &\qquad + \sum_{b \in \mathbb{F}\setminus\{0,a\}} \frac{\|\varphi\|_1}{q-1} \cdot 8q^{-3(n^2-1)/2} \\ &\leqslant (128q^{-\ell(k-\ell-m+1)/2} + 8q^{-(n^2-3)/2})\,\frac{\|\varphi\|_1}{q^{n^2}}, \quad (93)\end{aligned}$$

where the first step uses the triangle inequality, and the third step is a substitution from (90) and equation (83) of Lemma IV.2. To simplify this bound, recall from the theorem hypothesis that $n > k \geqslant 1$ and $\ell, m \geqslant 0$. Therefore, $\ell(k-\ell-m+1) \leqslant \ell(k-\ell+1) \leqslant (k+1)^2/4 \leqslant n^2/4 \leqslant n^2 - 3$. This results in $q^{-(n^2-3)/2} \leqslant q^{-\ell(k-\ell-m+1)/2}$, and thus (93) simplifies to

$$\|\Phi\| \leqslant 136q^{-\ell(k-\ell-m+1)/2} q^{-n^2}\,\|\varphi\|_1. \quad (94)$$

Next, we examine $\|\Phi\|_1$. We have

$$\begin{aligned}\sum_{\mathrm{rk}(A+B)=n} |\Phi_{A,B}| &= \sum_{\det(A+B)=a} |\Phi_{A,B}| \\ &= \sum_{\det(A+B)=a} \frac{|\varphi(n)|}{q^{n^2}|\mathrm{SL}(\mathbb{F}, n)|} \\ &= |\varphi(n)|,\end{aligned}$$

where the first and second steps are immediate from Claim IV.5, and the last step applies (80). Also,

$$\begin{aligned}\sum_{\mathrm{rk}(A+B)<n} |\Phi_{A,B}| &= \sum_{\mathrm{rk}(A+B)<n} |(E_\varphi)_{A,B}| \\ &= \|E_\varphi\| - \sum_{\mathrm{rk}(A+B)=n} |(E_\varphi)_{A,B}| \\ &= \|\varphi\|_1 - |\varphi(n)|,\end{aligned}$$

where the first step uses Claim IV.5, and the last step invokes Lemma III.11. These two equations yield

$$\|\Phi\|_1 = \|\varphi\|_1. \tag{95}$$

Continuing,

$$\begin{aligned}
\sum_{\operatorname{dom} F} F_{A,B}\Phi_{A,B} &= \sum_{\det(A+B)=a} \Phi_{A,B} - \sum_{\operatorname{rk}(A+B)=k} \Phi_{A,B} \\
&= \sum_{\det(A+B)=a} \frac{\varphi(n)}{q^{n^2}|\mathrm{SL}(\mathbb{F},n)|} \\
&\qquad - \sum_{\operatorname{rk}(A+B)=k} (E_\varphi)_{A,B} \\
&= \varphi(n) - \varphi(k) \\
&= |\varphi(n)| + |\varphi(k)| \\
&= \|\varphi\|_1 - \sum_{r\notin\{k,n\}} |\varphi(r)|,
\end{aligned} \tag{96}$$

where the second step uses Claim IV.5, the third step invokes Lemma III.11 and (80), and the fourth step is valid due to Lemma III.9(i), (ii). Finally,

$$\begin{aligned}
\sum_{\overline{\operatorname{dom} F}} |\Phi_{A,B}| &= \sum_{\operatorname{rk}(A+B)\notin\{n,k\}} |\Phi_{A,B}| + \sum_{\substack{\operatorname{rk}(A+B)=n \\ \det(A+B)\neq a}} |\Phi_{A,B}| \\
&= \sum_{\operatorname{rk}(A+B)\notin\{n,k\}} |\Phi_{A,B}| \\
&= \sum_{\operatorname{rk}(A+B)\notin\{n,k\}} |(E_\varphi)_{A,B}| \\
&= \sum_{r\notin\{n,k\}} |\varphi(r)|,
\end{aligned} \tag{97}$$

where the second and third steps use Claim IV.5, and the last step uses Lemma III.11. Now

$$\begin{aligned}
&\sum_{\operatorname{dom} F} F_{A,B}\Phi_{A,B} - \delta\|\Phi\|_1 - \sum_{\overline{\operatorname{dom} F}} |\Phi_{A,B}| \\
&\qquad = \|\varphi\|_1 - \delta\|\varphi\|_1 - 2\sum_{r\notin\{n,k\}} |\varphi(r)| \\
&\qquad \geqslant \left(1 - \delta - 2\sum_{r\notin\{n,k\}} |\varphi(r)|\right) \|\varphi\|_1,
\end{aligned} \tag{98}$$

where the first step uses (95)–(97), and the last step is legitimate by Lemma III.9(i).

Proposition II.9 implies, in view of (94) and (98), that

$$\begin{aligned}
\|F\|_{\Sigma,\delta} \geqslant \frac{1}{136} &\left(1 - \delta - 2\sum_{r\notin\{n,k\}} |\varphi(r)|\right) \\
&\times q^{\ell(k-\ell-m+1)/2}\, q^{n^2}.
\end{aligned} \tag{99}$$

Since $\sum_{r\notin\{n,k\}} |\varphi(r)| \leqslant 32q^{-m-1}$ by Lemma III.9(v), this proves (88). The alternate lower bound (89) follows by taking $\ell = k$ and $m = 0$ in (99) and noting that $\sum_{r\notin\{n,k\}} |\varphi(r)| = 0$ in this case (by Lemma III.9(iii)). ∎

By virtue of the approximate trace norm method, Theorem IV.4 yields the following tight lower bound on the communication complexity of the rank versus determinant problem.

Theorem (restatement of Theorem I.6)**.** *There is an absolute constant $c > 0$ such that for every finite field $\mathbb{F}$, every field element $a \in \mathbb{F} \setminus \{0\}$, and all integers $n > k \geqslant 0$,*

$$Q^*_{\frac{1}{2}-\frac{1}{4|\mathbb{F}|^{k/3}}}(\mathrm{RANKDET}^{\mathbb{F},n}_{k,a}) \geqslant c(1 + k^2 \log|\mathbb{F}|). \tag{100}$$

Proof: For $k = 0$, the claimed lower bound follows from the fact that $\mathrm{RANKDET}^{\mathbb{F},n}_{0,a}$ is nonconstant and hence has communication complexity at least 1 bit. For $k \geqslant 1$, our lower bounds on the approximate trace norm of $\mathrm{RANKDET}^{\mathbb{F},n}_{k,a}$ are identical to those for $\mathrm{RANK}^{\mathbb{F},n,n}_{k,n}$ (Theorems IV.4 and Theorem III.14, respectively). Accordingly, the proof here is identical to that of Theorem III.15, with equations (88) and (89) of Theorem IV.4 used in place of the corresponding equations (63) and (64) of Theorem III.14. ∎

As a consequence, we obtain an optimal communication lower bound for the unrestricted determinant problem.

Theorem (restatement of Theorem I.5)**.** *There is an absolute constant $c > 0$ such that for every finite field $\mathbb{F}$, every pair of distinct elements $a, b \in \mathbb{F}$, and all integers $n \geqslant 2$,*

$$Q^*_{\frac{1}{2}-\frac{1}{4|\mathbb{F}|^{(n-1)/3}}}(\mathrm{DET}^{\mathbb{F},n}_{a,b}) \geqslant cn^2 \log|\mathbb{F}|. \tag{101}$$

Proof: If $ab = 0$, then $\mathrm{DET}^{\mathbb{F},n}_{a,b}$ contains as a subproblem either $\mathrm{RANKDET}^{\mathbb{F},n}_{n-1,b}$ (when $a = 0$) or $\neg\,\mathrm{RANKDET}^{\mathbb{F},n}_{n-1,a}$ (when $b = 0$), and therefore (101) follows from Theorem I.6. If a and b are both nonzero, Theorem IV.3 gives

$$Q^*_{\frac{1}{2}-\frac{1}{4|\mathbb{F}|^{(n-1)/3}}}(\mathrm{DET}^{\mathbb{F},n}_{a,b}) \geqslant c'n^2 \log|\mathbb{F}| - \frac{1}{2}\log 24$$

for a small enough constant $c' > 0$. Taking a weighted average of this lower bound with the trivial lower bound of 1 bit settles (101). ∎

V. The subspace sum and intersection problems

As discussed in the introduction, our analysis of the subspace sum and subspace intersection problems has similarities with the rank problem but also diverges from it in important ways. Instead of additively composed matrices whose rows and columns are indexed by elements of $\mathbb{F}_q^{n\times n}$, we now have matrices with rows and columns indexed by subspaces, and each entry (A, B) depends solely on the dimension of $A \cap B$. While the construction of the univariate dual object is similar to that for the rank problem, its relation to the singular values of the dual matrix is significantly more intricate, and computing the spectral norm of the dual matrix is now a challenge. Our study of the spectral norm is based on ideas due to Knuth [25]. We start by formalizing the equivalence of the subspace sum and subspace intersection problems, which allows us to focus on the latter problem from then on.

Proposition V.1. *Let n, m, ℓ be nonnegative integers with $\max\{m, \ell\} \leqslant n$. Then for all integers d, D with $d \neq D$,*

$$\mathrm{SUM}_{d,D}^{\mathbb{F},n,m,\ell} = \mathrm{INTERSECT}_{m+\ell-d,m+\ell-D}^{\mathbb{F},n,m,\ell}, \tag{102}$$

$$\mathrm{SUM}_{d}^{\mathbb{F},n,m,\ell} = \mathrm{INTERSECT}_{m+\ell-d}^{\mathbb{F},n,m,\ell}. \tag{103}$$

Proof: Let $S, T \subseteq \mathbb{F}^n$ be arbitrary subspaces of dimension m and ℓ, respectively. Since $\dim(S+T) = m+\ell-\dim(S\cap T)$, we have

$$\mathrm{SUM}_{d,D}^{\mathbb{F},n,m,\ell}(S,T) = \mathrm{INTERSECT}_{m+\ell-d,m+\ell-D}^{\mathbb{F},n,m,\ell}(S,T),$$

settling (102). Analogously, for any subspaces $S, T \subseteq \mathbb{F}^n$ of dimension m and ℓ, respectively, we have $\dim(S+T) \leqslant d$ if and only if $\dim(S\cap T) \geqslant m+\ell-d$, which implies (103). ∎

We will now prove our main result on the subspace sum problem (stated in the introduction as Theorems I.7 and I.8) assuming our corresponding result on subspace intersection (Theorem I.9).

Proof of Theorems I.7 *and* I.8 *assuming Theorem* I.9*:* Recall that Theorem I.7 is a special case of Theorem I.8, corresponding to $\gamma = 1/3$. Therefore, it suffices to prove Theorem I.8. Define $r = m+\ell-D$ and $R = m+\ell-d$. Then the hypotheses $\max\{m, \ell\} \leqslant d < D \leqslant \min\{m+\ell, n\}$ and $\gamma \in [\frac{1}{3}q^{-(2d-m-\ell)/5}, \frac{1}{3}]$ of Theorem I.8 can be equivalently stated as

$$\max\{0, m+\ell-n\} \leqslant r < R \leqslant \min\{m, \ell\}, \tag{104}$$

$$\gamma \in [\tfrac{1}{3}q^{-(m+\ell-2R)/5}, \tfrac{1}{3}]. \tag{105}$$

Recall from Proposition V.1 that $\mathrm{SUM}_{d,D}^{\mathbb{F},n,m,\ell}$ is the same function as $\mathrm{INTERSECT}_{R,r}^{\mathbb{F},n,m,\ell}$, which in turn is the negation of $\mathrm{INTERSECT}_{r,R}^{\mathbb{F},n,m,\ell}$. Now the bounds for $\mathrm{SUM}_{d,D}^{\mathbb{F},n,m,\ell}$ claimed in Theorem I.8 follow from the bounds for $\mathrm{INTERSECT}_{r,R}^{\mathbb{F},n,m,\ell}$ in Theorem I.9, upon substituting $R = m+\ell-d$. This appeal to Theorem I.9 is legitimate due to (104) and (105).

Analogously, $\mathrm{SUM}_{d}^{\mathbb{F},n,m,\ell}$ is the same function as $\mathrm{INTERSECT}_{R}^{\mathbb{F},n,m,\ell}$ (Proposition V.1), and therefore the bounds claimed for $\mathrm{SUM}_{d}^{\mathbb{F},n,m,\ell}$ in Theorem I.8 follow from the bounds for $\mathrm{INTERSECT}_{R}^{\mathbb{F},n,m,\ell}$ in Theorem I.9, upon substituting $R = m+\ell-d$. ∎

The proof of Theorem I.9 is the focus of the remainder of this section, which we defer to the full version of our paper [29].

Acknowledgments

The authors are thankful to Alan Joel for useful discussions and his feedback on an earlier version of this manuscript.

References

[1] A. C.-C. Yao, "Some complexity questions related to distributive computing," in *Proceedings of the Eleventh Annual ACM Symposium on Theory of Computing* (STOC), 1979, pp. 209–213.

[2] ——, "Quantum circuit complexity," in *Proceedings of the Thirty-Fourth Annual IEEE Symposium on Foundations of Computer Science* (FOCS), 1993, pp. 352–361.

[3] J. I. Chu and G. Schnitger, "The communication complexity of several problems in matrix computation," *J. Complex.*, vol. 7, no. 4, pp. 395–407, 1991.

[4] ——, "Communication complexity of matrix computation over finite fields," *Math. Syst. Theory*, vol. 28, no. 3, pp. 215–228, 1995.

[5] X. Sun and C. Wang, "Randomized communication complexity for linear algebra problems over finite fields," in *Proceedings of the Twenty-Ninth International Symposium on Theoretical Aspects of Computer Science* (STACS), vol. 14, 2012, pp. 477–488.

[6] Y. Li, X. Sun, C. Wang, and D. P. Woodruff, "On the communication complexity of linear algebraic problems in the message passing model," in *Proceedings of the Twenty-Eighth International Symposium on Distributed Computing* (DISC), vol. 8784, 2014, pp. 499–513.

[7] K. L. Clarkson and D. P. Woodruff, "Numerical linear algebra in the streaming model," in *Proceedings of the Forty-First Annual ACM Symposium on Theory of Computing* (STOC), 2009, pp. 205–214.

[8] M. Bury and C. Schwiegelshohn, "Sublinear estimation of weighted matchings in dynamic data streams," in *Proceedings of the Twenty-Third Annual European Symposium on Algorithms* (ESA), 2015, pp. 263–274.

[9] S. Assadi, S. Khanna, and Y. Li, "On estimating maximum matching size in graph streams," in *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms* (SODA), 2017, pp. 1723–1742.

[10] S. Assadi, G. Kol, R. R. Saxena, and H. Yu, "Multi-pass graph streaming lower bounds for cycle counting, MAX-CUT, matching size, and other problems," in *Proceedings of the Sixty-First Annual IEEE Symposium on Foundations of Computer Science* (FOCS), 2020, pp. 354–364.

[11] L. Chen, G. Kol, D. Paramonov, R. R. Saxena, Z. Song, and H. Yu, "Almost optimal super-constant-pass streaming lower bounds for reachability," in *Proceedings of the Fifty-Third Annual ACM Symposium on Theory of Computing* (STOC), 2021, pp. 570–583.

[12] L. Lovász and M. E. Saks, "Lattices, Möbius functions and communication complexity," in *Proceedings of the Twenty-Ninth Annual IEEE Symposium on Foundations of Computer Science* (FOCS), 1988, pp. 81–90.

[13] P. B. Miltersen, N. Nisan, S. Safra, and A. Wigderson, "On data structures and asymmetric communication complexity," *J. Comput. Syst. Sci.*, vol. 57, no. 1, pp. 37–49, 1998.

[14] I. Kremer, "Quantum communication," Master's thesis, Hebrew University, Computer Science Department, 1995.

[15] A. A. Razborov, "Quantum communication complexity of symmetric predicates," *Izvestiya: Mathematics*, vol. 67, no. 1, pp. 145–159, 2003.

[16] N. Linial and A. Shraibman, "Lower bounds in communication complexity based on factorization norms," *Random Struct. Algorithms*, vol. 34, no. 3, pp. 368–394, 2009.

[17] A. A. Sherstov, "The pattern matrix method," *SIAM J. Comput.*, vol. 40, no. 6, pp. 1969–2000, 2011, preliminary version in *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing* (STOC), 2008.

[18] ——, "Communication lower bounds using dual polynomials," *Bulletin of the EATCS*, vol. 95, pp. 59–93, 2008.

[19] T. Lee and A. Shraibman, "Lower bounds in communication complexity," *Foundations and Trends in Theoretical Computer Science*, vol. 3, no. 4, pp. 263–398, 2009.

[20] R. P. Stanley, *Enumerative Combinatorics*, 2nd ed. Cambridge University Press, 2012, vol. I.

[21] P. Delsarte, "Association schemes and t-designs in regular semilattices," *J. Comb. Theory, Ser. A*, vol. 20, no. 2, pp. 230–243, 1976.

[22] J. Eisfeld, "The eigenspaces of the Bose-Mesner-algebras of the association schemes corresponding to projective spaces and polar spaces," *Des. Codes Cryptogr.*, vol. 17, no. 1-3, pp. 129–150, 1999.

[23] A. E. Brouwer, S. M. Cioaba, F. Ihringer, and M. McGinnis, "The smallest eigenvalues of Hamming graphs, Johnson graphs and other distance-regular graphs with classical parameters," *J. Comb. Theory, Ser. B*, vol. 133, pp. 88–121, 2018.

[24] S. M. Cioaba and H. Gupta, "On the eigenvalues of Grassmann graphs, bilinear forms graphs and Hermitian forms graphs," *Graphs Comb.*, vol. 38, no. 2, p. 30, 2022.

[25] D. E. Knuth, *Selected Papers on Discrete Mathematics*. CSLI Publications, 2001.

[26] A. A. Sherstov, "Strong direct product theorems for quantum communication and query complexity," *SIAM J. Comput.*, vol. 41, no. 5, pp. 1122–1165, 2012, preliminary version in *Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing* (STOC), 2011.

[27] E. Kushilevitz and N. Nisan, *Communication complexity*. Cambridge University Press, 1997.

[28] R. de Wolf, "Quantum computing and communication complexity," Ph.D. dissertation, University of Amsterdam, 2001.

[29] A. A. Sherstov and A. A. Storozhenko, "The communication complexity of approximating matrix rank," in *Electronic Colloquium on Computational Complexity (ECCC)*, September 2024.