Quantum Time-Space Tradeoffs for Matrix Problems

Paul Beame*
Computer Science & Engineering
University of Washington
Seattle, WA, USA
beame@cs.washington.edu

Niels Kornerup†
Computer Science
The University of Texas at Austin
Austin, TX, USA
nielskornerup@utexas.edu

Michael Whitmeyer‡
Computer Science & Engineering
University of Washington
Seattle, WA, USA
mdwhit@cs.washington.edu

ABSTRACT

We prove lower bounds on the time and space required for quantum computers to solve a wide variety of problems involving matrices, many of which have only been analyzed classically in prior work.

Using a novel way of applying recording query methods we show that for many linear algebra problems—including matrix-vector product, matrix inversion, matrix multiplication and powering—existing classical time-space tradeoffs also apply to quantum algorithms with at most a constant factor loss. For example, for almost all fixed matrices A, including the discrete Fourier transform (DFT) matrix, we prove that quantum circuits with at most T input queries and S qubits of memory require $T = \Omega(n^2/S)$ to compute matrix-vector product Ax for $x \in \{0,1\}^n$. We similarly prove that matrix multiplication for $n \times n$ binary matrices requires $T = \Omega(n^3/\sqrt{S})$. Because many of our lower bounds are matched by deterministic algorithms with the same time and space complexity, our results show that quantum computers cannot provide any asymptotic advantage for these problems at any space bound.

We also improve the previous quantum time-space tradeoff lower bounds for $n \times n$ Boolean (i.e. AND-OR) matrix multiplication from $T = \Omega(n^{2.5}/S^{1/2})$ to $T = \Omega(n^{2.5}/S^{1/4})$ which has optimal exponents for the powerful query algorithms to which it applies. Our method also yields improved lower bounds for classical algorithms.

CCS CONCEPTS

• **Theory of computation** → *Quantum complexity theory*; **Quantum query complexity**.

KEYWORDS

quantum time-space tradeoffs, query complexity, matrix-problems

ACM Reference Format:
Paul Beame, Niels Kornerup, and Michael Whitmeyer. 2024. Quantum Time-Space Tradeoffs for Matrix Problems. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing (STOC '24), June 24–28, 2024, Vancouver, BC, Canada.* ACM, New York, NY, USA, 12 pages. https://doi.org/10.1145/3618260.3649700

*Research supported by NSF grant CCF-2006359
†Research supported by Schmidt Sciences Polymath award to David Soloveichik
‡Research supported by NSF grant CCF-2006359 and Simons Foundation grant 928589

STOC '24, June 24–28, 2024, Vancouver, BC, Canada

ACM ISBN 979-8-4007-0383-6/24/06
https://doi.org/10.1145/3618260.3649700

1 INTRODUCTION

Matrix computations are fundamental and critically important in scientific computing, optimization, and machine learning among others. If quantum computers have a significant advantage over classical devices for these computations then it would yield a wide range of applications for them.

This prospect has seemed promising in that. with input presented an unconventional form, the HHL algorithm of Harrow, Hassidim, and Lloyd [14, 19] can efficiently ϵ-approximate properties of solutions of well-conditioned linear systems with exponential quantum advantage over the best current classical algorithms. Many extensions of the HHL algorithm have also been proposed using the quantum singular value transform (qSVT) framework [17, 21]. Despite initial hope of exponential speed-up for this framework, a series of papers by Tang and co-authors, and others (e.g. [12, 13, 16, 27]) has shown that, given a comparable input format to the HHL algorithm, these quantum algorithms can be replaced by classical ones with only a polynomial blowup in the running time.

This begs the question: What is the conventional quantum complexity of standard classical problems like explicitly computing the linear-system solutions, multiplying or inverting matrices, computing matrix-vector products, and computing the low rank approximation of a matrix? By the polynomial method, we know that computing a single inner product of n-bit vectors requires $\Omega(n)$ quantum queries [6] but linear algebra computations generally involve $\Omega(n)$ or $\Omega(n^2)$ such computations. Sherstov [24], generalizing results of Klauck, Špalek, and de Wolf [20] for the OR function, gave a strong direct product lower bound for quantum query complexity proved using the polynomial method, which yields strong lower bounds for inner products involving many *disjoint* input vectors. However, the matrix problems in linear algebra are very far from direct product problems since the vectors involved are highly correlated with each other so this does not help.

We resolve these questions for quantum computation of a wide array of linear algebra problems, proving lower bounds for quantum computation that are asymptotically the same as the best classical lower bounds. That is, for each fixed bound on the amount of memory allowed, we derive asymptotically the same time lower bound for the quantum algorithm as the time lower bound on classical algorithms with the same number of classical bits. Quantum memory is an even more critical resource than classical memory since it is a measure of the maximum number of qubits that maintain coherence at any time during the algorithm's execution; the first general-purpose fault-tolerant quantum computers will likely have very limited memory. Since many of the problems also have deterministic algorithms whose time and memory usage matches our lower bounds, our results show that there is provably no asymptotic quantum advantage at all in solving these linear algebra problems!

Our lower bounds hold in a query model where algorithms perform arbitrary input-independent unitary transformations on their state between quantum queries. This is sufficiently general that our lower bounds also apply to any reasonable model of quantum computation—including quantum circuits with a (classical) input is stored in quantum-readable read only memory (QROM). The keys to proving our bounds are new results improving on strong direct product theorems for matrix-vector products and matrix multiplication. While our bounds have the form of such theorems (success probability decays exponentially with the number of outputs), they also apply with almost completely overlapping sets of inputs, in contrast to the disjoint inputs necessary for direct product theorems.

While there is a large body of work proving strong classical time-space tradeoffs (e.g. [2, 3, 7, 11, 22, 28]) and a large body of work analyzing unrestricted quantum query algorithms versus their classical randomized counterparts (e.g [4, 6, 10, 15, 23, 25, 26]), there are just a few previous papers that analyze the quantum memory required to make use of these quantum queries. Klauck, Špalek, and de Wolf [20] extended the classical method of Borodin and Cook [11] for proving time-space tradeoffs to quantum circuits using their strong direct production theorem for OR. They showed that algorithms making T quantum queries and using S qubits of quantum memory require $T = \Theta(n^{1.5}/S^{1/2})$ to sort lists of length n, and require $T = \Omega(n^{2.5}/S^{1/2})$ to compute $n \times n$ Boolean matrix product. Ambainis, Špalek, and de Wolf [5] extended this direct product approach to 2-sided error algorithms computing k-threshold functions producing similar tradeoff lower bounds for systems of linear inequalities/equalities. After a long gap in results because of the difficulty of applying these methods Hamoudi and Magniez [18], used a refinement of the *compressed oracle* method of Zhandry [29], which uses a *recording query* basis that allows one to keep track of a quantum query algorithm as a superposition of basis states that have a natural classical query interpretation, to prove that any quantum algorithm that finds K disjoint collisions in an input of length n with T quantum queries and S qubits of memory must have $T = \Omega(KN^{1/3}/S^{1/3})$ and reproved the sorting lower bound of [20] using this method.

Our linear algebra lower bounds and methods. The strongest classical time-space tradeoff lower bounds for linear algebra problems are due to Abrahamson [3] who developed a powerful general method based on matrix rigidity. This yields state-of-the-art lower bounds for computation of Fourier transforms, convolution, matrix-vector products, matrix multiplication, matrix inversion, matrix powering, and linear system solving. The lack of any analogous results for quantum computation has been a substantial gap in our understanding. Our results show that all of the linear algebraic time-space tradeoff lower bounds shown by Abrahamson [3] also apply to quantum computers that can adaptively decide when to produce output values. Since many of these classical lower bounds are tight, there is hybrid classical-quantum algorithms also provide no advantage. In the full paper [9] we use results of [8] to derive asymptotically equivalent lower bounds for the stronger model of quantum cumulative memory complexity. A summary of our time-space tradeoff lower bounds is in Table 1.

Table 1: Summary of our quantum lower bounds, along with prior work. Bounds apply for input elements coming from any fixed subset D of a field with $d = |D|$.

Problem	Quantum Bound	Source
Matrix-Vector Product	$T = \Omega(n^2 \log d\ /S)$	Cor 3.3
Discrete Fourier Transform	$T = \Omega(n^2 \log d\ /S)$	Cor 3.4
Convolution	$T = \Omega(n^2 \log d\ /S)$	Cor 3.4
Binary Integer Multiplication	$T = \Omega(n^2/(S \log^2 n))$	Cor 3.5
Matrix Triple Product $f(A, B, C) = ABC$	$T = \Omega(n^4 \log d\ /S)$	Cor 3.6
Matrix Cubing	$T = \Omega(n^4 \log d\ /S)$	Cor 3.6
Matrix Inversion	$T = \Omega(n^4 \log d\ /S)$	Cor 3.6
System of Linear Equations	$T = \Omega(n^3 \log d\ /S)$	Cor 3.7
Matrix Multiplication	$T = \Omega(n^3 \sqrt{\log d\ /S})$	Thm 4.1
Matrix Squaring	$T = \Omega(n^3 \sqrt{\log d\ /S})$	Full paper
Boolean Matrix Multiplication	$T = \Omega(n^{2.5}/S^{1/2})$ $T = \Omega(n^{2.5}/S^{1/4})$	[20] Thm 5.4
Boolean Matrix Squaring	$T = \Omega(n^{2.5}/S^{1/4})$	Full paper

So far, compressed oracle arguments have followed a two-step pattern: First, show that the total amplitude of states with some form of unusual progress (i.e., the partial information so far is unexpectedly determinative of the answer) is small, Then show that without unusual progress the total amplitude of quantum states where many outputs are produced is small by breaking the algorithm's final state into mutually orthogonal components, each with small amplitude on the correct answers, typically using classical ideas.

However, with linear algebra problems, there is no form of unusual progress and no clear way to break states up into mutually orthogonal basis states. Instead, for most of our bounds, we use the recording query framework to allocate portions of the algorithm's state to a small number of non-orthogonal components (or buckets) that share some set of inputs that they know nothing about. We can then apply a classical argument to show that each component must have small amplitude on correct answers obtain the overall bound using the triangle inequality. For matrix multiplication, we need an exponential number of buckets but use flexibility in bucket allocation to show that only a small number can be used to cover almost all the amplitude.

Improved bounds for Boolean matrix operations. We improve the previous lower bound for quantum algorithms computing Boolean matrix multiplication given in [20] from $T = \Omega(n^{2.5}/S^{1/2})$ to $T = \Omega(n^{2.5}/S^{1/4})$. We also improve the classical lower bound tradeoff of $T = \Omega(n^3/S)$ for circuits (where T is circuit size and S is the circuit width) shown in [20] to $T = \Omega(n^3/S^{1/2})$. This answers a question of Klauck, Špalek, and de Wolf [20] who ventured that this was the likely tight tradeoff for classical computation of Boolean matrix multiplication. Our classical bound also dominates Abrahamson's $T = \Omega(n^{3.5}/S$ bound for $S \geq n$ in the more general branching

program model [2]. The exponents of n and S in our bounds are optimal for the general circuit models to which they apply. We prove these bounds using a more sophisticated embedding of k-fold direct product of OR into an arbitrary subset of k outputs of Boolean matrix multiplication. This embedding hinges on the number of colors needed for a certain kind of coloring of subsets of the $n \times n$ grid.

Finally, we convert lower bounds for Boolean matrix-vector products and systems of inequalities given in [5, 20] to yield instances hard for space S that do not depend on S.

2 PRELIMINARIES

We define the binary entropy function $H_2 : [0,1] \to \mathbb{R}$, by $H_2(p) = -p\log_2 p - (1-p)\log_2(1-p)$.

Proposition 2.1 (Shannon). *The number of subsets of $[k]$ of size at most αk is at most $2^{H_2(\alpha)\,k}$.*

Definition 2.2. An $m \times n$ matrix is (g,h,c)-*rigid* iff every $k \times w$ submatrix where $k \le g$ and $w \ge n-h$ has rank at least ck. We call $(g,h,1)$-rigid matrices (g,h)-rigid.

Matrix rigidity is a robust notion of rank and is an important property for proving time-space and cumulative complexity lower bounds for linear algebra. Fortunately, Yesha gives an explicit example of such a matrix and Abrahamson proved that there are many rigid square matrices.

Proposition 2.3 (Lemma 3.2 in [28]). *The $n \times n$ Discrete Fourier Transform (DFT) matrix is $(n/4, n/4, 1/2)$ rigid.*

Proposition 2.4 (Lemma 4.3 in [3]). *There is a constant $\gamma \in (0, \frac{1}{2})$ such that at least a $1 - d^{-1}(2/3)^{\gamma n}$ fraction of the matrices over $D^{n\times n}$ with $|D| = d$ are $(\gamma n, \gamma n)$-rigid.*

Unitary quantum circuits with oracle states. Throughout this paper, we consider quantum circuits that seek to compute target functions $f : D^n \to R^m$. Let $d = |D|$ and assume the existence of a bijective map $v : D \to \{0,\ldots,d-1\}$ that gives us an ordering on the elements of D. A T query quantum circuit is specified using input independent unitaries $U_0,\ldots,U_T$. These unitaries define a sequence of quantum states $|\psi_1\rangle_C,\ldots|\psi_T\rangle_C$ that an algorithm enters during its execution. We can think of each of these state $|\psi_t\rangle_C$ as a linear combination of basis vectors $|i,p,w\rangle$ where i represents an index to query, p represents a phase for the query, and w contains all the remaining qubits of the state.

Similar to [4, 18, 29], we define a general oracle operator O that interacts with an input register that start in a state $|\psi_0\rangle_O$. Given a distribution $\mathcal{D}$ over D^n, we can make $|\psi_0\rangle_O = \sum_{X\in D^n}\sqrt{\Pr_{X'\sim\mathcal{D}}[X=X']}\,|X\rangle$ to represent an input sampled from $\mathcal{D}$. We define our oracle operator O as $O\,|i,p,w\rangle\,|X\rangle = \omega_d^{pv(x_i)}\,|i,p,w\rangle\,|X\rangle$. Thus the joint state of the input and quantum circuit at the end of the computation is given by $|\psi_T\rangle = U_T O \ldots O U_0\,|0\rangle_C\,|\psi_0\rangle_O$.

The output of the quantum circuit is determined by measuring the work register of $|\psi_T\rangle_C$ in the standard basis and applying some input-independent post-processing function q to interpret the result as an output $\tau \in R^J$ where $J \subseteq [m]$. The correctness of these outputs is then determined by measuring the input registers in the standard basis to obtain the input X and evaluating whether τ is consistent with $f(X)$ which we denote by writing $\tau||f(X)$. In general we can define the projector Π_k where:

$$\Pi_k = \sum_{\substack{i,p,w,x_1,\ldots,x_n \\ \text{s.t. } q(w)||f(x_1,\ldots,x_n) \\ \text{and } |q(w)|\ge k}} |i,p,w,x_1,\ldots,x_n\rangle\langle i,p,w,x_1,\ldots,x_n|$$

The probability that the circuit produces a correct partial assignment of at least k outputs is given by $\|\Pi_k\,|\psi_T\rangle\|^2$. For a given partial assignment $q(w)$ to some outputs, we can define $\Pi_{q(w)}$ to be the projection onto the values of $|X\rangle$ where $q(w)||f(X)$. More specifically we have that:

$$\Pi_{q(w)} = \sum_{\substack{x_1,\ldots,x_n \\ \text{s.t. } q(w)||f(x_1,\ldots,x_n)}} |x_1,\ldots,x_n\rangle\langle x_1,\ldots,x_n| \tag{1}$$

By construction when q always produces a partial assignment of at least k elements we have that $\Pi_k = \sum_{i,p,w}|i,p,w\rangle\langle i,p,w|\otimes\Pi_{q(w)}$.

Space Bounded Quantum Computation. Without loss of generality, we think of quantum circuits as starting in the all $|0\rangle$ state and cycling between applying input queries O, arbitrary input-independent computation U_t, and intermediate measurements as in Figure 1. Adopting the notation of [8], we will consider the set of consecutive O, U_t, and measurement gates as layer L_t. The space of layer L_t is the number of qubits that are passed from layer L_t to L_{t+1} and is denoted S_t. We define the space of a circuit as the maximum space of any layer, the time as the total number of layers, and the cumulative memory as the sum over all the S_t.

Intermediate measurements enable circuits to produce parts of their output early and discard unnecessary ancillary qubits. Some prior quantum time-space tradeoff lower bounds required the quantum circuit to declare which outputs are produced at each layer (e.g. sorting, Boolean matrix multiplication, and systems of linear inequalities [5, 20]); however the recent collision-finding bounds in [18] extend the output model for quantum circuits to include indicator qubits specifying which (if any) outputs are being produced at each layer. This allows them to prove lower bounds against quantum algorithms that dynamically decide when they want to produce outputs based on their observed inputs. While our Boolean matrix bounds build on those in [20] and thus require a fixed time for each output bit, our linear algebra bounds work with this dynamic output model.

Our time-space tradeoffs follow the Borodin-Cook method, and thus rely on dividing a quantum circuit into blocks that each are unlikely to produce many correct outputs. We prove that these blocks cannot produce many outputs in the unitary quantum circuits model and then apply the results to our space bounded model using the differed measurement principle. After the first block, a quantum circuit will have some input-dependent state that can help it produce more outputs. Fortunately, a result by Aaronson (modified to work on mixed states) lets us bound how much this initial state can amplify the success probability.

Proposition 2.5 ([1]). *Let C be a quantum circuit, ρ be an S-qubit (possibly mixed) state, and π_{mix} be the S-qubit maximally mixed state. If C starting in initial state ρ produces some output z with probability*

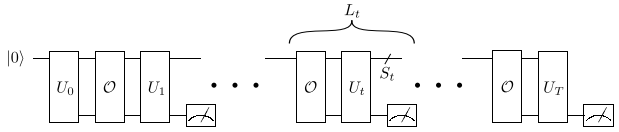

Figure 1: A general quantum circuit with T queries.

p, then C starting in state π_{mix} will produce z with probability at least $p/2^{2S}$.

The recording query technique and quantum lower bounds. Here we review the methods developed in [18, 29] that allow us to analyze what a quantum circuit learns about its input by making quantum queries. We will assume that $|\psi_0\rangle_O$ is the equal superposition state over all inputs. We can exchange the general query operator O with a recording query operator $\mathcal{R}$ defined as follows:

Definition 2.6 (adapted from [18]). Let $\mathcal{S}_1$ be the unitary operator where: $\mathcal{S}_1|\perp\rangle = \frac{1}{\sqrt{d}}\sum_{y\in D}|y\rangle$, $\mathcal{S}_1\frac{1}{\sqrt{d}}\sum_{y\in D}|y\rangle = |\perp\rangle$, and $\mathcal{S}_1\frac{1}{\sqrt{d}}\sum_{y\in D}\omega_d^{p\,\nu(y)}|y\rangle = \frac{1}{\sqrt{d}}\sum_{y\in D}\omega_d^{p\,\nu(y)}|y\rangle$ for all $p\neq 0$.

Let $\mathcal{S} = (I)_{i,p,w}\otimes(\mathcal{S}_1^{\otimes n})_{x_1,\ldots,x_n}$ and O be the standard oracle operator $O|i,p,w,x_1,\ldots,x_n\rangle = \omega_d^{p\,\nu(x_i)}|i,p,w,x_1,\ldots,x_n\rangle$. Then the *recording query oracle operator* $\mathcal{R}$ is defined as $\mathcal{SOS}$.

$\mathcal{S}_1$ introduces $\perp$ as a new value for the input registers. The $\perp$ symbol indicates that the algorithm knows nothing about that register of the oracle. By adding and correctly manipulating the $\perp$ symbols in the oracle's registers, we can record what the algorithm knows about the input. Since $\mathcal{S}^2 = I$, we can exactly characterize how the states of quantum circuits with oracles O and $\mathcal{R}$ relate.

Proposition 2.7 ([18]). *Let C be a quantum circuit that for each $j\le t$ applies unitary U_j after the j-th query. Let $\mathcal{S}$ be the unitary operation and $\mathcal{R}$ be the recording query oracle from Definition 2.6. Let*

$$|\psi_t\rangle = U_t O U_{t-1}\ldots U_1 O U_0\left(|0\rangle_{i,p,w}\otimes\frac{1}{d^{n/2}}\sum_{x_1,\ldots,x_n\in D}|x_1,\ldots,x_n\rangle_{x_1,\ldots,x_n}\right)$$

$$|\phi_t\rangle = U_t \mathcal{R} U_{t-1}\ldots U_1\mathcal{R}U_0\left(|0\rangle_{i,p,w}\otimes|\perp\rangle_{x_1,\ldots,x_n}\right)$$

be the states of C with oracle O or $\mathcal{R}$ respectively. Then $|\psi_t\rangle = \mathcal{S}|\phi_t\rangle$.

In other words, it is impossible to distinguish the final state $|\psi_T\rangle$ of a circuit with standard oracle O from the output with recording oracle $\mathcal{R}$ if we apply $\mathcal{S}$ to the registers of $\mathcal{R}$ after the final query. Thus the success probability of a quantum circuit with T queries is given by $\||\Pi_{succ}|\psi_T\rangle\|^2 = \||\Pi_{succ}\mathcal{S}|\phi_T\rangle\|^2$. Note that while $|\phi_T\rangle$ may have inputs in the $\perp$ state, Proposition 2.7 tells us that $\mathcal{S}|\phi_T\rangle$ will never have an input that is $\perp$. Thus when considering recording query oracles, it is safe to keep our current definitions of Π_{succ} and $\Pi_{q(w)}$ which will always project out any basis state where an input is assigned to $\perp$. We leverage the following property of $|\phi_T\rangle$ to bound the success probability of short quantum circuits.

Proposition 2.8 (Fact 3.2 in [18]). *The state $|\phi_t\rangle$ from Proposition 2.7 is a linear combination of basis states $|i,p,w,x_1,\ldots,x_n\rangle$ where at most t of the x_i are different from $\perp$.*

3 QUANTUM MATRIX-VECTOR PRODUCTS

In this section, we consider the task of — for a fixed matrix $A\in\mathbb{F}^{m\times n}$ — computing the function $f(x) = Ax$ for inputs $x\in D^m$ using a quantum circuit. We note that this is a fundamentally harder task than is considered in many quantum machine learning papers (for example [19]) as we require the circuit to output a classical vector $y\in\mathbb{F}^n$ rather than either a quantum state encoding the entries of y in the amplitudes or an estimate of $y^\dagger My$.

Also unlike many prior quantum time-space tradeoffs, including sorting [8, 18, 20] and Boolean matrix multiplication [20] (and our Theorem 5.4), but like those of [18] for disjoint collisions, our matrix-vector product and matrix multiplication lower bounds apply to circuits that can adaptively decide when to produce each output based on the observed inputs.

Theorem 3.1. *Let m be $n^{O(1)}$. Let A be an $m\times n$ matrix over a field $\mathbb{F}$ that is $(g(m),h(n),c)$-rigid for $c\in(0,1/2]$. Then any quantum circuit using time T and space S that computes a function $f: D^n\to\mathbb{F}^m$ for $D\subseteq\mathbb{F}$ with $d = |D|$ given by $f(x) = Ax$ with success probability larger than 2^{-S} requires that T is $\Omega(g(m)\,h(n)\log d\,/S)$; more precisely, T must be $\Omega(\min\{g(m)\,n\log d,\ m\,h(n)\log d\}/S)$.*

This theorem follows from the following key lemma, proven in the next subsection, which lets us bound the number of correct outputs produced by a shallow quantum circuit.

Lemma 3.2. *Let A be any (k,h,c)-rigid $m\times n$ matrix over a finite field $\mathbb{F}$ and let $f: D^n\to\mathbb{F}^m$ for $D\subseteq\mathbb{F}$ be defined by $f(x) = Ax$. Then for $\alpha > 0$ and for input x sampled uniformly from D^n and any quantum circuit C with at most αh queries to x, the probability that C produces k correct output values of $f(x)$ is at most $\lceil h/(ck)\rceil\,(2^{H_2(\alpha)}/|D|^{1-\alpha})^{ck}$.*

Note: For $\alpha\le 0.1717$ we have $1-\alpha-H_2(\alpha) > 1/6$ and hence the bound is at most $\lceil h/(ck)\rceil|D|^{-ck/6}$ for $d\ge 2$.

Proof of Theorem 3.1 from Lemma 3.2. First observe that as $S\ge\log_2 n$ and $T\ge n$ we know that $T\cdot S$ is $\Omega(n\log n)$ which is $\Omega(g(m)\,n\log|D|)$ if $g(m) < (12/c)\log_d n$. Therefore we can assume without loss of generality that $g(m)\ge(12/c)\log_d n$.

Let C be a quantum circuit with T queries and space S, write $h = h(n)$, $g = g(m)$, and let $\alpha = 0.1717$. We partition C into $\lceil T/(\alpha h)\rceil$ sub-circuits that each have at most αh queries. By combining Proposition 2.5 and Lemma 3.2, we know that each sub-circuit can produce $k\le g$ correct outputs with probability at most $2^{2S}\lceil h/(ck)\rceil\,d^{-ck/6}\le h\,2^{2S}d^{-ck/6}$.

Now suppose that $h\,2^{2S}d^{-cg/6} > 2^{-S}/T$. Then $T2^{3S} > d^{cg/6}/h\ge d^{cg/6}/n\ge d^{cg/12}$ by the assumption on g. Since $S\ge\log_2 n$ and T is at most polynomial in n (or the bound applies already), $T2^{3S}$ is at most $2^{c'S}$ for some constant $c' > 0$. This implies that S is $\Omega(g(m)\log d)$ and since $T\ge n$, we get that $T\cdot S$ is $\Omega(g(m)\,n\log|D|)$.

Otherwise set $k\le g$ to be the smallest integer such that $h\,2^{2S}d^{-ck/6}\le 2^{-S}/T$. Then the probability that a sub-circuit produces k correct outputs is at most $2^{-S}/T$. This gives $k = \lceil[6\log_2(hT)+18S]/(c\log_2 d)\rceil$, which is at most $c^*S/\log_2 d$ for some constant $c^* > 0$ since S is $\Omega(\log n)$ which is $\Omega(\log(hT))$.

Taking a union bound over the sub-circuits, the probability that any of them produces k correct outputs is at most 2^{-S}. Since f has m outputs, this means that $\lceil T/(\alpha h)\rceil\,(k-1)\ge m$ Since $T\ge n\ge\alpha h$,

we have $2Tk \geq \alpha mh$. Plugging in our upper bound on k we have that $2c^*TS/\log_2 d \geq \alpha mh$ and hence $T \cdot S$ is $\Omega(mh \log d)$ which is $\Omega(m\, h(n) \log |D|)$ as claimed. □

Applying Theorem 3.1 with Proposition 2.4 we obtain:

Corollary 3.3. *Let $\mathbb{F}$ be a field and $D \subseteq \mathbb{F}$ such that $d = |D| \geq 2$. For all but a $2^{-\Theta(n)}$ fraction of $A \in D^{n\times n}$, quantum circuits using space S require $\Omega(n^2 \log d\ /S)$ queries to compute Ax for $x \in D^n$.*

Using Theorem 3.1 we many time-space lower bounds for related problems. See the full paper [9] for the full details, which are identical to their classical counterparts proven in [3].

Since the discrete Fourier transform (DFT) matrix is rigid [28], Theorem 3.1 yields a time-space tradeoff for computing the DFT. Likewise since convolution between random vectors can be expressed as a matrix-vector product with a random Toeplitz matrix — which is rigid with high probability [3] — we also get a time-space tradeoff for convolution.

Corollary 3.4. *Let $\mathbb{F}$ be a field and $D \subseteq \mathbb{F}$ such that $d = |D|$. Any quantum circuit that computes the discrete Fourier transform of vectors in D^n (or the convolution of pairs of vectors in D^n) in time T and space S with probability at least 2^{-S} requires T to be $\Omega(n^2 \log d\ /S)$.*

Since convolution of vectors in $\{0,1\}^n$ is a sub-function of multiplication between $2n$ $\lceil \log n \rceil$ bit binary numbers we obtain:

Corollary 3.5. *A quantum circuit that multiplies two n bit binary numbers in time T and space S with probability at least 2^{-S} requires T to be $\Omega(n^2/(S \log^2 n))$.*

Let $B(\mathcal{Y})$ denote the vector formed by stacking the transposes of the rows of matrix $B(\mathcal{Y})$. Then $Y = ABC$ iff $\mathcal{Y} = (A \otimes C^T)\mathcal{B}$ [3]. When A and C are random matrices, $(A \otimes C^T)$ is a rigid matrix [3]. Thus we can use Theorem 3.1 to get a time-space tradeoff for computing the product of three matrices. Since ABC can be embedded as a sub-function of $f(X) = X^3$ and $f(X) = X^{-1}$, this also gives us an equivalent tradeoff for matrix cubing and matrix inversion [3].

Corollary 3.6. *Let $\mathbb{F}$ be a field and $D \subseteq \mathbb{F}$ such that $d = |D|$. Any quantum circuit that computes the product ABC on inputs $A, B, C \in D^{n\times n}$, or the cube A^3 or inverse A^{-1} in time T and space S with probability at least 2^{-S} requires T that is $\Omega(n^4 \log d\ /S)$.*

Since it is possible to invert a matrix by solving n systems of n linear equations, one of those systems must take $\Omega(n^3 \log d\ /S)$ time to solve.

Corollary 3.7. *Let $\mathbb{F}$ be a field and $D \subseteq \mathbb{F}$ such that $d = |D|$. Any quantum circuit that solves any $n \times n$ system of linear equations over D in time T and space S with probability at least 2^{-S} requires T that is $\Omega(n^3 \log d\ /S)$*

Following the methods of [8], all of the quantum time-space product lower bounds proven in this section can be extended to asymptotically equivalent lower bounds on the stronger notion of cumulative memory complexity.

3.1 Bounding the Success Probability of Small Depth Quantum Circuits

Proof of Lemma 3.2. Let $d = |D|$. For simplicity we will assume that $q(w)$—the output as a function of the measured value of the work register—always produces k outputs.[1] Let A be a (k, h, c)-rigid matrix. By Proposition 2.8 after $t \leq \alpha h$ queries in the recording query oracle model, we can write the state as:

$$|\phi_t\rangle = \sum_{i,p,w} \sum_{I\subseteq[n],\ |I|\leq t,\ y\in D^I} \alpha_{i,p,w,I,y} \,|i,p,w\rangle\, |y\rangle_I\, |\perp\rangle_{[n]\setminus I} \quad (2)$$

for some $\alpha_{i,p,w,I,y}$ with $\sum_{i,p,w,I,y} |\alpha_{i,p,w,I,y}|^2 = 1$. Thus by Proposition 2.7, the final state of the algorithm in the non-recording query oracle setting is given by: $|\psi_t\rangle = \mathcal{S}\,|\phi_t\rangle$ and since $\mathcal{S}$ behaves as the identity on $|\psi\rangle_C$ and the $|i,p,w\rangle$ are orthogonal basis states, we can rewrite $|\psi_t\rangle$ as:

$$\sum_{i,p,w} \beta_{i,p,w} \,|i,p,w\rangle \otimes \left[\mathcal{S}_1^{\otimes n} \sum_{I\subseteq[n],\ |I|\leq t,\ y\in D^I} \beta_{I,y}^{i,p,w}\, |y\rangle_I\, |\perp\rangle_{[n]\setminus I}\right]$$

for some $\beta_{i,p,w}$ and $\beta_{I,y}^{i,p,w}$ such that $\alpha_{i,p,w,I,y} = \beta_{i,p,w}\beta_{I,y}^{i,p,w}$, $\sum_{i,p,w} |\beta_{i,p,w}|^2 = 1$ and for each choice of i,p,w, we have that $\sum_{I,y} |\beta_{I,y}^{i,p,w}|^2 = 1$. With this decomposition, the success probability, which is $\left\|\Pi_k \mathcal{S}\,|\phi_t\rangle\right\|^2$ where $\Pi_{q(w)}$ is defined as in Equation (1) and is the projection of Π_k onto fixed values of $q(w)$, equals

$$\left\| \sum_{i,p,w} \beta_{i,p,w} \,|i,p,w\rangle \otimes \left[\Pi_{q(w)}\mathcal{S}_1^{\otimes n} \sum_{\substack{I\subseteq[n],\ |I|\leq t\\ y\in D^I}} \beta_{I,y}^{i,p,w}\, |y\rangle_I\, |\perp\rangle_{[n]\setminus I}\right]\right\|^2$$

Since $\sum_{i,p,w} |\beta_{i,p,w}|^2 = 1$ and the basis states $|i,p,w\rangle$ are orthogonal, we have

$$\left\|\Pi_k \mathcal{S}\,|\phi_t\rangle\right\|^2 \leq \max_{i,p,w} \left\|\Pi_{q(w)}\mathcal{S}_1^{\otimes n} \sum_{\substack{I\subseteq[n],\ |I|\leq t\\ y\in D^I}} \beta_{I,y}^{i,p,w}\, |y\rangle_I\, |\perp\rangle_{[n]\setminus I}\right\|^2. \quad (3)$$

We now fix i,p,w and let $A_{q(w)}$ be the submatrix of A restricted to the rows defined by the set of the k output values U associated with $q(w)$. We can describe $\Pi_{q(w)}$ as a projection onto basis states $|x\rangle = |x_1, \ldots, x_n\rangle$ such that $A_{q(w)}x = q(w)$.

Since the basis states $|y\rangle_I\,|\perp\rangle_{[n]\setminus I}$ for distinct I are orthogonal in the recording query basis, they remain orthogonal in the standard basis after the $\mathcal{S}$ operator is applied. However, the subsequent application of the $\Pi_{q(w)}$ projector makes these vectors no longer orthogonal.

To handle this, we bucket the sets $I \subseteq [n]$ with $|I| \leq t$ into a small number of buckets, $\mathcal{B}_1, \ldots$, so that for each bucket $\mathcal{B}_\ell$ we can bound:

$$\mu_\ell = \left\|\Pi_{q(w)}\mathcal{S}_1^{\otimes n} \sum_{I\in\mathcal{B}_\ell, y\in D^I} \beta_{I,y}^{i,p,w}\, |y\rangle_I\, |\perp\rangle_{[n]\setminus I}\right\|^2$$

and then we can use Cauchy-Schwarz to bound the success probability as a sum of the μ_ℓ.

Our key observation is that if a bucket of recording query basis states completely misses querying a fixed set of input variables that could determine a set of r output values, then one cannot do better than randomly guess those output values and the total contribution from that bucket has amplitude at most $d^{-r/2}$.

[1] If in general $q(w)$ produces more than k outputs, we only consider its first k outputs.

Lemma 3.8. *Let $U \subseteq [m]$ be a set of output indices and $V \subseteq [n]$ be a set of input indices with $|V| = |U| = r$ such that the submatrix $A_{U,V}$ is full rank. Fix $q \in \mathbb{F}^U$ and define Π_q to be the projection map onto the span of the set of basis states $|x\rangle$ with $x \in D^n$ such that $A_U x = q$. Then for any collection $\mathcal{B}$ of sets $I \subseteq [n] \setminus V$ and any quantum state $\sum_{I\in\mathcal{B},\ y\in D^I} \eta_{I,y} |y\rangle_I |\perp\rangle_{[n]\setminus I}$ we have*

$$\Big\|\Pi_q \mathcal{S}_1^{\otimes n} \sum_{I\in\mathcal{B},\ y\in D^I} \eta_{I,y} |y\rangle_I |\perp\rangle_{[n]\setminus I}\Big\|^2 \le \frac{1}{d^r}.$$

Proof. By definition each $I \in \mathcal{B}$ satisfies $I \cap V = \emptyset$, so

$$\begin{aligned}
&\Pi_q \mathcal{S}_1^{\otimes n} \sum_{I\in\mathcal{B},\ y\in D^I} \eta_{I,y} |y\rangle_I |\perp\rangle_{[n]\setminus I} \\
&= \Pi_q \mathcal{S}_1^{\otimes n}\Big[|\perp\rangle_V \otimes \sum_{I\in\mathcal{B},\ y\in D^I} \eta_{I,y} |y\rangle_I |\perp\rangle_{[n]\setminus(I\cup V)} \Big] \\
&= \Pi_q\big[\mathcal{S}_1^{\otimes j} |\perp\rangle_V \otimes \mathcal{S}_1^{\otimes(n-j)} \sum_{I\in\mathcal{B},\ y\in D^I} \eta_{I,y} |y\rangle_I |\perp\rangle_{[n]\setminus(I\cup V)} \Big] \\
&= \Pi_q\Big[\sum_{y'\in D^V} \frac{1}{\sqrt{d^r}} |y'\rangle_V \otimes \mathcal{S}_1^{\otimes(n-j)} \sum_{I\in\mathcal{B},\ y\in D^I} \eta_{I,y} |y\rangle_I |\perp\rangle_{[n]\setminus(I\cup V)} \Big].
\end{aligned}$$

Now $\mathcal{S}_1^{\otimes(n-j)} \sum_{I\in\mathcal{B},\ y\in D^I} \eta_{I,y} |y\rangle_I |\perp\rangle_{[n]\setminus(I\cup V)}$ is equal to $\sum_{z\in(D\cup\{\perp\})^{[n]\setminus V}} \delta_z |z\rangle_{n\setminus V}$ for some amplitudes δ_z satisfying $\sum_{z\in(D\cup\{\perp\})^{[n]\setminus V}} |\delta_z|^2 = 1$. For each value of $z \in D^{[n]\setminus V}$, since the sub-matrix $A_{U,V}$ is invertible, there is a unique value $y_z \in D^V$ such that $A_U(y_z \cup z) = q$ so we get that

$$\begin{aligned}
&\Big\|\Pi_q \mathcal{S}_1^{\otimes n} \sum_{I\in\mathcal{B},\ y\in D^I} \eta_{I,y} |y\rangle_I |\perp\rangle_{[n]\setminus I}\Big\|^2 \\
&= \Big\|\Pi_q\Big[\sum_{y'\in D^V} \frac{1}{\sqrt{d^r}} |y'\rangle_V \otimes \sum_{z\in(D\cup\{\perp\})^{n-j}} \delta_z |z\rangle_{[n]\setminus V} \Big]\Big\|^2 \\
&= \Big\| \frac{1}{\sqrt{d^r}} \cdot \Pi_q\Big[\sum_{z\in D^{[n]\setminus V}} \delta_z \sum_{y'\in D^V} |y'\rangle_V |z\rangle_{n\setminus V} \Big]\Big\|^2 \\
&= \Big\| \frac{1}{\sqrt{d^r}} \sum_{z\in D^{[n]\setminus V}} \delta_z |y_z\rangle_V |z\rangle_{n\setminus V})\Big\|^2 \le \frac{1}{d^r}
\end{aligned}$$

since $\sum_{z\in D^{[n]\setminus V}} |\delta_z|^2 \le 1$. □

Next we decompose the set of all I with $|I| \le t$ into buckets so that we can apply the above.

Lemma 3.9. *Let A be a (k,h,c)-rigid matrix and let $k' = \lceil ck \rceil$. Then for every subset U of k rows of A, there is a collection of disjoint k'-subsets of columns from $[n]$, $V_1, \ldots, V_\ell$ for $\ell = \lceil h/k' \rceil \le \lceil h/(ck) \rceil$ and corresponding sets of rows $U_1, \ldots, U_\ell \subseteq U$ such that for each $j \in [\ell]$, the $k' \times k'$ submatrix A_{U_j,V_j} is full rank. (In particular the union, W, of the sets V_j has size at least h.) If $c = 1$ then all $U_j = U$.*

Proof. Fix $U \in [m]$ with $|U| = k$. The following procedure constructs such a collection, one set at a time. We maintain a subset of W columns that is the union of the V_j constructed so far. Suppose that $|W| < h$. Then, by the (k,h,c)-rigidity of A, the submatrix $A_{U,[n]\setminus W}$ has rank at least k'. Hence there is a $k' \times k'$ submatrix A_{U_j,V_j} of $A_{U,[n]\setminus W}$ that has full rank k'. We now add V_j to the collection of k'-sets of columns, record its corresponding row set U_j, and set $W \leftarrow W \cup V_j$. This produces exactly $\lceil h/k' \rceil$ subsets. □

Fix the collection of sets $V_1, \ldots, V_\ell$ given by Lemma 3.9. Let $k'' = \lfloor \alpha k' \rfloor$. Suppose that $V_j = \{i_1, \ldots, i_{k'}\} \subseteq [n]$ with $i_1 \le \cdots \le i_{k'}$. For each $\lambda \in \binom{[k']}{k''}$, define the set V_j^λ to be the subset of V_j that has the k'' elements of V_j indexed by λ removed. (That is, $i_{j'} \notin V_j^\lambda$ iff $j' \in \lambda$.) Then $|V_j^\lambda| = k' - k'' \ge c(1-\alpha)k$. There are a total of $\binom{k'}{k''} \le 2^{H_2(\alpha)\,k'}$ possible values of λ and hence $\lceil h/k' \rceil \cdot 2^{H_2(\alpha)\,k'}$ sets of the form V_j^λ. These sets have two useful properties: first any subset of $[n]$ with size at most αh must miss some V_j^λ and second if the entries of x corresponding to some V_j^λ are uniformly random, then for any set of k indices in Ax, at least $c(1-\alpha)k$ of these values are also uniformly random.

Lemma 3.10. *For $t \le \alpha h$ and every $I \subseteq [n]$ with $|I| \le t$, there is some $j \le \lceil h/k' \rceil$ and $\lambda \in \binom{[k']}{k''}$ such that $I \subseteq [n] \setminus V_j^\lambda$.*

Proof. Fix such a set I with $|I| \le t$. Since $t \le \alpha h$, $|\bigcup_{j\in[\ell]} V_j| \ge h$, and the sets V_j are disjoint, by averaging there is some set V_j that has at most an α fraction of its elements in I. Hence V_j has at most $k'' \le \alpha k'$ elements of I. Choose a set $\lambda \in \binom{[k']}{k''}$ that contains the indices within V_j of all of the elements of $V_j \cap I$. Then by construction $I \cap V_j^\lambda = \emptyset$. □

By applying Lemma 3.10 we can associate each $I \subseteq [n]$ with $|I| \le t$ with a pair (j,λ) such that $I \in [n] \setminus V_j^\lambda$ and define bucket $\mathcal{B}_j^\lambda$ to consist of all such sets I associated with pair (j,λ).(Though some sets I could be associated with multiple pairs (j,λ), we choose only one such pair for each I.) Further, define a set $U_j^\lambda \subseteq U_j \subseteq [m]$ of the rows of $A_{q(w)}$ with $|U_j^\lambda| = k' - k''$ such that the submatrix $A_{U_j^\lambda,V_j^\lambda}$ is full rank. Such a subset of rows must exist since A_{U_j,V_j^λ} is a full rank matrix. Then let $q_j^\lambda = q(w)|_{U_j^\lambda}$ be the portion of the assignment $q(w)$ on the rows of U_j^λ.

We are now ready to provide an upper bound on the success probability from Equation (3).

$$\begin{aligned}
&\Big\|\Pi_{q(w)} \mathcal{S}_1^{\otimes n} \sum_{I\subseteq[n],\ |I|\le t,\ y\in D^I} \beta_{I,y}^{i,p,w} |y\rangle_I |\perp\rangle_{[n]\setminus I}\Big\|^2 \\
&= \Big\|\Pi_{q(w)} \mathcal{S}_1^{\otimes n} \sum_{j\in[\ell]} \sum_{\lambda\in\binom{[k']}{k''}} \sum_{I\in\mathcal{B}_j^\lambda,\ y\in D^I} \beta_{I,y}^{i,p,w} |y\rangle_I |\perp\rangle_{[n]\setminus I}\Big\|^2 \\
&\le \Big\| \sum_{j\in[\ell]} \sum_{\lambda\in\binom{[k']}{k''}} \Pi_{q_j^\lambda} \mathcal{S}_1^{\otimes n} \sum_{I\in\mathcal{B}_j^\lambda,\ y\in D^I} \beta_{I,y}^{i,p,w} |y\rangle_I |\perp\rangle_{[n]\setminus I}\Big\|^2. \quad (4)
\end{aligned}$$

Applying Lemma 3.8 with $r = k' - k''$, $q = q_j^\lambda$, $U = U_j^\lambda$, $V = V_j^\lambda$, and $\mathcal{B} = \mathcal{B}_j^\lambda$, we have that

$$\Big\|\Pi_{q_j^\lambda} \mathcal{S}_1^{\otimes n} \sum_{I\in\mathcal{B}_j^\lambda,\ y\in D^I} \beta_{I,y}^{i,p,w} |y\rangle_I |\perp\rangle_{[n]\setminus I}\Big\|^2 \le \frac{1}{d^{k'-k''}} \le \frac{1}{d^{(1-\alpha)\,k'}}.$$

and hence using Equation (4) we obtain that

$$\big\|\Pi_k \mathcal{S} |\phi_t\rangle\big\|^2 \le \ell \binom{k'}{k''} / d^{(1-\alpha)\,k'} \le \lceil h/k' \rceil\, (2^{H_2(\alpha)}/d^{(1-\alpha)})^{k'}.$$

Without loss of generality in our desired bound we can assume that $2^{H_2(\alpha)}/d^{(1-\alpha)} < 1$. Therefore the bound still applies when we replace k' by the potentially smaller ck as required. □

4 QUANTUM MATRIX MULTIPLICATION

While our matrix-vector product lower bound has led to all the applications so far, including the matrix triple product lower bound, our matrix multiplication lower bound requires a separate argument using ideas from the classical lower bound in [3]. Implementing this requires a much more subtle way of applying our bucketing method for states that allows us to concentrate on just a subset of the buckets containing most of the total amplitude.

Theorem 4.1. *Let $\mathbb{F}$ be a field and $D \subseteq \mathbb{F}$ with $d = |D|$. Then any quantum circuit C that uses time T and space S and computes the function $f : D^{2n^2} \to \mathbb{F}^{n^2}$ given by $f(A,B) = AB$ with success probability larger than $1/T$ must have T that is $\Omega(n^3\sqrt{\log d\,/S})$.*

Lemma 4.2. *Let $\gamma \in (0,1/2)$ and $f : D^{n^2} \times D^{n^2} \to \mathbb{F}^{n^2}$ for $D \subseteq \mathbb{F}$ with $|D| = d$ be defined by $f(A,B) = AB$. Then for any constant $\beta > 0$ and quantum circuit C with at most $h = \beta\gamma n\sqrt{k/2}$ queries to input matrices A, B sampled uniformly from D^{n^2}, the probability that A and B are $(\gamma n, \gamma n)$-rigid and C produces k correct output values of $f(A,B)$ is at most $16\min(k,n)^{\sqrt{k/2}}(2^{H_2(4\beta)}/d^{1-4\beta})^{k/4}$*

Note that for $\beta \le 0.0429$ we have $1 - 4\beta - H_2(4\beta) > 1/6$ so the bound is at most $16\min(k,n)^{\sqrt{k/2}}d^{-k/24}$.

Proof of Theorem 4.1 from Lemma 4.2. Let $\gamma \in (0,1/2)$ be the constant given by Proposition 2.4. By that proposition, the probability that either of two matrices A and B chosen uniformly randomly from D^{n^2} is not $(\gamma n, \gamma n)$-rigid is at most $2d^{-1}(2/3)^{\gamma n}$. Let C be a quantum circuit with T queries and space S. Let $\beta = 0.0429$, $d = |D|$, and set $k = \lceil 48(5S+5)/\log_2 d \rceil$. We partition C into $\left\lceil T/(\beta\gamma n\sqrt{k/2})\right\rceil$ sub-circuits that have at most $\beta\gamma n\sqrt{k/2}$ queries each. Without loss of generalities there are at most n^2 such sub-circuits. By combining Proposition 2.5 with Lemma 4.2, we know that for a uniformly random input, the probability that A and B are $(\gamma n, \gamma n)$-rigid matrices and a fixed sub-circuit can produce k outputs is at most $16k^{\sqrt{k/2}}2^{2S}d^{-k/24}$. Therefore the probability that A and B are $(\gamma n, \gamma n)$-rigid matrices and one of the sub-circuits produces k correct outputs is at most $16k^{\sqrt{k/2}}2^{2S}d^{-k/24}n^2$. Combining this with the probability that one of A or B is not $(\gamma n, \gamma n)$-rigid, the probability that there is a sub-circuit that produces k correct outputs is at most $16k^{\sqrt{k/2}}2^{2S}d^{-k/24}n^2 + 2d^{-1}(2/3)^{2\gamma n}$. Since we can assume without loss of generality that $T \le n^3$, for sufficiently large n, $2d^{-1}(2/3)^{2\gamma n} \le 1/(2T)$ and $k^{\sqrt{k/2}} \le 2^{k/48} \le d^{k/48}$. Plugging in our value of k and the fact that $S \ge \log_2 n$ without loss of generality gives a probability of at most

$$16k^{\sqrt{k/2}}2^{2S}d^{-k/24}n^2 + 2d^{-1}(2/3)^{2\gamma n} \le 16^{2^{2S}}d^{-k/48}n^2 + 1/(2T) \le 1/(2T) + 1/(2T) = 1/T.$$

Since C must be correct with probability larger than $1/T$, this implies that $(k-1)\left\lceil T/(\beta\gamma n\sqrt{k/2})\right\rceil \ge n^2$. Plugging in our value of k gives us that T is $\Omega(n^3\sqrt{\log d}/\sqrt{S+\log T})$. Since $S \ge \log_2 n$ and our bound trivially holds when T is $\omega(n^3\sqrt{\log d})$ there is a constant $c > 0$ such that $cS \ge \log_2 T$. Thus T is $\Omega(n^3\sqrt{\log d/S})$. □

4.1 Bounding Success

Proof of Lemma 4.2. Let $C = AB$, $\Pi_{\text{rigid}(A)}$ ($\Pi_{\text{rigid}(B)}$) be the projection onto inputs where A and B are $(\gamma n, \gamma n)$-rigid matrices, and define $\Pi_{\text{rigid}} = \Pi_{\text{rigid }A}\Pi_{\text{rigid }B}$. Assume that $q(w)$—the output as a function of the measured value of the work register—produces exactly k outputs; we ignore anything it produces after the first k. We will use $[A]$ to denote the set of indices of elements in A and likewise for $[B]$ and $[C]$. By Proposition 2.8, after $t \le h$ queries in the recording query basis, our state can be written as:

$$|\phi_t\rangle = \sum_{\substack{i,p,w \\ E\subseteq[A],F\subseteq[B],\ |E|+|F|\le t \\ x\in D^E, y\in D^F}} \alpha_{i,p,w,E,F,x,y}\,|i,p,w\rangle\,|x\rangle_E\,|\perp\rangle_{[A]\setminus E}\,|y\rangle_F\,|\perp\rangle_{[B]\setminus F}$$

for some $\alpha_{i,p,w,E,F,x,y}$ with $\sum_{i,p,w,E,F,x,y}|\alpha_{i,p,w,E,F,x,y}|^2 = 1$. We first apply an analogous series observations and decompositions to those that allowed us to derive (3) from (2) in the case of matrix-vector product: By Proposition 2.7, the final state in the standard oracle setting $|\psi_t\rangle = \mathcal{S}\,|\phi_t\rangle$. Because $\mathcal{S}$ behaves as the identity on $|\psi\rangle_C$ and each distinct choice of $|i,p,w\rangle$ gives an orthogonal basis state, $|\psi_t\rangle$ equals $\sum_{i,p,w}\beta_{i,p,w}\,|i,p,w\rangle \otimes \Psi_{i,p,w}$ for some $\beta_{i,p,w}$ with $\sum_{i,p,w}|\beta_{i,p,w}|^2 = 1$ and

$$\Psi_{i,p,w} = S_1^{\otimes 2n^2}\sum_{\substack{E\subseteq[A],F\subseteq[B] \\ |E|+|F|\le t \\ x\in D^E, y\in D^F}} \beta^{i,p,w}_{E,F,x,y}\,|x\rangle_E\,|\perp\rangle_{[A]\setminus E}\,|y\rangle_F\,|\perp\rangle_{[B]\setminus F}$$

for some $\beta^{i,p,w}_{E,F,x,y}$ such that $\sum_{E,F,x,y}|\beta^{i,p,w}_{E,F,x,y}|^2 = 1$ for each i,p,w. Now the probability over the choices of the input matrices and the result of the quantum algorithm making t queries that the matrices A and B are both $(\gamma n, \gamma n)$-rigid and the algorithm produces k correct output values from $C = AB$ is at most:

$$\begin{aligned}\left\|\Pi_k\Pi_{\text{rigid}}\,\mathcal{S}\,|\phi_t\rangle\right\|^2 &= \left\|\Pi_k\Pi_{\text{rigid}}\sum_{i,p,w}\beta_{i,p,w}\,|i,p,w\rangle\otimes\Psi_{i,p,w}\right\|^2 \\ &= \sum_{i,p,w}|\beta_{i,p,w}|^2\cdot\left\|\Pi_{q(w)}\Pi_{\text{rigid}}\Psi_{i,p,w}\right\|^2 \\ &\le \max_{i,p,w}\left\|\Pi_{q(w)}\Pi_{\text{rigid}}\Psi_{i,p,w}\right\|^2. \qquad (5)\end{aligned}$$

For the rest of the proof we fix an i,p,w to achieve the maximum value in (5) and prove an upper bound on the resulting probability. This fixes the output values $q(w)$; we write $G \subseteq [C]$ with $|G| = k$ for the set of indices of the outputs given by $q(w)$. To keep notations simpler in the remainder of the proof we observe that (5) is upper bounded by the maximum of

$$\left\|\Pi_{q(G)}\Pi_{\text{rigid}}S_1^{\otimes 2n^2}\sum_{\substack{E\subseteq[A],F\subseteq[B] \\ |E|+|F|\le t \\ x\in D^E, y\in D^F}}\beta_{E,F,x,y}\,|x\rangle_E\,|\perp\rangle_{[A]\setminus E}\,|y\rangle_F\,|\perp\rangle_{[B]\setminus F}\right\|^2 \qquad (6)$$

over all $\beta_{E,F,x,y}$ with $\sum_{E,F,x,y}|\beta_{E,F,x,y}|^2 = 1$, all sets $G \subseteq [C]$ with $|G| = k$ and all assignments $q(G)$ to G.

We will split the sum in (6) over the different sets E and F of queried input indices depending on how they relate to the set of output indices given by G. Let $r(G)$ be the set of rows containing elements of G and $c(G)$ be the set of columns containing elements of G.

We define a *light row of E* to be an element of $r(G)$ that contains at most $\beta\gamma n$ elements of E and define a *light column of F* to be an element of $c(G)$ that contains at most $\beta\gamma n$ elements of F. Since $|E| + |F| \leq t \leq \beta\gamma n\sqrt{k/2}$ we have $\leq \sqrt{k/2}$ rows of E in $r(G)$ and $\leq \sqrt{k/2}$ columns of F in $c(G)$ that are not light. We define $\mathcal{L}(E) \subseteq r(G)$, to be any set of $|r(G)| - \left\lfloor\sqrt{k/2}\right\rfloor$ light rows of E and $\mathcal{L}'(F) \subseteq c(G)$ to be any set of $c(G) - \left\lfloor\sqrt{k/2}\right\rfloor$ light columns of F. Therefore $|\{(i', j') \in G \mid i' \notin \mathcal{L}(E),\ j' \notin \mathcal{L}'(F)\}| \leq k/2$ so at least $k/2$ elements of G are in light rows of E or in light columns of F. Therefore for every pair (E, F) at least one of the sets of outputs $G^r_{\mathcal{L}(E)} = \{(i', j') \in G \mid i' \in \mathcal{L}(E)\}$ or $G^c_{\mathcal{L}(F)} = \{(i', j') \in G \mid j' \in \mathcal{L}'(F)\}$ has size $\geq k/4$.

Let $\mathcal{E}$ be the set of all $E \subseteq [A]$ with $|E| \leq t$ such that G has $\geq k/4$ outputs in light rows and $\mathcal{F}$ be the set of all $F \subseteq [B]$ with $|F| \leq t$ such that G has $\geq k/4$ outputs in light columns. We separately bound the contribution to (6) from pairs (E, F) with $E \in \mathcal{E}$ and $F \in \mathcal{F}$. The analyses of the two cases are completely symmetric up to matrix transposition. It will be convenient to focus on the case $F \in \mathcal{F}$ that there are many outputs of G in light columns and compute an upper bound on

$$\left\|\Pi_{q(G)}\Pi_{\text{rigid}}\mathcal{S}_1^{\otimes 2n^2}\sum_{\substack{E\subseteq[A]\\|E|\leq t\\x\in D^E}}\sum_{\substack{F\in\mathcal{F}\\y\in D^F}}\beta_{E,F,x,y}\,|x\rangle_E\,|\perp\rangle_{[A]\setminus E}\,|y\rangle_F\,|\perp\rangle_{[B]\setminus F}\right\|^2. \tag{7}$$

The case that $E \in \mathcal{E}$ has exactly the same upper bound as (7) by applying the argument to the transposed product B^TA^T and corresponding transposed sets F^T, E^T, and G^T. Hence, the quantity in (6) is at most 4 times that of (7).

To upper bound (7), we first remove the projection operator $\Pi_{\text{rigid }B}$ from $\Pi_{q(G)}\Pi_{\text{rigid}} = \Pi_{q(G)}\Pi_{\text{rigid }A}\Pi_{\text{rigid }B}$ to get $\Pi_{q(G)}\Pi_{\text{rigid }A}$. We then rewrite this combined projection operator as $\Pi_{q(G)}\Pi_{\text{rigid }A} = \sum_{A\ (\gamma n,\gamma n)\text{-}\mathbf{rigid}}\Pi_A \otimes \Pi^A_{q(G)}$ where Π_A is the projection onto the specific matrix A and for each A, $\Pi^A_{q(G)}$ is the projection onto the choices for matrix B such that $C = AB$ agrees with $q(w)$. We therefore obtain that (7) is at most

$$\begin{aligned}
&\Big\|\sum_{A\ (\gamma n,\gamma n)\text{-rigid}}(\Pi_A\otimes\Pi^A_{q(G)})\\
&\qquad\mathcal{S}_1^{\otimes 2n^2}\sum_{\substack{E\subseteq[A]\\|E|\leq t\\x\in D^E}}\sum_{\substack{F\in\mathcal{F}\\y\in D^F}}\beta_{E,F,x,y}\,|x\rangle_E\,|\perp\rangle_{[A]\setminus E}\,|y\rangle_F\,|\perp\rangle_{[B]\setminus F}\Big\|^2\\
&=\Big\|\sum_{A\ (\gamma n,\gamma n)\text{-rigid}}(\Pi_A\otimes\Pi^A_{q(G)}\mathcal{S}_1^{\otimes n^2})\\
&\qquad\sum_{A'\in(D\cup\{\perp\})^{[A]}}\sum_{\substack{F\in\mathcal{F}\\y\in D^F}}\beta_{A'}\beta^{A'}_{F,y}\,|A'\rangle_{[A]}\,|y\rangle_F\,|\perp\rangle_{[B]\setminus F}\Big\|^2\\
&=\Big\|\sum_{A\ (\gamma n,\gamma n)\text{-rigid}}\beta_A\,|A\rangle_{[A]}\otimes[\Pi^A_{q(G)}\mathcal{S}_1^{\otimes n^2}\sum_{\substack{F\in\mathcal{F}\\y\in D^F}}\beta^A_{F,y}\,|y\rangle_F\,|\perp\rangle_{[B]\setminus F}]\Big\|^2
\end{aligned} \tag{8}$$

for some β_A and $\beta^A_{F,y}$ such that $\sum_{A\in(D\cup\{\perp\})^{n^2}}|\beta_A|^2 = 1$ and $\sum_{F\in\mathcal{F},y\in D^{F0}}|\beta^A_{F,y}|^2 = 1$ for each A. Since $\Pi_{q(G)^A}$ only projects onto the $[B]$ input registers, each distinct choice of $|A\rangle_{[A]}$ gives orthogonal states so (8) equals

$$\begin{aligned}
&\sum_{A\ (\gamma n,\gamma n)\text{-rigid}}|\beta_A|^2\Big\|\Pi^A_{q(G)}\mathcal{S}_1^{\otimes n^2}\sum_{\substack{F\in\mathcal{F}\\y\in D^F}}\beta^A_{F,y}\,|y\rangle_F\,|\perp\rangle_{[B]\setminus F}\Big\|^2\\
&\leq\max_{A\ (\gamma n,\gamma n)\text{-rigid}}\Big\|\Pi^A_{q(G)}\mathcal{S}_1^{\otimes n^2}\sum_{\substack{F\in\mathcal{F}\\y\in D^F}}\beta^A_{F,y}\,|y\rangle_F\,|\perp\rangle_{[B]\setminus F}\Big\|^2.
\end{aligned} \tag{9}$$

We fix a $(\gamma n, \gamma n)$-rigid matrix A that maximizes (9) and partition the set $\mathcal{F}$ based on the set $\mathcal{L}'(F)$ which contains all but precisely $\left\lfloor\sqrt{k/2}\right\rfloor$ columns in $c(G)$. Therefore we can rewrite (9) as

$$\Big\|\sum_{H\in\binom{c(G)}{\lfloor\sqrt{k/2}\rfloor}}\Pi^A_{q(G)}\mathcal{S}_1^{\otimes n^2}\sum_{\substack{F\in\mathcal{F},\ \mathcal{L}'(F)=c(G)\setminus H\\y\in D^F}}\beta^A_{F,y}\,|y\rangle_F\,|\perp\rangle_{[B]\setminus F}\Big\|^2. \tag{10}$$

Since the different choices of F, and hence different choices of H, correspond to orthogonal basis states, we can upper bound (10) by

$$\binom{|c(G)|}{\lfloor\sqrt{k/2}\rfloor}\max_{H\in\binom{c(G)}{\lfloor\sqrt{k/2}\rfloor}}\Big\|\Pi^A_{q(G)}\mathcal{S}_1^{\otimes n^2}\sum_{\substack{F\in\mathcal{F}\\\mathcal{L}'(F)=c(G)\setminus H\\y\in D^F}}\beta^A_{F,y}\,|y\rangle_F\,|\perp\rangle_{[B]\setminus F}\Big\|^2. \tag{11}$$

We fix the set H achieving the maximum value in (11) which fixes the value of $\mathcal{L}'(F) = c(G) \setminus H$. This fixes the set $G^c_{\mathcal{L}'(F)}$ of elements in G that are in light columns of F (equivalently, not in H) which, since $F \in \mathcal{F}$, contains at least $k/4$ elements of G. Let G' be a fixed subset of $k/4$ of the elements of $G^c_{\mathcal{L}'(F)}$. By construction we have $c(G') \subseteq \mathcal{L}'(F)$. By only requiring that the outputs in G' are correct and using the fact that $|c(G)| \leq \min(k, n)$, we therefore can upper bound $\left\|\Pi_k\Pi_{\text{rigid}}\mathcal{S}\,|\phi_t\rangle\right\|^2$ by the maximum value of

$$4\ \min(k,n)\sqrt{k/2}\Big\|\Pi^A_{q(G')}\mathcal{S}_1^{\otimes n^2}\sum_{\substack{F\subseteq[B]\\c(G')\ \subseteq\ \mathcal{L}'(F),\ y\in D^F}}\beta'_{F,y}\,|y\rangle_F\,|\perp\rangle_{[B]\setminus F}\Big\|^2 \tag{12}$$

over all $G' \subseteq [C]$ with $|G'| = k/4$ and $\beta'_{F,y}$ with $\sum_{F,y}|\beta'_{F,y}|^2 = 1$.

For each $j \in c(G')$, let k_j be the number of elements of G' in column j. Our overall strategy is to consider the $j \in c(G')$ one by one, and show that the total amplitude on states where these k_j outputs are correct conditioned on the success for previous values of j is of the form $d^{-\delta k_j}$ for some fixed constant $\delta > 0$. These are k_j outputs of the matrix-vector product Ay^j where y^j is the j-th column of B and the fact that $c(G') \subseteq \mathcal{L}'(F)$ implies that F has made at most $\beta\gamma n$ queries to $y^{(j)}$.

We could try to apply the ideas of Lemma 3.2 to this collection of matrix-vector problems and create a set of buckets that is the product of the sets of column buckets for each j and bound each bucket separately. However, unlike Lemma 3.2, the value of many of the k_j can be very small, as low as 1, and the upper bounds using Lemmas 3.8 and 3.9 would yield a probability bound larger than 1.

Instead, we need a stronger argument to show that, except for a portion that is exponentially small in k, all of the amplitude can be allocated to a very small number of buckets. The following lemma gives the inductive step that allows us to define those buckets. Rather than thinking about each column $j \in c(G')$ as separate matrix-vector problems, it works by considering all of the answers in G' at once.

Lemma 4.3. *Let $G' \subseteq [C]$ with $|G'| = k/4$ and $\mathcal{F}'$ be a set of $F \subseteq [B]$ such that $c(G') \subseteq \mathcal{L}'(F)$. Let $\delta_{F,y} \in \mathbb{C}$ satisfy $\sum_{F\in\mathcal{F}', y\in D^F} |\delta_{F,y}|^2 = 1$. Let $C' \geq 2$ and define $\alpha = C'\beta$. Then there is an $\mathcal{F}'' \subseteq \mathcal{F}'$ and $\delta'_{F,y}$ such that $\sum_{F\in\mathcal{F}'', y\in D^F} |\delta'_{F,y}|^2 = 1$ and*

$$\left\|\Pi^A_{q(G')}\mathcal{S}_1^{\otimes n^2}\sum_{F\in\mathcal{F}',\ y\in D^F}\delta_{F,y}\,|y\rangle_F\,|\perp\rangle_{[B]\setminus F}\right\|^2 \leq \frac{2^{1+H_2(\alpha)\,k/4}}{d^{(1-\alpha)\,k/4}} + \frac{2}{C'}\left\|\Pi^A_{q(G')}\mathcal{S}_1^{\otimes n^2}\sum_{F\in\mathcal{F}'',\ y\in D^F}\delta'_{F,y}\,|y\rangle_F\,|\perp\rangle_{[B]\setminus F}\right\|^2.$$

Proof. For each $j \in c(G')$, define U^j to be the set of row indices of G' in column j and let $k_j = |U_j|$. Define $\ell_j = \lceil \gamma n/k_j \rceil$,apply Lemma 3.9 for each j, and let $V_1^j, \ldots, V_{\ell_j}^j$ be the collection of disjoint subsets of $[n]$ of size k_j found for each j such that each $k_j \times k_j$ sub-matrix $A_{U^jV_i^j}$ has full rank.

For each $F \in \mathcal{F}'$ and $i \in c(G')$, define F^j to be the set of row indices of elements of F in column j; since $c(G') \subseteq \mathcal{L}'(F)$, we have $|F^j| \leq \beta\gamma n$. For each $i \in [\ell_j]$ define

$$m_i^j = \sum_{F\in\mathcal{F}',\ y\in D^F} |\delta_{F,y}|^2 \cdot |F^j \cap V_i^j|.$$

Since $\sum_{F,y} |\delta_{F,y}|^2 = 1$, m_i^j can be viewed as the expected size of the overlap between the recorded queries in the j-th column of the matrix B and each V_i^j. Since for each j, the sets V_i^j are disjoint and $|F^j| \leq \beta\gamma n$ we have $\sum_{i\in[\ell_j]} m_i^j \leq \beta\gamma n$. Therefore, for each j, we have some index $i_j \in [\ell_j]$ such that $m_{i_j}^j \leq \beta\gamma n/\ell_j \leq \beta k_j$.

Since $\sum_{j\in c(G')} k_j = |G'| = k/4$, the expected total overlap between the recorded queries in the columns of G and the chosen sets $V_{i_j}^j$ for those columns is $\sum_j m_{i_j}^j \leq \sum_j \beta k_j = \beta k/4$. Define $\mathcal{F}''$ to be the set of $F \in \mathcal{F}'$ such that $\sum_j |F^j \cap V_{i_j}^j| \geq \alpha k/4 = C'\beta k/4$. By Markov's inequality we have $\sum_{F\in\mathcal{F}'',\ y\in D^F} |\delta_{F,y}|^2 \leq \frac{\sum_j m_{i_j}^j}{C'\beta k/4} \leq 1/C'$. We split our analysis for $\mathcal{F}'$ into two parts due to sets F in $\mathcal{F}''$ and $\mathcal{F}' \setminus \mathcal{F}''$, respectively.

Let $F \in \mathcal{F}''$, write $\kappa = \sum_{F\in\mathcal{F}'',\ y\in D^F} |\delta_{F,y}|^2 \leq 1/C'$. For $F \in \mathcal{F}''$, define $\delta'_{F,y} = \delta_{F,y}/\sqrt{\kappa}$. Then $\sum_{F\in\mathcal{F}'', y\in D^F} |\delta'_{F,y}|^2 = 1$ and

$$\left\|\Pi^A_{q(G')}\mathcal{S}_1^{\otimes n^2}\sum_{F\in\mathcal{F}'',\ y\in D^F}\delta_{F,y}\,|y\rangle_F\,|\perp\rangle_{[B]\setminus F}\right\|^2 \leq \frac{1}{C'}\left\|\Pi^A_{q(G')}\mathcal{S}_1^{\otimes n^2}\sum_{F\in\mathcal{F}'',\ y\in D^F}\delta'_{F,y}\,|y\rangle_F\,|\perp\rangle_{[B]\setminus F}\right\|^2. \quad (13)$$

We now consider $\mathcal{F}' \setminus \mathcal{F}''$. By definition, for $F \in \mathcal{F}' \setminus \mathcal{F}''$, we have $\sum_j |F^j \cap V_{i_j}^j| < \alpha k/4$. By definition we have $\sum_j |V_{i_j}^j| = \sum_j k_j = k/4$ so F must miss more than $(1-\alpha)k/4$ elements of the set $V = \bigcup_j (V_{i_j}^j \times \{j\})$ of size $k/4$. For each subset V' of V of size $k/4 - \lfloor \alpha k/4 \rfloor$ we define a bucket $\mathcal{B}_{V'}$ that contains sets F that must the elements of V' and assign each $F \in \mathcal{F}' \setminus \mathcal{F}''$ to a unique bucket in an arbitrary fixed way. There are at most $2^{H_2(\alpha)k/4}$ such buckets so:

$$\begin{aligned}
&\left\|\Pi^A_{q(G')}\mathcal{S}_1^{\otimes n^2}\sum_{F\in\mathcal{F}'\setminus\mathcal{F}'',\ y\in D^F}\delta_{F,y}\,|y\rangle_F\,|\perp\rangle_{[B]\setminus F}\right\|^2 \\
&\leq \left(\sum_{|V'|=k/4-\lfloor\alpha k/4\rfloor}\ \sum_{\substack{V'\subseteq V\\ F\in\mathcal{B}_{V'}\\ y\in D^F}}\left\|\Pi^A_{q(G')}\mathcal{S}_1^{\otimes n^2}\delta_{F,y}\,|y\rangle_F\,|\perp\rangle_{[B]\setminus F}\right\|\right)^2 \\
&\leq 2^{H_2(\alpha)\,k/4}\sum_{|V'|=k/4-\lfloor\alpha k/4\rfloor}\ \sum_{\substack{V'\subseteq V\\ F\in\mathcal{B}_{V'}\\ y\in D^F}}\left\|\Pi^A_{q(G')}\mathcal{S}_1^{\otimes n^2}\,|\perp\rangle_{V'}\delta_{F,y}\,|y\rangle_F\,|\perp\rangle_{[B]\setminus(F\cup V')}\right\|^2 \quad (14)
\end{aligned}$$

using the triangle and Jensen's inequalities.

Now, applying the $\mathcal{S}_1^{\otimes n^2}$ operator in (14) will convert the $|\perp\rangle_{V'}$ to a uniform superposition of all $|y'\rangle_{V'}$ for all $y' \in D^{V'}$ and convert $\sum_{\substack{F\in\mathcal{B}_{V'}\\ y\in D^F}} \delta_{F,y}\,|y\rangle_F\,|\perp\rangle_{[B]\setminus(F\cup V')}$ to some superposition of $|y''\rangle \in D^{[B]\setminus V'}$ with amplitudes some $\delta_{V',y''}$ such that $\sum_{y''} |\delta_{V',y''}|^2 = \sum_{F\in\mathcal{B}_{V'},y\in D^F} |\delta_{F,y}|^2$. Therefore, we can rewrite (14) as

$$2^{\frac{H_2(\alpha)k}{4}}\sum_{|V'|=k/4-\lfloor\alpha k/4\rfloor}\ \sum_{V'\subseteq V}\left\|\Pi^A_{q(G')}\left[\sum_{y'\in D^{V'}}\frac{1}{\sqrt{d^{|V'|}}}\,|y'\rangle_{V'}\right]\otimes\sum_{y''\in D^{[n]\setminus V'}}\delta_{V',y}\,|y\rangle_{[B]\setminus V'}\right\|^2. \quad (15)$$

We now consider the application of $\Pi^A_{q(G')}$. Let $V'_j \subseteq V_{i_j}^j$ be the set of row indices in column j of $V' \subseteq [B]$ and consider the corresponding set of columns in A. Since $A_{U^jV_{i_j}^j}$ has full rank, there is a subset $U_0^j \subseteq U^j$ with $|U_0^j| = |V'_j|$ so that $A_{U_0^jV'_j}$ also has full rank. Now define $G'_0 \subseteq G'$ to be $\bigcup_{j\in c(G)}(U_j \times \{j\})$ with size $|V'|$.

For each j, the outputs in $U_j \times \{j\} \subset [C]$ can be expressed as the matrix-vector product $A_{U_0^jV'_j}y^j_{V'_j} + M$ for some $|V'_j| \times |V'_j|$ matrix M defined by the product of the $U_0^j \times ([n] \setminus V'_j)$ submatrix of the fixed matrix A and $y^j_{[n]\setminus V'_j}$. Since $A_{U_0^jV'_j}$ is full rank, for each value of M given by $y^j_{[n]\setminus V'_j}$, there is precisely one value of $y^j_{V'_j}$ that will yield the output values $q(U_j \times \{j\})$. Putting the properties for the columns of $c(G')$ together, there is precisely one value $y' \in D^{V'}$ that yields the output values $q(G'_0)$. So (15) is at most

$$\begin{aligned}
&2^{\frac{H_2(\alpha)k}{4}}\sum_{|V'|=k/4-\lfloor\alpha k/4\rfloor}\ \sum_{V'\subseteq V}\left\|\Pi^A_{q(G'_0)}\left[\sum_{y'\in D^{V'}}\frac{1}{\sqrt{d^{|V'|}}}\,|y'\rangle_{V'}\right]\otimes\sum_{y''\in D^{[n]\setminus V'}}\delta_{V',y}\,|y\rangle_{[B]\setminus V'}\right\|^2 \\
&= 2^{H_2(\alpha)\,k/4}\sum_{|V'|=k/4-\lfloor\alpha k/4\rfloor}\ \sum_{V'\subseteq V}\left\|\frac{1}{\sqrt{d^{|V'|}}}\sum_{y''\in D^{[n]\setminus V'}}\delta_{V',y}\,|y\rangle_{[B]\setminus V'}\right\|^2 \\
&= 2^{H_2(\alpha)\,k/4}\sum_{|V'|=k/4-\lfloor\alpha k/4\rfloor}\ \sum_{V'\subseteq V}\frac{1}{d^{|V'|}}\sum_{F\in\mathcal{B}_{V'},y\in D^F}|\delta_{F,y}|^2 \\
&= 2^{H_2(\alpha)\,k/4}\cdot\frac{1}{d^{|V'|}}\sum_{F\in\mathcal{F}'\setminus\mathcal{F}'',y\in D^F}|\delta_{F,y}|^2 \leq \frac{2^{H_2(\alpha)\,k/4}}{d^{(1-\alpha)\,k/4}} \quad (16)
\end{aligned}$$

since the buckets $\mathcal{B}_{V'}$ partition $\mathcal{F}' \setminus \mathcal{F}''$.

We now combine the contributions from $\mathcal{F}''$ and $\mathcal{F}' \setminus \mathcal{F}''$. Applying Jensen's inequality together with the bounds in (13) and (16)

we obtain that

$$\begin{aligned}
&\left\|\Pi^A_{q(G')}\mathcal{S}_1^{\otimes n^2}\sum_{F\in\mathcal{F}',\ y\in D^F}\delta_{F,y}\,|y\rangle_F\,|\perp\rangle_{[B]\setminus F}\right\|^2\\
&\le 2\Bigg(\left\|\Pi^A_{q(G')}\mathcal{S}_1^{\otimes n^2}\sum_{F\in\mathcal{F}'\setminus\mathcal{F}'',\ y\in D^F}\delta_{F,y}\,|y\rangle_F\,|\perp\rangle_{[B]\setminus F}\right\|^2\\
&\qquad+\left\|\Pi^A_{q(G')}\mathcal{S}_1^{\otimes n^2}\sum_{F\in\mathcal{F}'',\ y\in D^F}\delta_{F,y}\,|y\rangle_F\,|\perp\rangle_{[B]\setminus F}\right\|^2\Bigg)\\
&\le \frac{2^{1+H_2(\alpha)\,k/4}}{d^{(1-\alpha)\,k/4}}+\frac{2}{C'}\left\|\Pi^A_{q(G')}\mathcal{S}_1^{\otimes n^2}\sum_{F\in\mathcal{F}'',\ y\in D^F}\delta'_{F,y}\,|y\rangle_F\,|\perp\rangle_{[B]\setminus F}\right\|^2.\ \square
\end{aligned}$$

Corollary 4.4. *Let $G'\subseteq[C]$ with $|G'|=k/4$, $\mathcal{F}'$ be a set of $F\subseteq[B]$ such that $c(G')\subseteq\mathcal{L}'(F)$, and $\sum_{F\in\mathcal{F}',y\in D^F}|\delta_{F,y}|^2=1$ for some $\delta_{F,y}$. Then*

$$\left\|\Pi^A_{q(G')}\mathcal{S}_1^{\otimes n^2}\sum_{F\in\mathcal{F}',\ y\in D^F}\delta_{F,y}\,|y\rangle_F\,|\perp\rangle_{[B]\setminus F}\right\|^2\le\frac{2^{2+H_2(4\beta)\,k/4}}{d^{(1-4\beta)\,k/4}}.$$

Proof. Let M be the maximum value of

$$\left\|\Pi^A_{q(G')}\mathcal{S}_1^{\otimes n^2}\sum_{F\in\mathcal{F}',\ y\in D^F}\delta_{F,y}\,|y\rangle_F\,|\perp\rangle_{[B]\setminus F}\right\|^2$$

over all choices of $\mathcal{F}'$ and $\delta_{F,y}$ with the required properties. This corollary follows from Lemma 4.3 with $C'=4$ by observing that the term multiplied by $2/C'$ is also upper bounded by M and hence $M\le 2^{1+H_2(4\beta)\,k/4}/d^{(1-4\beta)\,k/4}+M/2$. $\square$

Finally, plugging the bound from Corollary 4.4 into (12), we obtain that the probability that A and B are both $(\gamma n,\gamma n)$-rigid and C produces k correct output values for $C=AB$, $\left\|\Pi_k\Pi_{\text{rigid}}\mathcal{S}\,|\phi_t\rangle\right\|^2$, is at most $16\ \min(k,n)\sqrt{k/2}\left(2^{H_2(4\beta)}/d^{(1-4\beta)}\right)^{k/4}$ as desired. $\square$

5 QUANTUM BOUNDS FOR BOOLEAN MATRIX OPERATIONS

In this section we focus on Boolean matrix operations, which use (AND, OR) inner product which we denote by $u\bullet v=\bigvee_i(u_i\wedge v_i)$ and we extend this $\bullet$ notation to matrices.

5.1 Boolean Matrix Multiplication

Unlike what we have shown for algebraic problems, one can apply Grover's algorithm to each output of Boolean matrix multiplication to obtain a quantum advantage. For any constant c this gives quantum circuits computing $n\times n$ Boolean matrix multiplication $A\bullet B$ with error at most n^{-c} using space $O(\log n)$ and $O(n^{2.5}\log n)$. This is in contrast to the following result of Abrahamson which shows that classical algorithms as fast as this quantum algorithm require space $\tilde{\Omega}(n^{0.5})$ rather than $O(\log n)$.

Proposition 5.1 ([2]). *There is a probability distribution on input matrices and constants $0<c_1<c_2$ under which classical algorithms (branching programs) for Boolean matrix multiplication $A\bullet B$ using space S require time T for which $T\cdot S$ is $\Theta(n^{3.5})$ for $T\le c_1n^{2.5}$ and $\Theta(n^3)$ for $T\ge c_2n^{2.5}$.*

For quantum circuits, Klauck, Špalek, and de Wolf [20] proved the following time-space tradeoff lower bound which nearly matches the Grover-based upper bound when the space S is $O(\log n)$.

Proposition 5.2 ([20]). *Any bounded error quantum circuit that computes the $n\times n$ Boolean matrix multiplication $A\bullet B$ with T queries and space S requires T that is $\Omega(n^{2.5}/S^{0.5})$.*

Unlike our results in linear algebra results and Abrahamson's bounds, Proposition 5.2 only applies to circuits, where it is natural that the set of output values produced in each part of the computation is fixed independent of the input. It uses an embedding of the direct product of OR functions into any fixed set of k outputs of the Boolean matrix multiplication problem together with part (b) of the following strong direct product theorem.

Proposition 5.3 ([20]). *There are positive constants ε and γ such that the following hold: (a) Any randomized algorithm making at most εkn queries has success probability at most $2^{-\gamma k}$ in computing OR_n^k. (b) Any quantum algorithm making at most $\varepsilon k\sqrt{n}$ queries has success probability at most $2^{-\gamma k}$ in computing OR_n^k.*

Using a more efficient embedding of OR computations into outputs of the matrix multiplication problem we obtain the following improved lower bound for quantum computation.

Theorem 5.4. *Any quantum circuit computing $n\times n$ Boolean matrix multiplication $A\bullet B$ with T queries and space S and success probability more than 2^{-S} must have T that is $\Omega(n^{2.5}/S^{1/4})$.*

Both the exponent of n and that of S in our bound are optimal: The Grover-based algorithm shows that exponent of n is optimal since there is only a gap of $O(\log^{5/4}n)$ for space $\Theta(\log n)$. At the other end of the scale, in our quantum query model, an algorithm with space $3n^2$ can query and completely remember both matrices in $2n^2$ time, after which a global unitary transformation will produce the n^2 bits of output needed in the remaining qubits of working memory; hence the exponent of $1/4$ on S cannot be reduced.

Via similar improvement we also obtain the following theorem for classical computation, which dominates the lower bound of Proposition 5.1 for all values of S.

Theorem 5.5. *Any classical circuit (or other sequential model in which each output value is produced at a fixed time step) computing $n\times n$ Boolean matrix-multiplication with T queries and space S with success probability more than 2^{-S} must have T that is $\Omega(n^3/\sqrt{S})$.*

This answers a question of Klauck, Špalek, and de Wolf [20] who ventured that this might be the likely tight tradeoff for classical computation of Boolean matrix multiplication. Like the quantum lower bound it has optimal exponents for the model to which it applies. Our full paper [9] includes the details of the proof.

Theorem 5.4 follows from the following key lemma.

Lemma 5.6. *There are constants $\varepsilon,c'>0$ such that for any integer $k<n^2/100$ and quantum circuit C with at most $\varepsilon k^{3/4}n^{1/2}$ queries to x, the probability that C produces k correct output values of $n\times n$ Boolean matrix multiplication $A\bullet B$ is at most $2^{-\gamma k}$.*

Proof of Theorem 5.4 via Lemma 5.6. By applying Lemma 5.6 with $k=n^2/101$, we see that T must be $\Omega(n^2)$ and hence without

loss of generality we can assume that $\sqrt{S} < \alpha n$ for some arbitrarily small constant $\alpha > 0$. Let ε and γ be the constants from Lemma 5.6. Let $c = 3/(2\gamma)$ and define $k = cS$. Therefore for $\alpha \le 1/(10\sqrt{c})$ we obtain that $5\sqrt{k} = 5\sqrt{cS} < n/2$. By Lemma 5.6, since $k < n^2/100$, any quantum query algorithm with at most $\varepsilon k^{3/4} n^{1/2}$ queries has success probability at most $2^{-\gamma k} = 2^{-3S}$ of producing k correct outputs.

We prove the contrapositive of the theorem statement: Suppose that $T \le \varepsilon n^{2.5}/(cS)^{1/4} = \varepsilon n^{2.5}/k^{1/4}$. When we divide C into layers with $\varepsilon k^{3/4} n^{1/2}$ quantum queries each, there are at most n^2/k layers. Since there are a total of n^2 outputs, there must be some layer i during which at least k outputs are produced. Let E be the set of the first k outputs produced in layer i. By the argument above since the space is at most S, by Proposition 2.5 the probability that these k outputs are correct given the S qubits of input-dependent initial state at the beginning of layer i is at most 2^{2S} times larger than that of a circuit without them and the same number of queries, which is at most $2^{2S} \cdot 2^{-3S} = 2^{-S}$ which is what we needed to show. □

The proof of this key lemma is based on our improved method for embedding the direct product of *OR* functions into outputs of the Boolean matrix multiplication problem; this uses the following definition of an L-coloring of subsets of $[n] \times [n]$.

Definition 5.7. For $E \subseteq [n] \times [n]$ an *L-coloring* of E is a map $\chi : E \to [L]$ such that (1) within each color class either all rows are distinct or all columns are distinct, and (2) for each color ℓ there is a rectangle given by sets $R_\ell \subseteq [n]$ of rows and $C_\ell \subseteq [n]$ of columns such that the set of points of color ℓ is precisely $E \cap (R_\ell \times C_\ell)$. (Note that the rectangles $R_\ell \times C_\ell$ may overlap, but their overlap must not contain any points in E.) We say that a rectangle $R \times C \in [n] \times [n]$ is *colorable* iff $E \cap (R \times C)$ either has all its elements in different rows or all its elements in different columns.

Lemma 5.8. *Let $E \subseteq [n] \times [n]$ with $|E| = k$ and $L \le n$ be an integer with $L \le n/2$. If E has an L-coloring then $OR^k_{\lfloor n/L \rfloor}$ is a sub-function of the function that produces the k outputs of $A \bullet B$ indexed by E for $n \times n$ Boolean matrices A and B.*

Proof Sketch. Write $E = \dot{\bigcup}_{\ell=1}^{L} E_\ell$ where E_ℓ is the set of (i, j) in E in color class ℓ. We now divide $[n]$ into L disjoint blocks $b_1, \ldots, b_L$ of at least $\lfloor n/L \rfloor \ge 2$ elements each. Given the coloring and division into blocks, we define a partial assignment to the matrices A and B:

- If color class ℓ consists of points that do not share a column, for each $(i, j) \in E_\ell$, we set all entries of A_{i,b_ℓ} to 1 and leave all entries of $B_{b_\ell,j}$ unset.
- If color class ℓ consists of points that do not share a row, for each $(i, j) \in E_\ell$, we set all entries of $B_{b_\ell,j}$ to 1 and leave all the entries of A_{i,b_ℓ} unset.
- All entries of A and B that are not defined by the above two cases are set to 0.

It is not hard to check that the subfunction property holds. □

The lower bound of [20] corresponds to the trivial k-coloring that colors each element of E differently. For integer $k > 0$ define $L(k)$ to be the minimum number of colors L such that for all subsets $E \subseteq [n] \times [n]$ with $|E| \le k$, there is an L-coloring of E.

Lemma 5.9. *There are constants $c, c' > 0$ such that the following holds. Let k be an integer such that $L(k) \le n/2$. For any quantum circuit C with at most $ckn^{1/2}/L(k)^{1/2}$ queries to x, the probability that C produces k correct output values of $n \times n$ Boolean matrix product $A \bullet B$ is at most $2^{-c'k}$.*

Proof. Let E be any fixed set of k output positions in $A \bullet B$. States with different choices of E are orthogonal to each other so we show that for each fixed value of E the probability that the algorithm is correct has the given bound. Let $L \le L(k)$ be such that there is an L-coloring of E. By Lemma 5.8, $OR^k_{\lfloor n/L \rfloor}$ is a sub-function of the k outputs indexed by the set E. Since $L \le n/2$, $\lfloor n/L \rfloor \ge 2n/(3L)$ and $\sqrt{\lfloor n/L \rfloor} \ge 4\sqrt{n/L}/5$. Choose $c = 4\varepsilon/5$ and $c' = \gamma$ for ε and γ given in Proposition 5.3. By that proposition, the probability that C produces these k outputs correctly is at most $2^{-\gamma k} = 2^{-c'k}$. □

Then Lemma 5.6 follows from Lemma 5.9 and a bound on $L(k)$.

Lemma 5.10 (Coloring Lemma). $\sqrt{2k} \le L(k) \le 2\sqrt{6k} < 5\sqrt{k}$.

Proof. The lower bound follows from a set E consisting of a grid of side L with the lower triangular part removed which has $k = L(L+1)/2$ points, has two trivial L-colorings which are optimal since all L diagonal points must have different colors or they would violate the coloring conditions.

We now prove the upper bound on $L(k)$. Suppose that for some c with $0 < c \le \sqrt{k}$, we can always find a colorable rectangle $R \times C$ containing $r \ge c\sqrt{k}$ elements of E. Then we claim that $L(k) \le \frac{2}{c}\sqrt{k}$ as follows: First color that set with one color and apply induction to color the remaining $k' = k - r$ elements of E'. By induction there will be at most $\frac{2}{c}\sqrt{k'} = \frac{2}{c}\sqrt{k - r}$ colors needed to color E'. Now $k - r \le k - c\sqrt{k} \le k - c\sqrt{k} + c^4/2 = (\sqrt{k} - c/2)^2$. Therefore, $\sqrt{k - r} \le \sqrt{k} - c/2$ and hence the number of colors needed to color E', $\frac{2}{c}\sqrt{k - r} \le \frac{2}{c}\sqrt{k} - 1$. It follows that at most $\frac{2}{c}\sqrt{k}$ colors are needed to color E as required.

In the following we prove that we can always find a colorable rectangle $R \times C$ containing at least $\sqrt{k/6}$ elements of E, which implies the statement of the lemma by the above argument.

For any column j we write E^j for the set of i such that $(i, j) \in E$. We will have two candidates for the color class. The first candidate is given by the points in some row i with the largest number of elements of E. The second candidate is the colorable rectangle $R \times C$ given by the following procedure. This maintains a colorable rectangle, initially empty, that contains a large portion of the rows where the elements of E occur in the columns in C.

$R \leftarrow \emptyset$; $C \leftarrow \emptyset$; $D \leftarrow \emptyset$
While there is a j such that $|E^j \setminus (R \cup D)| \ge \frac{2}{3}|E^j|$
    $C \leftarrow C \cup \{j\}$
    $R \leftarrow (R \setminus E_j) \cup (E^j \setminus (R \cup D))$
    $D \leftarrow D \cup (R \cap E^j)$

Observe that the rectangle $R \times C$ contains exactly one element of E in every row, every row of $D \times C$ contains at least two elements of E, and there are no elements of E in $([n] \setminus (R \cup D)) \times C$. Also, when we add j to C in the loop, we have $|E^j \setminus (R \cup D)| \ge \frac{2}{3}|E^j|$, and therefore have $|R \cap E^j| \le \frac{1}{3}|E^j|$. It follows that $|R|$ increases by at least $\frac{1}{3}|E^j|$ during that iteration and $|D|$ increases by at most $\frac{1}{3}|E^j|$ and hence we have $|D| \le |R|$.

We let s be the larger of $|R|$, which is the size of this second candidate for the color class, and the length of the longest row in E. For convenience, write $Z = R \cup D$, $\overline{Z} = [n] \setminus Z$, and $\overline{C} = [n] \setminus C$.

We have $|Z| \leq 2|R| \leq 2s$ and $E \cap (\overline{Z} \times C) = \varnothing$. When the procedure finishes, for every column $j \in \overline{C}$, fewer than 2/3 of its points are in rows of $\overline{Z}$ and hence more than 1/3 of its points are in rows of Z. That is, we must have $|E^j \setminus Z| < \frac{2}{3}|E^j|$ and $|E^j \cap Z| > \frac{1}{3}|E^j|$ so $|E \cap (Z \times \overline{C})| > \frac{1}{2}|E \cap (\overline{Z} \times \overline{C})|$. As $\overline{Z} \times C$ has no points of E and each row has at most s points of E, the total number of points is

$$\begin{aligned} k &= |E \cap ([n] \times [n])| = |E \cap (Z \times [n])| + |E \cap (\overline{Z} \times [n])| \\ &\leq |Z|s + |E \cap (\overline{Z} \times [n])| = |Z|s + |E \cap (\overline{Z} \times \overline{C})| \\ &\leq |Z|s + 2|E \cap (Z \times \overline{C})| < |Z|s + 2|Z|s = 3|Z|s \leq 6s^2. \end{aligned}$$

Therefore $s \geq \sqrt{k/6}$. □

Lemma 5.10 also holds with $2\sqrt{2} + 2 = 4.828437..$ in place of $2\sqrt{6}$. Lemma 5.6 is a immediate corollary of Lemmas 5.9 and 5.10 which completes the proof of Theorem 5.4.

5.2 Boolean Matrix-Vector Products

Klauck, Špalek, and de Wolf [20, Theorem 23] proved that for every space bound S in $o(n/\log n)$, there is an $n \times n$ Boolean matrix $A^{(S)}$ such that every bounded-error quantum circuit with space S computing Boolean matrix-vector product $A^{(S)} \bullet x$ in T queries requires that T is $\Omega(\sqrt{n^3/S})$. Though the bound itself is good this does not yield a single function that is hard for all space bounds, as the matrix $A^{(S)}$ changes depending on the value of S. The same space dependent matrix $A^{(S)}$ was also used in [5] to prove similar bounds for problems involving linear inequalities.

In our full paper using the matrices $A^{(S)}$, we remove the dependence of the matrix on S with only a small loss in parameters (and we show a similar improvement on [5] for linear inequalities):

Theorem 5.11. *There is a fixed $m \times n$ Boolean matrix A with $m \leq n \log_2 n$ such that for every S that is $o(n/\log n)$ every bounded-error quantum circuit with space at most S that computes Boolean matrix-vector product $A \bullet x$ in T queries requires that T is $\Omega(\sqrt{n^3/S})$.*

REFERENCES

[1] Scott Aaronson. 2005. Limitations of Quantum Advice and One-Way Communication. *Theory of Computing* 1, 1 (2005), 1–28. https://doi.org/10.4086/toc.2005.v001a001

[2] Karl R. Abrahamson. 1990. A Time-Space Tradeoff for Boolean Matrix Multiplication. In *31st Annual Symposium on Foundations of Computer Science, Volume I.* IEEE Computer Society, St. Louis, MO, USA, 412–419. https://doi.org/10.1109/FSCS.1990.89561

[3] Karl R. Abrahamson. 1991. Time-Space Tradeoffs for Algebraic Problems on General Sequential Machines. *J. Comput. System Sci.* 43, 2 (1991), 269–289. https://doi.org/10.1016/0022-0000(91)90014-v

[4] Andris Ambainis. 2002. Quantum Lower Bounds by Quantum Arguments. *J. Comput. System Sci.* 64, 4 (2002), 750–767. https://doi.org/10.1006/jcss.2002.1826

[5] Andris Ambainis, Robert Špalek, and Ronald de Wolf. 2009. A New Quantum Lower Bound Method, with Applications to Direct Product Theorems and Time-Space Tradeoffs. *Algorithmica* 55, 3 (Nov 2009), 422–461. https://doi.org/10.1007/s00453-007-9022-9

[6] Robert Beals, Harry Buhrman, Richard Cleve, Michele Mosca, and Ronald de Wolf. 2001. Quantum Lower Bounds by Polynomials. *J. ACM* 48, 4 (jul 2001), 778–797. https://doi.org/10.1145/502090.502097

[7] Paul Beame. 1991. A General Sequential Time-Space Tradeoff for Finding Unique Elements. *SIAM J. Comput.* 20, 2 (1991), 270–277. https://doi.org/10.1137/0220017

[8] Paul Beame and Niels Kornerup. 2023. Cumulative Memory Lower Bounds for Randomized and Quantum Computation. In *50th International Colloquium on Automata, Languages, and Programming (ICALP 2023)*, Vol. 261. LIPIcs, Dagstuhl, Germany, 17:1–17:20. https://doi.org/10.4230/LIPIcs.ICALP.2023.17

[9] Paul Beame, Niels Kornerup, and Michael Whitmeyer. 2024. Quantum Time-Space Tradeoffs for Matrix Problems. *CoRR* abs/2401.05321 (2024). https://doi.org/10.48550/arxiv.2401.05321

[10] Ethan Bernstein and Umesh V. Vazirani. 1997. Quantum Complexity Theory. *SIAM J. Comput.* 26, 5 (1997), 1411–1473. https://doi.org/10.1137/S0097539796300921

[11] Allan Borodin and Stephen A. Cook. 1982. A Time-Space Tradeoff for Sorting on a General Sequential Model of Computation. *SIAM J. Comput.* 11, 2 (1982), 287–297. https://doi.org/10.1137/0211022

[12] Nadiia Chepurko, Kenneth L. Clarkson, Lior Horesh, Honghao Lin, and David P. Woodruff. 2022. Quantum-Inspired Algorithms from Randomized Numerical Linear Algebra. In *International Conference on Machine Learning, ICML 2022 (Proceedings of Machine Learning Research, Vol. 162).* PMLR, Baltimore, MD, USA, 3879–3900. https://proceedings.mlr.press/v162/chepurko22a.html

[13] Nai-Hui Chia, András Pal Gilyén, Tongyang Li, Han-Hsuan Lin, Ewin Tang, and Chunhao Wang. 2022. Sampling-based Sublinear Low-rank Matrix Arithmetic Framework for Dequantizing Quantum Machine Learning. *J. ACM* 69, 5 (2022), 33:1–33:72. https://doi.org/10.1145/3549524

[14] Andrew M. Childs, Robin Kothari, and Rolando D. Somma. 2015. Quantum Algorithm for Systems of Linear Equations with Exponentially Improved Dependence on Precision. *SIAM J. Comput.* 46 (2015), 1920–1950. https://api.semanticscholar.org/CorpusID:3834959

[15] David Deutsch and Richard Jozsa. 1992. Rapid solution of problems by quantum computation. *Proceedings of the Royal Society of London A* 439 (1992), 553–558. https://doi.org/10.1098/rspa.1992.0167

[16] András Gilyén, Zhao Song, and Ewin Tang. 2022. An improved quantum-inspired algorithm for linear regression. *Quantum* 6 (2022), 754. https://doi.org/10.22331/q-2022-06-30-754

[17] András Gilyén, Yuan Su, Guang Hao Low, and Nathan Wiebe. 2019. Quantum Singular Value Transformation and Beyond: Exponential Improvements for Quantum Matrix Arithmetics. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing* (Phoenix, AZ, USA) *(STOC 2019).* ACM, New York, NY, USA, 193–204. https://doi.org/10.1145/3313276.3316366

[18] Yassine Hamoudi and Frédéric Magniez. 2021. Quantum Time-Space Tradeoff for Finding Multiple Collision Pairs. In *16th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2021)*, Vol. 197. LIPIcs, Dagstuhl, Germany, 1:1–1:21. https://doi.org/10.4230/LIPIcs.TQC.2021.1

[19] Aram W. Harrow, Avinatan Hassidim, and Seth Lloyd. 2009. Quantum Algorithm for Linear Systems of Equations. *Physical Review Letters* 103, 15 (2009). https://doi.org/10.1103/physrevlett.103.150502

[20] Hartmut Klauck, Robert Špalek, and Ronald de Wolf. 2007. Quantum and Classical Strong Direct Product Theorems and Optimal Time-Space Tradeoffs. *SIAM J. Comput.* 36, 5 (2007), 1472–1493. https://doi.org/10.1137/05063235x arXiv:https://doi.org/10.1137/05063235X

[21] Guang Hao Low and Isaac L. Chuang. 2019. Hamiltonian Simulation by Qubitization. *Quantum* 3 (2019), 163. https://doi.org/10.22331/q-2019-07-12-163

[22] Yishay Mansour, Noam Nisan, and Prasoon Tiwari. 1993. The Computational Complexity of Universal Hashing. *Theor. Comput. Sci.* 107, 1 (1993), 121–133. https://doi.org/10.1016/0304-3975(93)90257-T

[23] Alexander A. Sherstov. 2011. Strong direct product theorems for quantum communication and query complexity. In *Proceedings of the 43rd ACM Symposium on Theory of Computing, STOC 2011.* ACM, San Jose, CA, USA, 41–50. https://doi.org/10.1145/1993636.1993643

[24] Alexander A. Sherstov. 2012. Strong Direct Product Theorems for Quantum Communication and Query Complexity. *SIAM J. Comput.* 41, 5 (2012), 1122–1165. https://doi.org/10.1137/110842661

[25] Robert Spalek. 2008. The Multiplicative Quantum Adversary. In *Proceedings of the 23rd Annual IEEE Conference on Computational Complexity, CCC 2008.* IEEE Computer Society, College Park, MD, USA, 237–248. https://doi.org/10.1109/CCC.2008.9

[26] Robert Spalek and Mario Szegedy. 2006. All Quantum Adversary Methods are Equivalent. *Theory Comput.* 2, 1 (2006), 1–18. https://doi.org/10.4086/TOC.2006.V002A001

[27] Ewin Tang. 2019. A quantum-inspired classical algorithm for recommendation systems. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC 2019.* ACM, Phoenix, AZ, USA, 217–228. https://doi.org/10.1145/3313276.3316310

[28] Yaacov Yesha. 1984. Time-space tradeoffs for matrix multiplication and the discrete Fourier transform on any general sequential random-access computer. *J. Comput. System Sci.* 29, 2 (1984), 183–197. https://doi.org/10.1016/0022-0000(84)90029-1

[29] Mark Zhandry. 2019. How to Record Quantum Queries, and Applications to Quantum Indifferentiability. In *Advances in Cryptology – CRYPTO 2019.* Springer International Publishing, Cham, 239–268.

Received 13-NOV-2023; accepted 2024-02-11