

An access model for quantum encoded data

Miguel Murça^{1,2,3}, Paul K. Faehrmann⁴, and Yasser Omar^{1,2,3}

¹Physics of Information and Quantum Technologies group, Centro de Física e Engenharia de Materiais Avançados (CeFEMA), Lisbon, Portugal

²Instituto Superior Técnico, Universidade de Lisboa, Lisbon, Portugal

³PQI — Portuguese Quantum Institute, Lisbon, Portugal

⁴Dahlem Center for Complex Quantum Systems, Freie Universität Berlin, 14195 Berlin, Germany

We introduce and investigate a data access model (*approximate sample and query*) that is satisfiable by the preparation and measurement of block encoded states, as well as in contexts such as classical quantum circuit simulation or Pauli sampling. We illustrate that this abstraction is compositional and has some computational power. We then apply these results to obtain polynomial improvements over the state of the art in the sample and computational complexity of distributed inner product estimation. By doing so, we provide a new interpretation for why Pauli sampling is useful for this task. Our results partially characterize the power of time-limited fault-tolerant quantum circuits aided by classical computation. They are a first step towards extending the classical data Quantum Singular Value Transform dequantization results to a quantum setting.

1 Introduction

Quantum computation subsumes classical computation. This is well-known, but it is stated under the assumption that the quantum computer is operating fault-tolerantly, i.e., that noise in the computation can be effectively suppressed. Despite ongoing and advancing efforts towards reaching this regime [1–7], to do so seems practically challenging. Thus, the current reality of quantum and classical computation is at odds with the first statement. This motivated a pragmatic exploration of the power of the existing limited computers [8–10], typically designated Noisy Intermediate-Scale Quantum (NISQ) devices, seeking to use them to go beyond classically tractable problems (or “achieve quantum advantage”). Careful analysis of these approaches revealed subtle caveats on the power of NISQ computers and algorithms: they must contend with optimization issues [11, 12], ingenious classical simulation techniques that exploit the presence of noise [13–16], and questions regarding the role of loading in and reading out of data, and whether that data is intrinsically quantum or not [9, 17–20].

These considerations motivate a related, if more theoretical, question: even if adopting a fault-tolerant quantum circuit setting, how powerful are hybrid algorithms: those that conjugate limited quantum computation (for some notion of limited) with classical computation? How strong can these limitations be before the model becomes at most as powerful as fully classical computation? Answers to these questions have the potential to impact the practical reality of NISQ: if a hybrid algorithm that does not have to contend with noisy operations can be proven to be unable to achieve advantage in some task, then surely a hybrid algorithm that furthermore must deal with noise cannot perform that same task. An instance of this argument is found in Ref. [20]. Another example is found in quantum query complexity results: a significant argument for the usefulness of quantum computation is the proven separation between classical and quantum computation under an oracle model [21]. If this separation were to not hold under some limitation on the number of queries performed coherently, this would argue against the usefulness of NISQ algorithms and computation. (This is not the case: Ref. [22] shows the existence of a problem for which a single

Miguel Murça: miguel.murca@tecnico.ulisboa.pt

quantum query is sufficient for an exponential query quantum advantage, and Ref. [23] shows that this separation also holds for total functions.)

1.1 Prior art

The line of work started by Van den Nest [24] and developed by Tang, Chia *et al.* [25–27] arguably addressed the power of quantum computation from this perspective. Ref. [24] showed that incorporating the readout and classical post-processing into the access model of a quantum algorithm resulted in a model that could be (efficiently) satisfied classically for some nontrivial cases. These included highly entangled states, classes of states beyond Clifford [28] and match-gate [29] states, the Quantum Fourier Transform, and even quantum subroutines of Simon’s [30] and Shor’s [31] algorithms. Therefore Ref. [24] evidenced, early on, that having access to limited fault-tolerant quantum computation with classical post-processing *could* be no better than purely classical processing: most strikingly illustrated by the results on Shor’s and Simon’s algorithms, this foreshadowed the relevance of data loading and readout in computational advantage.¹

Refs. [25, 26] leveraged this line of thought (combining it with the classical approach of Ref. [32]) to address a complementary question: generalizing the access model of Ref. [24] to matrices, what was the computational power of this kind of access with post-processing? Given that this access model could be taken to reflect measurements on a quantum RAM, this complementary question arguably addressed the power of a classical agent with access to limited fault-tolerant quantum computers.² Their results (where Ref. [26] strengthened and extended those of Ref. [25]) were surprisingly strong: such access to classical data, together with classical post-processing, could be enough to simulate a Quantum Singular Value Transformation (QSVT) [33] on the data, for sufficiently well behaved (for example, of low rank) data. Note that the QSVT framework is known to encompass a large number of significant quantum algorithms, including amplitude estimation [34] and matrix inversion [35], justifying the strength of this statement.

To better understand our contributions, it is useful to discuss in more detail how Chia *et al.* [26] obtained this result. Informally speaking, these authors showed the dequantization of QSVT on classical data by first abstracting the power of a quantum RAM (plus classical access to the loaded data) into a small set of promises on how this data could be accessed: the *sample and query* (SQ) access model. For a vector, these promises were that:

Definition 1. (*Sample and query* access to vectors — Informal statement of Def. 11)

1. One may sample an index by the corresponding entry’s ℓ_2 weight (suggesting a quantum measurement), taking time $\mathbf{s}$.
2. One may read exactly the ℓ_2 norm of the vector, taking time $\mathbf{n}$.
3. One may read exactly any given entry of the vector, taking time $\mathbf{q}$.

Chia *et al.* then showed that SQ access *composed*, in the sense that if one had SQ access to vectors x and y , one could compose this into SQ access to a linear combination of x and y ; or, extending the SQ access model to matrices, that one could compose SQ access to X and Y into SQ access to the inner or outer product, or to the Kronecker product, of X and Y . This sufficed to show that one could start with SQ access to, for example, matrix A and vector x , and “build up” to a QSVT $f(A^\dagger A)x$. It is worth noting that, to show composability, the authors had to also show that SQ *computed*, in the sense that, for example, having SQ access to x and y sufficed to approximate the inner product of x and y (a perspective more prevalent in Ref. [25]).

More rigorously, one finds that speaking only of the *sample and query* model is insufficient to make the above argument, and rather that the notion of SQ access needs to be relaxed to a more general ϕ -oversample and query (SQ $_\phi$) access model. (It suffices to think of ϕ as some relaxation

¹It also provided further evidence that the power of quantum algorithms (with respect to classical algorithms) was rooted beyond entanglement or an exponentially large representation space, although this was already known from results such as those in Refs. [28, 29].

²It is worth noting that the authors of Ref. [26] take a different point of view: given that there is not a known efficient way to prepare a quantum RAM, the authors interpret their results as evidence that a classical agent has no use for quantum power in the studied setting. This is because if the agent is willing to be inefficient in preparing a quantum RAM with their data, then there are classical (inefficient) procedures that suffice to satisfy the access model proposed by the authors.




Figure 1: The *approximate sample and query* (ASQ) access model restricts access to a vector to three rules (Fig. 1a). These can be satisfied by quantum or classical agents in the right conditions. By developing algorithms that operate over this abstraction, their runtime can be described in terms of the time necessary to carry out each access operation. If a classical agent can satisfy each operation fast enough, a dequantization result follows. In this paper, we establish some of the power of ASQ (Fig. 1b): we show that an agent can *compose* ASQ into linear combinations, and that they can use ASQ access to *compute* the inner product.

parameter, where the greater ϕ is than 1, the more imperfect the access is.) Under this relaxation, Ref. [26] showed that SQ_ϕ is approximately closed under arithmetic operations.

Since the SQ model was defined to abstract quantum RAM access to data, one might expect that the ability to prepare and measure a given state (possibly performing some auxiliary computation before measurement) would suffice to satisfy SQ access to the state’s representation in the computational basis (for example). But the consequences of this would be remarkable: because the results discussed above rely only on the data access promises of *sample and query*, all of the results of Ref. [26] would immediately apply to the scenario of preparation and measurement. The composability of the *sample and query* model would immediately imply that something akin to a linear combination of unitaries [36] could be attained at a known overhead, without having to prepare the various states coherently; and the dequantization results would imply that QSVTs could be carried out classically for sufficiently well conditioned states. Indeed, as noted by the authors of Ref. [26] themselves, and later in Ref. [37], the *sample and query* access model is *not* compatible with the case where one is given the ability to prepare and measure a general quantum state; in other words, when the input data is not classical. Thus, while applying to the particular case of quantum RAM with classical computation, the results of *sample and query* access do not immediately apply to all limited fault-tolerant quantum computation aided by classical post-processing.

1.2 Our contributions

In this work, we adapt the *sample and query* access model, such that it *can* be satisfied by preparation and measurement of a quantum state (with some auxiliary computation before measurement). We call this modified access model *approximate sample and query* (or ASQ). This modified model addresses the following points, which the original SQ model cannot accommodate:

1. Computing amplitudes or norms of quantum states to exponential precision by preparation and measurement is expected to be inefficient [35];
2. Estimating a certain amplitude to varying precision results in varying runtime (and, conversely, a given task may only require a certain degree of precision);
3. Quantum operations may succeed only probabilistically, and this success may not necessarily be signaled.

The *approximate sample and query* model addresses these points by incorporating the notion of precision (and its relationship with runtime), and admitting that the promises are satisfied only with some probability. Thus, informally, the promises of ASQ for a certain vector x become:

Definition 2. (*Approximate sample and query* for a vector x — Informal statement of Def. 20)

1. One may sample an index of x by the corresponding entry’s ℓ_2 weight (suggesting a quantum measurement), or fail with probability at most $1/3$ by raising an error flag. This action is assumed to take some time s .

2. One may estimate the ℓ_2 norm of x to some absolute error ϵ , or fail undetectably with probability at most $1/3$. This action takes time that varies with the choice of ϵ , and so is denoted $\mathbf{n}(\epsilon)$.
3. One may estimate a certain entry of x to some absolute error ϵ , or fail undetectably with probability at most $1/3$. This action also takes time varying with the choice of ϵ , thus denoted $\mathbf{q}(\epsilon)$.

Note that the model considers different admissible modes of failure for each operation. This is because if one can produce an estimation of a value to error ϵ with some probability bounded away from $1/2$, then one may show (by a Hoeffding bound) that a logarithmic overhead suffices to “boost” the correctness probability.³ However, the same cannot be done for the index sampling procedure, and so, in that case, undetectable failures become problematic.

We find that, nonetheless, these demands on how data may be accessed nicely conform to access to a quantum state: sampling may fail if, for example, post-selection for the sampled state fails. But, in that case, the failure is evident. Contrarily, the estimation of an amplitude by, for example, quantum phase estimation returns a value that is correct with high probability, but without any indication of success. Leveraging existing results on quantum tomography [38], we characterize the runtimes $\mathbf{s}$, $\mathbf{n}(\epsilon)$, and $\mathbf{q}(\epsilon)$ for the case where one may prepare quantum states that encode a (subnormalized) vector in a given subspace of their computational basis representation, without performing much more quantum computation before measuring. [In particular, for an $\mathbf{n}$ qubits state, we show how to satisfy ASQ access to the subnormalized vector with $O(\mathbf{n})$ quantum gates over $O(\log \mathbf{n})$ depth between the state preparation and measurement.] This is given in Remark 23. We show that these runtimes do not scale, for example, more than logarithmically with the dimension of the vector.

Perhaps surprisingly, we show that ASQ access can also be satisfied in other contexts. In particular, we characterize the runtimes $\mathbf{s}$, $\mathbf{n}(\epsilon)$, and $\mathbf{q}(\epsilon)$ for two setups other than preparation and measurement. First, we find that the promises of ASQ may be satisfied by a classical agent for a given state x if they are given a description of x by means of a quantum circuit that prepares it, and that circuit has few T gates and otherwise uses Clifford gates. This leverages existing results on the classical simulation of such circuits [39], and is given in Remark 25. Secondly, we find that if an agent has the ability to perform *Pauli sampling* of a certain quantum state x , as well as the ability to prepare and measure x , then one can satisfy ASQ access to the *representation of x in a Pauli basis*. (Pauli sampling here means sampling a Pauli string with probability proportional to its expectation value.) This is given in Remark 26.

We summarize these results in Fig. 1a.

Our motivation for investigating *approximate sample and query* is threefold. First: by defining a model that describes (to some extent) the ability to measure a quantum state, results on the computational power of this model (plus classical computation) establish results on the computational power of time-limited fault-tolerant quantum computation, when aided by classical computation. For example: a result on the complexity of computing the inner product of x and y given ASQ access to x and y establishes an upper bound to the computational complexity of the distributed computation of the inner product of two states. This is because if Alice and Bob can prepare, locally and respectively, states $|\psi\rangle$ and $|\phi\rangle$, then Alice and Bob have ASQ access to the representation of $|\psi\rangle$ and $|\phi\rangle$ in the computational basis (as stated above). Thus, these can be taken to be x and y , from whence using the ASQ result, Alice and Bob can compute the inner product between the two states with local operations and classical communication between them. Secondly: the compositional aspect of SQ is conceptually attractive, and we wish to find a similar property for the case of quantum data. As mentioned in the previous section, the power to prepare a linear combination of states is a significant quantum operation [36]. If one admits that the ASQ model captures (to some extent, and as established with computational results) the power of being able to prepare and measure a state, then a composition result on linear combinations shows that it is possible to (and effectively characterizes the overhead of) “faking” access to a linear combination of states, when given access to each state individually and complementing this with classical computation. Thirdly: given that Chia *et al.* [26] could dequantize QSVT based on the language of

³This is achieved by taking the median of several estimates logarithmic on the desired failure probability, a procedure typically referred to as “median of means”.

sample and query, this raises the question of whether an analogous result can be obtained for a sufficiently similar model that nonetheless reflects quantum access. Beyond the evident practical significance of such a quantization, as argued in the previous section, such an analogous result would provide much insight into the question: “in what conditions can a quantum computation be halted, and the remaining computation be carried out classically, without too much overhead?”

We address the first two points directly, and, with this, take a first step in addressing the third point. Namely, we find that the modifications of ASQ (with respect to SQ and as discussed above) mean the results of Chia *et al.* no longer apply. Thus, we first show that the ASQ model remains compositional, not going as far as showing that it is closed under arithmetic, but rather taking a first step in this direction by focusing on composition by linear combinations (motivated by linear combinations of unitaries, as stated). This requires us to define a relaxed access model, ASQ_ϕ (in line with AS_ϕ), after which we show the following:

Lemma 3. (Linear combinations of ASQ — Informal statement of Lemmas 27 and 29.) *Given ASQ_ϕ access to complex vectors $x_1, \dots, x_\tau$, and complex coefficients $\lambda_1, \dots, \lambda_\tau$, one may obtain $\text{ASQ}_{\phi'}$ access (with certainty) to the linear combination $\lambda_1 x_1 + \dots + \lambda_\tau x_\tau$, for suitable ϕ . If the vectors are approximately orthogonal and of constant and same size, then $\phi' \sim \phi\tau^2$. If, instead, one is willing to fail the construction with some probability, then $\phi' \sim \phi\tau$, matching that of Chia *et al.*’s Lemma 3.6.*

We also show that the ASQ model still holds some computational power. Again the results of Chia *et al.* no longer apply, so we focus on the task of inner product estimation. We show the following:

Lemma 4. (Inner product estimation with ASQ — Informal statement of Lemmas 30 and 32.) *Suppose (without loss of generality, normalized) complex vectors x and y of exponentially large dimension. Given efficient ASQ_ϕ access to x and the ability to compute entries of y , one may compute $x^\dagger y$ to error $\epsilon \|y\|_1$ efficiently with a probability of at least $2/3$. If one is given efficient ASQ_ϕ access to both x and y , then $x^\dagger y$ is efficiently computable to error $\epsilon \min(\|x\|_1, \|y\|_1)$ with probability at least $2/3$.*

The reader may be alarmed by the presence of the 1-norm: it may be of the order of the dimension of the vectors, which is taken to be exponentially large. However, we later find this dependence to be sufficient. We furthermore interpret this quantity as a concrete measure of “peakedness”, which may be controlled by a suitable choice of representation basis.

We summarize these results in Fig. 1b.

Finally, we show that the ASQ model has explanatory power and that the promises encoded by this model suffice to obtain improved results on the topic of distributed inner product estimation. We show that, under the light of the ASQ model, Pauli sampling is useful to compute the distributed inner product [40, 41] because, for sufficiently well-conditioned states, their representation in the Pauli basis has reasonably small 1-norm. We do so by first specializing Lemma 4:

Lemma 5. (Inner product estimation with exactly ASQ for real vectors — Informal statement of Lemma 34.) *If x and y are (without loss of generality, normalized) real vectors, and one has ASQ access to them (not ASQ_ϕ access), and if κ is an upper bound to $\max(\|x\|_1, \|y\|_1)$, then the dependence on κ of Lemma 4 may be improved. This requires that the runtime of estimating an entry $[\mathbf{q}(\epsilon)]$ is not too large.*

From here, we combine three observations: the first is that the inner product of two states may be computed in a Pauli basis (where the representations are real). The second is that the 1-norm of a state’s representation in the Pauli basis is directly connected to certain measures of “stabilizerness” [42, 43]. The third observation is that prior art on Pauli sampling [41] gives explicit procedures and runtimes for Pauli sampling, and thus gives $\mathbf{s}$, $\mathbf{n}(\epsilon)$ and $\mathbf{q}(\epsilon)$ for ASQ access to the states’ representation in the Pauli basis. By adapting this prior art somewhat (Corollary 37) and directly computing overlaps between Pauli amplitudes, we establish ASQ runtimes for Pauli sampling under a constraint in the “stabilizerness” measure corresponding to the 1-norm of the Pauli basis representation of the states.

This allows us to state a distributed inner product estimation algorithm that polynomially improves over the runtime and sample complexity of the state of the art (Theorem 35).

2 Notation and preliminaries

2.1 Notation

For some integer n , $[n]$ denotes the set $\{1, 2, \dots, n\}$. $\mathbf{i}$ denotes the imaginary unit. We denote vectors by lowercase letters. The i th entry of some vector x is denoted $x(i)$. $\|\cdot\|_p$ refers to the p -norm for vectors, and to the induced p -norm for matrices (thus, for some matrix A , $\|A\|_2$ refers to the induced two norm, or the Schatten ∞ -norm). $\|A\|_F$ refers to the Frobenius norm (or the Schatten 2-norm) of A . When unspecified, $\|\cdot\|$ refers to $\|\cdot\|_2$. We denote probability distributions by upright lowercase letters, e.g., p . We denote the probability of a certain event (for example, $X = x$) as $\mathbb{P}[X = x]$. Given a probability distribution p , we denote that a random variable X follows that distribution as $X \sim p$. We likewise denote the probability of observing value x under this distribution as $p(x)$ (i.e., this is equivalent to denoting $\mathbb{P}[X = x]$ for $X \sim p$). We will conflate the random variable with the value it takes where it is clear from context. The expectation of X is denoted $\mathbb{E}X$, and its variance $\text{VAR} X$. For an arbitrary vector x of dimension d , we often refer to a corresponding distribution, denoted p_x , to be a probability distribution with support $[d]$, satisfying $p_x(i) \equiv |x(i)|^2/\|x\|^2$ (i.e., p_x is the ℓ_2 sampling distribution of x). I_d denotes the d -dimensional identity matrix. Approximate quantities or estimates are generally denoted with a hat. All logarithms ($\log$) are natural unless otherwise noted. $\text{poly}(x)$ denotes some polynomial over x , and likewise $\text{polylog}(x)$ denotes some polynomial over $\log x$. $\log^k x$ denotes the k th power of $\log x$. The Fraktur font is used to make clear when referring to a number of qubits (thus, the base-2 logarithm of the dimension of the underlying state vector). E.g., an $\mathfrak{n}$ qubits pure quantum state has a $d \equiv 2^{\mathfrak{n}}$ -dimension underlying quantum state vector. “ x is defined as y ” is denoted $x := y$, and “ x taken to be the definition of y ” is denoted $x =: y$. Moreover, $x \equiv y$ emphasizes a coincidence by the definitions of x and y . “Big O notation” $[O(\cdot)]$ is used with its usual definition, with the use of “big omega” $[\Omega(\cdot)]$ and “big theta” $[\Theta(\cdot)]$ as popularized by Knuth [44]. The notation $\tilde{O}(\cdot)$ indicates that logarithmic terms on the quantities involved in the expression are suppressed. E.g., $\tilde{O}(xy + z)$ indicates terms of $\text{polylog}(x)$, $\text{polylog}(y)$, and $\text{polylog}(z)$ are suppressed in the argument.

2.2 Block encodings and projective encodings

Quantum circuits encode, forcibly, unitary maps. However, many mappings of interest are non-norm preserving. To deal with this, as well as to more generally encode non-unit norm vectors and matrices, one possible strategy is to encode the object or map in a suitable subspace, i.e., perform a block encoding:

Definition 6. (Block encoding [33]) *For A an $\mathfrak{n}$ qubit operator, if $\alpha, \epsilon > 0$ and $\mathfrak{m} \in \mathbb{N}$, then an $(\mathfrak{n} + \mathfrak{m})$ -qubits unitary U is an $(\alpha, \mathfrak{m}, \epsilon)$ -block encoding of A if*

$$\|A - \alpha(\langle 0|^{\otimes \mathfrak{m}} \otimes I_{2^{\mathfrak{n}}})U(|0\rangle^{\otimes \mathfrak{m}} \otimes I_{2^{\mathfrak{n}}})\| \leq \epsilon. \quad (1)$$

Informally, U is a block encoding of A because the block of U spanned by the $|0\rangle^{\otimes \mathfrak{m}}$ subspace of the first $\mathfrak{m}$ qubits is approximately proportional to A . This scaling is necessary because a block of U can never have norm greater than U itself, which cannot have norm greater than 1. Therefore, it is only possible to block encode A if $\alpha \geq \|A\|$.

It is possible to define an analogous notion for vectors:

Definition 7. (Block encoding — vectors) *For v a $2^{\mathfrak{n}}$ dimensional complex vector, defined up to a global phase, and if $\alpha, \epsilon > 0$ and $\mathfrak{m} \in \mathbb{N}$, then an $(\mathfrak{n} + \mathfrak{m})$ -qubits quantum state $|\psi\rangle$ is an $(\alpha, \mathfrak{m}, \epsilon)$ -block encoding of v if*

$$\|v - \alpha(\langle 0|^{\otimes \mathfrak{m}} \otimes I_{2^{\mathfrak{n}}})|\psi\rangle\| \leq \epsilon. \quad (2)$$

An $\mathfrak{m}$ -block encoding of v refers to a $(1, \mathfrak{m}, 0)$ -block encoding of v . Likewise an $(\mathfrak{m}, \epsilon)$ -block encoding of v refers to a $(1, \mathfrak{m}, \epsilon)$ -block encoding of v .

Block encodings may be further generalized to the notion of “projected encodings”, whereby there is a more generic projection operator, such that the operator or vector of interest is defined in the projector’s image:

Definition 8. (Projected encoding [33]) Let Π and $\tilde{\Pi}$ be projection operators, and U a unitary operator. Then $(U, \Pi, \tilde{\Pi})$ form a projected encoding of an operator A if $\tilde{\Pi}U\Pi \equiv A$.

This is straightforwardly generalized to approximations and operators of norm greater than one.

2.3 One-sided and two-sided error

A probabilistic algorithm (as are usually quantum algorithms) may fail in multiple, operationally different ways. For this reason, we wish to differentiate these failure modes in the analysis below. Borrowing from the language of probabilistic Turing machines and decision problems [45], we will refer to “one-sided” and “two-sided” types of error:

Definition 9. (One-sided error) A probabilistic decision algorithm fails with one-sided error p if with probability $1 - p$ the algorithm correctly accepts or rejects the input and with probability p rejects valid input. An algorithm with one-sided error never accepts invalid input.

For a non-decision algorithm, suppose that the algorithm outputs the result and a flag indicating if the result is valid. A one-sided algorithm may falsely report that the provided output is wrong, but will never report bogus output as correct. For this reason, we refer to probabilistic algorithms that may detectably fail with some probability as “one-sided”. For example: state post-selection in block encodings (cf. section 2.2).

Definition 10. (Two-sided error) A probabilistic decision algorithm fails with two-sided error p if with probability $1 - p$ the algorithm correctly accepts or rejects the input and with probability p randomly accepts or rejects the input.

Again we extend this definition to non-decision algorithms in the same fashion: if a probabilistic algorithm fails in an undetectable way, i.e., returns bogus output with some error probability, then we refer to them as “two-sided”. For example: estimating a coin’s bias from sampling, since one may be particularly unlucky and draw only heads, regardless of the coin’s weight.

2.4 Sample and Query

As discussed in Section 1.1, the *sample and query* access model is a data access model defined to investigate the role of a quantum RAM in obtaining quantum advantage with classical data. The *sample and query* (SQ) access model is defined as follows:

Definition 11. (Sample and query access [24, 26]) For a vector $v \in \mathbb{C}^d$, one has sample and query access to v if one can:

SQ1. read, for fixed $i \in [d]$, $v(i)$ to arbitrary precision, in time $\mathbf{q}$;

SQ2. obtain independent samples $i \in [d]$, following a distribution with probability mass function $p(i) = |v(i)|^2 / \|v\|^2$, in time $\mathbf{s}$;

SQ3. read $\|v\|$ to arbitrary precision, in time $\mathbf{n}$.

In the context where v is known classically and can be loaded into a quantum RAM (i.e., if one may prepare the state $\sum_{i \in [d]} [v(i)/\|v\|] |i\rangle$), it is clear that the requirements of SQ access are satisfied. Furthermore, because every entry can be known to arbitrary precision (or, at least, to exponential precision in runtime [24]), the probability of any $i \in [d]$ can be known exactly (or to exponential precision) in time $O(\mathbf{q} + \mathbf{n})$. These requirements may be split into query access (SQ1) and sample access (SQ2 and SQ3).

Ref. [26] introduces also the notion of ϕ -oversampling. Say that, for $d \in \mathbb{N}$, p and q are distributions supported on $[d]$. For $\phi \geq 1$, we say that p ϕ -oversamples q if, for all $i \in [d]$,

$$p(i) \geq \frac{1}{\phi} q(i). \quad (3)$$

If, for some $i \in [d]$, $p(i)$, $q(i)$ and ϕ are known exactly, then a sample of p can be transformed into a sample of q by rejection sampling. This motivates the definition of ϕ -oversample and query access (SQ $_\phi$) to a vector:

Definition 12. (ϕ -Oversample and query access) For a vector $v \in \mathbb{C}^d$ and $\phi \geq 1$, one has ϕ -oversample and query access to v if one has query access to v and sample and query access to a vector $\tilde{v} \in \mathbb{C}^d$, satisfying:

1. for all $i \in [d]$, $|\tilde{v}(i)|^2 \geq |v(i)|^2$;
2. $\|\tilde{v}\|^2 \leq \phi \|v\|^2$.

Note that, with these conditions, p_v is ϕ -oversampled by $p_{\tilde{v}}$.

Ref. [26] generalizes the SQ model to matrices. Denoting by $A(i, *)$ the i th row of A :

Definition 13. (ϕ -Oversample and query access — matrices [26]) For a matrix $A \in \mathbb{C}^{m \times n}$, one has SQ access to A if one has SQ access to every row of A , and SQ access to the vector of row norms, i.e., SQ access to the vector $a \in \mathbb{R}^m$, defined by its entries as $a(i) := \|A(i, *)\|$. One has ASQ_ϕ access to A if one has query access to every row of A and SQ_ϕ to $\tilde{A} \in \mathbb{C}^{m \times n}$, satisfying

1. for all $(i, j) \in [m] \times [n]$, $|\tilde{A}(i, j)|^2 \geq |A(i, j)|^2$;
2. $\|\tilde{A}\|_F^2 \leq \phi \|A\|_F^2$.

The model of *oversample and query* is approximately closed under arithmetic operations [26]. The inclusion of *oversampling and querying* — rather than just *sampling and querying* — is important to ensure this property. We refer to Ref. [26] for details, but note that these closures include mapping from SQ_ϕ access to matrices into $SQ_{\phi'}$ access to approximately linear combinations thereof, approximately the inner product thereof, or approximately general Lipschitz functions (and, in particular, polynomials) thereof.

2.5 Singly coherent copy tomography

In Ref. [38], van Apeldoorn *et al.* investigate the feasibility of tomography of a pure quantum state (say, $|\psi\rangle$) given the ability to prepare and measure copies of $|\psi\rangle$, without computing more than one copy of $|\psi\rangle$ in coherence. Contrary to what one might expect, it is possible to obtain an estimate of $|\psi\rangle$'s state vector in the computational basis (up to a global phase) to ϵ error in ∞ -norm distance (i.e., up to ϵ error in each entry of the tomography) in an amount of time and space that do not depend on the dimension of the state. We restate here Ref. [38]'s results significant to this work. We rephrase the results to include the runtime requirements for each task (since the original statements are concerned only with the necessary number of state preparations and measurements, but provide explicit constructions).

Theorem 14. (Estimation of state vector absolute values [38, Proposition 13]) Let $\epsilon > 0$, and $\delta \in [1/2, 1)$. Let $|\psi\rangle = \sum_{j \in [d]} \alpha_j |j\rangle$ be a quantum state, with $\alpha_j \in \mathbb{C}$, and let p , defined by its probability mass function as $p(j) = |\alpha_j|^2$, $j \in [d]$, be the probability distribution of the outcomes of a computational basis measurement. Let T be the time necessary to prepare and measure $|\psi\rangle$. Then, $O[\log(d/\delta)\epsilon^{-2}T]$ time [and in particular $8\epsilon^{-2} \log(2d/\delta)$ measurements] are sufficient to learn the absolute values of α up to ϵ ∞ -norm error, with success probability at least $1 - \delta$, and using only T quantum computation time between measurements. These values are stored in $O[\log(d/\delta)\epsilon^{-2}]$ space, and each entry can be read in $O(\log d)$ time.

The above Theorem does not provide phase information on the entries of the state vector. On the other hand, a global phase on the state is physically irrelevant. If one admits a procedure performing tomography of the state up to a global phase, then just the ability to prepare the state suffices to perform such a tomography (up to ∞ -norm error).

Theorem 15. (Estimation of state vector with phases by prepare and measure [38, Proposition 19]) Let $\epsilon > 0$ and $\delta \in [1/2, 1)$. Let $|\psi\rangle$ be a d -dimensional quantum state, which is prepared in black-box fashion in time T . Then, $\tilde{O}[T\epsilon^{-4} \log^2(d/\delta)]$ time is sufficient to compute a sparse representation of the quantum state (in the computational basis), up to ϵ ∞ -norm error, with success probability at least $1 - \delta$, and using at most $O[T + \epsilon^{-2} \log(d/\delta) \log^3 d]$ quantum time between measurements. The quantum state is stored in $O[\epsilon^{-2} \log(d/\delta) \log^2 d]$ space, and each entry can be read in $O(\log d)$ time.

See also Ref. [46] for the time necessary to perform some of the steps outlined in the proof of [38, Proposition 19].

2.6 Simulation of low T-gate count states

Circuits dominated by Clifford operations (Hadamard, phase, and controlled NOT gates) admit efficient classical simulation, in the sense that one may approximately sample from the distribution of measurements in the computational basis, and estimate the probability of a selected output [39]. In particular:

Theorem 16. (Simulation of low T-gate count states [39]) *Suppose a state $|\psi\rangle$ of n qubits is (classically) given by a quantum circuit description of a unitary U , such that $U|0\rangle^{\otimes n} = |\psi\rangle$. This description is taken to be of size polynomial in n , and by the gates of the circuit, which are either Hadamard, phase, controlled NOT, or T quantum gates. Let x be the representation of $|\psi\rangle$ in the computational basis. Then, if there are t T-gates in this description, and c Clifford gates in this description, one may sample from a distribution q supported on $[n]$, satisfying*

$$\sum_{i \in [n]} |q(i) - p_x(i)| \leq \Delta \quad (4)$$

in time $\tilde{O}[n(n+t)(c+t) + n(n+t)^3 + 2^t t^3 n^3 \Delta^{-4}] = \tilde{O}[\text{poly}(n)(1 + 2^t \Delta^{-4})]$; and, for choice of $i \in [n]$, one may estimate $p_x(i)$ to error ϵ with failure probability δ in time $O[(1+t)(c+t) + (n+t)^3 + 2^t t^3 \epsilon^{-2} \log \delta^{-1}] = O[\text{poly}(n)(1 + 2^t \epsilon^{-2} \log \delta^{-1})]$.

2.7 Pauli sampling and nonstabilizerness

The 4^n -qubit phaseless Pauli string operators form a basis to the space of $\mathbb{C}^{2^n \times 2^n}$. That is, if X , Y , and Z denote the corresponding single-qubit Pauli operators, then for an arbitrary⁴ ordering of the operators

$$P_i \in \{I, X, Y, Z\}^{\otimes n}, \quad i = 1, 2, \dots, 4^n, \quad (5)$$

the set of operators $\{P_i\}_i$ forms a basis of $\mathbb{C}^{2^n \times 2^n}$. This implies that every density matrix ρ of dimension d admits a decomposition of the form

$$\rho = \sum_{i \in [d^2]} \frac{\text{Tr}(\rho P_i)}{\sqrt{d}} P_i \quad (6)$$

where the factor of $d^{-1/2}$ is included to ensure orthonormality. Henceforth we will denote the representation of ρ in this basis as π_ρ , such that

$$\pi_\rho(i) = \frac{\text{Tr}(\rho P_i)}{\sqrt{d}}, \quad i = 1, 2, \dots, d^2. \quad (7)$$

Note that π_ρ is a real-valued vector and that it has unit 2-norm.

Pauli sampling corresponds to sampling $i \in [d^2]$ with probability $\pi_\rho(i)^2$ [43]. While effectively performing this sampling can be challenging [41, 47], the ability to do so can be powerful, enabling, for example, fidelity estimation in a constant number of samples [48].

Furthermore, the representation of a pure state $|\psi\rangle\langle\psi|$ in a Pauli basis (π_ψ) relates closely to the “nonstabilizerness” of $|\psi\rangle$ (sometimes also referred to as “magic”). Recall that a state $|\phi\rangle$ is a stabilizer state (or Clifford state) if there exists a unitary U composed of controlled NOT, Hadamard and phase gates, such that $U|0\rangle = |\phi\rangle$. Equivalently, stabilizer states of n qubits are those stabilized by a group of 2^n Pauli strings. Stabilizer states can be classically simulated and are less powerful than classical circuits [28, 49]. If a state has a large overlap with few stabilizer states, then we say that it has low “nonstabilizerness”. This fact can be exploited to, for example, classically compute and simulate measurements of the state [39]. As shown in Ref. [43], the α -Rényi entropies of π_ψ directly relate to various measures of “nonstabilizerness”, called stabilizer α -Rényi entropies, or simply α -magic:

$$M_\alpha(|\psi\rangle) = (1 - \alpha)^{-1} \log \left[\sum_{i \in [d^2]} \pi_\psi(i)^{2\alpha} \right] - \log d. \quad (8)$$

⁴Without loss of generality, but for concreteness, the ordering of the operators may be fixed in a “tableau” format, as introduced in [28]. This method maps $(x_1, x_2, \dots, x_{2n}) \in \{0, 1\}^{2n}$ to $(X^{x_1} Z^{x_2}) \otimes (X^{x_3} Z^{x_4}) \otimes (\dots) \otimes (X^{x_{2n-1}} Z^{x_{2n}})$, modulo the resulting phases.

If $H_\alpha(\cdot)$ denotes the α -Rényi entropy, $M_\alpha(|\psi\rangle) \equiv H_\alpha(p_{\pi_\psi}) - \log d$. Furthermore, from the properties of the Rényi entropies, $M_\alpha(|\psi\rangle) \in [0, \log d]$, with zero magic if and only if $|\psi\rangle$ is a stabilizer state. Also, this definition unifies other definitions of nonstabilizerness in the literature. In particular we highlight the connection made in Ref. [43] regarding “stab-norm”:

$$M_{1/2}(|\psi\rangle) \equiv 2 \log \mathbf{D}(|\psi\rangle) \quad (9)$$

where $\mathbf{D}(|\psi\rangle) := d^{-1} \sum_{i \in [d^2]} |\text{Tr}(P_i |\psi\rangle\langle\psi|)|$ is the “stab-norm” of $|\psi\rangle$ as defined in Ref. [42].

In Ref. [41], Hinsche *et al.* describe how to perform approximate Pauli sampling on a pure state provided that its M_0 magic is sufficiently bounded. (Approximate Pauli sampling means that the sampling follows a distribution not too different from the Pauli sampling distribution, in ℓ_1 distance.) Their procedure requires Bell measurements on $\rho \otimes \rho$, and has the following performance guarantees:

Theorem 17. (Performance guarantee for approximate Pauli sampling [41, Theorem 8]) *Let $\Delta \in (0, 1)$, and $\delta \in (0, 1)$. Let ρ be a pure, $\mathbf{n}$ qubit state with magic M_0 , with no more than χ entanglement across any of its bipartitions. Then, there exists an algorithm sampling from a distribution q such that, with overall probability $1 - \delta$*

$$\sum_{i \in [4^n]} |q(i) - p_{\pi_\rho}(i)| \leq \Delta, \quad (10)$$

provided that one can perform N Bell measurements on $\rho \otimes \rho$, with

$$N = O\left(\frac{\mathbf{n}^3 2^{4\chi} 2^{2M_0}}{\Delta^4} \log \frac{1}{\delta}\right). \quad (11)$$

Producing a single sample takes time $O(N\mathbf{n}^2)$ time in classical postprocessing of the obtained Bell measurement data.

3 Approximate Sample and Query

We define the *approximate sample and query* model as a modification of the *sample and query* model (sec. 2.4). The main observations motivating these modifications are that, even given arbitrarily many copies of some general state, one cannot cheaply compute a given amplitude to exponential precision; that the required precision will dictate the cost of the operation; and that quantum operations may fail depending on the operation, with one-sided or two-sided error.

Definition 18. (Approximate query access) *We say that we have Approximate Query (AQ) access to a vector $x \in \mathbb{C}^d$ if, for any choice of $\epsilon \in (0, 1]$, and for any $i \in [d]$, one may obtain an estimate of $x(i)$ to absolute error ϵ , with two-sided error probability of at most $1/3$, and runtime $\mathbf{q}(\epsilon)$.*

Definition 19. (Approximate sample access) *We say that we have Approximate Sample (AS) access to a vector $x \in \mathbb{C}^d$ if we may sample from a probability distribution p_x , defined by its probability mass function as*

$$p_x(i) = \frac{|x(i)|^2}{\|x\|^2}, \quad i = 1, 2, \dots, d \quad (12)$$

with one-sided error probability at most $1/3$ and runtime $\mathbf{s}$, and $\|x\|^2$ may be obtained to additive error ϵ , with at most two-sided error probability $1/3$, and runtime $\mathbf{n}(\epsilon)$. (By one-sided error probability over the sampling, we mean that either a sample is successfully produced, or the algorithm emits a flag indicating failure.)

The definition of Approximate *sample and query* then combines these two forms of access:

Definition 20. (Approximate sample and query) *Say that we have approximate sample and query (ASQ) access to a vector $x \in \mathbb{C}^d$ if we have AS and AQ access to it.*

The corresponding definition for oversampling follows naturally:

Definition 21. (Approximate ϕ -oversample and query access) For $\phi \geq 1$, say that we have Approximate ϕ -Oversample and Query (ASQ_ϕ) access to a vector $x \in \mathbb{C}^d$ if we have AQ access to x (Definition 18) and we have AS access to $\tilde{x} \in \mathbb{C}^d$ (Definition 19) satisfying that, for all $i \in [d]$, $|\tilde{x}(i)| \geq |x(i)|$ and $\|\tilde{x}\|^2 \leq \phi \|x\|^2$. [Recall that the time to query an entry of x to error ϵ is $\mathbf{q}(\epsilon)$, the time to query $\|\tilde{x}\|^2$ to error ϵ is $\mathbf{n}(\epsilon)$, and that the time to sample from $p_{\tilde{x}}$ is $\mathbf{s}$.]

Compare this definition with that of Def. 12: we are careful to require only *sample* access to $\tilde{x}$, whereas *oversample* and *query* asks for *sample* and *query* access to $\tilde{x}$.

3.1 Satisfying approximate sample and query access

The results of van Apeldoorn *et al.* [38], discussed in section 2.5, together with the standard properties of quantum states and measurements suffice to see that access to single-copy preparations of a quantum state suffices to satisfy ASQ access to the underlying quantum state vector (when represented in the computational basis). Because we only require one amplitude at a time, we may reduce the necessary quantum time after the state preparation:

Lemma 22. Let $|\psi\rangle$ be an $\mathbf{n}$ qubits state, prepared in time T in black-box fashion. Then, $O[T\epsilon^{-2}\mathbf{n}]$ time suffices to estimate $\langle i|\psi\rangle$ for any choice of $i \in [2^n]$ (up to a global phase, but with phases consistent across estimates), to error ϵ and with two-sided error probability of at most $1/3$. Furthermore, at most $O(T + \mathbf{n})$ quantum time [and a $O(\log \mathbf{n})$ depth overhead in the quantum circuits] is needed between measurements.

Proof. The proof is a modification of that of Theorem 15. Employing Theorem 14, retrieve an estimate of $|\langle i|\psi\rangle|$ for all $i \in [2^n]$ up to ∞ -norm error $\epsilon/16$ and with error probability $1/9$. This takes time $O[T\epsilon^{-2}\mathbf{n}]$, and the resulting vector is k -sparse, where $k = O(\epsilon^{-2}\mathbf{n})$. Denote this vector by r . In time $O(k)$, round any entry of this vector below $\epsilon/2$ to zero. Then, in time $O(k)$, determine $m \in [2^n]$ such that $r(m)$ is the greatest entry of r . Now, for given $j \in [2^n]$, if $r(j) = 0$, estimate $\langle j|\psi\rangle$ as 0. If $j = m$, estimate $\langle j|\psi\rangle$ as $r(m)$. Otherwise, write unitaries V and V' that map $|m\rangle \mapsto |0\rangle$ and $|i\rangle \mapsto |1\rangle$ and that then apply a Hadamard gate on the least significant bit, in the case of V , or then apply a phase gate followed by a Hadamard gate on the least significant bit, in the case of V' . This is done in time $O(\mathbf{n})$ and produces a circuit with $O(\mathbf{n})$ gates. Measuring $V|\psi\rangle$ and $V'|\psi\rangle$ $O(\epsilon^{-2})$ times (for a total $O[(T + \mathbf{n})\epsilon^{-2}]$ runtime) suffices to estimate both $|\langle m|\psi\rangle - \langle j|\psi\rangle|$ and $|\langle m|\psi\rangle - \mathbf{i}\langle j|\psi\rangle|$ to error $\epsilon/16$ with joint success probability at least $2/9$. Let these estimates be denoted by s and t respectively. Then, $[r(m)^2 + r(j)^2 - s^2]/[2r(j)] + \mathbf{i}[r(m)^2 + r(j)^2 - t^2]/[2r(j)]$ is an ϵ -error estimate of $\langle j|\psi\rangle$ with two-sided error probability at most $1/3$. The depth overhead is given by the depth of the circuits implementing V and V' , which are the same up to a difference of 1. The depth of V is the same as that of a circuit mapping $|0\rangle \mapsto |0\rangle$, $|m \oplus i\rangle \mapsto |1\rangle$, up to a difference of 1 (and where $\oplus$ denotes bitwise XOR). Since $O(1)$ layers of controlled NOT gates can map $|u\rangle$ to $|v\rangle$ such that $|v| = |u|/2$ (while mapping $|0\rangle$ to $|0\rangle$), and since $|m \oplus i| \leq \mathbf{n}$, the depth of $O(\log \mathbf{n})$ follows. $\square$

Then, the following remark follows:

Remark 23. Given the ability to prepare a quantum state $|\psi\rangle$ of $\mathbf{n}$ qubits in time T (in black-box fashion), and the ability to run quantum circuits of up to time $O(T + \mathbf{n})$ [or the ability to run quantum circuits $O(\log \mathbf{n})$ deeper], one has ASQ access (Definition 20) to $x \in \mathbb{C}^{2^n}$, where $x(i) = \langle i|\psi\rangle$, up to a global phase on $|\psi\rangle$. The runtime costs are $\mathbf{n}(\epsilon) = 0$, $\mathbf{s} = T$, $\mathbf{q}(\epsilon) = O(T\epsilon^{-2}\mathbf{n})$. More generally, given the ability to prepare an $\mathbf{m}$ -block encoding of a (subnormalized) state $|\phi\rangle$ of $\mathbf{n}$ qubits in time T (in black-box fashion), and the ability to run circuits of up to time $O(T + \mathbf{n})$, one has ASQ access to $x \in \mathbb{C}^{2^n}$, where $x(i) = \langle i|\phi\rangle$, up to a global phase on $|\phi\rangle$. In this case, the runtime costs are $\mathbf{n} = O(\epsilon^{-2}T)$, $\mathbf{s} = O(T\|x\|^{-2})$, $\mathbf{q}(\epsilon) = O(T\epsilon^{-2}\mathbf{n})$.

Note that the procedures outlined in Theorem 14 and Lemma 22 apply unchanged to the block encoding setting, from whence the above remark specifies $\mathbf{q}(\epsilon)$ for block encodings. On the other hand, the sample time $\mathbf{s}$ for block encodings corresponds to a single successful post-selection. Because successful post-selection happens with probability $\|x\|^2$ (from the standard properties of quantum measurement), and from a Chebyshev inequality, the stated sample time $\mathbf{s}$ follows. (Further, these do not require that $\|x\|$ is known.) From this follows also the norm query time $\mathbf{n}$.

The remark above does not consider $(\mathbf{m}, \epsilon)$ -block encodings: clearly, access to a $(\mathbf{m}, \epsilon)$ -block encoding of a vector x gives *approximate sample and query* access to x' satisfying $\|x' - x\| \leq \epsilon$. While the consequences of this for query access are obvious — one has query access to x up to error ϵ — the effects on *approximate sample access* are made more meaningful by the following lemma (also observed in [25, Lemma 4.1] for the particular case of the reals, and for which the proof proceeds identically):

Lemma 24. *Let $x, y \in \mathbb{C}^d$ be two vectors such that $\|x - y\| \leq \epsilon$ for some $\epsilon > 0$. Let p_x and p_y be the distributions with probability mass function $p_x(i) = |x(i)|/\|x\|^2$, and $p_y(i) = |y(i)|^2/\|y\|^2$. Then, the distance between these two distributions is bounded as*

$$\sum_{i \in [d]} |p_x(i) - p_y(i)| \leq \frac{4\epsilon}{\|x\|}. \quad (13)$$

Thus, we have shown that *sample and query* access captures at least some of the power of the ability to prepare and measure (singly coherent) copies of a block encoding of a state.

However, the definition allows us to satisfy ASQ access in other contexts. For example, leveraging Theorem 16 (cf. section 2.6), we note that, given a low T -gate count circuit describing a quantum state x , we obtain a situation quite similar to that of $(\mathbf{m}, \epsilon)$ -block encodings:

Remark 25. *Given a classical description of a quantum state $|\psi\rangle$ in the terms of Theorem 16, thus by a circuit with at most t T -gates, and denoting x as the representation of $|\psi\rangle$ in the computational basis, one has approximate sample access to x' satisfying $\|x - x'\| \leq \Delta/4$ at runtime cost $\mathbf{n}(\epsilon) = 0$ and $\mathbf{s} = \tilde{O}[\text{poly}(n)(1 + 2^t \Delta^{-4})]$. Furthermore, for choice of $i \in [2^n]$, the circuit performing a swap test between $U|0\rangle^{\otimes n}$ and $|i\rangle$ has $O(\mathbf{n})$ qubits and $O(t)$ T -gates; thus also one has approximate query access to x at runtime cost $O[\text{poly}(n)(1 + 2^t \epsilon^{-2})]$. This does not require quantum computation.*

Finally, the *approximate sample and query* access model is sufficiently general that it also captures some of the power of measuring a quantum state in a completely different basis; namely, the Pauli string basis:⁵

Remark 26. *Given the ability to prepare copies of a quantum state $|\psi\rangle$ of dimension $2^n \equiv d$ in time T , and if performing Pauli sampling on $|\psi\rangle$ takes time T' , then one has ASQ access to $\pi_{|\psi\rangle\langle\psi|}$ [cf. eq. (7)], at runtime costs $\mathbf{n}(\epsilon) = 0$, $\mathbf{s} = T'$, and $\mathbf{q}(\epsilon) = O(Td^{-1}\epsilon^{-2})$.*

Note that the query runtime cost follows from the fact that, having chosen a given Pauli string P_i , rotated single-qubit measurements suffice to estimate $\langle\psi|P_i|\psi\rangle$. From Ref. [50], we may attain error ϵ with $O(\epsilon^{-2})$ samples; and, from eq. (7), error $\epsilon/\sqrt{d}$ suffices. This procedure requires at most $O(T + \mathbf{n})$ quantum time between measurements.

3.2 Composing approximate oversample and query access

One significant property of the *sample and query* model is that it can be composed, meaning that *sample and query* access to separate objects (say, two vectors x and y) can be leveraged to obtain (over-)sample and query access to arithmetic over those objects (say, $x + y$). This is the property enabling partial dequantization of linear algebra algorithms in Ref. [26]. While the proofs in Ref. [26] for composition do not translate to the *approximate oversample and query* model (as they rely on the ability to query to arbitrary precision), we now illustrate that *approximate oversample and query* access retains composability. We do not show a result as strong as those stated for *sample and query*, but take a first step in this direction. Since the ability to perform linear combinations of quantum states given the ability to prepare each state individually is already a significant operation in quantum computation [36], we show that ASQ_ϕ composes over linear combinations.

Lemma 27. (Linear combinations from ASQ_ϕ) *Let $\phi \geq 1$. Given ASQ_ϕ to τ vectors $x_1, x_2, \dots, x_\tau \in \mathbb{C}^d$, and τ nonzero complex coefficients $\lambda_1, \dots, \lambda_\tau \in \mathbb{C}$, one may construct $\text{ASQ}_{\phi'}$ access to the*

⁵We note here the parallel to Ref. [24]’s “rotated bases computationally tractable states”.

linear combination $u := \sum_{j \in [\tau]} \lambda_j x_j$, where $\phi' = \phi \tau^2 \kappa^2$, and κ is the spectral condition number of the matrix whose columns are the vectors $x_1, \dots, x_\tau$. The costs of access are:

$$\mathbf{q}(\epsilon) = O[\tau \log \tau \mathbf{q}_x(\epsilon/\|\lambda\|)] \quad (14)$$

$$\mathbf{n}(\epsilon) = O\{\tau \log \tau \mathbf{n}_x[\epsilon/(\tau^2 \|\lambda\|^2)]\} \quad (15)$$

$$\mathbf{s} = O(\tau + \mathbf{s}_x), \quad (16)$$

where $\lambda = (\lambda_1, \dots, \lambda_\tau)$, and $\mathbf{q}_x$, $\mathbf{n}_x$, and $\mathbf{s}_x$ are upper bounds to the costs defined in Definition 21 for all $x_1, \dots, x_\tau$.

Proof. Query access to u is attained simply by querying each entry to sufficient precision; namely, for $j \in [\tau]$, $i \in [n]$, if $x_j(i)$ is known to precision $\epsilon/|\lambda_j|$, $u(i)$ is known to precision ϵ . With the median of $18 \log(6\tau)$ copies of each estimate (considering the real and imaginary part of each estimate separately), one also ensures that all estimates are correct with overall probability $2/3$. [Note that this requires each $x_j(i)$ to be estimated to precision $\epsilon/(\sqrt{2}|\lambda_j|)$, but from this results only a constant factor.]

To show that oversampling in the stated conditions is also possible, we propose an explicit procedure and then argue that the oversampling factor ϕ' is as defined. Let $\tilde{u}$ be the “oversampling vector” for u , such that we should satisfy sample access to $\tilde{u}$ (cf. also Definition 21). Then sample from $\tilde{u}$ by first choosing j with probability $|\lambda_j|^2/\|\lambda\|^2$ and outputting i as sampled from approximate ϕ -oversampling of x_j . This corresponds to sampling from $p_{\tilde{u}}$, where

$$\tilde{u}(i) = \sqrt{\tau \left(\sum_{j \in [\tau]} \|\tilde{x}_j\|^2 \right) \sum_{k \in [\tau]} |\lambda_k|^2 \frac{|\tilde{x}_k(i)|^2}{\|\tilde{x}_k\|^2}}. \quad (17)$$

For all $i \in [d]$, $|\tilde{u}(i)| \geq |u(i)|$, by a triangle inequality and a geometric mean–harmonic mean inequality, and because $\tau \sum_{j \in [\tau]} \|\tilde{x}_j\|^2 \geq \tau \|\tilde{x}_k\|^2$ for all $k \in [\tau]$. Thus this procedure samples from $p_{\tilde{u}}$, which is oversampling p_u . The oversampling factor ϕ' is bounded, by definition, by an upper bound to $\|\tilde{u}\|^2/\|u\|^2$. Write X as the $(d \times \tau)$ matrix whose columns are $x_1, \dots, x_\tau$, and note that $\sum_{j \in [\tau]} \|x_j\|^2 \equiv \|X\|_F^2 \equiv \text{Tr}(X^\dagger X)$. Then, we have that

$$\phi' \leq \frac{\phi \tau \|X\|_F^2 \|\lambda\|^2}{\|X\lambda\|^2} \leq \frac{\phi \tau}{\sigma_{\min}(\frac{X}{\|X\|_F})^2} \leq \phi \tau^2 \kappa^2 \quad (18)$$

where $\sigma_{\min}(A)$ denotes the smallest singular value of A , and κ denotes the ratio of the largest to the smallest singular value of X . The first inequality follows from $\|\tilde{x}_j\|^2 \leq \phi \|x_j\|^2$, the second inequality by a “min-max” formulation of the singular values, and the third inequality follows from the identity $\text{Tr}(X^\dagger X) = \sum_{i \in [\text{Rank } X]} \sigma_i^2$, if σ_i is the i th singular value of X . Since the outlined procedure requires a single sample, it correctly fails with one-sided $1/3$ probability error.

Finally, estimating the norm $\|\tilde{u}\|^2$ requires only that every norm $\|\tilde{x}_j\|^2$ is known to error $\epsilon/(\tau^2 \|\lambda\|^2)$. With the median of $18 \log(3\tau)$ copies of each estimate for $\|\tilde{x}_j\|$ (as enabled by the ASQ_ϕ access to each $\tilde{x}_j$), one can ensure that every estimate is correct with two-sided $1/3$ probability error. $\square$

The dependence on the condition number κ in the lemma above is not a problem if the various $x_1, \dots, x_\tau$ are approximately orthogonal and not too different in magnitude. But this is a rather restrictive condition. Indeed, the conditions stated by the ASQ model regarding the error “sidedness” for each operation prevent an approach exactly like that of Ref. [26, Lemma 3.9].⁶ If we are willing to fail the construction with some probability, then an analogous result is recovered, and the dependence on such a κ vanishes.

We first require a lemma on obtaining relative error estimates from absolute error:

⁶One finds that many proofs regarding the SQ model rely on exactly querying a given entry or the norm of the vector to cancel out undesirable terms. The introduction of error and randomness impedes the same approach.

Algorithm 1 Calculating a relative error estimate via online access to randomized queries to absolute error estimates.

Require: relative precision ρ

Require: failure probability δ

Require: online access to $\hat{x}(\epsilon)$ such that $|\hat{x}(\epsilon) - x| \leq \epsilon$ with probability $2/3$

<pre> 1: $k \leftarrow 1$ loop: 2: $\epsilon \leftarrow 2^{-k}/\sqrt{2}$ 3: $l \leftarrow \text{[EMPTY LIST]}$ 4: do $18 \log[10^4 2^{k+1}/\delta]$ times: 5: Sorted insert $\hat{x}(\epsilon)$ into l 6: end 7: $\mu \leftarrow \text{MEDIAN}(l)$ 8: if $\epsilon \leq \frac{1}{2}\mu$ then 9: goto final 10: end if </pre>	<pre> 11: $k \leftarrow k + 1$ 12: goto loop final: 13: $\perp \leftarrow \mu - \epsilon$ 14: $r \leftarrow \text{[EMPTY LIST]}$ 15: $c \leftarrow \text{[EMPTY LIST]}$ 16: do $18 \log[8/\delta]$ times: 17: Sorted insert $\text{Re}[\hat{x}(\rho\perp)]$ into r 18: Sorted insert $\text{Im}[\hat{x}(\rho\perp)]$ into c 19: end 20: return $\text{MEDIAN}(r) + \mathbf{i} \text{MEDIAN}(c)$ </pre>
--	--

Lemma 28. For choice of $\rho \in (0, 1]$, and given the ability to estimate a quantity $x \in \mathbb{C}$ to absolute error ϵ in time $\mathbf{q}(\epsilon)$ with probability $2/3$, algorithm 1 computes an estimate of x to relative error ρ with probability $1 - \delta$. The algorithm halts with overwhelming probability in time $O[\mathbf{q}(\rho|x|)(1 + \log \frac{1}{|x|}) \log \frac{1}{\delta}]$ [under the assumption that $\mathbf{q}\{\Omega(\epsilon)\} = O\{\mathbf{q}(\epsilon)\}$].

Proof. Algorithm 1 computes increasingly accurate estimates of x , looking to halt when the estimate's absolute value (say, $|\hat{x}|$) is twice as large as the estimation error (say ϵ). This ensures that $\epsilon \leq |x|/3$, and that $\frac{2}{3}|x| \leq |\hat{x}| \leq 2|x|$, i.e., that $|\hat{x}|$ is of the order of $|x|$. Once this bound for $|x|$ is obtained (say, $\perp$), it suffices to output an estimate of error $\rho\perp$.

To deal with the randomization in the estimations of $|\hat{x}|$, we use the fact that, for $w \in \mathbb{R}$ and $\hat{w}$ such that $|w - \hat{w}| \leq \epsilon$ with probability $2/3$, the median of $18 \log(1/\delta)$ independent copies of $\hat{w}$ is within ϵ error of w with probability $1 - \delta$ (from Hoeffding's inequality). This allows the algorithm to guarantee that: false positives do not occur within the loop with more than (overall) probability $\delta/2$; and that the loop is not kept running by bogus estimates for too long. Let k_{MAX} be the maximum number of iterations of the loop needed to achieve a sufficiently small value of ϵ : from the considerations above, $k_{\text{MAX}} = \lceil \log_2(3/|x|) \rceil$. After k exceeds this value, the probability of repeatedly looping from bad estimates vanishes overwhelmingly: the probability of reaching the $(k_{\text{MAX}} + m)$ th iteration of the loop is $(\delta/2^{k_{\text{MAX}}+1})(\delta/2^{k_{\text{MAX}}+2}) \dots (\delta/2^{k_{\text{MAX}}+m})$, from which it is easy to see that with probability at least $1 - \omega$, the loop halts at the $(k_{\text{MAX}} + m)$ th iteration, where $m = \log_2(1/\omega)/[1 + k_{\text{MAX}} + \log_2(1/\delta)]$. Setting $\delta = \delta'\omega$, this ensures that the loop halts in k_{MAX} iterations with probability $1 - \omega$. (If the loop does not halt in k_{MAX} , it halts soon thereafter with overwhelming probability, from the same considerations.) From this observation, the explicit form of the algorithm, and the value of k_{MAX} follows the stated runtime. $\square$

With this lemma, we may conclude an alternative linear combination composition lemma that does not have a condition number dependence and is very similar to [26, Lemma 3.9], at the cost of succeeding only probabilistically:

Lemma 29. (Probabilistic linear combinations from ASQ $_\phi$) Let $\phi \geq 1$ and $\delta \in (0, 1)$. Given ASQ $_\phi$ to τ vectors $x_1, \dots, x_\tau \in \mathbb{C}^d$, and τ nonzero complex coefficients $\lambda_1, \dots, \lambda_\tau$, one may construct ASQ $_{\phi'}$ access to the linear combination $u := \sum_{j \in [\tau]} \lambda_j x_j$ with probability at least $1 - \delta$, where $\phi' = 4\tau\phi(\sum_{k \in [\tau]} \|\lambda_k x_k\|^2)/\|u\|^2$. The costs of access are:

$$\mathbf{q}(\epsilon) = O[\tau \log \tau \mathbf{q}_x(\epsilon/\|\lambda\|)] \quad (19)$$

$$\mathbf{n}(\epsilon) = O(1) \quad (20)$$

$$\mathbf{s} = O(\tau + \mathbf{s}_x), \quad (21)$$

where $\lambda = (\lambda_1, \dots, \lambda_\tau)$, and $\mathbf{q}_x$, $\mathbf{n}_x$, and $\mathbf{s}_x$ are upper bounds to the costs defined in Definition 21 for all $x_1, \dots, x_\tau$. For $\|x\|^2$ the smallest $\|x_j\|^2$, $j = 1, \dots, \tau$, this construction requires a one-time preliminary procedure that runs in time $O[\tau \mathbf{n}(\|x\|^2)(1 + \log \frac{1}{\|x\|}) \log \frac{\tau}{\delta}]$ [with overwhelming probability and under the assumption that $\mathbf{n}\{\Omega(\epsilon)\} = O\{\mathbf{n}(\epsilon)\}$].

Proof. Using Lemma 28, start by computing a $1/2$ -relative error estimate of each norm $\|\tilde{x}_j\|^2$, $j = 1, \dots, \tau$, with overall error probability δ , and store these estimates as $\hat{n}_1^2, \dots, \hat{n}_\tau^2$. The runtime for this follows directly from Lemma 28 and is as stated therein. The construction fails only if any of these estimates are wrong; henceforth we assume these are correct.

We otherwise proceed similarly to 27: estimating each $x_j(i)$ to error $\epsilon/|\lambda_j|$ suffices, with the median of means being used to ensure every estimate is correct with probability $2/3$.

Sampling is done by first choosing j with probability $|\lambda_j|^2 \hat{n}_j^2 / (\sum_{k \in [\tau]} |\lambda_k|^2 \hat{n}_k^2)$ and outputting i as sampled from approximate ϕ -oversampling of x_j . This corresponds to sampling from $p_{\tilde{u}}$, where $\tilde{u}$ is defined by its entries as

$$\tilde{u}(i) = \sqrt{2\tau \sum_{k \in [\tau]} |\lambda_k|^2 \hat{n}_k^2 \frac{|\tilde{x}_k(i)|^2}{\|\tilde{x}_k\|^2}}. \quad (22)$$

This is for all $i \in [d]$, $|\tilde{u}(i)|^2 \geq |u(i)|^2$, by a triangle and a geometric mean–harmonic mean inequality, and making use of the fact that, by hypothesis and for all $j \in [\tau]$, $\frac{1}{2}\|\tilde{x}\|^2 \leq \hat{n}_j \leq \frac{3}{2}\|\tilde{x}\|^2$. To bound ϕ' , note that $\|\tilde{u}\|^2 = 2\tau \sum_{k \in [\tau]} |\lambda_k|^2 \hat{n}_k^2$, and again make use of the bounds of $\hat{n}_j$, and the fact that $\|\tilde{x}_j\|^2 \leq \phi\|x_j\|^2$. Furthermore, this explicit norm can be calculated once as a preliminary step, after which computing it corresponds to a lookup. $\square$

3.3 Computation with approximate oversample and query access

An important characterization of the ASQ model is its computational power: it is only meaningful to say that one can satisfy the “promises” of ASQ if one knows how to perform useful computation by leveraging those promises. Here, we focus on the task of performing an inner product estimation between two vectors, given ASQ access to either one or both of the vectors. The computational cost of computing the inner product between two states, given the ability to prepare singly coherent copies of each, is recognized as significant (for example, for fidelity estimation) and is under ongoing study [40, 41, 51, 52]. On the other hand, the inner product can be regarded as the simplest nontrivial “reduction” of two vectors into a scalar to estimate and is also considered in Refs. [25–27]. This task will find immediate application in the application of ASQ to distributed inner product estimation (Sec. 4)

We begin with the asymmetric case, i.e., *approximate sample and query* access to one vector, and classical oracle access to the entries of another:

Lemma 30. (Asymmetric inner product estimation) Let $x, y \in \mathbb{C}^d$, and $\epsilon \in (0, 1]$. Suppose ASQ $_\phi$ access to x and query access to y (in the language of Ref. [26]; cf. Def. 11). Then, one may estimate $x^\dagger y$ to additive error $\epsilon\|y\|_1$ and success probability $2/3$ by algorithm 2. The runtime of the algorithm is

$$\tilde{O}[\phi\epsilon^{-2}\|x\|^2 \log d + \phi\epsilon^{-2}\|x\|^2 \mathbf{q}(\epsilon) + \phi^2\epsilon^{-4}\|x\|^4 \log d \mathbf{s} + \mathbf{n}(\|x\|^2)],$$

with very high probability, and under the assumption that $\mathbf{q}[\Omega(\epsilon)] = O[\mathbf{q}(\epsilon)]$ and likewise for $\mathbf{n}$.

Algorithm 2 Calculating an inner product estimate via ASQ_ϕ access to one of the vectors (x), and classical RAM access to the other (y), with success probability at least $2/3$, and absolute error at most $\epsilon \|y\|_1$.

Require: ASQ_ϕ access to $x \in \mathbb{C}^d$ Require: Query access to $y \in \mathbb{C}^d$ Require: Precision $\epsilon \in (0, 1]$ 1: function IMPROVE(x, ϵ, δ) 2: $r \leftarrow []$ 3: $c \leftarrow []$ 4: $\epsilon \leftarrow \epsilon/\sqrt{2}$ 5: do $\lceil 18 \log(1/\delta) \rceil$ times: 6: Sorted insert $\text{Re}[\hat{x}(\epsilon)]$ into r 7: Sorted insert $\text{Im}[\hat{x}(\epsilon)]$ into c 8: end 9: return $\text{MEDIAN}(r) + \mathbf{i} \text{MEDIAN}(c)$ 10: end function 11: $\hat{n}^2 \leftarrow [\text{ALGORITHM 1}](\ x\ ^2, \rho = \frac{1}{4}, \delta = \frac{1}{9})$ 12: $\gamma \leftarrow \epsilon^2/(135 \hat{n}^2)$	13: $h \leftarrow [\text{EMPTY HISTOGRAM}]$ 14: do $\lceil 512 \gamma^{-2} \log(18d) \rceil$ times: 15: repeat 16: $i \leftarrow \text{OVERSAMPLE}(x)$ 17: until sample is valid. 18: Insert i into h. 19: end 20: $m \leftarrow []$ 21: do $\lceil 7 \hat{n}^2 \epsilon^{-2} \rceil$ times: 22: repeat 23: $i \leftarrow \text{OVERSAMPLE}(x)$ 24: until sample is valid. 25: $\hat{p} \leftarrow \text{FREQUENCY}(h; i)$ 26: if $\hat{p} \geq 3\gamma/2$ then 27: $\hat{x} \leftarrow \text{IMPROVE}[x(i), \epsilon/4, \epsilon^2/(127 \hat{n}^2)]$ 28: Insert $(\hat{p} + \gamma/2)^{-1} \hat{x}^* y(i)$ into m. 29: end if 30: end 31: return $\text{AVERAGE}(m)$
--	---

Proof. Broadly speaking, Algorithm 2 computes an approximation to every probability using Theorem 14. Then, it filters the values of i , as sampled from the ASQ_ϕ access to the input vector x , excluding entries for which $x(i)$ is small. By using a Chebyshev inequality, and showing that the bias and variance of the considered point estimation are sufficiently bounded, one can ensure that the average of a sufficient number of samples is $\epsilon \|y\|_1$ close to the true inner product. This is in line with Refs. [25–27]; but one now needs to adapt the filtering technique and the bias and variance bounding to the presence of error in the knowledge of the entries of x .

We now discuss the proof in detail.

Denote the “oversampling vector” of x as $\tilde{x}$ (cf. Def. 21). Let $\hat{n}^2$ be an estimate of $\|\tilde{x}\|^2$ to relative error $1/4$ with probability $8/9$, as made possible by ASQ_ϕ to x and Lemma 28. This is attained, accordingly, in time $\mathcal{O}\{\mathbf{n}(\|\tilde{x}\|^2)[1 + \log(1/\|\tilde{x}\|^2)]\}$, and we henceforth assume the estimate is correct. Define $\gamma := \epsilon^2/(135 \hat{n}^2)$. Taking also $\mathcal{O}(\gamma^{-2} \log d)$ samples from $p_{\tilde{x}}$, and from Theorem 14, one knows also every $p_{\tilde{x}}(i)$ to error $\gamma/8$, with overall probability $8/9$. [Looking up each estimate of $p_{\tilde{x}}(i)$, say, $\hat{p}_i$, takes time $\mathcal{O}(\log d)$, per the Theorem.]

We also note the following preliminary observations:

- (a) $\|\tilde{x}\|/\hat{n} \leq 2$;
- (b) $|x(i)| \leq |\tilde{x}(i)| = \|\tilde{x}\| \sqrt{p_{\tilde{x}}(i)}$, for all $i \in [d]$; and
- (c) $|a/(\hat{a} + \epsilon) - 1| \leq 2\epsilon/a$, for all positive a and $\hat{a}$ satisfying $|a - \hat{a}| \leq \epsilon$, as a direct consequence of this fact.

Now, consider the following point estimator, if $i \sim p_{\tilde{x}}$, and $\hat{x}_i$ is an $\epsilon/4$ estimate of $x(i)$:

$$\widehat{x^\dagger y}(i) = \begin{cases} 0 & \text{if } \hat{p}_i \leq \frac{3}{2}\gamma \\ \frac{1}{\hat{p}_i + \gamma/8} \hat{x}_i^* y(i) & \text{otherwise.} \end{cases} \quad (23)$$

Let us say that the first case corresponds to the “rejection” of i , and the second corresponds to the “acceptance” of i . From the error bounds on $\hat{p}_i$, for all rejected i , one has that $p_{\tilde{x}}(i) \leq 2\gamma$, and for

all accepted i , $p_{\tilde{x}}(i) \geq \gamma$. This lets us bound the bias of the estimator as

$$|\mathbb{E}_{i \sim p_{\tilde{x}}} \widehat{x^\dagger y}(i) - x^\dagger y| \leq \left| \sum_{i: p_{\tilde{x}}(i) \leq 2\gamma} x(i)^* y(i) \right| + \left| \sum_{i: p_{\tilde{x}}(i) \geq \gamma} \frac{p_{\tilde{x}}(i)}{\hat{p}_i} \hat{x}_i^* y(i) - x(i)^* y(i) \right| \quad (24)$$

$$\leq \sqrt{2\gamma} \|\tilde{x}\| \|y\|_1 + \sum_{i: p_{\tilde{x}}(i) \geq \gamma} |y(i)| \left\{ |x(i)| \left| \frac{p_{\tilde{x}}(i)}{\hat{p}_i + \gamma/8} - 1 \right| + |\hat{x}_i - x(i)| \left[1 + \left| \frac{p_{\tilde{x}}(i)}{\hat{p}_i + \gamma/8} - 1 \right| \right] \right\} \quad (25)$$

$$\leq \frac{1}{4} \epsilon \|y\|_1 + \sum_{i: p_{\tilde{x}}(i) \geq \gamma} |y(i)| \left[\|\tilde{x}\| \frac{\gamma}{4\sqrt{p_{\tilde{x}}(i)}} + \frac{5}{16} \epsilon \right] \quad (26)$$

$$\leq \frac{2}{3} \epsilon \|y\|_1. \quad (27)$$

The first inequality follows from the definition of the estimator, the triangle inequality, and possibly double counting, the second inequality from observations (b) and (c), and the third inequality from the definition of γ and all of the preliminary observations.

The variance is bounded as

$$\text{VAR}_{i \sim p_{\tilde{x}}} \widehat{x^\dagger y}(i) \leq \mathbb{E}_{i \sim p_{\tilde{x}}} |\widehat{x^\dagger y}(i)|^2 \quad (28)$$

$$\leq \sum_{i: \hat{p}_i \geq \frac{3}{2}\gamma} \frac{p_{\tilde{x}}(i)}{\hat{p}_i + \gamma/8} \frac{1}{\hat{p}_i + \gamma/8} |\hat{x}_i|^2 |y(i)|^2 \quad (29)$$

$$\leq \sum_{i: \hat{p}_i \geq \frac{3}{2}\gamma} \frac{1}{p_{\tilde{x}}(i)} [|x(i)| + \frac{\epsilon}{4}]^2 |y(i)|^2 \quad (30)$$

$$= \sum_{i: \hat{p}_i \geq \frac{3}{2}\gamma} \frac{|x(i)|^2}{p_{\tilde{x}}(i)} |y(i)|^2 + \frac{(\epsilon/4)^2}{p_{\tilde{x}}(i)} |y(i)|^2 + \frac{\epsilon/2}{\sqrt{p_{\tilde{x}}(i)}} \frac{|x(i)|}{\sqrt{p_{\tilde{x}}(i)}} |y(i)|^2 \quad (31)$$

$$\leq \|\tilde{x}\|^2 \|y\|^2 + \frac{3}{2} 4\hat{n} \|\tilde{x}\| \|y\|^2 + 2\hat{n} \|\tilde{x}\| \|y\|^2 \quad (32)$$

$$\leq 10 \|\tilde{x}\|^2 \|y\|^2. \quad (33)$$

The first inequality follows from the definition of variance for complex random variables, the second inequality from a triangle inequality (and the definition of the estimator), and the third inequality from the observation that $\hat{p}_i \geq p_{\tilde{x}}(i) - \gamma/2$. The second to last inequality follows from the definition of $p_{\tilde{x}}(i)$ as $|\tilde{x}(i)|^2 / \|\tilde{x}\|^2$ and the imposed condition that $|\tilde{x}(i)| \geq |x(i)|$, for all i , as well as the summing index condition and the definition of γ .

With these bounds, it follows from the Chebyshev inequality that the average of $O(\frac{\|\tilde{x}\|^2 \|y\|^2}{\epsilon^2}) = O(\hat{n}^2 \epsilon^{-2})$ valid samples of this estimator suffice to estimate a value that is at most $\frac{1}{3} \epsilon \|y\|$ away from the estimator's mean — thus, at most $\epsilon \|y\|_1$ away from $x^\dagger y$ — with probability at least $1/18$.

From a Hoeffding inequality, taking the median of $O[\log(\hat{n}^2 \epsilon^{-2})]$ independent estimates ensures that every $\hat{x}_i$ is a correct estimate of $x(i)$ to error $\epsilon/4$, with overall error probability $1/18$. From a union bound, it follows that the whole procedure estimates $x^\dagger y$ to error $\epsilon \|y\|_1$ with probability $2/3$.

To find the stated runtime, we first note that with very high probability, sampling k valid samples from ϕ -oversampling access to x takes time $O(k\mathbf{s})$. This follows from a Hoeffding inequality, the one-sidedness of the error, and the fact that, by definition, the probability of a valid sample is at least $2/3$. Then it follows immediately that the total runtime of the procedure outlined above and in algorithm 2 is

$$O[\mathbf{n}(\|\tilde{x}\|^2/4)(1+\log \frac{1}{\|\tilde{x}\|^2})] + O(\mathbf{s}\gamma^{-2} \log d) + O(\mathbf{s}\hat{n}^2 \epsilon^{-2}) + O\{\lceil \mathbf{q}(\epsilon/4) \log(\hat{n}^2 \epsilon^{-2}) + \log d \rceil \hat{n}^2 \epsilon^{-2}\}. \quad (34)$$

Using the definition of γ , the fact that $\hat{n}^2 \leq \frac{5}{4} \|\tilde{x}\|^2$, that $\|\tilde{x}\|^2 \leq \phi \|x\|^2$ by definition, and the assumption that constant factors can be “moved out” of the cost runtime functions, the stated runtime follows. $\square$

Remark 31. A restricted form of this lemma extends to the case where sampling is not exact in the distribution. In particular: if $x \in \mathbb{C}^d$ such that $\|x\| \leq 1$, and $x' \in \mathbb{C}^d$ such that

$$\sum_{i \in [d]} |p_{\tilde{x}}(i) - p_{x'}(i)| \leq \frac{C\epsilon}{\sqrt{\phi}} \quad (35)$$

for $p_{\tilde{x}}$ a distributed ϕ -oversampling p_x , and for sufficient choice of constant C , then ϕ -Approximate Sample access to x' (Def. 19) and Approximate Query access to x (Def. 18) suffices to estimate $x^\dagger y$ to additive error $\epsilon\|y\|_1$, with runtime

$$\tilde{O}[\phi\epsilon^{-2} \log d + \phi\epsilon^{-2} \mathbf{q}(\epsilon) + \phi^2\epsilon^{-4} \log d \mathbf{s}] \quad (36)$$

with very high probability and under the same assumptions as those in the lemma above. This requires that ϕ or an upper bound thereof is known (whereas Lemma 30 does not). The proof proceeds like that of Lemma 30 with technical complications, and so is deferred to Appendix A.

While the dependence on the 1-norm of y provides for a weaker result than that of, for example, Ref. [25], where it is possible to attain absolute error in time completely independent of the dimension of the underlying vectors, we will nonetheless find that such a dependence is sufficient to prove, for example, Theorem 35. On the other hand, this 1-norm has a simple geometric interpretation: the more y is “aligned” with a single measurement basis, the easier it is to estimate the inner product. At the same time, we note that there is no dependence on the 1-norm of x , such that only the 1-norm of the classically queried vector plays a role in increasing the runtime. This means that even if x has maximal 1-norm (e.g., $|x(i)| = d^{-1/2}\|x\|$), the procedure is still efficient provided that y has a small 1-norm (for example, that y is known to be sufficiently sparse, even if which elements are nonzero is unknown).

Now, we analyze the symmetric case, in which one has *approximate sample and query* access to two vectors. In light of the remarks of section 3.1, this case can be taken to translate a restricted view of the case where one knows how to prepare and measure two states, and wishes to estimate the inner product of the underlying state vectors. We find that if, per the above Theorem, 1-norm reflects “peakedness” in the sampling, only one of the provided vectors needs to be sufficiently “peaked” in the symmetric case.

Lemma 32. (Symmetric inner product estimation) *Let $x, y \in \mathbb{C}^d$, and $\epsilon \in (0, 1]$. Suppose ASQ_ϕ access to x and y is given. One may estimate $x^\dagger y$ to absolute error $\epsilon[1 + \min(\|x\|_1/\|x\|, \|y\|_1/\|y\|)]$ with probability $2/3$ by algorithm 3. Assuming, for simplicity, that $\|x\|, \|y\| \leq 1$, the runtime of the algorithm is*

$$\tilde{O}[\phi^2\epsilon^{-2} \log d + \phi^2\epsilon^{-2} \mathbf{q}(\epsilon^2\|x\|\|y\|) + \epsilon^{-4} \log d \mathbf{s} + \mathbf{n}(\|x\|^2) + \mathbf{n}(\|y\|^2)]$$

with very high probability, and under the assumption that $\mathbf{q}[\Omega(\epsilon)] = O[\mathbf{q}(\epsilon)]$, and likewise for $\mathbf{n}$.

Proof. The proof proceeds similarly to that of Lemma 30.

Denote the oversampling distribution of $ASQ_\phi(x)$ by $p_{\tilde{x}}$ and that of $ASQ_\phi(y)$ by $p_{\tilde{y}}$, for “oversampling vectors” $\tilde{x}$ and $\tilde{y}$ for x and y , respectively. Let p be the distribution obtained by uniformly randomly outputting from either $p_{\tilde{x}}$ or $p_{\tilde{y}}$, from which we have the probability mass function $p(i) = [p_{\tilde{x}}(i) + p_{\tilde{y}}(i)]/2$, for all $i \in [d]$. Note that p is a 2-oversampling distribution for both $p_{\tilde{x}}$ and $p_{\tilde{y}}$. Thus also, for all $i \in [d]$, $p_{\tilde{x}}(i) = |\tilde{x}(i)|^2/\|\tilde{x}\|^2 \leq 2p(i)$, implying $|x(i)| \leq |\tilde{x}(i)| \leq \sqrt{2p(i)}\|\tilde{x}\|$. The analogous statements for y also hold.

Let $\hat{n}_x^2$ and $\hat{n}_y^2$ be relative error $1/2$ estimates of $\|\tilde{x}\|^2$ and $\|\tilde{y}\|^2$, respectively, with joint probability $8/9$. These are obtained by algorithm 1 and the ASQ_ϕ access, and so with high probability in time $O\{\mathbf{n}(\|\tilde{x}\|^2)[1 + \log(1/\|\tilde{x}\|^2)]\}$ for $\hat{n}_x^2$, and $O\{\mathbf{n}(\|\tilde{y}\|^2)[1 + \log(1/\|\tilde{y}\|^2)]\}$ for $\hat{n}_y^2$. Note that, with these choices, $1/2\|\tilde{x}\|^2 \leq \hat{n}_x^2 \leq 3/2\|\tilde{x}\|^2$, from which follows $\hat{n}_x/\|\tilde{x}\| < \sqrt{2}$, and likewise for $\hat{n}_y, y$.

Define $\gamma := \Theta[\min(1, \hat{n}_x^2) \min(1, \hat{n}_y^2)\epsilon^2]$. From Theorem 14, $O(\gamma^{-2} \log d)$ samples of p are sufficient to know, for every $i = 1, \dots, d$, $p(i)$ to error $\gamma/2$ and with overall probability $8/9$. [Looking up each estimate of $p(i)$, say $\hat{p}_i$, takes time $O(\log d)$, per the Theorem.]

Algorithm 3 Calculating an inner product estimate via ASQ_ϕ access to both vectors (x and y), with success probability at least $2/3$ and absolute error at most $\epsilon[1 + \min(\frac{\|x\|_1}{\|x\|}, \frac{\|y\|_1}{\|y\|})]$.

Require: ASQ_ϕ access to $x, y \in \mathbb{C}^d$ 22: $\hat{n}_x^2 \leftarrow [\text{ALGORITHM 1}](\|x\|^2, \rho = \frac{1}{2}, \delta = \frac{1}{18})$
Require: Precision $\epsilon \in (0, 1]$ 23: $\hat{n}_y^2 \leftarrow [\text{ALGORITHM 1}](\|y\|^2, \rho = \frac{1}{2}, \delta = \frac{1}{18})$

1: **function** $\text{IMPROVE}(x, \epsilon, \delta)$ 24: $\gamma \leftarrow \min(1, \hat{n}_x^{-2}) \min(1, \hat{n}_y^{-2}) \epsilon^2 / 100$
2: $r \leftarrow []$ 25: $\epsilon_x \leftarrow \epsilon \sqrt{\gamma} \min(1, \hat{n}_y^{-1}) / 100$
3: $c \leftarrow []$ 26: $\epsilon_y \leftarrow \epsilon \sqrt{\gamma} \min(1, \hat{n}_x^{-1}) / 100$
4: $\epsilon \leftarrow \epsilon / \sqrt{2}$
5: **do** $\lceil 18 \log(1/\delta) \rceil$ **times:** 27: $l \leftarrow [\text{EMPTY HISTOGRAM}]$
6: Sorted insert $\text{Re}[\hat{x}(\epsilon)]$ into r 28: **do** $32\gamma^{-2} \log(18d)$ **times:**
7: Sorted insert $\text{Im}[\hat{x}(\epsilon)]$ into c 29: $i \leftarrow \text{SAMPLE}()$
8: **end** 30: Insert i into h .
9: **return** $\text{MEDIAN}(r) + \mathbf{i} \text{MEDIAN}(c)$ 31: **end**

10: **end function**

11: **function** SAMPLE 32: $m \leftarrow 864(1 + 2\hat{n}_x^2)(1 + 2\hat{n}_y^2) \epsilon^{-2}$
12: $z \leftarrow \text{FAIR COIN}$ 33: $l \leftarrow []$
13: **if** z is HEADS **then** 34: **do** m **times:**
14: **repeat** $i \leftarrow \text{OVERSAMPLE}(x)$ 35: $i \leftarrow \text{SAMPLE}()$
15: **until** sample is valid. 36: $\hat{p}_i \leftarrow \text{FREQUENCY}(h; i)$
16: **else** 37: **if** $\hat{p}_i \leq \frac{3}{2}\gamma$ **then**
17: **repeat** $i \leftarrow \text{OVERSAMPLE}(y)$ 38: $v \leftarrow 0$
18: **until** sample is valid. 39: **else**
19: **end if** 40: $\hat{x}_i \leftarrow \text{IMPROVE}(x(i), \epsilon_x, 1/(18m))$
20: **return** i 41: $\hat{y}_i \leftarrow \text{IMPROVE}(y(i), \epsilon_y, 1/(18m))$
21: **end function** 42: $v \leftarrow (\hat{p}_i + \gamma/2)^{-1} \hat{x}_i^* \hat{y}_i$
43: **end if**
44: Sorted insert v into l
45: **end**
46: **return** $\text{MEDIAN}(l)$

Furthermore, for $i \sim p$, consider $\hat{x}_i$ and $\hat{y}_i$ to be estimates of, respectively, $x(i)$ and $y(i)$, to absolute error, respectively, $\Theta[\sqrt{\gamma}\epsilon/\max(1, \hat{n}_y)]$ and $\Theta[\sqrt{\gamma}\epsilon/\max(1, \hat{n}_x)]$. With these, we define the point estimator

$$\widehat{x^\dagger y}(i) = \begin{cases} 0 & \text{if } \hat{p}_i \leq \frac{3}{2}\gamma \\ \frac{1}{\hat{p}_i + \gamma/2} \hat{x}_i^* \hat{y}_i & \text{otherwise.} \end{cases} \quad (37)$$

This estimator has bias (with respect to $x^\dagger y$) bounded as (from a triangle inequality and double counting)

$$|\mathbb{E}_{i \sim p} \widehat{x^\dagger y} - x^\dagger y| \leq \sum_{i: p(i) \leq 2\gamma} |x(i)||y(i)| + \sum_{i: p(i) \geq \gamma} \left| \frac{p(i)}{\hat{p}_i + \gamma/2} \hat{x}_i^* \hat{y}_i - x(i)^* y(i) \right|. \quad (38)$$

The first term is further bounded as

$$\begin{aligned} \sum_{i: p(i) \leq 2\gamma} |x(i)||y(i)| &\leq 2\sqrt{\gamma} \min(\|\tilde{x}\| \|y(i)\|, |x(i)| \|\tilde{y}\|) = \frac{1}{5} \frac{\|\tilde{x}\| \|\tilde{y}\|}{\hat{n}_x \hat{n}_y} \epsilon \min\left(\frac{\|x\|_1}{\|x\|}, \frac{\|y\|_1}{\|y\|}\right) \\ &\leq \frac{2}{5} \epsilon \min\left(\frac{\|x\|_1}{\|x\|}, \frac{\|y\|_1}{\|y\|}\right) \end{aligned} \quad (39)$$

where we have used $|x(i)| \leq \sqrt{2p(i)}$ and the condition on the summing index. The second term is further bounded as

$$\begin{aligned} &\sum_{i: p(i) \geq \gamma} \left| \frac{p(i)}{\hat{p}_i + \gamma/2} \hat{x}_i^* \hat{y}_i - x(i)^* y(i) \right| \\ &\leq \sum_{i: p(i) \geq \gamma} \left| \frac{p(i)}{\hat{p}_i + \gamma/2} - 1 \right| |x(i)||y(i)| + |\hat{x}_i^* \hat{y}_i - x(i)^* y(i)| + \left| \frac{p(i)}{\hat{p}_i + \gamma/2} - 1 \right| |\hat{x}_i^* \hat{y}_i - x(i)^* y(i)|. \end{aligned} \quad (40)$$

which can then be further bounded term-by-term:

$$\sum_{i: p(i) \geq \gamma} \left| \frac{p(i)}{\hat{p}_i + \gamma/2} \|x(i)||y(i)| \right| \leq \sum_{i: p(i) \geq \gamma} \frac{\gamma}{p(i)} |x(i)||y(i)| \quad (41)$$

$$\leq \frac{\sqrt{2}}{10} \frac{\|\tilde{x}\| \|\tilde{y}\|}{\hat{n}_x \hat{n}_y} \epsilon \min\left(\frac{\|x\|_1}{\|x\|}, \frac{\|y\|_1}{\|y\|}\right) \leq \frac{\sqrt{2}}{5} \epsilon \min\left(\frac{\|x\|_1}{\|x\|}, \frac{\|y\|_1}{\|y\|}\right); \quad (42)$$

$$\sum_{i: p(i) \geq \gamma} |\hat{x}_i^* \hat{y}_i - x(i)^* y(i)| \leq \sum_{i: p(i) \geq \gamma} |\hat{x}_i - x(i)| |\hat{y}_i - y(i)| + |x(i)| |\hat{y}_i - y(i)| + |\hat{x}_i - x(i)| |y(i)| \quad (43)$$

$$\leq \sum_{i: p(i) \geq \gamma} \left(\frac{1}{100} \epsilon^2 \gamma + \frac{\sqrt{2}}{100} \sqrt{p(i)} \gamma \left[\frac{\|\tilde{x}\|}{\max(1, \hat{n}_x)} + \frac{\|\tilde{y}\|}{\max(1, \hat{n}_y)} \right] \right) \leq \frac{1}{4} \epsilon; \quad (44)$$

$$\begin{aligned} \sum_{i: p(i) \geq \gamma} \left| \frac{p(i)}{\hat{p}_i + \gamma/2} - 1 \right| |\hat{x}_i^* \hat{y}_i - x(i)^* y(i)| &\leq \sum_{i: p(i) \geq \gamma} \frac{\gamma}{p(i)} |\hat{x}_i^* \hat{y}_i - x(i)^* y(i)| \\ &\leq \sum_{i: p(i) \geq \gamma} |\hat{x}_i^* \hat{y}_i - x(i)^* y(i)| \leq \frac{1}{4} \epsilon. \end{aligned} \quad (45)$$

Above, we've used the aforementioned fact that $|\hat{a} - a| \leq \epsilon$ implies $|a/\hat{a} - 1| \leq 2\epsilon/a$, the relation between $|x(i)|$ and $p(i)$ (and between $|y(i)|$ and $p(i)$), the condition on the summing index (to upper bound γ by $p(i)$ where necessary), and the various error bounds.

Thus we conclude that the bias of the estimator is bounded as $O\{\epsilon[1 + \min(\frac{\|x\|_1}{\|x\|}, \frac{\|y\|_1}{\|y\|})]\}$.

Likewise, we bound the variance of the estimator as

$$\text{VAR}_{i \sim \mathbf{p}} \widehat{x^\dagger y} \leq \mathbb{E}_{i \sim \mathbf{p}} [\widehat{x^\dagger y}]^2 \leq \sum_{i: \mathbf{p}(i) \geq \gamma} \frac{1}{\mathbf{p}(i)} [|x(i)| + \frac{1}{10} \gamma \max(1, \hat{n}_x)]^2 [|y(i)| + \frac{1}{10} \gamma \max(1, \hat{n}_y)]^2 \quad (46)$$

$$\leq \sum_{i: \mathbf{p}(i) \geq \gamma} \frac{4}{\mathbf{p}(i)} [|x(i)|^2 + \frac{1}{100} \gamma^2 \max(1, \hat{n}_x^2)] [|y(i)|^2 + \frac{1}{100} \gamma^2 \max(1, \hat{n}_y^2)] \quad (47)$$

$$\leq 10(1 + \|\tilde{x}\|^2)(1 + \|\tilde{y}\|^2). \quad (48)$$

The first inequality uses the observations that the error bounds for $\hat{x}_i$ may be written as $\gamma \max(1, \hat{n}_x)/10$ (similarly for $\hat{y}_i$), and that the error bounds on $\hat{p}_i$ imply that $(\hat{p}_i + \gamma/2)^{-1} \leq \mathbf{p}(i)^{-1}$. The second inequality uses a geometric mean–quadratic mean inequality. The third inequality uses a Cauchy-Schwarz inequality, the summing index’s condition, and the inequality $\|x\|^2 \leq \|\tilde{x}\|^2$ (likewise for y).

Then it follows from a Chebyshev bound that $\text{O}[\epsilon^{-2}(1 + \|\tilde{x}\|^2)(1 + \|\tilde{y}\|^2)] = \text{O}[\epsilon^{-2}(1 + \hat{n}_x^2)(1 + \hat{n}_y^2)]$ valid samples have their average at most $\epsilon/2$ away from the estimator’s mean — thus are at most $\epsilon[1 + \max(\frac{\|x\|_1}{\|x\|}, \frac{\|y\|_1}{\|y\|})]$ away from $x^\dagger y$ — with probability $1/18$. From a Hoeffding bound, the error probability of each sample of $\hat{x}_i$ and $\hat{y}_i$ can be reduced to ensure that every sample is a correct estimate of $x(i)$, $y(i)$, respectively, with overall probability $1/18$. From a union bound, it follows that the final estimation is correct with probability $2/3$.

From the outlined procedure, and explicitly from algorithm 3, one finds that the runtime is, with high probability,

$$\begin{aligned} & \text{O}[\mathbf{n}(\|\tilde{x}\|^2/2)(1 + \log \frac{1}{\|\tilde{x}\|^2})] + \text{O}[\mathbf{n}(\|\tilde{y}\|^2/2)(1 + \log \frac{1}{\|\tilde{y}\|^2})] \\ & \quad + \text{O}(\mathbf{s} \gamma^{-2} \log d) + \text{O}[\epsilon^{-2}(1 + \|\tilde{x}\|^2)(1 + \|\tilde{y}\|^2) \log d] \\ & \quad + \text{O}\{\epsilon^{-2}(1 + \|\tilde{x}\|^2)(1 + \|\tilde{y}\|^2) \log[\epsilon^{-2}(1 + \|\tilde{x}\|^2)(1 + \|\tilde{y}\|^2)] \mathbf{q}[\epsilon^2 \min(1, \|\tilde{x}\|) \min(1, \|\tilde{y}\|)]\}. \end{aligned} \quad (49)$$

The stated complexity follows. $\square$

Remark 33. A restricted form of this lemma extends to the case where sampling is not exact in either distribution. In particular: let $\phi \geq 1$, $\epsilon \in (0, 1]$, and $x, y \in \mathbb{C}^d$, such that $\|x\|, \|y\| \leq 1$. Suppose $x', y' \in \mathbb{C}^d$ satisfy

$$\sum_{i \in [d]} |\mathbf{p}_{\tilde{x}}(i) - \mathbf{p}_{x'}(i)| \leq \frac{C\epsilon}{\phi}, \quad \sum_{i \in [d]} |\mathbf{p}_{\tilde{y}}(i) - \mathbf{p}_{y'}(i)| \leq \frac{C\epsilon}{\phi} \quad (50)$$

where $\mathbf{p}_{\tilde{x}}$ is a distribution ϕ -oversampling $\mathbf{p}_x$, $\mathbf{p}_{\tilde{y}}$ is a distribution ϕ -oversampling $\mathbf{p}_y$, and $C \leq 1$ is a suitably chosen constant. Then, Approximate Sample access to x' and y' and Approximate Query access to x and y suffices to estimate $x^\dagger y$ to additive error $\epsilon[1 + \min(\|x\|_1, \|y\|_1)]$ and success probability $2/3$, halting with high probability in time

$$\text{O}[\phi^2 \epsilon^{-2} \log d + \phi^2 \epsilon^{-2} \mathbf{q}(\|x\| \|y\| \epsilon^2) + \phi^2 \epsilon^{-4} \log d \mathbf{s}].$$

This requires that ϕ or an upper bound to ϕ is known (whereas the above lemma does not). The proof proceeds similarly to that of Lemma 32, and so is deferred to Appendix B.

The lemmas above have assumed the generality of complex vectors and oversampling access, such that if κ is an upper bound to $\min(\|x\|_1, \|y\|_1)$ is known, and one wishes to attain absolute error on the estimate of $x^\dagger y$, the runtime scales with at least κ^4 . If one is willing to restrict the considered setting, and if querying for entries of x and y is not much more expensive than sampling, then this dependence on κ can be improved:

Lemma 34. Let $x, y \in \mathbb{R}^d$ such that $\|x\| = \|y\| = 1$. Choose $\epsilon \in (0, 1]$ and $\Delta \in (0, 1]$. Assume Approximate Query access to x and y and Approximate Sample access to x' and y' , where $\|x - x'\| \leq \Delta/2$ and $\|y - y'\| \leq \Delta/2$. Let κ be a known upper bound to $\min(\|x\|_1, \|y\|_1)$. Then, $x^T y$ may be estimated to an absolute error of $\epsilon + \Delta$ with probability $2/3$ in time

$$\tilde{\text{O}}[\epsilon^{-2} \mathbf{s} + \mathbf{q}(\epsilon^2 \kappa^{-1})].$$

Proof. Observe that, for fixed $i \in [d]$, at cost $\tilde{O}[\mathbf{q}(\epsilon)]$, both $x(i)$ and $y(i)$ can be known, with high probability, to additive error ϵ . Define $\mathbf{p}$ as that resulting from uniformly randomly sampling from either $\mathbf{p}_x$ or $\mathbf{p}_y$. If $\hat{x}_i, \hat{y}_i$ are the estimates for, respectively, $x(i)$ and $y(i)$, then by a geometric argument one has that

$$\sqrt{\hat{p}_i} := \sqrt{\frac{\hat{x}_i^2 + \hat{y}_i^2}{2}} \implies |\sqrt{\hat{p}_i} - \sqrt{\mathbf{p}(i)}| \leq \epsilon. \quad (51)$$

Define, similarly, $\mathbf{q}$ as the distribution corresponding to uniformly randomly sampling from either $\mathbf{p}_{x'}$ or $\mathbf{p}_{y'}$. From Lemma 24, one has that

$$\sum_{i \in [d]} |\mathbf{q}(i) - \mathbf{p}(i)| \leq 2\Delta. \quad (52)$$

Let C_1 and C_2 be suitably chosen constants, and define $\gamma := C_1\epsilon/\kappa$ and $\epsilon_Q := C_2\epsilon^2/\kappa$. By eq. (51), take $\sqrt{\hat{p}_i}$ and $\hat{p}_i$ be estimates of, respectively, $\sqrt{\mathbf{p}(i)}$ and $\mathbf{p}(i)$ to error $\gamma/2$. Furthermore let $\hat{x}_i$ and $\hat{y}_i$ be estimates of, respectively, $x(i)$ and $y(i)$ to error ϵ_Q . Obtaining these estimates (with high probability) takes time $O[\mathbf{q}(\gamma) + \mathbf{q}(\epsilon_Q)]$, by the Approximate Query access.

Now, take a point estimator like that of lemmas 30 and 32, but “rejecting” samples based on the value of $\sqrt{\hat{p}_i}$:

$$\widehat{x^T y}(i) \leq \begin{cases} 0 & \text{if } \sqrt{\hat{p}_i} \leq \frac{3}{2}\gamma \\ \frac{\hat{x}_i \hat{y}_i}{\hat{p}_i} & \text{otherwise.} \end{cases} \quad (53)$$

We start by bounding the bias of this estimator under sampling of $\mathbf{p}$, rather than under sampling of $\mathbf{q}$:

$$|\mathbb{E}_{i \sim \mathbf{p}} \widehat{x^T y}(i) - x^T y| \leq \sum_{i: \mathbf{p}(i)^{1/2} \leq 2\gamma} |x(i)||y(i)| + \sum_{i: \mathbf{p}(i)^{1/2} \geq \gamma} |\mathbf{p}(i) \frac{\hat{x}_i \hat{y}_i}{\frac{\hat{x}_i^2 + \hat{y}_i^2}{2}} - x(i)y(i)|. \quad (54)$$

Since, from its definition, $2\mathbf{p}(i) \geq x(i)^2$ and $2\mathbf{p}(i) \geq y(i)^2$, and from the condition in the summation, it is straightforward to see that

$$\sum_{i: \mathbf{p}(i)^{1/2} \leq 2\gamma} |x(i)||y(i)| \leq 2\sqrt{2}\gamma \min(\|x\|_1, \|y\|_1) = 2\sqrt{2}C_1\epsilon. \quad (55)$$

To bound the right-hand side term of eq. (54), we use the fact that $\hat{p}_i$ is always “consistent” with $\hat{x}_i$ and $\hat{y}_i$, and that the sampling is “exact” (in the sense that we are not oversampling). Namely, we have that

$$\sum_{i: \mathbf{p}(i)^{1/2} \geq \gamma} |\mathbf{p}(i) \frac{\hat{x}_i \hat{y}_i}{\frac{\hat{x}_i^2 + \hat{y}_i^2}{2}} - x(i)y(i)| = \sum_{i: \mathbf{p}(i)^{1/2} \geq \gamma} 2\mathbf{p}(i) |f(\hat{x}_i, \hat{y}_i) - f(x(i), y(i))|, \quad (56)$$

where $f(a, b) := \frac{ab}{a^2 + b^2}$. On the other hand, for $\sqrt{a^2 + b^2} \geq \sqrt{2}\gamma$, f is Lipchitz continuous, with constant $(\sqrt{2}\gamma)^{-1}$, as may be seen:

$$\left| \frac{\partial f}{\partial a} \right|(a, b) = |a| \frac{|a^2 - b^2|}{(a^2 + b^2)^2} \leq \frac{|a|}{a^2 + b^2}, \quad \left| \frac{\partial f}{\partial b} \right|(a, b) = |b| \frac{|a^2 - b^2|}{(a^2 + b^2)^2} \leq \frac{|b|}{a^2 + b^2}; \quad (57)$$

$$\implies \|\nabla f\| \leq \frac{\sqrt{a^2 + b^2}}{a^2 + b^2} \leq \frac{1}{\sqrt{2}\gamma}. \quad (58)$$

Moreover, $\|(\hat{x}_i, \hat{y}_i) - (x(i), y(i))\| \leq \sqrt{2}\epsilon_Q$. Thus it follows that

$$\sum_{i: \mathbf{p}(i)^{1/2} \geq \gamma} |\mathbf{p}(i) \frac{\hat{x}_i \hat{y}_i}{\frac{\hat{x}_i^2 + \hat{y}_i^2}{2}} - x(i)y(i)| \leq \sum_{i: \mathbf{p}(i)^{1/2} \geq \gamma} 2\mathbf{p}(i) \frac{1}{\sqrt{2}\gamma} \sqrt{2}\epsilon_Q \leq 2 \frac{C_2}{C_1} \epsilon. \quad (59)$$

The variance of the estimator under $\mathbf{p}$ is also easily bounded by a geometric mean–arithmetic mean inequality:

$$\text{VAR}_{i \sim \mathbf{p}} \widehat{x^T y}(i) \leq \mathbb{E}_{i \sim \mathbf{p}} \widehat{x^T y}(i)^2 = \sum_{i: \mathbf{p}(i)^{1/2} \geq \gamma} \mathbf{p}(i) \left(\frac{\hat{x}_i \hat{y}_i}{\frac{\hat{x}_i^2 + \hat{y}_i^2}{2}} \right)^2 \leq \sum_{i: \mathbf{p}(i)^{1/2} \geq \gamma} \mathbf{p}(i) \leq 1. \quad (60)$$

Now we show that because $p(i)$ is not estimated by sampling, sampling from q does not affect the analysis substantially: the variance is still bounded by 1 by exactly the approach of eq. (60), and the bias is now bounded as

$$|\mathbb{E}_{i \sim q} \widehat{x^T y}(i) - x^T y| \leq |\mathbb{E}_{i \sim q} \widehat{x^T y}(i) - \mathbb{E}_{i \sim p} \widehat{x^T y}(i)| + (2\sqrt{2}C_1 + 2\frac{C_2}{C_1})\epsilon \leq (4\sqrt{2}C_1 + 2\frac{C_2}{C_1})\epsilon + \Delta, \quad (61)$$

since, using again a geometric mean–arithmetic mean inequality,

$$|\mathbb{E}_{i \sim q} \widehat{x^T y}(i) - \mathbb{E}_{i \sim p} \widehat{x^T y}(i)| \leq \sum_{i: p(i)^{1/2} \leq 2\gamma} |x(i)||y(i)| + \sum_{i: p(i)^{1/2} \geq \gamma} |q(i) - p(i)| \frac{\hat{x}_i \hat{y}_i}{\hat{p}_i} \quad (62)$$

$$\leq 2\sqrt{2}\gamma \min(\|x\|_1, \|y\|_1) + \frac{1}{2}2\Delta = 2\sqrt{2}C_1\epsilon + \Delta. \quad (63)$$

The stated runtime follows from an analysis analogous to that of the proof of Lemma 32. $\square$

4 Application to distributed inner product estimation

Finally, we show that the access model encoded by *approximate sample and query* and the computational results derived in Section 3.3 suffice to derive a result of independent interest. (We refer the reader to Section 2.7 and Ref. [41] for a review of the concepts considered in the following statement.)

Theorem 35. (Distributed inner product estimation under entanglement and stabilizer norm constraints) *Let $|\psi\rangle$ and $|\phi\rangle$ be pure states of n qubits, such that both have no more than χ bipartite entanglement across any bipartition of their qubits, and have stabilizer norm (at most) D . If preparing $|\psi\rangle$ or $|\phi\rangle$ takes time $O(T)$, and assuming Bell measurements on $|\psi\rangle$ or $|\phi\rangle$ take time $O(T)$, then there is an (explicit) procedure, which does not require simultaneously coherent copies of $|\psi\rangle$ and $|\phi\rangle$, and that estimates $|\langle\psi|\phi\rangle|^2$ to absolute error ϵ with high probability in time*

$$\tilde{O}[n^5 e^{4\chi} D^2 \epsilon^{-6} + T \epsilon^{-4} n^3 e^{4\chi} D^2 + T \epsilon^{-2} D^2]. \quad (64)$$

To prove Theorem 35, we begin by adapting [41, Lemmas 2 and 3] to the following:

Lemma 36. *Let ρ be a d -dimensional pure quantum state, and $\tau \in [0, 1]$. Denote $\alpha_\rho(i) := \text{Tr}(\rho P_i)$, and let $p_{\pi_\rho}(i) = \alpha_\rho(i)^2/d$ be the probability mass function associated to Pauli sampling from ρ . Let $F_\rho(\tau) := \mathbb{P}_{i \sim p_{\pi_\rho}}[\alpha_\rho(i)^2 < \tau]$, and $M_{1/2}(\rho)$ be the stabilizer 1/2-Rényi entropy, as defined in section 2.7 (using natural base logarithms). Then,*

$$F_\rho(\tau) \leq \sqrt{\tau} e^{\frac{1}{2} M_{1/2}(\rho)}. \quad (65)$$

Proof. We have that, by definition,

$$F_\rho(\tau) = \sum_{i: \alpha_\rho(i)^2 < \tau} \frac{\alpha_\rho(i)^2}{d}. \quad (66)$$

On the other hand,

$$M_{1/2}(\rho) = 2 \log \left[\sum_{i \in [d^2]} \frac{\alpha_\rho(i)}{d} \right] = 2 \log \left[\sum_{i \in [d^2]} \frac{\alpha_\rho(i)^2}{d} \frac{1}{\alpha_\rho(i)} \right] \quad (67)$$

implying that

$$e^{\frac{1}{2} M_{1/2}(\rho)} = \mathbb{E}_{i \sim p_{\pi_\rho}} \left[\frac{1}{\alpha_\rho(i)} \right]. \quad (68)$$

Let X be the random variable $X = \alpha_\rho(Y)$ for $Y \sim p_{\pi_\rho}$. The above implies that $\mathbb{E} X = 2^{\frac{1}{2} M_{1/2}(\rho)}$, and so we may apply Markov's inequality:

$$\mathbb{P}(X \geq \tau^{-1/2}) \leq \frac{e^{\frac{1}{2} M_{1/2}(\rho)}}{\tau^{-1/2}} \implies \mathbb{P}[\alpha_\rho(Y)^2 \leq \tau] \leq \sqrt{\tau} e^{\frac{1}{2} M_{1/2}(\rho)}. \quad (69)$$

But $\mathbb{P}_{i \sim p_{\pi_\rho}}[\alpha_\rho(i) \leq \tau] \equiv F_\rho(\tau)$, concluding the proof. $\square$

Then we may directly adapt the proof of [41, Theorem 8], making use of the above lemma:

Corollary 37. (Corollary of [41, Theorem 8]) *Let ρ be a d dimensional pure quantum state. Let $M_{1/2}(\rho)$ be its stabilizer $1/2$ -Rényi entropy, and χ be the maximum Rényi entanglement entropy across all possible bipartitions of ρ . Finally, let p_{π_ρ} be the distribution with support $\{1, \dots, d^2\}$ corresponding to Pauli sampling on ρ , as defined in section 2.7. For a choice of $\Delta, \delta \in (0, 1)$, one may generate $i \sim q$, where, with probability at least $1 - \delta$,*

$$\sum_{i \in [d^2]} |p_{\pi_\rho}(i) - q(i)| \leq \Delta, \quad (70)$$

requiring N copies of $\rho \otimes \rho$ Bell measurements, with

$$N = O(e^{4\chi} e^{2M_{1/2}(\rho)} \Delta^{-4} \log^3 d \log \frac{1}{\delta}). \quad (71)$$

Each generated sample requires $O(N \log^2 d)$ time in classical postprocessing of the obtained Bell measurement data.

Proof. From [41, Lemma 1] and the adapted Lemma 36, and for f, γ, ϵ as defined in Ref. [41],

$$f(\gamma) \leq F_\rho(\gamma e^{2\chi}) \leq \sqrt{\gamma} e^{\chi} e^{\frac{1}{2} M_{1/2}(\rho)}. \quad (72)$$

Then, from [41, Theorem 7], the algorithm therein proposed samples from a distribution q such that

$$\sum_{i \in [d^2]} |p_{\pi_\rho}(i) - q(i)| \leq f(\gamma) + \exp\left(\frac{4\epsilon \log d}{\gamma}\right) - 1. \quad (73)$$

Thus we wish to choose γ and ϵ such that

$$\begin{cases} f(\gamma) \leq \frac{\Delta}{2} \\ \exp\left(\frac{4\epsilon \log d}{\gamma}\right) \leq \exp\left(\frac{\Delta}{4}\right) \leq 1 + \frac{\Delta}{2} \end{cases} \iff \begin{cases} \gamma \leq \left(\frac{\Delta}{2}\right)^2 (e^\chi)^{-2} e^{-M_{1/2}(\rho)} \\ \epsilon = \frac{\Delta}{16} \frac{\gamma}{\log d} \end{cases} \quad (74)$$

ensuring that the total variation distance between p_{π_ρ} and q is guaranteed to be at most Δ . The rest of the proof proceeds exactly like in Ref. [41]. $\square$

But now note that, from section 2.7, $|\langle \psi | \phi \rangle|^2 \equiv \text{Tr}(|\psi\rangle\langle\psi| |\phi\rangle\langle\phi|) \equiv \pi_{|\psi\rangle\langle\psi|}^T \pi_{|\phi\rangle\langle\phi|}$, and

$$e^{\frac{1}{2} M_{1/2}(\rho)} \equiv \mathbf{D}(\rho) \equiv d^{-1} \sum_{i \in [d]^2} |\text{Tr}(P_i \rho)| \equiv d^{-1/2} \|\pi_\rho\|_1. \quad (75)$$

Thus Theorem 35 follows as an immediate corollary of Remark 26, Lemma 34, and Corollary 37.

We note that the approach we have outlined above differs from that of Ref. [41] in the way Pauli sampling is used to construct an inner product estimator. As opposed to their symmetric and antisymmetric protocols, we leverage ASQ to create an estimator that more directly computes an overlap between the “Pauli amplitudes” (square root of the Pauli distributions of the two states). From this results the polynomial improvements in sample and time complexity.

We also note that by proving Theorem 35 with the apparatus of *approximate sample and query* we provide an alternative interpretation for why Pauli sampling is useful for the task of distributed inner product estimation, and why only states with low nonstabilizerness admit efficient procedures. On one hand, we see that the Pauli representation is especially cheap to query from when one is allowed to measure the expectation value of Pauli strings. Due to the normalization factor of $d^{-1/2}$ in each entry of the Pauli representation of a state, one may learn it to error $\epsilon/d^{1/2}$ by learning the Pauli string’s expectation value to only error ϵ . On the other hand, we find that the classes of states with low nonstabilizerness measures correspond to states with low 1-norm when expressed in the Pauli basis. Thus, we may argue that for some states sampling in the Pauli basis is useful because the states’ representation in that basis is quite “peaked”.

5 Discussion and outlook

We have shown that the *sample and query* model can be adapted to more accurately reflect the ability to prepare and measure block encoded states, thereby addressing the comments of Ref. [37]. This adaptation preserves the (arguably) two defining characteristics of the original *sample and query* model: those of composability, and the ability to perform meaningful computation with this form of access. However, this adaptation is not without its drawbacks: the randomized nature of the model poses new problems in deriving results analogous to those of Ref. [26], and, in general, the proof techniques employed in Ref. [25] and onwards do not seem to translate well without careful adaptation.

The ultimate goal for such a framework would be to achieve a dequantization result akin to that of Ref. [26] for the ASQ model: not only because of the practical implications of such a result, but also because of the insight regarding quantum computational power that this would provide. While we hope to have established a first step in this direction, it is unclear if such a result should be expected. We can extend the *approximate sample and query* model to block encodings of matrices — we do so in Appendix C — but we observe that Ref. [26]’s results strongly rely on the classical technique of “sketching” [32]. This technique does not translate immediately to a randomized and finite error setting (as is the case for the other proof techniques, as stated), and so the next steps in this direction will require either a careful adaptation or a completely original approach.

Nonetheless, we have shown that the ASQ model is both expressive and satisfiable in varied setups. Faithfully to the kinds of results obtained with the SQ model, Theorem 35 establishes not only “worst-case scenario” computational complexities, but rather a full characterization with respect to the input data. This provides better bounds for “well behaved” data, as well as quantifiable overheads for “badly behaved” data.

Acknowledgements

The authors would like to thank M. Szegedy, D. Magano, S. Pratapsi, E. Trombetti, L. Bugalho, M. Casariego, M. Hinsche, S. Jerbi, J. Carrasco, M. Ioannou, A. Kontorovich, and D. Berend for their help and comments. We thank the support from FCT — Fundação para a Ciência e a Tecnologia (Portugal), namely through project UIDB/04540/2020, as well as from projects QuantHEP and HQCC supported by the EU QuantERA ERA-NET Cofund in Quantum Technologies and by FCT (QuantERA/0001/2019 and QuantERA/004/2021, respectively), from the EU Quantum Flagship project EuRyQa (101070144), from the BMWi (PlanQK), and the BMBF (QPIC-1, HYBRID). MM would like to further acknowledge and thank FCT — Fundação para a Ciência e a Tecnologia (Portugal) for its support through the scholarship with identifier 2021.05528.BD and DOI 10.54499/2021.05528.BD.

References

- [1] Rajeev Acharya, Igor Aleiner, Richard Allen, Ningfeng Zhu, et al. “Suppressing quantum errors by scaling a surface code logical qubit”. *Nature* **614**, 676–681 (2023).
- [2] M. DeCross, R. Hagshenas, M. Foss-Feig, et al. “Computational power of random quantum circuits in arbitrary geometries”. *Phys. Rev. X* **15**, 021052 (2025).
- [3] Craig Gidney, Noah Shutty, and Cody Jones. “Magic state cultivation: growing T states as cheap as CNOT gates”. Unpublished. (2024) [arXiv:2409.17595](https://arxiv.org/abs/2409.17595).
- [4] Benjamin L. Brock, Shraddha Singh, Devoret, Michel H., et al. “Quantum error correction of qudits beyond break-even”. *Nature* **641**, 612–618 (2025).
- [5] Rajeev Acharya, Dmitry A. Abanin, Zobrist, Nicholas, et al. “Quantum error correction below the surface code threshold”. *Nature* **638**, 920–926 (2024).
- [6] Adam Wills, Min-Hsiu Hsieh, and Hayata Yamasaki. “Constant-overhead magic state distillation”. *Nature Physics* **21**, 1842–1846 (2025).
- [7] Hengyun Zhou, Chen Zhao, Mikhail D. Lukin, et al. “Low-overhead transversal fault tolerance for universal quantum computation”. *Nature* **646**, 303–308 (2025).

- [8] Abhinav Kandala, Antonio Mezzacapo, Jay M. Gambetta, et al. “Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets”. *Nature* **549**, 242–246 (2017).
- [9] M. Cerezo, Guillaume Verdon, Hsin-Yuan Huang, Lukasz Cincio, and Patrick J. Coles. “Challenges and opportunities in quantum machine learning”. *Nature Computational Science* **2**, 567–576 (2022).
- [10] Kishor Bharti, Alba Cervera-Lierta, Alán Aspuru-Guzik, et al. “Noisy intermediate-scale quantum algorithms”. *Reviews of Modern Physics* **94** (2022).
- [11] Jarrod R. McClean, Sergio Boixo, Hartmut Neven, et al. “Barren plateaus in quantum neural network training landscapes”. *Nature Communications* **9** (2018).
- [12] Martín Larocca, Supanut Thanasilp, M. Cerezo, et al. “Barren plateaus in variational quantum computing”. *Nature Reviews Physics* **7**, 174–189 (2025).
- [13] Thomas Häner and Damian S. Steiger. “0.5 petabyte simulation of a 45-qubit quantum circuit”. In Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis. *SC ’17*. ACM (2017).
- [14] Cupjin Huang, Fang Zhang, Michael Newman, Mario Szegedy, Jianxin Chen, et al. “Classical simulation of quantum supremacy circuits”. Unpublished. (2020) [arXiv:2005.06787](#).
- [15] Igor L. Markov, Aneeqa Fatima, Sergei V. Isakov, and Sergio Boixo. “Massively parallel approximate simulation of hard quantum circuits”. In 2020 57th ACM/IEEE Design Automation Conference (DAC). *Pages 1–6*. IEEE (2020).
- [16] Begušić Tomislav and Garnet Kin-Lic Chan. “Fast classical simulation of evidence for the utility of quantum computing before fault tolerance”. Unpublished. (2023) [arXiv:2306.16372](#).
- [17] Hsin-Yuan Huang, Richard Kueng, and John Preskill. “Predicting many properties of a quantum system from very few measurements”. *Nature Physics* **16**, 1050–1057 (2020).
- [18] Hsin-Yuan Huang, Michael Broughton, Jarrod R. McClean, et al. “Power of data in quantum machine learning”. *Nature Communications* **12** (2021).
- [19] Hsin-Yuan Huang, Michael Broughton, Jarrod R. McClean, et al. “Quantum advantage in learning from experiments”. *Science* **376**, 1182–1186 (2022).
- [20] Sitan Chen, Jordan Cotler, Hsin-Yuan Huang, and Jerry Li. “The complexity of *nisq*”. *Nature Communications* **14** (2023).
- [21] Ethan Bernstein and Umesh Vazirani. “Quantum complexity theory”. *SIAM Journal on Computing* **26**, 1411–1473 (1997).
- [22] Scott Aaronson and Andris Ambainis. “Forrelation: A problem that optimally separates quantum from classical computing”. In Proceedings of the forty-seventh annual ACM symposium on Theory of Computing. *STOC ’15*. ACM (2015).
- [23] Xiaoming Sun and Yufan Zheng. “Hybrid decision trees: Longer quantum time is strictly more powerful”. Unpublished. (2019) [arXiv:1911.13091](#).
- [24] Maarten Van Den Nest. “Simulating quantum computers with probabilistic methods”. *Quantum Info. Comput.* **11**, 784–812 (2011).
- [25] Ewin Tang. “A quantum-inspired classical algorithm for recommendation systems”. In Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing. *STOC ’19*. ACM (2019).
- [26] Nai-Hui Chia, András Gilyén, Chunhao Wang, et al. “Sampling-based sublinear low-rank matrix arithmetic framework for dequantizing quantum machine learning”. In Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing. ACM (2020).
- [27] François Le Gall. “Robust dequantization of the quantum singular value transformation and quantum machine learning algorithms”. *computational complexity* **34** (2025).
- [28] Scott Aaronson and Daniel Gottesman. “Improved simulation of stabilizer circuits”. *Physical Review A* **70** (2004).
- [29] Leslie G. Valiant. “Quantum circuits that can be simulated classically in polynomial time”. *SIAM Journal on Computing* **31**, 1229–1254 (2002).
- [30] Daniel R. Simon. “On the power of quantum computation”. *SIAM Journal on Computing* **26**, 1474–1483 (1997).
- [31] P.W. Shor. “Algorithms for quantum computation: discrete logarithms and factoring”. In Proceedings 35th Annual Symposium on Foundations of Computer Science. *Pages 124–134*. SFCS-94. IEEE Comput. Soc. Press (1994).

- [32] P. Dagum, R. Karp, M. Luby, and S. Ross. “An optimal algorithm for monte carlo estimation”. In Proceedings of IEEE 36th Annual Foundations of Computer Science. Pages 142–149. (1995).
- [33] András Gilyén, Yuan Su, Guang Hao Low, and Nathan Wiebe. “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. STOC ’19. ACM (2019).
- [34] Gilles Brassard, Peter Høyer, Michele Mosca, and Alain Tapp. “Quantum amplitude amplification and estimation” (2002).
- [35] Aram W. Harrow, Avinandan Hassidim, and Seth Lloyd. “Quantum algorithm for linear systems of equations”. *Physical Review Letters***103** (2009).
- [36] Andrew M. Childs and Nathan Wiebe. “Hamiltonian simulation using linear combinations of unitary operations”. *Quantum Information and Computation***12** (2012).
- [37] Jordan Cotler, Hsin-Yuan Huang, and Jarrod R. McClean. “Revisiting dequantization and quantum advantage in learning tasks”. Unpublished. (2021) [arXiv:2112.00811](#).
- [38] Joran van Apeldoorn, Arjan Cornelissen, András Gilyén, and Giacomo Nannicini. “Quantum tomography using state-preparation unitaries”. Pages 1265–1318. Society for Industrial and Applied Mathematics. (2023).
- [39] Sergey Bravyi and David Gosset. “Improved Classical Simulation of Quantum Circuits Dominated by Clifford Gates”. *Physical Review Letters* **116**, 250501 (2016).
- [40] Anurag Anshu, Zeph Landau, and Yunchao Liu. “Distributed quantum inner product estimation”. In Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing. STOC ’22. ACM (2022).
- [41] M. Hinsche, M. Ioannou, S. Jerbi, L. Leone, J. Eisert, and J. Carrasco. “Efficient distributed inner-product estimation via pauli sampling”. *PRX Quantum* **6**, 030354 (2025).
- [42] Earl T. Campbell. “Catalysis and activation of magic states in fault-tolerant architectures”. *Physical Review A***83** (2011).
- [43] Lorenzo Leone, Salvatore F. E. Oliviero, and Alioscia Hamma. “Stabilizer rényi entropy”. *Physical Review Letters***128** (2022).
- [44] Donald E. Knuth. “Big omicron and big omega and big theta”. *ACM SIGACT News* **8**, 18–24 (1976).
- [45] Sanjeev Arora and Boaz Barak. “Computational complexity: A modern approach”. *Cambridge University Press*. USA (2009). 1st edition.
- [46] Vivek V. Shende, Aditya K. Prasad, Igor L. Markov, and John P. Hayes. “Reversible logic circuit synthesis”. In Proceedings of the 2002 IEEE/ACM international conference on Computer-aided design - ICCAD ’02. Pages 353–360. ICCAD ’02. ACM Press (2002).
- [47] Sitan Chen, Jordan Cotler, Hsin-Yuan Huang, and Jerry Li. “Exponential separations between learning with and without quantum memory”. In 2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS). Volume 169, pages 574–585. IEEE (2022).
- [48] Steven T. Flammia and Yi-Kai Liu. “Direct fidelity estimation from few pauli measurements”. *Physical Review Letters***106** (2011).
- [49] M. Van den Nest. “Classical simulation of quantum computation, the Gottesman-Knill theorem, and slightly beyond”. *Quantum Information and Computation* **10**, 258–271 (2010).
- [50] Masashi Okamoto. “Some inequalities relating to the partial sum of binomial probabilities”. *Annals of the Institute of Statistical Mathematics* **10**, 29–35 (1959).
- [51] Srinivasan Arunachalam and Louis Schatzki. “Distributed inner product estimation with limited quantum communication”. Unpublished. (2024) [arXiv:2410.12684](#).
- [52] Weiyuan Gong, Jonas Haferkamp, Qi Ye, and Zhihan Zhang. “On the sample complexity of purity and inner product estimation”. Unpublished. (2024) [arXiv:2410.12712](#).

A Proof of Remark 31

Let $\epsilon \in (0, 1]$, $\phi \geq 1$, and $x, x', \tilde{x} \in \mathbb{C}^d$, such that for a suitable constant $C_1 > 0$,

$$\|\tilde{x} - x'\| \leq \frac{C_1 \epsilon}{4\sqrt{\phi}} \|\tilde{x}\|, \quad (76)$$

$$\|\tilde{x}\|^2 \leq \phi \|x\|^2, \quad (77)$$

$$\|x\| \leq 1, \quad (78)$$

and $p_{\tilde{x}}$ is a ϕ -oversampling distribution for p_x .

Assume Approximate Sample access to x' and Approximate Query access to x . From Lemma 24, the ℓ_1 distance between $p_{\tilde{x}}$ and $p_{x'}$ is at most $C_1 \epsilon / \sqrt{\phi} =: \Delta$. Furthermore, from a triangle inequality, $\|x'\| \leq (1 + 2\Delta) \|\tilde{x}\|$, and from the definition of $p_{\tilde{x}}$, $|x'(i)| = \|x'\| \sqrt{p_{x'}(i)}$.

From the Approximate Query access to x , one has that for a suitable constant C_2 and in time $\mathbf{q}(C_2 \epsilon)$, and with probability $2/3$,

$$\forall_{i \in [d]} \quad \hat{x}_i : |\hat{x}_i - x(i)| \leq C_2 \epsilon. \quad (79)$$

Let, for convenience of notation, $\epsilon_Q := C_2 \epsilon$.

Finally, for a choice of constant C_3 , one has from Theorem 14 that taking $O[\epsilon^{-2} \phi^2 \log(d/\delta)]$ samples of $p_{x'}$ is sufficient to estimate

$$\forall_{i \in [d]} \quad \hat{p}_i : |\hat{p}_i - p_{x'}(i)| \leq C_3 \frac{\epsilon^2}{2\phi} \quad (80)$$

with overall probability $1 - \delta$. Again for convenience of notation we define $\epsilon_P := C_3 \epsilon^2 / \phi$.

With these choices, define the point estimator

$$\widehat{x^\dagger y}(i) = \begin{cases} 0 & \text{if } \hat{p}_i \leq \frac{3}{2} \epsilon_P, \\ \frac{1}{\hat{p}_i + \epsilon_P/2} \hat{x}_i^* y(i) & \text{otherwise.} \end{cases} \quad (81)$$

Using the definitions above, the triangle inequality (possibly with double counting), and Hölder inequalities, as well as the fact that $|x(i)| \leq |\tilde{x}(i)|$, and ℓ_p norm inequalities, the bias of this estimator is bounded as

$$\left| \mathbb{E}_{i \sim p_{x'}} \widehat{x^\dagger y}(i) - x^\dagger y \right| \quad (82)$$

$$\leq \sum_{i: p_{x'}(i) \leq 2\epsilon_P} |x(i)| |y(i)| + \sum_{i: p_{x'}(i) \geq \epsilon_P} \left| \frac{p_{x'}(i)}{\hat{p}_i + \epsilon_P/2} \hat{x}_i^* y(i) - x(i)^* y(i) \right| \quad (83)$$

$$\leq \left[\sum_{i: p_{x'}(i) \leq 2\epsilon_P} |\tilde{x}(i) - x'(i)| |y(i)| + |x'(i)| |y(i)| \right] \quad (84)$$

$$\begin{aligned} &+ \left[\sum_{i: p_{x'}(i) \geq \epsilon_P} \left| \frac{p_{x'}(i)}{\hat{p}_i + \epsilon_P/2} - 1 \right| |x(i)| |y(i)| + |\hat{x}_i - x(i)| |y(i)| \right. \\ &\quad \left. + \left| \frac{p_{x'}(i)}{\hat{p}_i + \epsilon_P/2} - 1 \right| |\hat{x}_i - x(i)| |y(i)| \right] \\ &\leq \left[\Delta \|y\|_\infty + \sum_{i: p_{x'}(i) \leq 2\epsilon_P} \|x'\| \sqrt{p_{x'}(i)} |y(i)| \right] \end{aligned} \quad (85)$$

$$\begin{aligned} &+ \left[\sum_{i: p_{x'}(i) \geq \epsilon_P} \frac{\epsilon_P}{p_{x'}(i)} |x(i)| |y(i)| + \epsilon_Q |y(i)| + \frac{\epsilon_P}{p_{x'}(i)} \epsilon_Q |y(i)| \right] \\ &\leq \left[\Delta \|y\|_1 + 2\sqrt{\epsilon_P} (1 + 2\Delta) \|\tilde{x}\| \|y\|_1 \right] \end{aligned} \quad (86)$$

$$\begin{aligned} &+ \left[\sum_{i: p_{x'}(i) \geq \epsilon_P} |\tilde{x}(i) - x'(i)| |y(i)| + |x'(i)| |y(i)| + \epsilon_Q \|y\|_1 + \epsilon_Q \|y\|_1 \right] \\ &\leq [C_1 + \sqrt{C_3} (1 + 2C_1)] \epsilon \|y\|_1 + \Delta \|y\|_1 + \sum_{i: p_{x'}(i) \geq \epsilon_P} \|x'\| \sqrt{p_{x'}(i)} |y(i)| + 2\epsilon_Q \|y\|_1 \end{aligned} \quad (87)$$

$$\leq [C_1 + (1 + 2C_1) \sqrt{C_3} + 2C_2] \epsilon \|y\|_1. \quad (88)$$

Likewise, we bound the variance of this estimator with the same techniques and a geometric mean–quadratic mean inequality:

$$\text{VAR}_{i \sim p_{x'}} \widehat{x^\dagger y}(i) \leq \mathbb{E}_{i \sim p_{x'}} |\widehat{x^\dagger y}(i)|^2 \quad (89)$$

$$\leq \sum_{i: p_{x'}(i) \geq \epsilon_P} \frac{p_{x'}(i)}{\hat{p}_i + \epsilon_P/2} \frac{1}{\hat{p}_i + \epsilon_P/2} |\hat{x}_i|^2 |y(i)|^2 \quad (90)$$

$$\leq \sum_{i: p_{x'}(i) \geq \epsilon_P} \frac{1}{p_{x'}(i)} [|x(i)| + \epsilon_Q]^2 |y(i)|^2 \quad (91)$$

$$\leq \sum_{i: p_{x'}(i) \geq \epsilon_P} \frac{1}{p_{x'}(i)} [|x'(i)| + |\tilde{x}(i) - x'(i)| + \epsilon_Q]^2 |y(i)|^2 \quad (92)$$

$$\leq 3\|x'\|^2 \|y\|^2 + 3 \frac{1}{\epsilon_P} \Delta^2 \|y\|^2 + 3 \frac{\epsilon_Q^2}{\epsilon_P} \|y\|^2 \quad (93)$$

$$\leq [3(1 + 2C_1)^2 + 3 \frac{C_1^2}{C_3} + 3 \frac{C_2^2}{C_3}] \phi \|y\|^2. \quad (94)$$

Therefore we have a bias bounded as $O(\epsilon \|y\|_1)$, and a variance bounded as $O(\phi \|y\|^2)$. It follows from the Chebyshev inequality that by taking the average of $O(\phi \epsilon^{-2} \delta^{-1})$ valid samples of the point estimator, the resulting value deviates from $x^\dagger y$ at most $\epsilon \|y\|_1$ with probability at least $1 - \delta$. The rest of the proof follows identically to that of Lemma 30.

B Proof of Remark 33

Let $\epsilon \in [0, 1]$, and $x, y, \tilde{x}, \tilde{y}$ such that for suitable constants C_1, C_2, C_3 : $\|x\|, \|y\| \leq 1$, $\|\tilde{x}\|^2 \leq \phi \|x\|^2$, $\|\tilde{y}\|^2 \leq \phi \|y\|^2$, and

$$\|x' - \tilde{x}\| \leq \frac{C_1 \epsilon}{4\sqrt{\phi}} \|\tilde{x}\|, \quad \|y' - \tilde{y}\| \leq \frac{C_1 \epsilon}{4\sqrt{\phi}} \|\tilde{y}\|. \quad (95)$$

Define, for convenience of notation, $\Delta := C_1 \epsilon / \sqrt{\phi}$. Furthermore assume that $p_{\tilde{x}}$ is a ϕ -oversampling distribution for p_x , and $p_{\tilde{y}}$ is a ϕ -oversampling distribution for p_y . Note that the conditions imposed on $\|\tilde{x} - x'\|$ and $\|\tilde{y} - y'\|$ imply that $\|x'\| \leq (1 + \Delta/4) \|\tilde{x}\|$ and likewise $\|y'\| \leq (1 + \Delta/4) \|\tilde{y}\|$.

Let p and q be defined by their probability mass functions; for all $i \in [d]$,

$$p(i) = \frac{p_{\tilde{x}}(i) + p_{\tilde{y}}(i)}{2}, \quad q(i) = \frac{p_{x'}(i) + p_{y'}(i)}{2}. \quad (96)$$

One may sample from p by outputting uniformly randomly from either $p_{\tilde{x}}$ or $p_{\tilde{y}}$, and from q by outputting uniformly randomly from either $p_{x'}$ or $p_{y'}$.

From Theorem 14, one has that $O(\epsilon^{-4} \phi^2 \log d)$ samples of q suffice to estimate $q(i)$ to error $C_2 \epsilon^2 / (2\phi)$ with probability 8/9. [After this, per the Theorem, reading an estimate takes time $O(\log d)$.] Let, for ease of notation $\gamma := C_2 \epsilon^2 / \phi$, such that each estimate of $q(i)$, say, $\hat{q}_i$ is accurate to error $\gamma/2$ (with error probability 1/9).

Likewise, if $\epsilon_Q := C_3 \epsilon^2 / \phi$, one has from the Approximate Query access that, for all $i = 1, \dots, d$, it is possible to compute an estimate of $x(i)$ and $y(i)$ — say, $\hat{x}_i$ and $\hat{y}_i$, respectively — to absolute error ϵ_Q and with error probability 1/3 in time $\mathbf{q}(\epsilon_Q)$.

To relate $|x(i)|$ to $q(i)$, we will make use of the following bound, following from the definitions above and a triangle inequality:

$$|x(i)| \leq |\tilde{x}(i)| \leq |x'(i)| + |x'(i) - \tilde{x}(i)| \leq \|x'\| \sqrt{p_{x'}(i)} + |x'(i) - \tilde{x}(i)| \quad (97)$$

$$\leq (1 + \Delta/4) \|\tilde{x}\| \sqrt{2q(i)} + |x'(i) - \tilde{x}(i)| \leq (1 + 2\Delta) \sqrt{2\phi} \sqrt{q(i)} + |x'(i) - \tilde{x}(i)|. \quad (98)$$

An analogous bound holds for $|y(i)|$. We will also make use of the following fact, which makes use of an arithmetic mean–geometric mean inequality:

$$|x(i)| |y(i)| \leq \frac{|\tilde{x}(i)|}{\|\tilde{x}\|} \frac{|\tilde{y}(i)|}{\|\tilde{y}\|} \|\tilde{x}\| \|\tilde{y}\| \leq \phi p(i) \leq \phi [q(i) + |p(i) - q(i)|]. \quad (99)$$

Finally, as elsewhere, we will make use of the fact that if $\hat{q}_i$ is an estimate of $q(i)$ to error $\epsilon_P/2$, then $(\hat{q}_i + \epsilon_P/2)^{-1} \leq q(i)^{-1}$.

Now, define the estimator

$$\widehat{x^\dagger y}(i) = \begin{cases} 0 & \text{if } \hat{q}_i \leq \frac{3}{2}\gamma \\ \frac{1}{\hat{q}_i + \gamma/2} \hat{x}_i^* \hat{y}_i & \text{otherwise.} \end{cases} \quad (100)$$

The bias of this estimator (with respect to $x^\dagger y$) is bounded as

$$|\mathbb{E}_{i \sim q} \widehat{x^\dagger y} - x^\dagger y| \leq \sum_{i: q(i) \leq 2\gamma} |x(i)||y(i)| + \sum_{i: q(i) \geq \gamma} \left| \frac{q(i)}{\hat{q}_i + \gamma/2} \hat{x}_i^* \hat{y}_i - x(i)^* y(i) \right| \quad (101)$$

as follows from a triangle inequality and some potential double counting. We further bound these two terms separately, making use of the Cauchy-Schwarz inequality and the identities above:

$$\sum_{i: q(i) \leq 2\gamma} |x(i)||y(i)| \leq (1 + \frac{\Delta}{4}) 2\sqrt{\phi\gamma} \|y\|_1 + \sum_{i: q(i) \leq 2\gamma} |x'(i) - \tilde{x}(i)||y(i)| \quad (102)$$

$$\leq 2(1 + \frac{\Delta}{4}) \sqrt{\gamma\phi} \|y\|_1 + \|\tilde{x} - x'\| \|y\| \leq 2(1 + \frac{\Delta}{4}) \sqrt{\gamma\phi} \|y\|_1 + \frac{\Delta}{4} \sqrt{\phi}. \quad (103)$$

But, by a symmetry argument, we have

$$\sum_{i: q(i) \leq 2\gamma} |x(i)||y(i)| \leq 2(1 + \frac{\Delta}{4}) \sqrt{\gamma\phi} \min(\|x\|_1, \|y\|_1) + \frac{\Delta}{4} \sqrt{\phi}. \quad (104)$$

The second term is bounded as

$$\sum_{i: q(i) \geq \gamma} \left| \frac{q(i)}{\hat{q}_i + \gamma/2} \hat{x}_i^* \hat{y}_i - x(i)^* y(i) \right| \leq \sum_{i: q(i) \geq \gamma} \left| \frac{q(i)}{\hat{q}_i + \gamma/2} - 1 \right| |x(i)||y(i)| + |\hat{x}_i^* \hat{y}_i - x(i)^* y(i)| \quad (105)$$

$$+ \left| \frac{q(i)}{\hat{q}_i + \gamma/2} - 1 \right| |\hat{x}_i^* \hat{y}_i - x(i)^* y(i)|$$

$$\leq \sum_{i: q(i) \geq \gamma} \frac{\gamma}{q(i)} |x(i)||y(i)| + 2[\epsilon_Q |x(i)| + \epsilon_Q |y(i)| + \epsilon_Q^2] \quad (106)$$

$$\leq (1 + \frac{\Delta}{4}) \sqrt{2\phi\gamma} \min(\|x\|_1, \|y\|_1) + \frac{\Delta}{4} \sqrt{\phi} + \left[4\sqrt{\phi} \frac{\epsilon_Q}{\sqrt{\gamma}} + 2\frac{\epsilon_Q}{\gamma} \epsilon_Q \right]. \quad (107)$$

Therefore, the bias is overall bounded as

$$(2\sqrt{2C_2})\epsilon \min(\|x\|_1, \|y\|_1) + \left(\frac{C_1}{4} + \frac{4C_3}{\sqrt{C_2}} + \frac{2C_3^2}{C_2} \right) \epsilon = O\{\epsilon [1 + \min(\|x\|_1, \|y\|_1)]\}. \quad (108)$$

We now bound the variance of the estimator:

$$\text{VAR}_{i \sim q} \widehat{x^\dagger y}(i) \leq \mathbb{E}_{i \sim q} |\widehat{x^\dagger y}(i)|^2 \leq \sum_{i: q(i) \geq \gamma} \frac{1}{q(i)} |\hat{x}_i|^2 |\hat{y}_i|^2 \leq 4 \sum_{i: q(i) \geq \gamma} \frac{1}{q(i)} (|x(i)|^2 + \epsilon_Q^2)(|y(i)|^2 + \epsilon_Q^2) \quad (109)$$

$$\leq 4 \sum_{i: q(i) \geq \gamma} \frac{1}{q(i)} |x(i)|^2 |y(i)|^2 + 4 \sum_{i: q(i) \geq \gamma} \frac{1}{q(i)} \epsilon_Q^2 (|x(i)|^2 + |y(i)|^2) + 4 \sum_{i: q(i) \geq \gamma} \frac{1}{q(i)} \epsilon_Q^4. \quad (110)$$

Again further bounding term by term:

$$\sum_{i:q(i) \geq \gamma} \frac{1}{q(i)} |x(i)|^2 |y(i)|^2 \leq 4\phi \sum_{i:q(i) \geq \gamma} \frac{p(i)}{q(i)} |x(i)| |y(i)| \leq 4\phi + 4\phi \sum_{i:q(i) \geq \gamma} \frac{|p(i) - q(i)|}{q(i)} |x(i)| |y(i)| \quad (111)$$

$$\leq 4\phi + 4\phi^2 \Delta + 4\phi^2 \sum_{i:q(i) \geq \gamma} |p(i) - q(i)|^2 \frac{1}{q(i)} \leq 4\phi + 4\phi^2 \Delta + 4\phi^2 \frac{\Delta^2}{\gamma} \quad (112)$$

$$\leq 4(1 + C_1 + \frac{C_1^2}{C_2}) \phi^2; \quad (113)$$

$$4 \sum_{i:q(i) \geq \gamma} \frac{1}{q(i)} \epsilon_Q^2 (|x(i)|^2 + |y(i)|^2) \leq 8\phi \sum_{i:q(i) \geq \gamma} \frac{p(i)}{q(i)} \epsilon_Q^2 \leq 8\phi \frac{\epsilon_Q^2}{\gamma} = \frac{8C_3^2}{C_2}; \quad (114)$$

$$4 \sum_{i:q(i) \geq \gamma} \frac{1}{q(i)} \epsilon_Q^4 \leq 4 \frac{\epsilon_Q^4}{\gamma^2} = \frac{256 C_3^4}{C_2^2}. \quad (115)$$

$$(116)$$

Thus, the variance is bounded as $O(\phi^2)$, and the rest of the proof proceeds identically to that of Lemma 32.

C Approximate sample and query for block encoded matrices

Definition C.1. (Approximate ϕ -oversample and query access — matrices) For $\phi \geq 1$, say that one has Approximate ϕ -Oversample and Query (ASQ_ϕ) access to a matrix $A \in \mathbb{C}^{n \times m}$ if, for a matrix $\tilde{A}$, satisfying $\|\tilde{A}\|_F^2 \leq \phi \|A\|_F^2$ and $|\tilde{A}(i, j)|^2 \geq |A(i, j)|$, for all $i, j \in [m] \times [n]$, one has query access to every column of A ; AS access (cf. Def. 19) to every column of $\tilde{A}$; and one has AS access to the vector of column norms of $\tilde{A}$, i.e., $\tilde{a} \in \mathbb{R}^m$ with entries

$$\tilde{a}(j) = \sqrt{\sum_{i \in [n]} |\tilde{A}(i, j)|^2}$$

with one-sided error probability at most $1/3$ and runtime s_A .

Let U be a m -block encoding of a $2^n \times 2^n$ matrix A (cf. section 2.2). Given access to singly coherent applications of both U and $U^\dagger$, we claim that it is possible to satisfy ASQ_ϕ access to A . The first two conditions, as outlined above, are easy to satisfy: if one prepares $U|0\rangle^{\otimes m}|j\rangle$, then one has a block encoding of the j th column of A , and we have seen that this is sufficient to satisfy ASQ access to the j th column of A . Therefore, it remains to show that the third condition can also be satisfied. Consider the following procedure:

1. Choose uniformly randomly a string of n bits.
2. Prepare the corresponding computational basis state.
3. Run the computational basis state through $U^\dagger$ (with the block register suitably initialized to all-zeroes).
4. Postselecting for the correct block, output the readout of the remaining register.

One finds that the output is the index of a column of A (say, k), with probability $\|A(*, k)\|^2 / \|A\|_F^2$, where $A(*, k)$ is the k th column of A . Moreover, for the procedure to succeed with probability $2/3$, one finds that $O(d\|A\|_F^2)$ rounds of post-selection are expected to be required.