

Forging quantum data: classically defeating an IQP-based quantum test

Gregory D. Kahanamoku-Meyer

Department of Physics, University of California at Berkeley, Berkeley, CA 94720

Recently, quantum computing experiments have for the first time exceeded the capability of classical computers to perform certain computations—a milestone termed “quantum computational advantage.” However, verifying the output of the quantum device in these experiments required extremely large classical computations. An exciting next step for demonstrating quantum capability would be to implement tests of quantum computational advantage with efficient classical verification, such that larger system sizes can be tested and verified. One of the first proposals for an efficiently-verifiable test of quantumness consists of hiding a secret classical bitstring inside a circuit of the class IQP, in such a way that samples from the circuit’s output distribution are correlated with the secret. The classical hardness of this protocol has been supported by evidence that directly simulating IQP circuits is hard, but the security of the protocol against other (non-simulating) classical attacks has remained an open question. In this work we demonstrate that the protocol is not secure against classical forgery. We describe a classical algorithm that can not only convince the verifier that the (classical) prover is quantum, but can in fact can extract the secret key underlying a given protocol instance. Furthermore, we show that the key extraction algorithm is efficient in practice for problem sizes of hundreds of qubits. Finally, we provide an implementation of the algorithm, and give the secret vector underlying the “\$25 challenge” posted online by the authors of the original paper.

1 Introduction

Recent experiments have demonstrated groundbreaking quantum computational power in the laboratory, showing *quantum computational advantage* [1–4]. In

Gregory D. Kahanamoku-Meyer: gkm@berkeley.edu

the past decade, much theoretical work has gone into designing experimental protocols expressly for this purpose, and providing evidence for the classical hardness of reproducing the experimental results [5–18]. A difficulty with many of them, however, is that the quantum machine’s output is hard to *verify*. In many cases, the best known algorithm for directly checking the solution is equivalent to classically performing the computational task itself. This presents challenges for validation of the test’s results, because an ideal demonstration of quantum advantage occurs in the regime where a classical solution is not just difficult, but impossible with current technology. In that regime, experiments have had to resort to indirect methods to demonstrate that their devices are producing correct results [1–4].

In 2009, an efficiently-verifiable test of quantum computational advantage was proposed based on “instantaneous quantum polynomial-time” (IQP) circuits—quantum circuits in which all operations commute [19]. The protocol places only moderate requirements on the quantum device, making it potentially a good candidate for near-term hardware. Furthermore, later papers showed based on reasonable assumptions that classically sampling from the resulting distribution should be hard [20, 21]. This suggests that a “black-box” approach to cheating classically (by simply simulating the quantum device) is indeed computationally hard, and only a couple hundred qubits would be required to make a classical solution intractable.

Importantly, however, the classical verifier of the efficiently-verifiable protocol does not explicitly check whether the prover’s samples come from the correct distribution (in fact, doing such a check efficiently is probably not possible [20]). Instead, the sampling task is designed such that bitstrings from its distribution will be orthogonal to some secret binary vector \mathbf{s} with high probability, and it is this property that is checked by the verifier. A question that has remained open is whether a classical machine can efficiently generate samples satisfying the orthogonality check, without necessarily approximating the actual circuit’s distribution. In this




Figure 1: Mean time to extract the secret vector \mathbf{s} from X -programs constructed as described in [19]. Shaded region is the first to third quartile of the distribution of runtimes. We observe that the time is polynomial and fast in practice even up to problem sizes of hundreds of qubits. See Section 3.2 for a discussion of the $\mathcal{O}(n^2)$ scaling. The data points were computed by applying the algorithm to 1000 unique X -programs at each problem size. The secret vector was successfully extracted for every X -program tested. Experiments were completed using one thread on an Intel 8268 “Cascade Lake” processor.

work we show that the answer to this question is yes. We give an explicit algorithm that can extract the secret bistring \mathbf{s} underlying any instance of the protocol, thus making it trivial to generate orthogonal samples that pass the verifier’s test. The main results described here are a statement of the algorithm, a proof that a single iteration of it will extract the secret vector \mathbf{s} with probability $1/2$ (which can be made arbitrarily close to 1 by repetition), and empirical results demonstrating that the algorithm is efficient in practice (summarized in Figure 1).

The following is a summary of the paper’s structure. In Section 2, we review the protocol’s construction and some relevant analysis from the original paper. In Section 3 we describe the algorithm to extract the secret key, and therefore break the protocol’s security against classical provers. There we also discuss briefly our implementation of the algorithm. In Section 4 we discuss the results, and provide the secret key underlying the “\$25 challenge” that accompanied the publication of the protocol.

2 Background

Overview of protocol Here we summarize the IQP-based protocol for quantum advantage, in the standard

cryptographic terms of an ostensibly quantum *prover* attempting to prove its quantum capability to a classical *verifier*. We refer the reader to the work that proposed the protocol for any details not covered here [19]. The core of the protocol is a sampling problem. The verifier generates a Hamiltonian H_P consisting of a sum of products of Pauli X operators, and asks the quantum prover to generate samples by measuring the state $e^{iH_P\theta} |0^{\otimes n}\rangle$ for some value of the “action” θ . The Hamiltonian H_P is designed such that the measured bitstrings $\{\mathbf{x}_i\}$ are biased with respect to a secret binary vector \mathbf{s} , so that $\mathbf{x}_i \cdot \mathbf{s} = 0$ with high probability (where (\cdot) represents the binary inner product, modulo 2). The classical verifier, with knowledge of \mathbf{s} , can quickly check that the samples have such a bias. Since \mathbf{s} should be only known to the verifier, it was conjectured that the only efficient way to generate such samples is by actually computing and measuring the quantum state [19]. However, in Section 3 we show that it is possible to extract \mathbf{s} classically from just the description of the Hamiltonian.

X-programs A Hamiltonian of the type used in this protocol can be described by a rectangular matrix of binary values, for which each row corresponds to a term of the Hamiltonian. Given such a binary matrix P (called an “ X -program”), the Hamiltonian is

$$H_P = \sum_i \prod_j X^{P_{ij}} \quad (2.1)$$

In words, a 1 in P at row i and column j corresponds to the inclusion of a Pauli X operator on the j^{th} site in the i^{th} term of the Hamiltonian. The X -program also has one additional parameter θ , which is the “action”—the integrated energy over time for which the Hamiltonian will be applied. For the protocol relevant to this work, the action is set to $\theta = \pi/8$ (see below).

Embedding a bias and verifying the output

In order to bias the output distribution along \mathbf{s} , a submatrix with special properties is embedded within the matrix P . Notationally, for a vector \mathbf{s} and matrix P , let the submatrix $P_{\mathbf{s}}$ be that which is generated by deleting all rows of P that are orthogonal to \mathbf{s} . Letting \mathbf{X} represent the distribution of measurement results for a given X -program, it can be shown that the probability that a measurement outcome is orthogonal to the vector \mathbf{s} , $\Pr[\mathbf{X} \cdot \mathbf{s} = 0]$, depends only on the submatrix $P_{\mathbf{s}}$. The rows of P that are orthogonal to \mathbf{s} are irrelevant. The protocol uses that fact to attempt to hide $P_{\mathbf{s}}$ (and thus \mathbf{s}): starting with a matrix $P_{\mathbf{s}}$ that produces a bias, we may attempt to hide it in a larger matrix P by appending rows that are random aside from having

$\mathbf{p} \cdot \mathbf{s} = 0$, and then scrambling the new, larger matrix in a way that preserves the bias.

But what matrix $P_{\mathbf{s}}$ should one start with? In the protocol, the verifier sets $P_{\mathbf{s}}$ to the generator matrix for a binary code of block length $q \equiv 7 \pmod{8}$ whose codewords \mathbf{c} have $\text{wt}(\mathbf{c}) \in \{-1, 0\} \pmod{4}$ (and both those weights are represented, that is, the codewords do not all have weight $0 \pmod{4}$). In [19], the authors suggest specifically using a binary quadratic residue (QR) code because it has the desired codeword weights. The action θ is set to $\pi/8$. As described in Facts 2.1 and 2.2 below, this configuration leads to a gap between the quantum and classical probabilities of generating samples orthogonal to \mathbf{s} (for the best known classical strategy before this work). The verifier's check is then simply to request a large number of samples, and determine if the fraction orthogonal to \mathbf{s} is too large to have likely been generated by any classical strategy.

In the two Facts below, we recall the probabilities corresponding to the quantum strategy and previously best-known classical strategy [19]. The reasoning behind the classical strategy (Fact 2.2) forms the setup for the new algorithm described in this paper; it is worth understanding its proof before moving on to the algorithm in Section 3.

Fact 2.1. Quantum strategy

Let P be an X -program constructed by embedding a submatrix with the properties described above. Let \mathbf{X} be a random variable representing the distribution of bitstrings from an n -qubit quantum state $e^{iH_P\pi/8}|0\rangle$ measured in the Z basis, where H_P is defined as in Equation 2.1. Then,

$$\Pr[\mathbf{X} \cdot \mathbf{s} = 0] = \cos^2\left(\frac{\pi}{8}\right) \approx 0.85 \dots \quad (2.2)$$

Proof. The entire proof is contained in [19]. To summarize, it is shown that for any string \mathbf{z} and corresponding submatrix $P_{\mathbf{z}}$, the probability is

$$\Pr[\mathbf{X} \cdot \mathbf{z} = 0] = \mathbb{E}_{\mathbf{c}} [\cos^2(\theta \cdot (q - 2\text{wt}(\mathbf{c})))] \quad (2.3)$$

where θ is the action, q is the number of rows in $P_{\mathbf{z}}$ and the expectation is taken over the codewords \mathbf{c} of the code generated by the submatrix $P_{\mathbf{z}}$. When the values of $\theta = \pi/8$, $q \equiv 7 \pmod{8}$ and $\text{wt}(\mathbf{c}) \in \{-1, 0\} \pmod{4}$ corresponding to the specific submatrix $P_{\mathbf{s}}$ are substituted into this expression, the result is Equation 2.2. □

Fact 2.2. Classical strategy of [19]

Again let P be an X -program constructed by embedding a submatrix with the properties described above. Let \mathbf{d}, \mathbf{e} be two bitstrings of length n (the length

of a row of P). Define $P_{\mathbf{d}, \mathbf{e}}$ as the matrix generated by deleting the rows of P orthogonal to \mathbf{d} or \mathbf{e} .¹ Let $\mathbf{y} = \sum_{\mathbf{p}_i \in P_{\mathbf{d}, \mathbf{e}}} \mathbf{p}_i$ be the vector sum of the rows of $P_{\mathbf{d}, \mathbf{e}}$. Letting \mathbf{Y} be the random variable representing the distribution of \mathbf{y} when \mathbf{d} and \mathbf{e} are chosen uniformly at random, then

$$\Pr[\mathbf{Y} \cdot \mathbf{s} = 0] = 3/4 \quad (2.4)$$

Proof. (From [19]) With \mathbf{y} defined as above, we have

$$\mathbf{y} \cdot \mathbf{s} = \sum_{\mathbf{p}_i \in P_{\mathbf{d}, \mathbf{e}}} \mathbf{p}_i \cdot \mathbf{s} \quad (2.5)$$

By definition, $\mathbf{p}_i \cdot \mathbf{s} = 1$ if $\mathbf{p}_i \in P_{\mathbf{s}}$. Therefore $\mathbf{y} \cdot \mathbf{s}$ is equivalent to simply counting the number of rows common to $P_{\mathbf{s}}$ and $P_{\mathbf{d}, \mathbf{e}}$, or equivalently, counting the rows in $P_{\mathbf{s}}$ for which $\mathbf{p} \cdot \mathbf{d}$ and $\mathbf{p} \cdot \mathbf{e}$ are both 1. We can express this using the matrix-vector products of $P_{\mathbf{s}}$ with \mathbf{d} and \mathbf{e} :

$$\mathbf{y} \cdot \mathbf{s} = \sum_{\mathbf{p}_i \in P_{\mathbf{s}}} (\mathbf{p} \cdot \mathbf{d})(\mathbf{p} \cdot \mathbf{e}) \quad (2.6)$$

$$= (P_{\mathbf{s}} \mathbf{d}) \cdot (P_{\mathbf{s}} \mathbf{e}) \quad (2.7)$$

Considering that $P_{\mathbf{s}}$ is the generator matrix for an error correcting code, denote $\mathbf{c}_{\mathbf{d}} = P_{\mathbf{s}} \mathbf{d}$ as the encoding of \mathbf{d} under $P_{\mathbf{s}}$. Then we have

$$\mathbf{y} \cdot \mathbf{s} = \mathbf{c}_{\mathbf{d}} \cdot \mathbf{c}_{\mathbf{e}} \quad (2.8)$$

Now, note that if a code has $\text{wt}(\mathbf{c}) \in \{-1, 0\} \pmod{4}$ for all codewords \mathbf{c} , the extended version of that code (created by adding a single parity bit) is doubly even, that is, has all codeword weights exactly $0 \pmod{4}$. A doubly even binary code is necessarily self-dual, meaning all its codewords are orthogonal. This implies that any two codewords $\mathbf{c}_{\mathbf{d}}$ and $\mathbf{c}_{\mathbf{e}}$ of the original (non-extended) code have $\mathbf{c}_{\mathbf{d}} \cdot \mathbf{c}_{\mathbf{e}} = 0$ iff either $\mathbf{c}_{\mathbf{d}}$ or $\mathbf{c}_{\mathbf{e}}$ has even parity. Half of our code's words have even parity and $\mathbf{c}_{\mathbf{d}}$ and $\mathbf{c}_{\mathbf{e}}$ are random codewords, so the probability that either of them has even parity is $3/4$. Thus, the probability that $\mathbf{y} \cdot \mathbf{s} = 0$ is $3/4$, proving the fact. □

In the next section, we show that the classical strategy just described can be improved.

3 Algorithm

The classical strategy described in Fact 2.2 above generates vectors that are orthogonal to \mathbf{s} with probability $3/4$. The key to classically defeating the protocol is that it is possible to *correlate* the vectors

¹In [19], $P_{\mathbf{d}, \mathbf{e}}$ is written as $P_{\mathbf{d}} \cap P_{\mathbf{e}}$.

generated by that strategy, such that there is a non-negligible probability of generating a large set of vectors that *all* are orthogonal to \mathbf{s} . These vectors form a system of linear equations that can be solved to yield \mathbf{s} . Finally, with knowledge of \mathbf{s} it is trivial to generate samples that pass the verifier's test.

We follow a modified version of the classical strategy of Fact 2.2 to generate each vector in the correlated set. Crucially, instead of choosing random bitstrings for both \mathbf{d} and \mathbf{e} each time, we generate a single random bitstring \mathbf{d} and hold it constant, only choosing new random values for \mathbf{e} with each iteration. If the encoding $\mathbf{c}_{\mathbf{d}}$ of \mathbf{d} under $P_{\mathbf{s}}$ has even parity, *all* of the generated vectors \mathbf{m}_i will have $\mathbf{m}_i \cdot \mathbf{s} = 0$ (see Theorem 3.1 below). This occurs with probability $1/2$ over our choice of \mathbf{d} .

In practice, it is more convenient to do the linear solve if all $\mathbf{m}_i \cdot \mathbf{s} = 1$ instead of 0. This can be easily accomplished by adding to each \mathbf{m}_i a vector \mathbf{m}^* with $\mathbf{m}^* \cdot \mathbf{s} = 1$. It turns out that $\mathbf{m}^* = \sum_{\mathbf{p} \in \text{rows}(P)} \mathbf{p}$ has this property; see proof of Theorem 3.1.

The explicit algorithm for extracting the vector \mathbf{s} is given in Algorithm 1.

3.1 Analysis

In this section we present a theorem and an empirical claim which demonstrate together that Algorithm 1 can be used to efficiently extract the key from any X -program constructed according to the protocol described in Section 2. The theorem shows that with probability $1/2$ a single iteration of the algorithm finds the vector \mathbf{s} . The empirical claim is that Algorithm 1 is efficient.

Theorem 3.1. *On input an X -program P containing a unique submatrix $P_{\mathbf{s}}$ with the properties described in Section 2, a single iteration of Algorithm 1 will output the vector \mathbf{s} corresponding to $P_{\mathbf{s}}$ with probability $\frac{1}{2}$.*

Proof. If \mathbf{s} is contained in the set $\{\mathbf{s}_i\}$ generated in step 4 of the algorithm, the correct vector \mathbf{s} will be output via the check in step 5 because there is a unique submatrix $P_{\mathbf{s}}$ with codewords having $\text{wt}(\mathbf{c}) \in \{-1, 0\} \pmod{4}$. \mathbf{s} will be contained in $\{\mathbf{s}_i\}$ as long as M satisfies the equation $M\mathbf{s} = \mathbf{1}$. Thus the proof reduces to showing that $M\mathbf{s} = \mathbf{1}$ with probability $1/2$.

Each row of M is

$$\mathbf{m}_i = \mathbf{m}^* + \bar{\mathbf{m}}_i \quad (3.1)$$

for a vector $\bar{\mathbf{m}}_i$ defined as

$$\bar{\mathbf{m}}_i = \sum_{\substack{\mathbf{p} \in \text{rows}(P) \\ \mathbf{p} \cdot \mathbf{d} = \mathbf{p} \cdot \mathbf{e} = 1}} \mathbf{p} \quad (3.2)$$

Algorithm1 EXTRACTKEY(P)

The algorithm to extract the secret vector \mathbf{s} from an X -program P . n is the number of columns in the X -program, and $\overset{\$}{\leftarrow}$ means “select uniformly from the set.”

1. Let $\mathbf{m}^* = \sum_{\mathbf{p} \in \text{rows}(P)} \mathbf{p}$.
 2. Pick $\mathbf{d} \overset{\$}{\leftarrow} \{0, 1\}^n$.
 3. Generate a large number (say $2n$) of vectors \mathbf{m}_i via the following steps, collecting the results into the rows of a matrix M .
 - (a) Pick $\mathbf{e} \overset{\$}{\leftarrow} \{0, 1\}^n$
 - (b) Let $\mathbf{m}_i = \mathbf{m}^* + \sum_{\substack{\mathbf{p} \in \text{rows}(P) \\ \mathbf{p} \cdot \mathbf{d} = \mathbf{p} \cdot \mathbf{e} = 1}} \mathbf{p}$
 4. Via linear solve, find the set of vectors $\{\mathbf{s}_i\}$ satisfying $M\mathbf{s}_i = \mathbf{1}$, where $\mathbf{1}$ is the vector of all ones.
 5. For each candidate vector \mathbf{s}_i :
 - (a) Extract $P_{\mathbf{s}_i}$ from P by deleting the rows of P orthogonal to \mathbf{s}_i
 - (b) If adding a parity bit to each of the columns \mathbf{c} of $P_{\mathbf{s}_i}$ yields the generator matrix for a code that is doubly even (all basis codewords are doubly even and mutually orthogonal), return \mathbf{s} and exit.
 6. No candidate vector \mathbf{s} was found; return to step 2.
-

Here we will show that $\mathbf{m}^* \cdot \mathbf{s} = 1$ always and $\bar{\mathbf{m}}_i \cdot \mathbf{s} = 0$ for all i with probability $1/2$, implying that $M\mathbf{s} = \mathbf{1}$ with probability $1/2$.

First we show that $\mathbf{m}^* \cdot \mathbf{s} = 1$. \mathbf{m}^* is the sum of all rows of P , so we have

$$\mathbf{m}^* \cdot \mathbf{s} = \sum_{\mathbf{p} \in \text{rows}(P)} \mathbf{p} \cdot \mathbf{s} = \sum_{\mathbf{p} \in \text{rows}(P_s)} 1 \quad (3.3)$$

We see that the inner product is equal to the number of rows in the submatrix $P_s \pmod{2}$. This submatrix is a generator matrix for a code of block size $7 \pmod{8}$; thus the number of rows is odd and

$$\mathbf{m}^* \cdot \mathbf{s} = 1 \quad (3.4)$$

Now we turn to showing that $\bar{\mathbf{m}}_i \cdot \mathbf{s} = 0$ for all i with probability $1/2$. In the proof of Fact 2.2, it was shown that for any two vectors \mathbf{d} and \mathbf{e} , vectors $\bar{\mathbf{m}}_i$ generated by summing rows \mathbf{p}_i of P for which $\mathbf{d} \cdot \mathbf{p}_i = \mathbf{e} \cdot \mathbf{p}_i = 1$ have

$$\bar{\mathbf{m}}_i \cdot \mathbf{s} = 0 \text{ iff } \mathbf{c}_d \text{ or } \mathbf{c}_e \text{ has even parity} \quad (3.5)$$

where \mathbf{c}_d and \mathbf{c}_e are the encodings under P_s of \mathbf{d} and \mathbf{e} respectively. If \mathbf{d} is held constant for all i , and \mathbf{d} happened to be chosen such that $\mathbf{c}_d = P_s \mathbf{d}$ has even parity, then $\bar{\mathbf{m}}_i \cdot \mathbf{s} = 0$ for all i by Equation 3.5. Because half of the codewords have even parity, for \mathbf{d} selected uniformly at random we have $\bar{\mathbf{m}}_i \cdot \mathbf{s} = 0$ for all i with probability $1/2$.

We have shown that $\mathbf{m}^* \cdot \mathbf{s} = 1$ always and $\bar{\mathbf{m}}_i \cdot \mathbf{s} = 0$ for all i with probability $1/2$. Therefore we have

$$\Pr_{\mathbf{d}} [\mathbf{m}_i \cdot \mathbf{s} = 1 \ \forall i] = 1/2$$

Thus $M\mathbf{s} = \mathbf{1}$ with probability $1/2$. The algorithm will output \mathbf{s} whenever $M\mathbf{s} = \mathbf{1}$, proving the theorem. \square

Before we move on, we remark that while Theorem 3.1 treats X -programs containing a single unique submatrix with the relevant properties, the algorithm can easily be modified to return the vectors \mathbf{s} corresponding to *all* such submatrices, if more exist, by simply accumulating all vectors \mathbf{s} for which the check in Step 5(b) succeeds. We do note, however, that for the protocol described in Section 2, the probability of “extraneous” submatrices other than the one intentionally built into the matrix arising by chance is vanishingly small—corresponding to the probability that a random binary linear code happens to be doubly even and self-dual, which is bounded from above by $1/4^n$.

Now, having established that each iteration of the algorithm outputs \mathbf{s} with probability $1/2$, we now turn to analyzing its runtime.




Figure 2: **(a)** The average number of candidate vectors checked before the secret vector \mathbf{s} was found, when the algorithm was applied to 1000 unique X -programs at each problem size tested. We observe that the number of vectors to check is qualitatively constant in n . **(b)** The number of unconstrained degrees of freedom $n - \text{rank}(M)$ for matrices M generated in step 3 of Algorithm 1, for “good” choices of \mathbf{d} such that $M\mathbf{s} = \mathbf{1}$. The rapidly decaying tail qualitatively implies that it is rare for any more than a few degrees of freedom to remain unconstrained. The blue bars represent the distribution over 1000 unique X -programs of size $n = 245$. The algorithm was then re-run on the X -programs that had $n - \text{rank}(M) > 4$ to generate the orange bars.

Claim 3.1. (empirical) Algorithm 1 halts in $\mathcal{O}(n^3)$ time on average.

All steps of the algorithm except for step 5 have $\mathcal{O}(n^3)$ scaling by inspection. The obstacle preventing Claim 3.1 from trivially holding is that it is hard to make a rigorous statement about how large the set of candidate vectors $\{\mathbf{s}_i\}$ is. Because $|\{\mathbf{s}_i\}| = 2^{n - \text{rank}(M)}$, we’d like to show that on average, the rank of M is close to or equal to n . It seems reasonable that this would be the case: we are generating the rows of M by summing rows from P , and P must have full rank because it contains a rank- n error correcting code. But the rows of P summed into each \mathbf{m}_i are not selected independently—they are always related via their connection to the vectors \mathbf{d} and \mathbf{e} , and it’s not clear how these correlations affect the linear independence of the resulting \mathbf{m}_i .

Despite the lack of a proof, empirical evidence

supports Claim 3.1 when the algorithm is applied to X -programs generated in the manner described in Section 2. Figure 2(a) shows the average number of candidate keys checked by the algorithm before \mathbf{s} is found, as a function of problem size. The value is constant, demonstrating that the average size of the set $\{\mathbf{s}_i\}$ does not scale with n . Furthermore, the value is small—only about 4. This implies that M usually has high rank. In Figure 2(b) we plot explicitly the distribution of the rank of the matrix M over 1000 runs of the algorithm on unique X -programs of size $n = 245$. The blue bars (on the left of each pair) show the distribution over all X -programs tested, and the sharply decaying tail supports the claim that low-rank M almost never occur.

A natural next question is whether there is some feature of the X -programs in that tail that causes M to be low rank. To investigate that question, the algorithm was re-run 100 times on each of the X -programs that had $n - \text{rank}(M) > 4$ in the blue distribution. The orange bars of Figure 2(b) (on the right of each pair) plot the distribution of $n - \text{rank}(M)$ for that second run. The similarity of the blue and orange distributions suggests that the rank of M is not correlated between runs; that is, the low rank of M in the first run was not due to any feature of the input X -programs. From a practical perspective, this data suggests that if the rank of M is found to be unacceptably low, the algorithm can simply be re-run with new randomness and the rank of M is likely to be higher the second time.

3.2 Implementation

An implementation of Algorithm 1 in the programming language Julia (along with the code to generate the figures in this manuscript) is available online [22]. Figure 1 shows the runtime of this implementation for various problem sizes. Experiments were completed using one thread on an Intel 8268 “Cascade Lake” processor.

Note that Figure 1 shows $\mathcal{O}(n^2)$ scaling, rather than $\mathcal{O}(n^3)$ from Claim 3.1. This is due to data-level parallelism in the implementation. \mathbb{Z}_2^n vectors are stored as the bits of 64-bit integers, so operations like vector addition can be performed on 64 elements at once via bitwise operations. Furthermore, with AVX SIMD CPU instructions, those operations can be applied to multiple 64-bit integers in one CPU cycle. Thus, for n of order 100, the ostensibly $\mathcal{O}(n)$ vector inner products and vector sums are performed in constant time, removing one factor of n from the runtime. The tests in Figure 1 were performed on a CPU with 512 bit vector units.

4 Discussion

Modifications to the protocol A natural question is whether it is possible to modify the original protocol such that this attack is not successful. Perhaps P can be engineered such that either 1) it is not possible to generate a large number of vectors that all have a known inner product with \mathbf{s} , or 2) the rank of the matrix M formed by these generated vectors will never be sufficiently high to allow solution of the linear system.

For 1), our ability to generate many vectors orthogonal to \mathbf{s} relies on the fact that the code generated by the hidden submatrix $P_{\mathbf{s}}$ has codewords \mathbf{c} with $\text{wt}(\mathbf{c}) \in \{-1, 0\} \pmod{4}$, as shown in the proof of Theorem 3.1. Unfortunately, this property regarding the weights of the codewords is precisely what gives the quantum sampling algorithm its bias toward generating vectors with $\mathbf{x} \cdot \mathbf{s} = 0$ (see Fact 2.1). This fact seems to preclude the possibility of simply removing the special property of the submatrix $P_{\mathbf{s}}$ to prevent the attack.

For 2), the main obstacle is that the matrix P *must* have rank n because embedded in it is a code of rank n . The only hope is to somehow engineer the matrix such that linear combinations generated in the specific way described above will not themselves be linearly independent. It is not at all clear how one would do that, and furthermore, adding structure to the previously-random extra rows of P runs the risk of providing even more information about the secret vector \mathbf{s} . Perhaps one could *prove* that the rank of M will be large even for worst-case inputs P —this could be an interesting future direction.

Protocols with provable hardness The attack described in this paper reiterates the value of building protocols for which passing the test itself, rather than just simulating the quantum device, can be shown to be hard under well-established cryptographic assumptions. In the past few years, a number of new trapdoor claw-free function based constructions have been proposed for demonstrating quantum computational advantage [16–18, 23], as well as some based on other types of cryptography [24, 25]. Unfortunately, such rigorous results come with a downside, which is an increase in the size and complexity of circuits that must be run on the quantum device. Exploring simplified protocols that are provably secure is an exciting area for further research.

The \$25 challenge When the protocol was first proposed in [19], it was accompanied by an internet challenge. The authors posted a specific instance of the matrix P , and offered \$25 to anyone who could send them samples passing the verifier’s check. The secret

vector s corresponding to their challenge matrix P is (encoded as a base-64 string):

```
BilbHzjYxrOHYH401EJFB0XZbps4a54kH8flrRgo/g==
```

The key was extracted using the implementation of Algorithm 1 described in Section 3.2.

Shepherd and Bremner, the authors of the challenge, have graciously confirmed that this indeed is the correct key.

Summary and outlook Here we have described a classical algorithm that passes the interactive quantum test described in [19]. We have proven that a single iteration of the algorithm will return the underlying secret vector with probability $1/2$, and empirically shown that it is efficient. The immediate implication of this result is that the protocol in its original form is no longer effective as a test of quantum computational power. While it may be possible to reengineer that protocol to thwart this attack, this paper reiterates the value of proving the security of the verification step. Furthermore, while protocols for quantum advantage with provable classical hardness are valuable in their own right, they can also be used as building blocks for achieving new, more complex cryptographic tasks, like certifiable random number generation, secure remote state preparation, and even the verification of arbitrary quantum computations [16, 26, 27]. As quantum hardware continues to improve and to surpass the abilities of classical machines, quantum cryptographic tools will play an important role in making quantum computation available as a service. Establishing the security of these protocols is an important first step.

Acknowledgements The author was supported in this work by the National Defense Science and Engineering Graduate Fellowship (NDSEG).

References

- [1] Frank Arute et al. “Quantum supremacy using a programmable superconducting processor”. *Nature* **574**, 505–510 (2019).
- [2] Han-Sen Zhong et al. “Quantum computational advantage using photons”. *Science* **370**, 1460–1463 (2020).
- [3] Yulin Wu, Wan-Su Bao, Sirui Cao, Fusheng Chen, Ming-Cheng Chen, Xiawei Chen, Tung-Hsun Chung, Hui Deng, Yajie Du, Daojin Fan, Ming Gong, Cheng Guo, Chu Guo, Shaojun Guo, Lianchen Han, Linyin Hong, He-Liang Huang, Yong-Heng Huo, Liping Li, Na Li, Shaowei Li, Yuan Li, Futian Liang, Chun Lin, Jin Lin, Haoran Qian, Dan Qiao, Hao Rong, Hong Su, Lihua Sun, Liangyuan Wang, Shiyu Wang, Dachao Wu, Yu Xu, Kai Yan, Weifeng Yang, Yang Yang, Yangsen Ye, Jianghan Yin, Chong Ying, Jiale Yu, Chen Zha, Cha Zhang, Haibin Zhang, Kaili Zhang, Yiming Zhang, Han Zhao, Youwei Zhao, Liang Zhou, Qingling Zhu, Chao-Yang Lu, Cheng-Zhi Peng, Xiaobo Zhu, and Jian-Wei Pan. “Strong Quantum Computational Advantage Using a Superconducting Quantum Processor”. *Physical Review Letters* **127**, 180501 (2021).
- [4] Qingling Zhu, Sirui Cao, Fusheng Chen, Ming-Cheng Chen, Xiawei Chen, Tung-Hsun Chung, Hui Deng, Yajie Du, Daojin Fan, Ming Gong, Cheng Guo, Chu Guo, Shaojun Guo, Lianchen Han, Linyin Hong, He-Liang Huang, Yong-Heng Huo, Liping Li, Na Li, Shaowei Li, Yuan Li, Futian Liang, Chun Lin, Jin Lin, Haoran Qian, Dan Qiao, Hao Rong, Hong Su, Lihua Sun, Liangyuan Wang, Shiyu Wang, Dachao Wu, Yulin Wu, Yu Xu, Kai Yan, Weifeng Yang, Yang Yang, Yangsen Ye, Jianghan Yin, Chong Ying, Jiale Yu, Chen Zha, Cha Zhang, Haibin Zhang, Kaili Zhang, Yiming Zhang, Han Zhao, Youwei Zhao, Liang Zhou, Chao-Yang Lu, Cheng-Zhi Peng, Xiaobo Zhu, and Jian-Wei Pan. “Quantum computational advantage via 60-qubit 24-cycle random circuit sampling”. *Science Bulletin* **67**, 240–245 (2022).
- [5] Scott Aaronson and Alex Arkhipov. “The computational complexity of linear optics”. In Proceedings of the forty-third annual ACM symposium on Theory of computing. Pages 333–342. STOC ’11New York, NY, USA (2011). Association for Computing Machinery.
- [6] Edward Farhi and Aram W. Harrow. “Quantum Supremacy through the Quantum Approximate Optimization Algorithm” (2019). [arXiv:1602.07674](https://arxiv.org/abs/1602.07674).
- [7] A. P. Lund, Michael J. Bremner, and T. C. Ralph. “Quantum sampling problems, BosonSampling and quantum supremacy”. *npj Quantum Information* **3**, 1–8 (2017).
- [8] Aram W. Harrow and Ashley Montanaro. “Quantum computational supremacy”. *Nature* **549**, 203–209 (2017).
- [9] Sergio Boixo, Sergei V. Isakov, Vadim N. Smelyanskiy, Ryan Babbush, Nan Ding, Zhang Jiang, Michael J. Bremner, John M. Martinis, and Hartmut Neven. “Characterizing quantum supremacy in near-term devices”. *Nature Physics* **14**, 595–600 (2018).
- [10] Barbara M. Terhal. “Quantum supremacy, here we come”. *Nature Physics* **14**, 530–531 (2018).
- [11] C. Neill, P. Roushan, K. Kechedzhi, S. Boixo,

- S. V. Isakov, V. Smelyanskiy, A. Megrant, B. Chiaro, A. Dunsworth, K. Arya, R. Barends, B. Burkett, Y. Chen, Z. Chen, et al. “A blueprint for demonstrating quantum supremacy with superconducting qubits”. *Science* **360**, 195–199 (2018).
- [12] Sergey Bravyi, David Gosset, and Robert König. “Quantum advantage with shallow circuits”. *Science* **362**, 308–311 (2018).
- [13] Sergey Bravyi, David Gosset, Robert König, and Marco Tomamichel. “Quantum advantage with noisy shallow circuits”. *Nature Physics* **16**, 1040–1045 (2020).
- [14] Adam Bouland, Bill Fefferman, Chinmay Nirkhe, and Umesh Vazirani. “On the complexity and verification of quantum random circuit sampling”. *Nature Physics* **15**, 159–163 (2019).
- [15] Scott Aaronson and Sam Gunn. “On the Classical Hardness of Spoofing Linear Cross-Entropy Benchmarking”. *Theory of Computing* **16**, 1–8 (2020).
- [16] Zvika Brakerski, Paul Christiano, Urmila Mahadev, Umesh Vazirani, and Thomas Vidick. “A Cryptographic Test of Quantumness and Certifiable Randomness from a Single Quantum Device”. *Journal of the ACM (JACM)* (2021).
- [17] Zvika Brakerski, Venkata Koppula, Umesh Vazirani, and Thomas Vidick. “Simpler Proofs of Quantumness”. In Steven T. Flammia, editor, 15th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2020). Volume 158 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 8:1–8:14. Dagstuhl, Germany (2020). Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.
- [18] Gregory D. Kahanamoku-Meyer, Soonwon Choi, Umesh V. Vazirani, and Norman Y. Yao. “Classically-Verifiable Quantum Advantage from a Computational Bell Test”. *Nature Physics* **18**, 918–924 (2022).
- [19] Dan Shepherd and Michael J. Bremner. “Temporally unstructured quantum computation”. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* **465**, 1413–1439 (2009).
- [20] Michael J. Bremner et al. “Classical simulation of commuting quantum computations implies collapse of the polynomial hierarchy”. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* **467**, 459–472 (2011).
- [21] Michael J. Bremner et al. “Average-Case Complexity Versus Approximate Simulation of Commuting Quantum Computations”. *Physical Review Letters* **117**, 080501 (2016).
- [22] Gregory D. Kahanamoku-Meyer (2023). code: <https://doi.org/10.5281/zenodo.7545881>.
- [23] Yusuf Alnawakhtha, Atul Mantri, Carl A. Miller, and Daochen Wang. “Lattice-Based Quantum Advantage from Rotated Measurements” (2022). [arXiv:2210.10143](https://arxiv.org/abs/2210.10143).
- [24] Takashi Yamakawa and Mark Zhandry. “Verifiable Quantum Advantage without Structure”. In 2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS). Pages 69–74. (2022).
- [25] Yael Kalai, Alex Lombardi, Vinod Vaikuntanathan, and Lisa Yang. “Quantum Advantage from Any Non-local Game”. In Proceedings of the 55th Annual ACM Symposium on Theory of Computing. Pages 1617–1628. STOC 2023 New York, NY, USA (2023). Association for Computing Machinery.
- [26] Alexandru Gheorghiu and Thomas Vidick. “Computationally-Secure and Composable Remote State Preparation”. In 2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS). Pages 1024–1033. (2019).
- [27] Urmila Mahadev. “Classical Verification of Quantum Computations”. In 2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS). Pages 259–267. (2018).