Classically-Verifiable Quantum Advantage from a Computational Bell Test

Gregory D. Kahanamoku-Meyer,[1] Soonwon Choi,[1] Umesh V. Vazirani,[2] and Norman Y. Yao[1]

[1]*Department of Physics, University of California at Berkeley, Berkeley, CA 94720*

[2]*Department of Electrical Engineering and Computer Science, University of California at Berkeley, Berkeley, CA 94720*

We propose and analyze a novel interactive protocol for demonstrating quantum computational advantage, which is efficiently classically verifiable. Our protocol relies upon the cryptographic hardness of trapdoor claw-free functions (TCFs). Through a surprising connection to Bell's inequality, our protocol avoids the need for an adaptive hardcore bit, with essentially no increase in the quantum circuit complexity and no extra cryptographic assumptions. Crucially, this expands the set of compatible TCFs, and we propose two new constructions: one based upon the decisional Diffie-Hellman problem and the other based upon Rabin's function, $x^2 \bmod N$. We also describe two unique features of our interactive protocol: (i) it allows one to discard so-called "garbage bits", thereby removing the need for reversibility in the quantum circuits, and (ii) there exists a natural post-selection scheme, which significantly reduces the fidelity needed to demonstrate quantum advantage. Finally, we design several efficient circuits for $x^2 \bmod N$ and describe a blueprint for their implementation on a Rydberg-atom-based quantum computer.

I. INTRODUCTION

The development of large-scale programmable quantum hardware has opened the door to testing a fundamental question in the theory of computation: can quantum computers outperform classical ones for certain tasks? This idea, termed quantum computational advantage, has motivated the design of novel algorithms and protocols to demonstrate advantage with minimal quantum resources such as qubit number and gate depth [1–10]. Such protocols are naturally characterized along two axes: the computational speedup and the ease of verification. The former distinguishes whether a quantum algorithm exhibits a polynomial or super-polynomial speedup over the best known classical one. The latter classifies whether the correctness of the quantum computation is *efficiently verifiable* by a classical computer. Along these axes lie three broad paths to demonstrating advantage: 1) sampling from entangled quantum many-body wavefunctions, 2) solving a deterministic problem, e.g. prime factorization, via a quantum algorithm, and 3) proving quantumness through interactive protocols.

Sampling-based protocols directly rely on the classical hardness of simulating quantum mechanics [1, 3, 7–10]. The "computational task" is to prepare and measure a generic complex many-body wavefunction with little structure. As such, these protocols typically require minimal resources and can be implemented on near-term quantum devices [11, 12]. The correctness of the sampling results, however, is exponentially difficult to verify. This has an important consequence: in the regime beyond the capability of classical computers, the sampling results cannot be explicitly checked, and quantum computational advantage can only be inferred (e.g. extrapolated from simpler circuits).

Algorithms in the second class of protocols are naturally broken down by whether they exhibit polynomial or super-polynomial speed-ups. In the case of polynomial speed-ups, there exist notable examples that are *provably* faster than any possible classical algorithm [13, 14]. However, polynomial speed-ups are tremendously challenging to demonstrate in practice, due to the slow growth of the separation between classical and quantum run-times [15]. Accordingly, the most attractive algorithms for demonstrating advantage tend to be those with a super-polynomial speed-up, including Abelian hidden subgroup problems such as factoring and discrete logarithms [16]. The challenge is that for all known protocols of this type, the quantum circuits required to demonstrate advantage are well beyond the capabilities of near-term experiments.

The final class of protocols demonstrates quantum advantage through an *interactive proof* [17–24]. At a high level, this type of protocol involves multiple rounds of communication between the classical verifier and the quantum prover; the prover must give self-consistent responses despite not knowing what the verifier will ask next. This requirement of self-consistency rules out a broad range of classical cheating strategies and can imbue "hardness" into questions that would otherwise be easy to answer. To this end, interactive protocols expand the space of computational problems that can be used to demonstrate quantum advantage; from a more pragmatic perspective, this can enable the realization of efficiently verifiable quantum advantage on near-term quantum hardware.

Recently, a beautiful interactive protocol was introduced that can operate both as a test for quantum advantage and as a generator of certifiable quantum randomness [17]. The core of the protocol is a two-to-one function, f, built on the computational problem known as learning with errors (LWE) [25]. The demonstration of advantage leverages two important properties of the function: first, it is *claw-free*, meaning that it is computationally hard to find a pair of inputs (x_0, x_1) such that $f(x_0) = f(x_1)$. Second, there exists a *trapdoor*: given some secret data t, it becomes possible to efficiently invert f and reveal the pair of inputs mapping to any out-

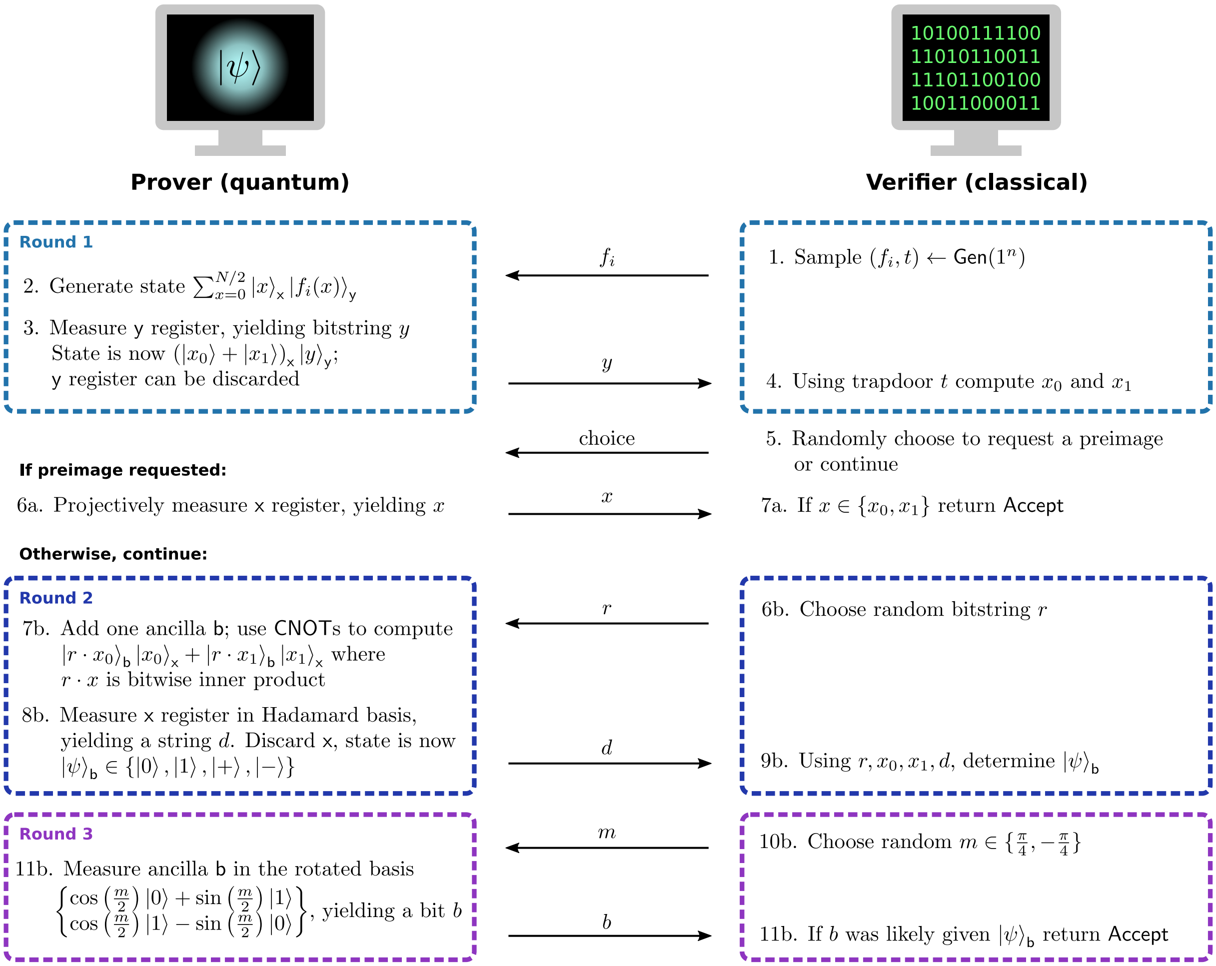

FIG. 1. Schematic representation of the interactive quantum advantage protocol. In the first round of interaction, the classical verifier (right) selects a specific function from a trapdoor claw-free family and the quantum prover (left) evaluates it over a superposition of inputs. The goal of the second round is to condense the information contained in the prover's superposition state onto a single ancilla qubit. The final round of interaction effectively performs a Bell inequality measurement, whose outcome is cryptographically protected.

put. However, to fully protect against cheating provers, the protocol requires a stronger version of the claw-free property called the *adaptive hardcore bit*, namely, that for any input x_0 (which may be chosen by the prover), it is computationally hard to find even a single bit of information about x_1 [26]. The need for an adaptive hardcore bit within this protocol severely restricts the class of functions that can operate as verifiable tests of quantum advantage.

Here, we propose and analyze a novel interactive quantum advantage protocol that removes the need for an adaptive hardcore bit, with essentially zero overhead in the quantum circuit and no extra cryptographic assumptions [27]. We present four main results. First, we demonstrate how an idea from tests of Bell's inequality can serve the same cryptographic purpose as the adaptive hardcore bit [28]. In essence, our interactive protocol is a variant of the CHSH (Clauser, Horne, Shimony, Holt) game [29] in which one player is replaced by a cryptographic construction. Normally, in CHSH, two quantum parties are asked to produce correlations that would be impossible for classical devices to produce. If space-like separation is enforced to rule out communication between the two parties, then the correlations constitute a proof of quantumness. In our case, the space-like separation is replaced by the computational hardness of a cryptographic problem. In particular, the quantum prover holds a qubit whose state depends on the cryptographic secret in the same way that the state of one CHSH player's qubit depends on the secret measurement basis of the other player. As in CHSH, a quantum device can pass the verifier's test with probability $\sim 85\%$, but a classical device can only succeed with probability at most 75%. This finite gap in success probabilities is precisely what enables a verifiable test of quantum advantage.

Second, by removing the need for an adaptive hardcore bit, our protocol broadens the landscape of functions that can be used for interactive tests of quantum advantage (see Table I and Methods). We populate this list with two new constructions. The first is based on the decisional Diffie-Hellman problem (DDH) [30–32]; the second utilizes the function $f_N(x) = x^2 \bmod N$ with N the product of two primes, which forms the backbone of the Rabin cryptosystem [33, 34]. On the one hand, DDH is appealing because the elliptic-curve version of the problem is particularly hard for classical computers [35–37]. On the other hand, $x^2 \bmod N$ can be implemented significantly more efficiently, while its hardness is equivalent to factoring. We hope that these two constructions will provide a foundation for the search for more TCFs with desirable properties (small key size and efficient quantum circuits).

Third, we describe two key features of our protocol that facilitate its use in practice: a way to significantly reduce overhead arising from the reversibility requirement of quantum circuits, and a scheme for increasing noisy devices' probability of passing the test. Normally, quantum implementations of classical functions like the TCFs used in this protocol have some overhead, due to the need to make the circuit reversible in order to be consistent with unitarity [38–42]. Our protocol exhibits the surprising property that it permit a measurement scheme to discard so-called "garbage bits" that arise during the computation, allowing classical circuits to be converted into quantum ones with essentially zero overhead. In the case of a noisy quantum device, the protocol also enables an inherent post-selection scheme for detecting and removing certain types of quantum errors. With this scheme it is possible for quantum devices to trade off low quantum fidelities for an increase in the overall runtime, while still passing the cryptographic test.

Finally, focusing on the TCF $x^2 \bmod N$, we provide explicit quantum circuits—both asymptotically optimal (requiring only $\mathcal{O}(n \log n)$ gates and $\mathcal{O}(n)$ qubits), as well as those aimed for near-term quantum devices. We show that a verifiable test of quantum advantage can be achieved with $\sim 10^3$ qubits and a gate depth $\sim 10^5$ (see Methods). We also co-design a specific implementation of $x^2 \bmod N$ optimized for a programmable Rydberg-based quantum computing platform. The native physical interaction corresponding to the Rydberg blockade mechanism enables the direct implementation of multi-qubit-controlled arbitrary phase rotations without the need to decompose such gates into universal two-qubit operations [43–47]. Access to such a native gate immediately reduces the gate depth for achieving quantum advantage by an order of magnitude.

II. INTERACTIVE PROTOCOL FOR QUANTUM ADVANTAGE

Our full protocol is shown diagrammatically in Figure 1. It consists of three rounds of interaction between the prover and verifier (with a "round" being a challenge from the verifier, followed by a response from the prover). The first round generates a multi-qubit superposition over two bit strings that would be cryptographically hard to compute classically. The second round maps this superposition onto the state of one ancilla qubit, retaining enough information to ensure that the resulting single-qubit state is also hard to compute classically. The third round takes this single qubit as input to a CHSH-type measurement, allowing the prover to generate a bit of data that is correlated with the cryptographic secret in a way that would not be possible classically. Having described the intuition behind the protocol, we now lay out each round in detail.

Problem	Trap door	Claw-free	Adaptive hard-core bit	Asymptotic complexity (gate count)
LWE [17]	✓	✓	✓	$n^2 \log^2 n$
$x^2 \bmod N$	✓	✓	✗	$n \log n$
Ring-LWE [18]	✓	✓	✗	$n \log^2 n$
Diffie-Hellman	✓	✓	✗	$n^3 \log^2 n$
Shor's alg.	—	—	—	$n^2 \log n$

TABLE I. Cryptographic constructions for interactive quantum advantage protocols. Big-$\mathcal{O}$ notation is implied and factors of $\log \log n$ and smaller are dropped. For references and derivations of the circuit complexities, see supplementary information [48].

A. Description of the protocol

The goal of the first round is to generate a superposition over two colliding inputs to the trapdoor claw-free function (TCF). It begins with the verifier choosing an instance f_i of the TCF along with the associated trapdoor data t; f_i is sent to the prover. As an example, in the case of $x^2 \bmod N$, the "index" i is the modulus N, and the trapdoor data is its factorization, p, q. The prover now initializes two registers of qubits, which we denote as the x and y registers. On these registers, they compute the entangled superposition $|\psi\rangle = \sum_x |x\rangle_\mathsf{x} |f_i(x)\rangle_\mathsf{y}$, over all x in the domain of f_i. The prover then measures the y register in the standard basis, collapsing the state to $(|x_0\rangle + |x_1\rangle)_\mathsf{x} |y\rangle_\mathsf{y}$, with $y = f(x_0) = f(x_1)$. The measured bitstring y is then sent to the verifier, who uses the secret trapdoor to compute x_0 and x_1 in full.

At this point, the verifier randomly chooses to either request a projective measurement of the x register, ending the protocol, or to continue with the second and third rounds. In the former case, the prover communicates the result of that measurement, yielding either x_0 or x_1, and the verifier checks that indeed $f(x) = y$. In the latter case, the protocol proceeds with the final two rounds.

The second round of interaction converts the many-qubit superposition $|\psi\rangle = |x_0\rangle_\mathsf{x} + |x_1\rangle_\mathsf{x}$ into a single-qubit state $\{|0\rangle_{\mathsf{b}}, |1\rangle_{\mathsf{b}}, |+\rangle_{\mathsf{b}}, |-\rangle_{\mathsf{b}}\}$ on an ancilla qubit b. The final state of b depends on the values of both x_0 and x_1. The round begins with the verifier choosing a random bitstring r of the same length as x_0 and x_1, and sending it to the prover. Using a series of CNOT gates from the x register to b, the prover computes the state $|r \cdot x_0\rangle_{\mathsf{b}} |x_0\rangle_{\mathsf{x}} + |r \cdot x_1\rangle_{\mathsf{b}} |x_1\rangle_{\mathsf{x}}$, where $r \cdot x$ denotes the binary inner product. Finally, the prover measures the x register in the Hadamard basis, storing the result as a bitstring d which is sent to the verifier. This measurement disentangles x from b without collapsing b's superposition. At the end of the second round, the prover's state is $(-1)^{d \cdot x_0} |r \cdot x_0\rangle_{\mathsf{b}} + (-1)^{d \cdot x_1} |r \cdot x_1\rangle_{\mathsf{b}}$, which is one of $\{|0\rangle, |1\rangle, |+\rangle, |-\rangle\}$. Crucially, it is cryptographically hard to predict whether this state is one of $\{|0\rangle, |1\rangle\}$ or $\{|+\rangle, |-\rangle\}$.

The final round of our protocol can be understood in analogy to the CHSH game [29]. While the prover cannot extract the polarization axis from their single qubit (echoing the no-signaling property of CHSH), they can make a measurement that is *correlated* with it. This measurement outcome ultimately constitutes the proof of quantumness. In particular, the verifier requests a measurement in an intermediate basis, rotated from the Z axis around Y, by either $m = \pi/4$ or $-\pi/4$. Because the measurement basis is never perpendicular to the state, there will always be one outcome that is more likely than the other (specifically, with probability $\cos^2(\pi/8) \approx 0.85$). The verifier returns Accept if this "more likely" outcome is the one received.

In the next section, we prove that a quantum device can cause the verifier to Accept with substantially higher probability than any classical prover. A full test of quantum advantage would consist of running the protocol many times, until it can be established with high statistical confidence that the device has exceeded the classical probability bound.

B. Completeness and soundness

We now prove completeness (the noise-free quantum success probability) and soundness (an upper bound on the classical success probability). Recall that after the first round of the protocol, the verifier chooses to either request a standard basis measurement of the first register, or to continue with the second and third rounds. In the proofs below, we analyze the prover's success probability across these two cases separately. We denote the probability that the verifier will accept the prover's string x in the first case as p_x, and the probability that the verifier will accept the single-qubit measurement result in the second case as p_m.

1. Perfect quantum prover (completeness)

Theorem 1. *An error-free quantum device honestly following the interactive protocol will cause the verifier to return* Accept *with* $p_x = 1$ *and* $p_m = \cos^2(\pi/8) \approx 0.85$.

Proof. If the verifier chooses to request a projective measurement of x after the first round, an honest quantum prover succeeds with probability $p_x = 1$ by inspection.

If the verifier chooses to instead perform the rest of the protocol, the prover will hold one of $\{|0\rangle, |1\rangle, |+\rangle, |-\rangle\}$ after round 2. In either measurement basis the verifier may request in round 3, there will be one outcome that occurs with probability $\cos^2(\pi/8)$, which is by construction the one the verifier accepts. Thus, an honest quantum prover has $p_m = \cos^2(\pi/8) \approx 0.85$. □

2. Classical prover (soundness)

Theorem 2. *Assume the function family used in the interactive protocol is claw-free. Then, p_x and p_m for any classical prover must obey the relation*

$$p_x + 4p_m - 4 < \epsilon(n) \tag{1}$$

where ϵ is a negligible function of n, the security parameter (i.e. problem size).

Proof. We prove by contradiction. Assume that there exists a classical machine $\mathcal{A}$ for which $p_x + 4p_m - 4 \geq \mu(n)$, for a non-negligible function μ. We show that there exists another algorithm $\mathcal{B}$ that uses $\mathcal{A}$ as a subroutine to find a pair of colliding inputs to the claw-free function, a contradiction.

Given a claw-free function instance f_i, $\mathcal{B}$ acts as a simulated verifier for $\mathcal{A}$. $\mathcal{B}$ begins by supplying f_i to $\mathcal{A}$, after which $\mathcal{A}$ returns a value y, completing the first round of interaction. $\mathcal{B}$ now chooses to request the projective measurement of the x register, and stores the result as x_0. This will be a valid preimage with probability p_x.

Next, $\mathcal{B}$ *rewinds* the execution of $\mathcal{A}$, to its state before x_0 was requested. Crucially, rewinding is possible because $\mathcal{A}$ is a classical algorithm. $\mathcal{B}$ now proceeds by running $\mathcal{A}$ through the second and third rounds of the protocol for many different values of the bitstring r (Fig. 1), rewinding each time.

We now show that, for r selected uniformly at random, $\mathcal{B}$ can extract the value of the inner product $r \cdot x_1$ with probability $p_{r \cdot x_1} \geq 1 - 2(1 - p_m)$. $\mathcal{B}$ begins by sending r to $\mathcal{A}$, and receiving the bitstring d. $\mathcal{B}$ then requests the measurement result in both the $m = \pi/4$ and $m = -\pi/4$ bases, by rewinding in between. If both the received values are "correct" (i.e. would be accepted by the real verifier), they uniquely determine the single-qubit state $|\psi\rangle \in \{|0\rangle, |1\rangle, |+\rangle, |-\rangle\}$ that would be held by an honest quantum prover. This state reveals whether $r \cdot x_0 = r \cdot x_1$, and because $\mathcal{B}$ already holds x_0, $\mathcal{B}$ can compute $r \cdot x_1$. By definition, each "measurement" from the prover is correct with probability p_m. Thus, via union bound, the probability that both are correct is $p_{r\cdot x_1} \geq 1-2(1-p_m)$.

Now, we show that extracting $r \cdot x_1$ in this way allows x_1 to be determined in full even in the presence of noise, by rewinding many times and querying for specific (correlated) choices of r. In particular, the above construction is a noisy oracle to the encoding of x_1 under the Hadamard code. By the Goldreich-Levin theorem [49], list decoding applied to such an oracle will generate a polynomial-length list of candidates for x_1. If the noise rate of the oracle is noticeably less than 1/2, x_1 will be contained in that list; $\mathcal{B}$ can iterate through the candidates until it finds one for which $f(x_1) = y$.

By Lemma 1 in the Methods, for a particular iteration of the protocol, the probability that list decoding succeeds is bounded by $p_{x_1} > 2p_{r\cdot x_1} - 1 - 2\mu'(n)$, for a noticeable function $\mu'(n)$ of our choice [50]. Setting $\mu'(n) = \mu(n)/4$ and combining with the previous result yields $p_{x_1} > 1 - 4(1-p_m) - \mu(n)/2$.

Finally, via union bound, the probability that $\mathcal{B}$ returns a claw is

$$P_{\mathcal{B}} \geq 1 - (1-p_{x_0}) - (1-p_{x_1}) > p_x + 4p_m - 4 - \mu(n)/2$$

and via the assumption that $p_x + 4p_m - 4 > \mu(n)$ we have

$$P_{\mathcal{B}} > \mu(n)/2$$

a contradiction. □

If we let $p_x = 1$, the bound requires that $p_m < 3/4 + \epsilon(n)$ for a classical device, while $p_m \approx 0.85$ for a quantum device, matching the classical and quantum success probabilities of the CHSH game. In the supplementary information [48], we provide an example of a classical algorithm saturating the bound with $p_x = 1$ and $p_m = 3/4$.

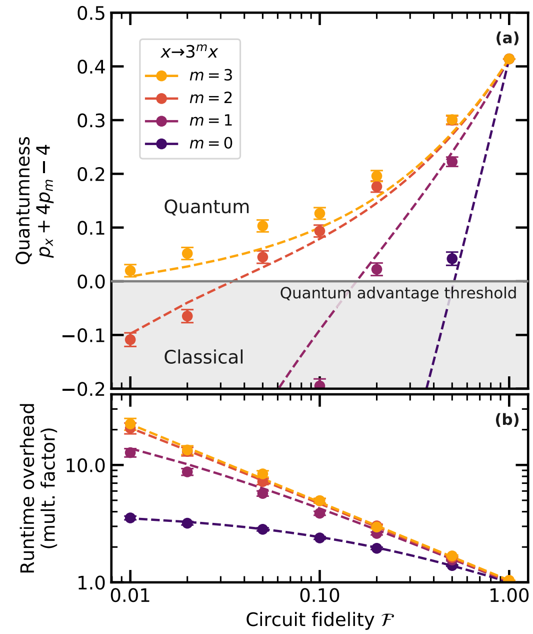

FIG. 2. Performance of our post-selection scheme when redundancy is added to the function $x^2 \bmod N$ by mapping it to $(3^m x)^2 \bmod 3^{2m} N$. Numerical simulations are performed on a quantum circuit implementing the Karatsuba algorithm for $m = \{0, 1, 2, 3\}$ (see supplemental information [48]). **(a)** "Quantumness" measured in terms of the classical bound from Eqn. 1 as a function of the total circuit fidelity. With $m = 3$, even a quantum device with only 1% circuit fidelity can demonstrate quantum advantage. **(b)** Depicts the increased runtime associated the post-selection scheme, which arises from a combination of slightly larger circuit sizes and the need to re-run the circuit multiple times. The latter is by far the dominant effect. Dashed lines are a theory prediction with no fit parameters; points are the result of numerical simulations at $n = 512$ bits and error bars depict 2σ uncertainty.

C. Robustness: Error mitigation via postselection

The existence of a finite gap between the classical and quantum success probabilities implies that our protocol can tolerate a certain amount of noise. A direct implementation of our interactive protocol on a noisy quantum device would require an overall fidelity of $\sim 83\%$ in order to exceed the classical bound [51]. To allow devices with lower fidelities to demonstrate quantum advantage, our protocol allows for a natural tradeoff between fidelity and runtime, such that the classical bound can, in principle, be exceeded with only a small [e.g. $1/\mathrm{poly}(n)$] amount of coherence in the quantum device [52].

The key idea is based upon postselection. For most TCFs, there are many bitstrings of length n that are not valid outputs of f. Thus, if the prover detects such a y value in step 3 (Fig. 1), they can simply discard it and try again [53]. In principle, the verifier can even use their trapdoor data to silently detect and discard iterations of the protocol with invalid y [54]. Since y is a function of x_0 and x_1, one might hope that this postselection scheme also rejects states where x_0 or x_1 has become corrupt. Although this may not always be the case, we demonstrate numerically that this assumption holds for a specific implementation of $x^2 \bmod N$ in the following subsection. One could also compute a classical checksum of x_0 and x_1 before and after the main circuit to ensure that they have not changed during its execution. Assuming that such bit-flip errors are indeed rejected, the possibility remains of an error in the phase between $|x_0\rangle$ and $|x_1\rangle$. In the supplementary information [48], we demonstrate that a prover holding the correct bitstrings but with an error in the phase can still saturate the classical bound; if the prover can avoid phase errors even a small fraction of the time, they will push past the classical threshold.

1. Numerical analysis of the postselection scheme for $x^2 \bmod N$

Focusing on the function $f(x) = x^2 \bmod N$, we now explicitly analyze the effectiveness of the postselection scheme. For this function, approximately 1/4 of the bitstrings of length n are valid outputs, so one would naively expect to reject about 3/4 of corrupted bitstrings. By introducing additional redundancy into the outputs of f, one can further decrease the probability that a corrupted y will incorrectly be accepted. As an example, let us consider mapping $x^2 \bmod N$ to the function $(kx)^2 \bmod k^2 N$ for some integer k. This is particularly convenient because the prover can validate y by simply checking whether it is a multiple of k^2. Moreover, the mapping adds only $\log k$ bits to the size of the problem, while rejecting a fraction $1-1/k^2$ of corrupted bitstrings.

We perform extensive numerical simulations demonstrating that postselection allows for quantum advantage to be achieved using noisy devices with low circuit fidelities (Fig. 2). We simulate quantum circuits for $(kx)^2 \bmod k^2 N$ at a problem size of $n = 512$ bits. Assuming a uniform gate fidelity across the circuit, we analyze the success rate of a quantum prover for $k = 3^m$ and $m = \{0, 1, 2, 3\}$. The choice of $k = 3^m$, and details of the simulation, are explained in the supplementary information [48].

For $m = 0$, the circuit implements our original function $x^2 \bmod N$, where in the absence of postselection, an overall circuit fidelity of $\mathcal{F} \sim 0.83$ is required to achieve quantum advantage. As depicted in Fig. 2(a), even for $m = 0$, our postselection scheme improves the advantage threshold down to $\mathcal{F} \sim 0.51$. For $m = 2$, circuit fidelities with $\mathcal{F} \lesssim 0.1$ remain well above the quantum advantage threshold, while for $m = 3$ the required circuit fidelity drops below 1%.

However, there is a tradeoff. In particular, one expects the overall runtime to increase for two reasons: (i) there will be a slight increase in the circuit size for $m > 0$ and (ii) one may need to re-run the quantum circuit many times until a valid y is measured. Somewhat remarkably, a runtime overhead of only 4.7x already enables quantum advantage to be achieved with an overall circuit fidelity of 10% [Fig. 2(b)]. Crucially, this increase in runtime is overwhelmingly due to re-running the quantum circuit and does not imply the need for longer experimental coherence times.

D. Efficient quantum evaluation of irreversible classical circuits

The central computational step in our interactive protocol (i.e. step 2, Fig. 1) is for the prover to apply a unitary of the form:

$$\mathcal{U}_{f_i} \sum_x |x\rangle_{\mathsf{x}} |0^{\otimes m}\rangle_{\mathsf{y}} = \sum_x |x\rangle_{\mathsf{x}} |f_i(x)\rangle_{\mathsf{y}}, \tag{2}$$

where $f_i(x)$ is a classical function. This type of unitary operation is ubiquitous across quantum algorithms, and a common strategy for its implementation is to convert the gates of a classical circuit into quantum gates. Generically, this process induces substantial overhead in both time and space complexity owing to the need to make the circuit reversible to preserve unitarity [38, 39]. This reversibility is often achieved by using an additional register, g, of so-called "garbage bits" and implementing: $\mathcal{U}'_{f_i} \sum_x |x\rangle_{\mathsf{x}} |0^{\otimes m}\rangle_{\mathsf{y}} |0^{\otimes l}\rangle_{\mathsf{g}} = \sum_x |x\rangle_{\mathsf{x}} |f_i(x)\rangle_{\mathsf{y}} |g_i(x)\rangle_{\mathsf{g}}$. For each gate in the classical circuit, enough garbage bits are added to make the operation injective. In general, to maintain coherence, these bits cannot be discarded but must be "uncomputed" later, adding significant complexity to the circuits.

A particularly appealing feature of our protocol is the existence of a measurement scheme to discard garbage bits, allowing for the direct mapping of classical to quantum circuits with no overhead. Specifically, we envision the prover measuring the qubits of the g register in the Hadamard basis and storing the results as a bitstring h, yielding the state,

$$|\psi\rangle = \sum_x (-1)^{h \cdot g_i(x)} |x\rangle_{\mathsf{x}} |f_i(x)\rangle_{\mathsf{y}}. \tag{3}$$

The prover has avoided the need to do any uncomputation of the garbage bits, at the expense of introducing phase flips onto some elements of the superposition. These phase flips do not affect the protocol, so long as the verifier can determine them. While classically computing $h \cdot g_i(x)$ is efficient for any x, computing it for *all* terms in the superposition is infeasible for the verifier. However, our protocol provides a natural way around this. The verifier can wait until the prover has collapsed the superposition onto x_0 and x_1, before evaluating $g_i(x)$ only on those two inputs [55].

Crucially, the prover can measure away garbage qubits as soon as they would be discarded classically, instead of waiting until the computation has completed. If these qubits are then reused, the quantum circuit will use no more space than the classical one. This feature allows for significant improvements in both gate depth and qubit number for practical implementations of the protocol (see last rows of Table I in Methods).

III. QUANTUM CIRCUITS FOR TRAPDOOR CLAW-FREE FUNCTIONS

While all of the trapdoor, claw-free functions listed in Table I can be utilized within our interactive protocol, each has its own set of advantages and disadvantages. For example, the TCF based on the Diffie-Hellman problem (described in the Methods) already enables a demonstration of quantum advantage at a key size of 160 bits (with a hardness equivalent to 1024 bit integer factorization [37]); however, building a circuit for this

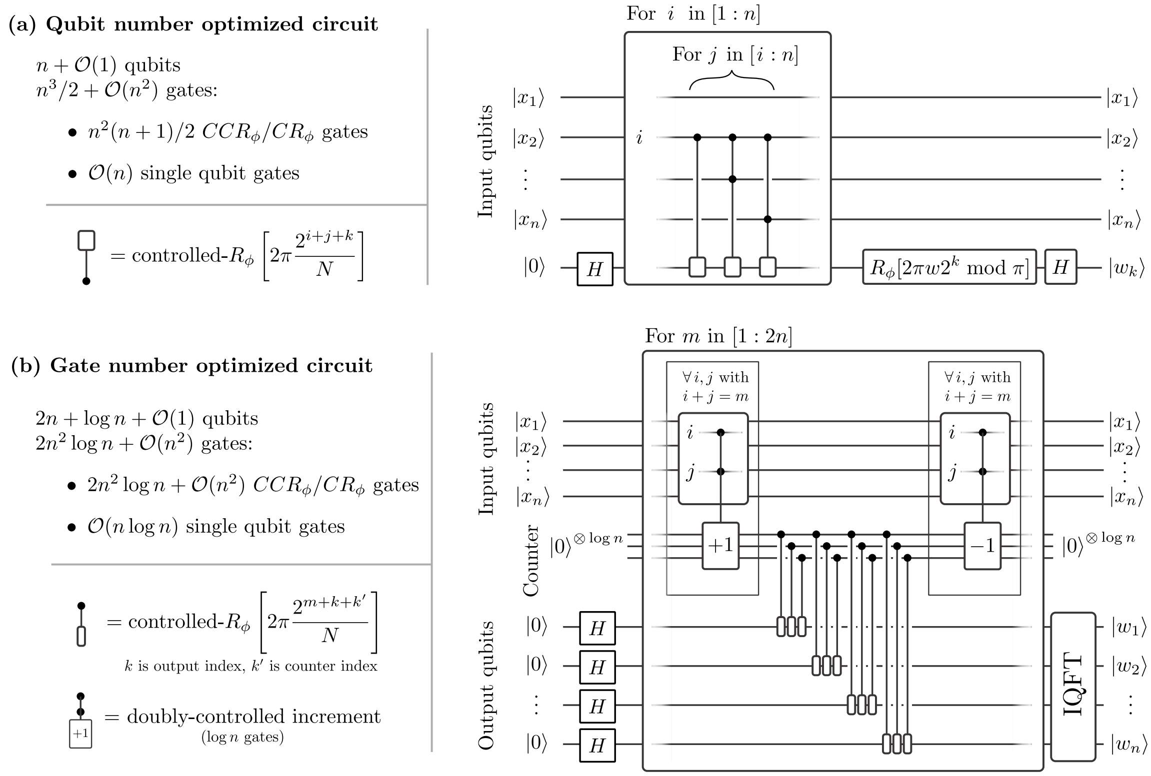

FIG. 3. Quantum circuits implementing step 2 of our interactive protocol for $f(x) = x^2 \bmod N$. **(a)** Depicts a quantum circuit optimized for qubit number. The circuit shown computes the k^{th} bit of $w = x^2/N$ and should be iterated for k. This iteration should begin at the least significant bit to ensure that the final phase rotation can be estimated classically. Note that the only entangling operations necessary for the circuit are doubly-controlled gates, which can be natively implemented using the Rydberg blockade (see Section III B). **(b)** Depicts a quantum circuit optimized for gate number. By combining gates of equal phase, one can reduce the overall circuit complexity to $\mathcal{O}(n^2 \log n)$ gates. We note that neither circuit generates garbage bits, since they are not built by converting from irreversible classical gates.

TCF requires a quantum implementation of Euclid's algorithm, which is challenging [56]. Thus, we focus on designing quantum circuits implementing Rabin's function, $x^2 \bmod N$.

A. Quantum circuits for $x^2 \bmod N$

We explore four different circuits (implementations of these algorithms in `Python` using the `Cirq` library are included as supplementary files [57]). The first two are quantum implementations of the classical Karatsuba and "schoolbook" integer multiplication algorithms, where we leverage the reversibility optimizations described in Section II D (see supplementary information [48]). The latter pair, described below, are intrinsically quantum algorithms that use Ising interactions to directly compute $x^2 \bmod N$ in the phase. Using those circuits, we propose a near-term demonstration of our interactive protocol on a Rydberg-based quantum computer [44, 47]; crucially, the so-called "Rydberg blockade" interaction *natively* realizes multi-qubit controlled phase rotations, from which the entire circuits shown in Figure 3 are built (up to single qubit rotations). A comparison of approximate gate counts for each of the four circuits can be seen in Table I in the Methods.

We now describe the two circuits, amenable to near-term quantum devices, that utilize quantum phase estimation to implement the function $f(x) = x^2 \bmod N$. The intuition behind our approach is as follows: we will compute x^2/N in the phase and transfer it to an output register via an inverse quantum Fourier transform [58, 59]; the modulo operation occurs automatically as the phase wraps around the unit circle, avoiding the need for a separate reduction step.

In order to implement $\sum_x |x\rangle_{\mathsf{x}} |x^2 \bmod N\rangle_{\mathsf{y}}$, we design a circuit to compute:

$$(\mathbb{I} \otimes \mathrm{IQFT})\, \tilde{\mathcal{U}}_{w_N}\, (\mathbb{I} \otimes \mathrm{H}^{\otimes \mathrm{m}})\, |x\rangle\, |0^{\otimes m}\rangle = |x\rangle\, |w\rangle \tag{4}$$

where H is a Hadamard gate, IQFT represents an inverse quantum Fourier transform, $w \equiv x^2/N = 0.w_1 w_2 \cdots w_m$ is an m-bit binary fraction [60], and $\tilde{\mathcal{U}}_{w_N}$ is the diagonal unitary,

$$\tilde{\mathcal{U}}_{w_N}\, |x\rangle\, |z\rangle = \exp\left(2\pi i \frac{x^2}{N} z\right) |x\rangle\, |z\rangle\, . \tag{5}$$

By performing a binary decomposition of the phase in

Eqn. 5:

$$\exp\left(2\pi i \frac{x^2}{N} z\right) = \prod_{i,j,k} \exp\left(2\pi i \frac{2^{i+j+k}}{N} x_i x_j z_k\right), \quad (6)$$

one immediately finds that $\tilde{\mathcal{U}}_{w_N}$ is equivalent to applying a series of controlled-controlled-phase rotation gates of angle,

$$\phi_{ijk} = \frac{2\pi 2^{i+j+k}}{N} \pmod{2\pi}. \quad (7)$$

Here, the control qubits are i, j in the x register, while the target qubit is k in the y register. Crucially, the value of this phase for any i, j, k can be computed classically when the circuit is compiled.

As depicted in Figure 3, we propose two explicit circuits to implement $\tilde{\mathcal{U}}_{w_N}$, one optimizing for qubit count, and the other optimizing for gate count. The first circuit [Fig. 3(a)] takes advantage of the fact that the output register is measured immediately after it is computed; this allows one to replace the m output qubits with a single qubit that is measured and reused m times. Moreover, by replacing groups of doubly-controlled gates with a Toffoli and a series of singly-controlled gates, one ultimately arrives at an implementation, which requires $n^3/2 + \mathcal{O}(n^2)$ gates, but only $n + \mathcal{O}(1)$ qubits.

Our second circuit [Fig. 3(b)], which optimizes for gate count, leverages the fact that ϕ_{ijk} (Eqn. 7) only depends on $i + j + k$, allowing one to combine gates with a common sum. In this case, one can define $m = i + j$ and then, for each value of m, simply "count" the number of values of i, j for which both control qubits are 1. By then performing controlled gates off of the qubits of the counter register, one can reduce the total gate complexity by a factor of $n/\log n$, leading to a implementation with $2n^2 \log n + \mathcal{O}(n^2)$ gates.

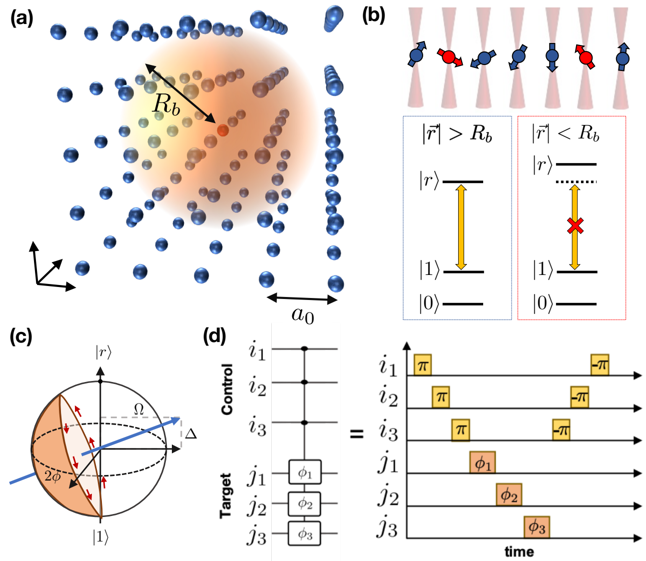

FIG. 4. **a,** Schematic illustration of a three dimensional array of neutral atoms with Rydberg blockade interactions. The blockade radius can be significantly larger than the inter-atom spacing, enabling multi-qubit entangling operations. **b,** As an example, Rydberg atoms can be trapped in an optical tweezer array. The presence of an atom in a Rydberg excited state (red) shifts the energy levels of nearby atoms (blue), preventing the driving field (yellow arrow) from exciting them to their Rydberg state, $|r\rangle$. **c,** A single qubit phase rotation can be implemented by an off-resonant Rabi oscillation between one of the qubit states, e.g., $|1\rangle$, and the Rydberg excited state. This imprints a tunable, geometric phase ϕ, which is determined by the detuning Δ and Rabi frequency Ω. **d,** Multi-qubit controlled-phase rotations are implemented via a sequence of π-pulses between the $|0\rangle \leftrightarrow |r\rangle$ transition of control atoms (yellow) and off-resonant Rabi oscillations on the target atoms (orange).

B. Experimental implementation

Motivated by recent advances in the creation and control of many-body entanglement in programmable quantum systems [11, 61–63], we propose an experimental implementation of our interactive protocol based upon neutral atoms coupled to Rydberg states [47]. We envision a three dimensional system of either alkali or alkaline-earth atoms trapped in an optical lattice or optical tweezer array [Fig. 4(a)] [64–66]. To be specific, we consider ^{87}Rb with an effective qubit degree of freedom encoded in hyperfine states: $|0\rangle = |F = 1, m_F = 0\rangle$ and $|1\rangle = |F = 2, m_F = 0\rangle$. Gates between atoms are mediated by coupling to a highly-excited Rydberg state $|r\rangle$, whose large polarizability leads to strong van der Waals interactions. This microscopic interaction enables the so-called Rydberg "blockade" mechanism—when a single atom is driven to its Rydberg state, all other atoms within a blockade radius, R_b, become off-resonant from the drive, thereby suppressing their excitation [Fig. 4(a,b)] [43].

Somewhat remarkably, this blockade interaction enables the *native* implementation of all multi-qubit-controlled phase gates depicted in the circuits in Figure 3. In particular, consider the goal of applying a $C^k R_\phi^\ell$ gate; this gate applies phase rotations, $\{\phi_1, \phi_2, \ldots, \phi_\ell\}$, to target qubits $\{j_1, j_2, \ldots j_\ell\}$ if all k control qubits $\{i_1, i_2, \ldots i_k\}$ are in the $|1\rangle$ state [Fig. 4(d)]. Experimentally, this can be implemented as follows: (i) sequentially apply (in any order) resonant π-pulses on the $|0\rangle \leftrightarrow |r\rangle$ transition for the k desired control atoms, (ii) off-resonantly drive the $|1\rangle \leftrightarrow |r\rangle$ transition of each target atom with detuning Δ and Rabi frequency Ω for a time duration $T = 2\pi/(\Omega^2 + \Delta^2)^{1/2}$ [Fig. 4(c)], (iii) sequentially apply [in the opposite order as in (i)] resonant $-\pi$-pulses (i.e. π-pulses with the opposite phase) to the k control atoms to bring them back to their original state. The intuition for why this experimental sequence implements the $C^k R_\phi^\ell$ gate is straightforward. The first step creates a blockade if any of the control qubits are in the $|0\rangle$ state, while the second step imprints a phase, $\phi = \pi(1 - \Delta/\sqrt{\Delta^2 + \Omega^2})$, on the $|1\rangle$ state, only in the absence of a blockade. Note that tuning the values of ϕ_i for each of the target qubits simply corresponds to adjusting the detuning and Rabi frequency of the off-resonant drive in the second step [Fig. 4(c,d)].

Demonstrations of our protocol can already be implemented in current generation Rydberg experiments, where a number of essential features have recently been shown, including: 1) the coherent manipulation of individual qubits trapped in a 3D tweezer array [64, 65], 2) the deterministic loading of atoms in a 3D optical lattice [66], and 3) fast entangling gate operations with fidelities, $F \geq 0.974$ [44–46]. In order to estimate the number of entangling gates achievable within decoherence time scales, let us imagine choosing a Rydberg state with a principal quantum number $n \approx 70$. This yields a strong van der Waals interaction, $V(\vec{r}) = C_6/r^6$, with a C_6 coefficient $\sim (2\pi)\,880$ GHz$\cdot\mu$m^6 [67]. Combined with a coherent driving field of Rabi frequency $\Omega \sim (2\pi)\,1-10$ MHz, the van der Waals interaction can lead to a blockade radius of up to, $R_b = (C_6/\Omega)^{1/6} \sim 10\,\mu$m. Within this radius, one can arrange $\sim 10^2$ all-to-all interacting qubits, assuming an atom-to-atom spacing of approximately, $a_0 \approx 2\mu$m [68]. In current experiments, the decoherence associated with the Rydberg transition is typically limited by a combination of inhomogeneous Doppler shifts and laser phase/intensity noise, leading to $1/T_2 \sim 10-100$ kHz [44, 69, 70]. Taking everything together, one should be able to perform $\sim 10^3$ entangling gates before decoherence occurs (this is comparable to the number of two-qubit entangling gates possible in other state-of-the-art platforms [11, 71]). While this falls short of enabling an immediate full-scale demonstration of classically verifiable quantum advantage, we hasten to emphasize that the ability to directly perform multi-qubit entangling operations significantly reduces the cost of implementing our interactive protocol. For example, the standard decomposition of a Toffoli gate uses 6 CNOT gates and 7 T and $T^\dagger$ gates, with a gate depth of 12 [72–74]; an equivalent three qubit gate can be performed in a single step via the Rydberg blockade mechanism.

IV. CONCLUSION AND OUTLOOK

The interplay between classical and quantum complexities ultimately determines the threshold for any quantum advantage scheme. Here, we have proposed a novel interactive protocol for classically verifiable quantum advantage based upon trapdoor claw-free functions; in addition to proposing two new TCFs [Table I], we also provide explicit quantum circuits that leverage the microscopic interactions present in a Rydberg-based quantum computer. Our work allows near-term quantum devices to move one step closer toward a loophole-free demonstration of quantum advantage and also opens the door to a number of promising future directions.

First, our proof of soundness only applies to classical adversaries; whether it is possible to extend our protocol's security to quantum adversaries remains an open question [75]. A quantum-secure proof could enable our protocol's use in a number of applications, such as certifiable random number generation [17] and the verification of arbitrary quantum computations [76]. Second, our work motivates the search for new trapdoor claw-free functions, which can be evaluated in the smallest possible quantum volume. Cryptographic primitives such as Learning Parity with Noise (LPN), which are designed for use in low-power devices such as RFID cards, represent a promising path forward [77]. More broadly, one could also attempt to build modified protocols, which simplify either the requirements on the cryptographic function or the interactions; interestingly, recent work has demonstrated the use of a cryptographic hash function to remove the need for interactions in a TCF-based proof of quantumness [18]. Finally, while we have focused our experimental discussions on Rydberg atoms, a number of other platforms also exhibit features that facilitate the protocol's implementation. For example, both trapped ions and cavity-QED systems can allow all-to-all connectivity, while superconducting qubits can be engineered to have biased noise [78]. This latter feature would allow noise to be concentrated into error modes detectable by our proposed post-selection scheme.

We gratefully acknowledge the insights of and discussions with A. Bouland, A. Gheorghiu, Z. Landau, L. Lewis, and T. Vidick. We are particularly indebted to Joonhee Choi for insights about Rydberg-based quantum computing as well as Sanjam Garg for his input regarding trapdoor function constructions. This work was supported by the NSF QLCI program through grant number OMA-2016245, and the DOD through MURI grant number FA9550-18-1-0161. NYY acknowledges support from the David and Lucile Packard foundation and a Google research award. GDKM acknowledges support from the Department of Defense (DOD) through the National Defense Science & Engineering Graduate Fellowship (NDSEG) Program. SC acknowledges support from the Miller Institute for Basic Research in Science.

[1] S. Aaronson and A. Arkhipov, in *Proceedings of the forty-third annual ACM symposium on Theory of computing*, STOC '11 (Association for Computing Machinery, New York, NY, USA, 2011) pp. 333–342.
[2] E. Farhi and A. W. Harrow, arXiv:1602.07674 (2016).
[3] M. J. Bremner, A. Montanaro, and D. J. Shepherd, Physical Review Letters **117**, 080501 (2016).
[4] A. P. Lund, M. J. Bremner, and T. C. Ralph, npj Quantum Information **3**, 1 (2017).
[5] A. W. Harrow and A. Montanaro, Nature **549**, 203 (2017).
[6] B. M. Terhal, Nature Physics **14**, 530 (2018).

[7] S. Boixo, S. V. Isakov, V. N. Smelyanskiy, R. Babbush, N. Ding, Z. Jiang, M. J. Bremner, J. M. Martinis, and H. Neven, Nature Physics **14**, 595 (2018).
[8] A. Bouland, B. Fefferman, C. Nirkhe, and U. Vazirani, Nature Physics **15**, 159 (2019).
[9] S. Aaronson and L. Chen, in *32nd Computational Complexity Conference (CCC 2017)*, Leibniz International Proceedings in Informatics (LIPIcs), Vol. 79, edited by R. O'Donnell (Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 2017) pp. 22:1–22:67.
[10] 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.*, Science **360**, 195 (2018).
[11] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. S. L. Brandao, D. A. Buell, B. Burkett, Y. Chen, Z. Chen, B. Chiaro, *et al.*, Nature **574**, 505 (2019).
[12] H.-S. Zhong, H. Wang, Y.-H. Deng, M.-C. Chen, L.-C. Peng, Y.-H. Luo, J. Qin, D. Wu, X. Ding, Y. Hu, P. Hu, X.-Y. Yang, W.-J. Zhang, H. Li, *et al.*, Science **370**, 1460 (2020).
[13] S. Bravyi, D. Gosset, and R. König, Science **362**, 308 (2018).
[14] S. Bravyi, D. Gosset, R. Koenig, and M. Tomamichel, arXiv:1904.01502 [quant-ph] (2019).
[15] They also have some other caveats: a provable speedup of $\mathcal{O}(1)$ quantum complexity over $\mathcal{O}(n)$ classical complexity is promising, but just reading the input may require $\mathcal{O}(n)$ time, hiding the computational speedup in practice.
[16] P. W. Shor, SIAM Journal on Computing **26**, 1484 (1997).
[17] Z. Brakerski, P. Christiano, U. Mahadev, U. Vazirani, and T. Vidick, arXiv:1804.00640 [quant-ph] (2019).
[18] Z. Brakerski, V. Koppula, U. Vazirani, and T. Vidick, arXiv:2005.04826 [quant-ph] (2020).
[19] D. Aharonov, M. Ben-Or, E. Eban, and U. Mahadev, arXiv:1704.04487 (2017).
[20] J. Watrous, arXiv:cs/9901015 (1999).
[21] A. Kitaev and J. Watrous, in *Proceedings of the thirty-second annual ACM symposium on Theory of computing* (2000) pp. 608–617.
[22] H. Kobayashi and K. Matsumoto, Journal of Computer and System Sciences **66**, 429 (2003).
[23] J. Fitzsimons and T. Vidick, in *Proceedings of the 2015 Conference on Innovations in Theoretical Computer Science* (2015) pp. 103–112.
[24] I. L. Markov, A. Fatima, S. V. Isakov, and S. Boixo, arXiv:1807.10749 (2018).
[25] O. Regev, in *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, STOC '05 (Association for Computing Machinery, New York, NY, USA, 2005) pp. 84–93.
[26] To be precise, it is hard to find both x_0 and the parity of any subset of the bits of x_1.
[27] One other proposal exists for removing the need for an adaptive hardcore bit [18]; it applies a cryptographic hash function on the superposition in addition to the TCF. This requires larger quantum circuits, and the extra cryptographic assumption of security of the hash function. On the other hand, that construction reduces the protocol to only one round of interaction, making it *non-interactive* and thus potentially asynchronous, which is useful in some practical settings.
[28] J. S. Bell, Physics Physique Fizika **1**, 195 (1964).
[29] J. F. Clauser, M. A. Horne, A. Shimony, and R. A. Holt, Physical Review Letters **23**, 880 (1969).
[30] W. Diffie and M. Hellman, IEEE Transactions on Information Theory **22**, 644 (1976).
[31] C. Peikert and B. Waters, in *Proceedings of the fortieth annual ACM symposium on Theory of computing*, STOC '08 (Association for Computing Machinery, New York, NY, USA, 2008) pp. 187–196.
[32] D. M. Freeman, O. Goldreich, E. Kiltz, A. Rosen, and G. Segev, in *Public Key Cryptography – PKC 2010*, Lecture Notes in Computer Science, edited by P. Q. Nguyen and D. Pointcheval (Springer, Berlin, Heidelberg, 2010) pp. 279–295.
[33] M. O. Rabin, *Digitalized signatures and public-key functions as intractable as factorization*, Technical Report (Massachusetts Institute of Technology, USA, 1979).
[34] S. Goldwasser, S. Micali, and R. L. Rivest, SIAM Journal on Computing **17**, 281 (1988).
[35] V. S. Miller, in *Advances in Cryptology — CRYPTO '85 Proceedings*, Lecture Notes in Computer Science, edited by H. C. Williams (Springer, Berlin, Heidelberg, 1986) pp. 417–426.
[36] N. Koblitz, Mathematics of Computation **48**, 203 (1987).
[37] E. Barker, *Recommendation for Key Management Part 1: General*, Tech. Rep. NIST SP 800-57pt1r4 (National Institute of Standards and Technology, 2016).
[38] C. H. Bennett, SIAM Journal on Computing **18**, 766 (1989).
[39] R. Y. Levine and A. T. Sherman, SIAM Journal on Computing **19**, 673 (1990).
[40] D. Aharonov, A. Kitaev, and N. Nisan, in *Proceedings of the thirtieth annual ACM symposium on Theory of computing* (1998) pp. 20–30.
[41] H. M. H. Babu, M. R. Islam, S. M. A. Chowdhury, and A. R. Chowdhury, in *17th International Conference on VLSI Design. Proceedings.* (IEEE, 2004) pp. 757–760.
[42] S. Kotiyal, H. Thapliyal, and N. Ranganathan, in *2014 27th international conference on VLSI design and 2014 13th international conference on embedded systems* (IEEE, 2014) pp. 545–550.
[43] M. Saffman, Journal of Physics B: Atomic, Molecular and Optical Physics **49**, 202001 (2016).
[44] H. Levine, A. Keesling, G. Semeghini, A. Omran, T. T. Wang, S. Ebadi, H. Bernien, M. Greiner, V. Vuletić, H. Pichler, and M. D. Lukin, Phys. Rev. Lett. **123**, 170503 (2019).
[45] T. Graham, M. Kwon, B. Grinkemeyer, Z. Marra, X. Jiang, M. Lichtman, Y. Sun, M. Ebert, and M. Saffman, Physical Review Letters **123**, 230501 (2019).
[46] I. S. Madjarov, J. P. Covey, A. L. Shaw, J. Choi, A. Kale, A. Cooper, H. Pichler, V. Schkolnik, J. R. Williams, and M. Endres, Nature Physics **16**, 857 (2020).
[47] A. Browaeys and T. Lahaye, Nature Physics **16**, 132 (2020).
[48] See Supplementary Information for additional details and supporting derivations.
[49] O. Goldreich and L. A. Levint, in *In Proceedings of the Twenty First Annual ACM Symposium on Theory of Computing* (1989) pp. 25–32.
[50] The oracle's noise rate is *not* simply $p_{r \cdot x_1}$: that is the probability that any single value $r \cdot x_1$ is correct, but all of the queries to the oracle are correlated (they are for the same iteration of the protocol, and thus the same

value of y).
[51] This number comes from solving the classical bound (Equation 1) for circuit fidelity $\mathcal{F}$, with $p_x = \mathcal{F}$ and $p_m = 1/2 + \mathcal{F}/2$.
[52] This is true even if the coherence is exponentially small in n. Of course, with arbitrarily low coherence the runtime may become excessively large such that quantum advantage cannot be demonstrated—the point is that regardless of runtime, the classical probability bound can be exceeded with a device that has arbitrarily low circuit fidelity.
[53] This scheme will only remove errors in the first round of the protocol, but fortunately, one expects the overwhelming majority of the quantum computation, and thus also the majority of errors, to occur in that round.
[54] This procedure does not leak data to a classical cheater, because the verifier does not communicate which runs were discarded. Furthermore, it does not affect the soundness of Theorem 2, because the machine $\mathcal{B}$ in that theorem's proof can simply iterate until it encounters a valid y.
[55] This is true because $g_i(x)$ is the result of adding extra output bits to the gates of a classical circuit, which is efficient to evaluate on any input.
[56] T. Häner, S. Jaques, M. Naehrig, M. Roetteler, and M. Soeken, arXiv:2001.09580 [quant-ph] (2020).
[57] Code is available at `https://github.com/GregDMeyer/quantum-advantage` and is archived on Zenodo [79].
[58] T. G. Draper, arXiv:quant-ph/0008033 (2000).
[59] S. Beauregard, arXiv:quant-ph/0205095 (2003).
[60] We must take $m > n + \mathcal{O}(1)$ to sufficiently resolve the value $x^2 \bmod N$ in post-processing.
[61] J. Zhang, G. Pagano, P. W. Hess, A. Kyprianidis, P. Becker, H. Kaplan, A. V. Gorshkov, Z.-X. Gong, and C. Monroe, Nature **551**, 601 (2017).
[62] P. Scholl, M. Schuler, H. J. Williams, A. A. Eberharter, D. Barredo, K.-N. Schymik, V. Lienhard, L.-P. Henry, T. C. Lang, T. Lahaye, A. M. Läuchli, and A. Browaeys, arXiv:2012.12268 [cond-mat, physics:physics, physics:quant-ph] (2020).
[63] S. Ebadi, T. T. Wang, H. Levine, A. Keesling, G. Semeghini, A. Omran, D. Bluvstein, R. Samajdar, H. Pichler, W. W. Ho, S. Choi, S. Sachdev, M. Greiner, V. Vuletic, and M. D. Lukin, arXiv:2012.12281 [cond-mat, physics:physics, physics:quant-ph] (2020).
[64] Y. Wang, X. Zhang, T. A. Corcovilos, A. Kumar, and D. S. Weiss, Phys. Rev. Lett. **115**, 043003 (2015).
[65] Y. Wang, A. Kumar, T.-Y. Wu, and D. S. Weiss, Science **352**, 1562 (2016).
[66] A. Kumar, T.-Y. Wu, F. Giraldo, and D. S. Weiss, Nature **561**, 83 (2018).
[67] R. Löw, H. Weimer, J. Nipper, J. B. Balewski, B. Butscher, H. P. Büchler, and T. Pfau, Journal of Physics B: Atomic, Molecular and Optical Physics **45**, 113001 (2012).
[68] We note that this spacing is ultimately limited by a combination of the optical diffraction limit and the orbital size of $n \approx 70$ Rydberg states.
[69] S. de Léséleuc, D. Barredo, V. Lienhard, A. Browaeys, and T. Lahaye, Physical Review A **97**, 053803 (2018).
[70] Y. Liu, Y. Sun, Z. Fu, P. Xu, X. Wang, X. He, J. Wang, and M. Zhan, arXiv:2012.12589 [quant-ph] (2020).
[71] V. M. Schäfer, C. J. Ballance, K. Thirumalai, L. J. Stephenson, T. G. Ballance, A. M. Steane, and D. M. Lucas, Nature **555**, 75 (2018).
[72] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information: 10th Anniversary Edition*, 10th ed. (Cambridge University Press, USA, 2011).
[73] V. V. Shende and I. L. Markov, Quantum Information & Computation **9**, 461 (2009).
[74] A. Barenco, C. H. Bennett, R. Cleve, D. P. DiVincenzo, N. Margolus, P. Shor, T. Sleator, J. A. Smolin, and H. Weinfurter, Physical Review A **52**, 3457 (1995).
[75] By "secure against quantum adversaries" we mean that it is possible to ensure that the quantum prover is actually following the prescribed protocol, to whatever extent is necessary for the desired cryptographic task.
[76] U. Mahadev, arXiv:1804.01082 [quant-ph] (2018).
[77] K. Pietrzak, in *SOFSEM 2012: Theory and Practice of Computer Science*, Lecture Notes in Computer Science, edited by M. Bieliková, G. Friedrich, G. Gottlob, S. Katzenbeisser, and G. Turán (Springer, Berlin, Heidelberg, 2012) pp. 99–114.
[78] S. Puri, L. St-Jean, J. A. Gross, A. Grimm, N. E. Frattini, P. S. Iyer, A. Krishna, S. Touzard, L. Jiang, A. Blais, S. T. Flammia, and S. M. Girvin, Science Advances **6**, eaay5901 (2020).
[79] G. D. Kahanamoku-Meyer, Gregdmeyer/quantum-advantage: v1.0 (2021), doi:10.5281/zenodo.4641454.

Methods and Extended Data: Classically-Verifiable Quantum Advantage from a Computational Bell Test

Gregory D. Kahanamoku-Meyer,[1] Soonwon Choi,[1] Umesh V. Vazirani,[2] and Norman Y. Yao[1]

[1] *Department of Physics, University of California at Berkeley, Berkeley, CA 94720*

[2] *Department of Electrical Engineering and Computer Science, University of California at Berkeley, Berkeley, CA 94720*

LIST DECODING LEMMA

In this section we prove a bound on the probability that list decoding will succeed for a particular value of y, given an oracle's noise rate over *all* values of y. Recall that by the Goldreich-Levin theorem [1], list decoding of the Hadamard code is possible if the noise rate is noticeably less than 1/2.

Lemma 1. *Consider a binary-valued function over two inputs $g : Y \times \{0,1\}^n \to \{0,1\}$, and a noisy oracle $\mathcal{G}$ to that function. Assuming some distribution of values $y \in Y$ and $r \in \{0,1\}^n$, define $\epsilon \equiv \Pr_{y,r}[\mathcal{G}(y,r) \neq g(y,r)]$ as the "noise rate" of the oracle. Now define the conditional noise rate for a particular $y \in Y$ as*

$$\epsilon_y \equiv \Pr_r\left[\mathcal{G}(y,r) \neq g(y,r)\right] \tag{1}$$

Then, the probability that ϵ_y is less than $1/2 - \mu(n)$ for any positive function μ, over randomly selected y, is

$$p_{\text{good}} \equiv \Pr_y\left[\epsilon_y < 1/2 - \mu(n)\right] \geq 1 - 2\epsilon - 2\mu(n). \tag{2}$$

Proof. Let $S \subseteq Y$ be the set of y values for which $\epsilon_y < 1/2 - \mu(n)$. Then by definition we have

$$\epsilon = p_{\text{good}} \cdot \epsilon_{y \in S} + (1 - p_{\text{good}}) \cdot \epsilon_{y \notin S} \tag{3}$$

Noting that we must have $\epsilon_y \geq 1/2 - \mu(n)$ for $y \notin S$ by definition, we may minimize the right hand side of Equation 3, yielding the bound

$$\epsilon > p_{\text{good}} \cdot 0 + (1 - p_{\text{good}}) \cdot (1/2 - \mu(n)) \tag{4}$$

Rearranging this expression we arrive at

$$p_{\text{good}} > 1 - 2\epsilon - 2\mu(n)$$

which is what we desired to show. □

TRAPDOOR CLAW-FREE FUNCTION CONSTRUCTIONS

Here we present two trapdoor claw-free function families (TCFs) for use in the protocol of this paper. These families are defined by three algorithms: Gen, a probabilistic algorithm which selects an index i specifying one function in the family and outputs the corresponding trapdoor data t; f_i, the definition of the function itself; and T, a trapdoor algorithm which efficiently inverts f_i for any i, given the corresponding trapdoor data t. Here we provide the definitions of the function families; proofs of their cryptographic properties are included in the supplementary information. In these definitions we use a security parameter λ following the notation of cryptographic literature; λ is informally equivalent to the "problem size" n discussed in the main text.

TCF from Rabin's function $x^2 \bmod N$

"Rabin's function" $f_N(x) = x^2 \bmod N$, with N the product of two primes, was first used in the context of public-key cryptography and digital signatures [2, 3]. We use it to define the trapdoor claw-free function family $\mathcal{F}_{\text{Rabin}}$, as follows.

Function generation

$\mathsf{Gen}\left(1^\lambda\right)$

1. Randomly choose two prime numbers p and q of length $\lambda/2$ bits, with $p \bmod 4 \equiv q \bmod 4 \equiv 3 \bmod 4$ [4].
2. Return $N = pq$ as the function index, and the tuple (p, q) as the trapdoor data.

Function definition

$f : [N/2] \to [N]$ is defined as

$$f_N(x) = x^2 \bmod N \tag{5}$$

The domain is restricted to $[N/2]$ to remove extra trivial collisions of the form $(x, -x)$.

Trapdoor

The trapdoor algorithm is the same as the decryption algorithm of the Rabin cryptosystem [2]. On input y and key (p, q), the Rabin decryption algorithm returns four integers $(x_0, x_1, -x_0, -x_1)$ in the range [0, N). x_0 and x_1 can then be selected by choosing the two values that are smaller than $N/2$. See proof in supplementary information for an overview of the algorithm.

TCF from Decisional Diffie-Hellman

We now present a trapdoor claw-free function family $\mathcal{F}_{\mathrm{DDH}}$ based on the decisional Diffie-Hellman problem (DDH). DDH is defined for a multiplicative group $\mathbb{G}$; informally, the DDH assumption states that for a group generator g and two integers a and b, given g, g^a, and g^b it is computationally hard to distinguish g^{ab} from a random group element. We expand on a known DDH-based trapdoor one-way function construction [5, 6], adding the claw-free property to construct a TCF.

Function generation

$\mathsf{Gen}\left(1^\lambda\right)$

1. Choose a group $\mathbb{G}$ of order $q \sim \mathcal{O}(2^\lambda)$, and a generator g for that group.
2. For dimension $k \sim \mathcal{O}(\lambda)$ choose a random invertible matrix $\mathbf{M} \in \mathbb{Z}_q^{k\times k}$.
3. Compute $g^{\mathbf{M}} = \left(g^{\mathbf{M}_{ij}}\right) \in \mathbb{G}^{k\times k}$ (element-wise exponentiation).
4. Choose a secret vector $\mathbf{s} \in \{0,1\}^k$; compute the vector $g^{\mathbf{Ms}}$ (where $\mathbf{Ms}$ is the matrix-vector product, and again the exponentiation is element-wise).
5. Publish the pair $\left(g^{\mathbf{M}}, g^{\mathbf{Ms}}\right)$, retain $(\mathbf{M}, \mathbf{s})$ as the trapdoor data.

Function definition

Let m be a power of two with $m \sim \mathcal{O}\left(k^2\right)$. We define the function f as $f(b||\mathbf{x}) := f_b(\mathbf{x})$, where $||$ denotes concatenation, for a pair of functions $f_b : \mathbb{Z}_m^k \to \mathbb{G}^k$:

$$f_0\left(\mathbf{x}\right) = g^{\mathbf{Mx}} \tag{6}$$
$$f_1\left(\mathbf{x}\right) = g^{\mathbf{Mx}} g^{\mathbf{Ms}} = g^{\mathbf{M}(\mathbf{x}+\mathbf{s})} \tag{7}$$

Trapdoor

The algorithm takes as input the trapdoor data $(\mathbf{M}, \mathbf{s})$ and a value $y = g^{\mathbf{Mx}_0} = g^{\mathbf{M}(\mathbf{x}_1+\mathbf{s})}$, and returns the claw $(\mathbf{x}_0, \mathbf{x}_1)$.

$T((\mathbf{M}, \mathbf{s}), y)$

1. Compute $\mathbf{M}^{-1}$ using $\mathbf{M}$.
2. Compute $g^{\mathbf{M}^{-1}\mathbf{M}\mathbf{x}_0} = g^{\mathbf{x}_0}$.
3. Take the discrete logarithm of each element of $g^{\mathbf{x}_0}$, yielding $\mathbf{x}_0$. Crucially, this is possible because the elements of $\mathbf{x}$ are in $\mathbb{Z}_m$ and $m = \mathsf{poly}(n)$, so the discrete logarithm can be computed in polynomial time by brute force.
4. Compute $\mathbf{x}_1 = \mathbf{x}_0 - \mathbf{s}$
5. Return $(\mathbf{x}_0, \mathbf{x}_1)$

TABLE OF CIRCUIT SIZES

A comparison of the resource requirements for computing $x^2 \bmod N$, for various problem sizes and circuit designs, is presented in Table I. These counts are generated in the "abstract circuit" model, in which error correction, qubit routing, and other practical considerations are not included. For schoolbook and Karatsuba circuits, circuits are decomposed into a Clifford+T gate set. For the "phase" circuits, we allow controlled arbitrary phase rotations, as we expect these circuits to be appropriate for hardware (physical) qubits where these gates are native. Accordingly, we do not provide T gate counts for those circuits.

TABLE I. Circuit sizes for various values of $n = \log N$. Values may vary for different N of the same length. *From analytic estimate rather than building explicit circuit. †Reversible circuits constructed using Q# implementation of Ref. [7], and scaled to include Montgomery reduction. ‡Estimate from [8].

Circuit	Qubits	Gates	T Gates	Depth
$n = 128$ (takes seconds on a desktop [9])				
Phase circuit 1	128	1.1×10^6	—	1.1×10^6
Phase circuit 2	264	4.3×10^5	—	6.3×10^4
Schoolbook	515	9.1×10^5	3.9×10^5	1.9×10^4
Karatsuba	942	7.7×10^5	3.3×10^5	2.0×10^3
$n = 400$ (takes hours on a desktop [9])				
Phase circuit 1	400	$3.3 \times 10^{7*}$	—	$3.3 \times 10^{7*}$
Phase circuit 2	812	$4.2 \times 10^{6*}$	—	$6.2 \times 10^{5*}$
Schoolbook	1603	8.7×10^6	3.6×10^6	5.9×10^4
Karatsuba	3051	5.4×10^6	2.3×10^6	5.3×10^4
$n = 829$ (record for factoring [10])				
Phase circuit 1	829	$3.0 \times 10^{8*}$	—	$2.9 \times 10^{8*}$
Phase circuit 2	1671	$1.8 \times 10^{7*}$	—	$2.6 \times 10^{6*}$
Schoolbook	3319	3.8×10^7	1.6×10^7	$1.2 \times 10^{5*}$
Karatsuba	5522	1.8×10^7	7.7×10^6	$1.1 \times 10^{5*}$
$n = 1024$ (exceeds factoring record)				
Phase circuit 1	1024	$5.6 \times 10^{8*}$	—	$5.5 \times 10^{8*}$
Phase circuit 2	2061	$2.7 \times 10^{7*}$	—	$4.0 \times 10^{6*}$
Schoolbook	4097	5.7×10^7	2.4×10^7	$1.5 \times 10^{5*}$
Karatsuba	6801	2.6×10^7	1.1×10^7	$1.4 \times 10^{5*}$
Other algs. at $n = 1024$				
Rev. schoolbook †	8192	6.4×10^8	2.2×10^8	1.1×10^8
Rev. Karatsuba †	12544	5.7×10^8	1.9×10^8	2.4×10^7
Shor's alg. ‡	3100	$4.8 \times 10^{9*}$	$1.9 \times 10^{9*}$	—

[1] O. Goldreich and L. A. Levint, in *In Proceedings of the Twenty First Annual ACM Symposium on Theory of Computing* (1989) pp. 25–32.
[2] M. O. Rabin, *Digitalized signatures and public-key functions as intractable as factorization*, Technical Report (Massachusetts Institute of Technology, USA, 1979).

[3] S. Goldwasser, S. Micali, and R. L. Rivest, SIAM Journal on Computing **17**, 281 (1988).
[4] In practice, p and q must be selected with some care such that Fermat factorization and Pollard's $p-1$ algorithm [11] cannot be used to efficiently factor N classically. Selecting p and q in the same manner as for RSA encryption would be effective [12].
[5] C. Peikert and B. Waters, in *Proceedings of the fortieth annual ACM symposium on Theory of computing*, STOC '08 (Association for Computing Machinery, New York, NY, USA, 2008) pp. 187–196.
[6] D. M. Freeman, O. Goldreich, E. Kiltz, A. Rosen, and G. Segev, in *Public Key Cryptography – PKC 2010*, Lecture Notes in Computer Science, edited by P. Q. Nguyen and D. Pointcheval (Springer, Berlin, Heidelberg, 2010) pp. 279–295.
[7] C. Gidney, arXiv:1904.07356 [quant-ph] (2019).
[8] C. Gidney and M. Ekerå, arXiv:1905.09749 [quant-ph] (2019).
[9] CADO-NFS, `http://cado-nfs.gforge.inria.fr/`, accessed: 2020-06-27.
[10] P. Zimmermann, Factorization of RSA-250, `https://lists.gforge.inria.fr/pipermail/cado-nfs-discuss/2020-February/001166.html` (2020), accessed: 2020-06-27.
[11] J. M. Pollard, Mathematical Proceedings of the Cambridge Philosophical Society **76**, 521 (1974).
[12] R. L. Rivest, A. Shamir, and L. Adleman, Communications of the ACM **21**, 120 (1978).

Supplementary Information: Classically-Verifiable Quantum Advantage from a Computational Bell Test

Gregory D. Kahanamoku-Meyer,[1] Soonwon Choi,[1] Umesh V. Vazirani,[2] and Norman Y. Yao[1]

[1] *Department of Physics, University of California at Berkeley, Berkeley, CA 94720*
[2] *Department of Electrical Engineering and Computer Science, University of California at Berkeley, Berkeley, CA 94720*

arXiv:2104.00687v1 [quant-ph] 1 Apr 2021

CRYPTOGRAPHIC PROOFS OF TCF PROPERTIES

Here we prove the cryptographic properties of the trapdoor claw-free functions (TCFs) presented in the Methods section of the main text. We base our definitions on the Noisy Trapdoor Claw-free Function family (NTCF) definition given in Definition 3.1 of Ref. [1], with certain modifications such as removing the adaptive hardcore bit requirement and the "noisy" nature of the functions.

We emphasize that in the definitions below, we define security only against classical attackers. Both the $x^2 \bmod N$ and DDH constructions could be trivially defeated by a quantum adversary via Shor's algorithm; since the purpose of the protocol in this paper is to demonstrate quantum capability, this type of adversary is allowed.

TCF definition

We use the following definition of a Trapdoor Claw-free Function family:

Definition 1. *Let λ be a security parameter, K a set of keys, and X_k and Y_k finite sets for each $k \in K$. A family of functions*

$$\mathcal{F} = \{f_k : X_k \to Y_k\}_{k\in K}$$

is called a trapdoor claw free (TCF) family if the following conditions hold:

1. ***Efficient Function Generation.*** *There exists an efficient probabilistic algorithm* Gen *which generates a key $k \in K$ and the associated trapdoor data t_k:*

$$(k, t_k) \leftarrow \mathsf{Gen}(1^\lambda)$$

2. ***Trapdoor Injective Pair.*** *For all keys $k \in K$, the following conditions hold:*

   (a) *Injective pair: Consider the set R_k of all tuples (x_0, x_1) such that $f_k(x_0) = f_k(x_1)$. Let $X'_k \subseteq X_k$ be the set of values x which appear in the elements of R_k. For all $x \in X'_k$, x appears in exactly one element of R_k; furthermore, $\lim_{\lambda\to\infty} |X'_k|/|X_k| = 1$.*

   (b) *Trapdoor: There exists an efficient deterministic algorithm T such that for all $y \in Y_k$ and (x_0, x_1) such that $f_k(x_0) = f_k(x_1) = y$, $T(t_k, y) = (x_0, x_1)$.*

3. ***Claw-free.*** *For any non-uniform probabilistic polynomial time (nu-PPT) classical Turing machine $\mathcal{A}$, there exists a negligible function $\epsilon(\cdot)$ such that*

$$\Pr\left[f_k(x_0) = f_k(x_1) \wedge x_0 \neq x_1 | (x_0, x_1) \leftarrow \mathcal{A}(k)\right] < \epsilon(\lambda)$$

   where the probability is over both choice of k and the random coins of $\mathcal{A}$.

4. ***Efficient Superposition.*** *There exists an efficient quantum circuit that on input a key k prepares the state*

$$\frac{1}{\sqrt{|X_k|}} \sum_{x\in X_k} |x\rangle |f_k(x)\rangle$$

Proof of $x^2 \bmod N$ TCF

In this section we prove that the function family $\mathcal{F}_{\text{Rabin}}$ (defined in Methods) is a TCF by demonstrating each of the properties of Definition 1. Most of the properties follow directly from properties of the Rabin cryptosystem [2]; we reproduce several of the arguments here for completeness.

Theorem 1. *The function family $\mathcal{F}_{\text{Rabin}}$ is trapdoor claw-free, under the assumption of hardness of integer factorization.*

Proof. We demonstrate each of the properties of Definition 1:

1. **Efficient Function Generation.** Sampling large primes to generate p, q and N is efficient [2].

2. **Trapdoor Injective Pair.**

   (a) Injective pair: For any $y \in Y_k$, consider the two values $a < p/2$ and $b < q/2$ such that $a^2 \equiv y \bmod p$ and $b^2 \equiv y \bmod q$. These values exist because y is a quadratic residue modulo pq, therefore it is also a quadratic residue modulo p and q. Define $c \equiv 1 \bmod p \equiv 0 \bmod q$ and $d \equiv 0 \bmod p \equiv 1 \bmod q$. The following four values x in the range $[0, N)$ have $x^2 \equiv y \bmod N$: $ac+bd, ac-bd, -ac+bd, -ac-bd$. Exactly two of these values are in the range $[0, N/2)$ of the TCF, and constitute the injective pair.

   (b) Trapdoor: Because p and q were selected to have $p \equiv q \equiv 3 \bmod 4$, a and b in the expressions above can always be computed as $a = y^{(p+1)/4} \bmod p$ and $b = y^{(q+1)/4} \bmod q$.

3. **Claw-free.** We show that knowledge of a claw x_0, x_1 can be used directly to factor N. Writing the claw as $(ac+bd, ac-bd)$ using the values a, b, c, d from above, we have $x_0 + x_1 = 2ac$. Because $c = 0 \bmod q$, $\gcd(x_0 + x_1, N) = q$ can be efficiently computed, which then also yields $p = N/q$. Thus, an algorithm that could be used efficiently to find claws could be equally used to efficiently factor N, which we assume to be hard.

4. **Efficient Superposition.** The set of preimages X is the set of integers $[0, N/2)$. A uniform superposition $\sum_{x\in X} |x\rangle$ may be computed by generating a uniform superposition of all bitstrings of length n (via Hadamard gate on every qubit), and then evaluating a comparator circuit that generates the state $\sum |x\rangle |x < N/2\rangle$ where $|x < N/2\rangle$ is a bit on an ancilla. If this ancilla is then measured and the result is $|1\rangle$, the state is collapsed onto the superposition $\sum_{x\in X} |x\rangle$ (if the result is $|0\rangle$ the process should simply be repeated). Then a multiplication circuit to an empty register may be executed to generate the desired state $\sum_{x\in X} |x\rangle |x^2 \bmod N\rangle$.

□

Proof of Decisional Diffie-Hellman TCF

We now prove that $\mathcal{F}_{\text{DDH}}$ (defined in Methods) forms a trapdoor claw-free function family.

Theorem 2. *The function family $\mathcal{F}_{\text{DDH}}$ is trapdoor claw-free, under the assumption of hardness of the decisional Diffie-Hellman problem for the group $\mathbb{G}$.*

Proof. We demonstrate each of the properties of Definition 1:

1. **Efficient Function Generation.** Each step of Gen is efficient by inspection.

2. **Trapdoor Injective Pair.**

   (a) Injective pair: First we note that the matrix $\mathbf{M}$ is chosen to be invertible, thus f_0 and f_1 are one-to-one. Therefore for all $\mathbf{x}_0 \in X_k$, at most one other preimage $\mathbf{x}_1 \in X_k$ has $f_k(\mathbf{x}_0) = f_k(\mathbf{x}_1)$. Furthermore, since colliding pairs have the structure $(0||\mathbf{x}'_0), (1||\mathbf{x}'_1)$ with $\mathbf{x}'_0 = \mathbf{x}'_1 + \mathbf{s}$ and $\mathbf{s} \in \{0,1\}^k$, the only preimages that will *not* form part of a colliding pair are those where $\mathbf{x}'_0$ has a zero element at an index where $\mathbf{s}$ is nonzero, or $\mathbf{x}'_1$ has an element equal to $m-1$ where $\mathbf{s}$ is nonzero (the vector element will be outside of the range of vector elements for the other vector). The fraction of preimages for which this occurs is $\delta(\lambda) = 1 - (1 - 1/m)^k$. Since $m \sim \mathcal{O}(k^2)$ and $k \sim \mathcal{O}(\lambda)$, we have $\lim_{\lambda\to\infty} \delta(\lambda) = 0$.

(b) Trapdoor: The steps of the algorithm T are efficient by inspection. Crucially, the discrete logarithm of each vector element is possible by brute force, because the elements of $\mathbf{x}_0$ only take values up to polynomial in λ.

3. **Claw-free.** An algorithm which could efficiently compute a claw $(0||\mathbf{x}'_0, 1||\mathbf{x}'_1)$ could then trivially compute the secret vector $\mathbf{s} = \mathbf{x}'_0 - \mathbf{x}'_1$. For any matrix $\mathbf{M}'$, the existence of an algorithm to uniquely determine $\mathbf{s}$ from $(g^{\mathbf{M}'}, g^{\mathbf{M}'\mathbf{s}})$ would directly imply an algorithm for determining whether $\mathbf{M}'$ has full rank. But DDH implies it is computationally hard to determine whether a matrix $\mathbf{M}'$ is invertible given $g^{\mathbf{M}'}$ [3, 4]. Therefore DDH implies the claw-free property.

4. **Efficient Superposition.** Because m is a power of two, a superposition of all possible preimages $\mathbf{x}$ can be computed by applying Hadamard gates to every qubit in a register all initialized to $|0\rangle$. The function f can then be computed by a quantum circuit implementing a classical algorithm for the group operation of $\mathbb{G}$.

□

EXPLANATION OF CIRCUIT COMPLEXITIES

Here we describe each of the asymptotic circuit complexities listed in Table I of the main text. For these estimates we drop factors of $\log\log n$ or less. In all cases, we assume integer multiplication can be performed in time $\mathcal{O}(n \log n)$ using the Schonhage-Strassen algorithm.

We emphasize that the value of n necessary to achieve classical hardness in practice varies widely among these functions, and also that the asymptotic complexities here may not be applicable at practical values of n.

LWE [1, 5] The LWE cost is dominated by multiplying an $\mathcal{O}(n \log n) \times n$ matrix of integers by a length n vector. The integers are of length $\log n$, so each multiplication is expected to take approximately $\mathcal{O}(\log n)$ time. Thus, the evaluation of the entire function requires $\mathcal{O}(n^2 \log^2 n)$ operations.

$\boldsymbol{x^2}$ **mod** $\boldsymbol{N}$ [2] The function can be computed in time $\mathcal{O}(n \log n)$ using Schonhage-Strassen multiplication algorithm and Montgomery reduction for the modulus.

Ring-LWE [6–9] Ring-LWE is dominated by the cost of multiplying one polynomial by $\log n$ other polynomials. Through Number Theoretic Transform techniques similar to the Schonhage-Strassen algorithm, each polynomial multiplication can be performed in time $\mathcal{O}(n \log n)$, so the total runtime is $\mathcal{O}(n \log^2 n)$. We note that integer multiplication and polynomial multiplication can be mapped onto each other, so the runtimes for x^2 mod N and Ring-LWE scale identically except for the fact that Ring-LWE requires $\log n$ multiplications instead of $\mathcal{O}(1)$.

Diffie-Hellman [3, 4, 10] The Diffie-Hellman based construction defined in Methods requires performing multiplication of a $k \times k$ matrix by a vector, with $k \sim \mathcal{O}(n)$. However, the "addition" operation for the matrix-vector multiply is the group operation of $\mathbb{G}$; we expect this operation to have complexity at least $\mathcal{O}(n \log n)$ (for e.g. integer multiplication). The exponentiation operations have exponent at most $m \sim \mathcal{O}(k^2)$, so can be performed in $\mathcal{O}(\log n)$ group operations. So, for each of the k^2 matrix elements one must perform an operation of complexity $\mathcal{O}(n \log^2 n)$, yielding a total complexity of $\mathcal{O}(n^3 \log^2 n)$.

Shor's Algorithm [11] Allowing for the use of Schonhage-Strassen integer multiplication, Shor's algorithm requires $\mathcal{O}(n^2 \log n \log\log n)$ gates [12].

OPTIMAL CLASSICAL ALGORITHM

Here we provide an example of a classical algorithm that saturates the probability bound of Theorem 2 of the main text. It has $p_x = 1$ and $p_m = 3/4$.

For a TCF $f : X \to Y$, consider a classical prover that simply picks some value $x_0 \in X$, and then computes y as $f(x_0)$, without ever having knowledge of x_1. If the verifier requests a projective x measurement, they always return x_0, causing the verifier to accept with $p_x = 1$. In the other case (performing rounds 2 and 3 of the protocol), upon receiving r they compute $b_0 = x_0 \cdot r$. The cheating prover now simply assumes that $x_0 \cdot r = x_1 \cdot r$, and thus that the correct single-qubit state that would be held by a quantum prover is $|b_0\rangle$, and returns measurement outcomes accordingly. With probability $1/2$, $|b_0\rangle$ is in fact the correct single-qubit state; in this case they can always cause the verifier to accept. On the other hand, if $x_0 \cdot r \neq x_1 \cdot r$, the correct state is either $|+\rangle$ or $|-\rangle$. With probability $\frac{1}{2}$, the measurement outcome reported by the cheating prover will happen to be correct for this state too. Overall, this cheating prover will have $p_m = (1 + \frac{1}{2})/2 = \frac{3}{4}$.

Thus we see $p_x + 4p_m - 4 = 1 + 4 \cdot \frac{3}{4} - 4 = 0$ which saturates the bound.

QUANTUM CIRCUITS FOR KARATSUBA AND SCHOOLBOOK MULTIPLICATION

Classically, multiplication of large integers is generally performed using recursive algorithms such as Schonhage-Strassen [13] and Karatsuba which have complexity as low as $\mathcal{O}(n \log n)$. In the quantum setting, the need to store garbage bits at each level of recursion has limited their usefulness [14, 15]. There does exist a reversible construction of Karatsuba multiplication that uses a linear number of qubits [16], but due to overhead required for its implementation it does not begin to outperform schoolbook multiplication until the problem size reaches tens of thousands of bits.

Leveraging the irreversibility described in Section IID of the main text, we are able use these recursive algorithms directly, without needing to maintain garbage bits for later uncomputation. We implement both the $\mathcal{O}(n^{1.58})$ Karatsuba multiplication algorithm and the simple $\mathcal{O}(n^2)$ "schoolbook" algorithm. Due to efficiencies gained from discarding garbage bits, we find that the Karatsuba algorithm already begins to outcompete schoolbook multiplication at problem sizes of under 100 bits. Thus Karatsuba seems to be the best candidate for "full-scale" tests of quantum advantage at problem sizes of $n \sim 500 - 1000$ bits. We also note that the Schonhage-Strassen algorithm scales even better than Karatsuba as $\mathcal{O}(n \log n \log\log n)$. However, even in classical applications it has too much overhead to be useful at these problem sizes. We leave its potential quantum implementation to a future work.

The multiplication algorithms just described do not include the modulo N operation, it must be performed in a separate step. We implement the modulo using only two classical-quantum multiplications and one addition via Montgomery reduction [17]. Montgomery reduction does introduce a constant R' into the product, but this factor can be removed in classical post-processing after $y = x^2 R' \bmod N$ is measured.

Finally, we note that at the implementation level, optimizing *classical* circuits for modular integer multiplication has received significant study in the context of performing cryptography on embedded devices and FPGAs [18–20]. Mapping such optimized circuits into the quantum context may be a promising avenue for further research.

DETAILS OF POST-SELECTION SCHEME

In this section we describe several details of the post-selection scheme proposed in Section IIC of the main text.

Quantum prover with no phase coherence saturates the classical bound

Consider the two states $|\psi_\pm\rangle = (|x_0\rangle \pm |x_1\rangle)_{\mathsf{x}} |y\rangle_{\mathsf{y}}$ for some claw (x_0, x_1) with $y = f_k(x_0) = f_k(x_1)$. Note that $|\psi_+\rangle$ is the state that would be held by a noise-free prover. Suppose a noisy quantum prover is capable of generating the mixed state

$$\rho_\delta = (1/2 + \delta) |\psi_+\rangle \langle\psi_+| + (1/2 - \delta) |\psi_-\rangle \langle\psi_-| \,. \tag{1}$$

In words, they are able to generate a state that is a superposition of the correct bitstrings, but with the correct phase only $1/2 + \delta$ fraction of the time. Here we show that such a prover can exceed the classical threshold of Theorem 2 of the main text, whenever $\delta > 0$. We proceed by examining this prover's behavior during the protocol.

First, we note that if the verifier requests a projective x measurement after Round 1 of the protocol, this prover will always succeed—they simply measure the x register as instructed, and the phase is not relevant. Thus, using the notation of Theorem 2, $p_x = 1$. With this value set, to exceed the bound we must achieve $p_m > 3/4$. Naively performing the rest of the protocol as described in the main text does not exceed the bound when δ is small. However, the noisy prover can exceed the bound if they adjust the angle of their measurements in the third round of protocol (but preserve the sign of the measurement requested by the prover). We now demonstrate how.

Define $|\phi\rangle$ as the "correct" single-qubit state at the end of Round 2—one of $\{|0\rangle, |1\rangle, |+\rangle, |-\rangle\}$. Let $f_\updownarrow$ be the probability that our noisy prover holds the correct state when $|\phi\rangle \in \{|0\rangle, |1\rangle\}$, and $f_\leftrightarrow$ the corresponding probability when $|\phi\rangle \in \{|0\rangle, |1\rangle\}$. In the first case, the potential phase error of our prover does not affect the single-qubit state, so $f_\updownarrow = 1$. In the other case, the state is only correct when the phase is correct, so $f_\leftrightarrow = 1/2 + \delta$. We see that our prover will hold the correct single-qubit state with probability greater than 3/4. But, if they naively measure in the prescribed off-diagonal basis, for small δ their success probability will be less than 3/4. This can be rectified by adjusting the rotation angle of the measurement basis.

Letting $\pm\theta$ define the pair of measurement angles m used in step 3 of the protocol (nominally $\theta = \pi/4$), we can now express the prover's success probability p_m as

$$p_m = \frac{1}{2}\left[\cos^2\left(\frac{\theta}{2}\right) f_{\updownarrow} + \cos^2\left(\frac{\theta}{2} - \frac{\pi}{4}\right) f_{\leftrightarrow} + \sin^2\left(\frac{\theta}{2}\right)(1 - f_{\updownarrow}) + \sin^2\left(\frac{\theta}{2} - \frac{\pi}{4}\right)(1 - f_{\leftrightarrow})\right] \tag{2}$$

If the prover measures with $\theta = \pi/4$ as prescribed in the protocol, the success rate will be $p_m \approx 0.68 + \mathcal{O}(\delta) < 3/4$. However, if they instead adjust their measurement angle to $\theta = \delta$, they instead achieve $p_m = 3/4 + 3\delta^2/8 - \mathcal{O}(\delta^3)$, which exceeds the classical bound (provided that δ is large enough to be noticeable).

In practice, both $f_{\updownarrow}$ and $f_{\leftrightarrow}$ are likely to be less than one; the optimal measurement angle can be determined as

$$\theta_{\text{opt}} = \tan^{-1}\left(\frac{2f_{\leftrightarrow} - 1}{2f_{\updownarrow} - 1}\right) \tag{3}$$

which is the result of optimizing Equation 2 over θ. In a real experiment, it would be most effective to empirically determine $f_{\updownarrow}$ and $f_{\leftrightarrow}$ and then use Equation 3 to determine the optimal measurement angle.

Details of simulation and error model

We now describe the details of the numerical simulation that was used to generate Figure 2 of the main text. For several values of the overall circuit fidelity $\mathcal{F}$, we established a per-gate fidelity as $f = \mathcal{F}^{1/N_g}$ where N_g is the number of gates in the $x^2 \bmod N$ circuit. We then generated a new circuit to compute the function $(3^m x)^2 \bmod 3^{2m} N$ for various values of m (see next subsection for an explanation of the choice $k = 3^m$). For each gate in the new circuit, with probability $1 - f$ we added a Pauli "error" operator randomly chosen from $\{X, Y, Z\}$ to one of the qubits to which the gate was being applied.

For the simulation, we randomly chose two primes p and q that multiplied to yield an integer N of length 512 bits. We then randomly chose a large set of colliding preimage pairs, and simulated the circuit separately for each such preimage (which is classically efficient, since the circuits only consist of X, CNOT, and Toffoli gates). The relative phase between each pair of preimages (due to error gates) was tracked explicitly during the simulation. Finally, the expected success rate of the prover was determined by analyzing the correctness of the bitstrings and their relative phase at the end of the circuit.

Choice of $k = 3^m$ to improve postselection for $x^2 \bmod N$

In the previous subsection, we map the TCF $f_N = x^2 \bmod N$ to the function $f'_N = (kx)^2 \bmod k^2 N$. To achieve this at the implementation level, we may use essentially the same circuit for modular multiplication; the only new requirement is to efficiently generate a superposition of multiples kx in the x register. We generate this superposition by starting with a uniform superposition over values x and then multiplying by k.

Normally, quantum multiplication circuits (like those we use to evaluate $x^2 \bmod N$) perform an out-of-place multiplication, where the result is stored in a new register. In this case, however, it is preferable to do the multiplication "in-place," where the result is stored in the input register itself—this way the y value is computed directly from the input register and thus is more likely to reflect errors that may occur in the input.

In general, performing in-place multiplication is complicated, particularly on a quantum register, because the input is being modified as it is being consumed (not to mention concerns about reversibility). However, multiplication by small constants is much simpler to implement. By setting k to a power of three, we are able to implement the in-place multiplication by performing a sequence of in-place multiplications by 3, which can each be performed quite efficiently (see implementation in the attached `Cirq` code [21]).

[1] Z. Brakerski, P. Christiano, U. Mahadev, U. Vazirani, and T. Vidick, arXiv:1804.00640 [quant-ph] (2019).
[2] M. O. Rabin, *Digitalized signatures and public-key functions as intractable as factorization*, Technical Report (Massachusetts Institute of Technology, USA, 1979).

[3] C. Peikert and B. Waters, in *Proceedings of the fortieth annual ACM symposium on Theory of computing*, STOC '08 (Association for Computing Machinery, New York, NY, USA, 2008) pp. 187–196.
[4] D. M. Freeman, O. Goldreich, E. Kiltz, A. Rosen, and G. Segev, in *Public Key Cryptography – PKC 2010*, Lecture Notes in Computer Science, edited by P. Q. Nguyen and D. Pointcheval (Springer, Berlin, Heidelberg, 2010) pp. 279–295.
[5] O. Regev, in *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, STOC '05 (Association for Computing Machinery, New York, NY, USA, 2005) pp. 84–93.
[6] Z. Brakerski, V. Koppula, U. Vazirani, and T. Vidick, arXiv:2005.04826 [quant-ph] (2020).
[7] V. Lyubashevsky, C. Peikert, and O. Regev, *On Ideal Lattices and Learning with Errors Over Rings*, Tech. Rep. 230 (2012).
[8] R. de Clercq, S. S. Roy, F. Vercauteren, and I. Verbauwhede, in *Proceedings of the 2015 Design, Automation & Test in Europe Conference & Exhibition*, DATE '15 (EDA Consortium, San Jose, CA, USA, 2015) pp. 339–344.
[9] S. S. Roy, F. Vercauteren, N. Mentens, D. D. Chen, and I. Verbauwhede, in *Cryptographic Hardware and Embedded Systems – CHES 2014*, Lecture Notes in Computer Science, edited by L. Batina and M. Robshaw (Springer, Berlin, Heidelberg, 2014) pp. 371–391.
[10] W. Diffie and M. Hellman, IEEE Transactions on Information Theory **22**, 644 (1976).
[11] P. W. Shor, SIAM Journal on Computing **26**, 1484 (1997).
[12] C. Zalka, arXiv:quant-ph/9806084 (1998).
[13] A. Schönhage and V. Strassen, Computing **7**, 281 (1971).
[14] L. Kowada, R. Portugal, and C. Figueiredo, J. UCS **12**, 499 (2006).
[15] A. Parent, M. Roetteler, and M. Mosca, arXiv:1706.03419 [quant-ph] (2017).
[16] C. Gidney, arXiv:1904.07356 [quant-ph] (2019).
[17] P. L. Montgomery, Mathematics of Computation **44**, 519 (1985).
[18] K. Javeed, D. Irwin, and X. Wang, in *Cloud Computing and Security*, Lecture Notes in Computer Science, edited by X. Sun, A. Liu, H.-C. Chao, and E. Bertino (Springer International Publishing, Cham, 2016) pp. 251–260.
[19] M. Morales-Sandoval and A. Diaz-Perez, IET Computers & Digital Techniques **10**, 102 (2016).
[20] Y. Yang, C. Wu, Z. Li, and J. Yang, Microprocessors and Microsystems **47**, 209 (2016).
[21] Code is available at `https://github.com/GregDMeyer/quantum-advantage` and is archived on Zenodo [22].
[22] G. D. Kahanamoku-Meyer, Gregdmeyer/quantum-advantage: v1.0 (2021), doi:10.5281/zenodo.4641454.