

Interactive Proofs of Quantumness via Mid-Circuit Measurements

Daiwei Zhu^{1,2,9,12†*}, Gregory D. Kahanamoku-Meyer^{3,4†}, Laura Lewis^{5,6}, Crystal Noel^{1,7,8}, Or Katz^{7,8}, Bahaa Harraz¹, Qingfeng Wang^{1,2,11}, Andrew Risinger^{1,2}, Lei Feng^{1,2}, Debopriyo Biswas^{1,2}, Laird Egan^{1,2}, Alexandru Gheorghiu^{5,10}, Yunseong Nam⁹, Thomas Vidick⁵, Umesh Vazirani^{3,4}, Norman Y. Yao^{3,4}, Marko Cetina^{1,7}, Christopher Monroe^{1,2,7,8,9}

¹Joint Quantum Institute, Departments of Physics and Electrical and Computer Engineering, University of Maryland, College Park, MD 20742, USA

²Joint Center for Quantum Information and Computer Science, NIST/University of Maryland, College Park, MD 20742, USA

³Department of Physics, University of California, Berkeley, CA 94720, USA

⁴Materials Sciences Division, Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA

⁵Department of Computing and Mathematical Sciences, California Institute of Technology, CA 91125, USA

⁶Division of Physics, Mathematics, and Astronomy, California Institute of Technology, CA 91125-0001, USA

⁷Duke Quantum Center and Department of Physics, Duke University, Durham, NC 27708, USA

⁸Department of Electrical and Computer Engineering, Duke University, Durham, NC 27708, USA

⁹IonQ, Inc., College Park, MD 20740, USA

¹⁰Institute for Theoretical Studies, ETH Zürich, CH 8001, Switzerland

¹¹Chemical Physics Program and Institute for Physical Science and Technology, University of Maryland, College Park, MD 20742, USA ¹²Departments of Electrical and Computer Engineering, University of Maryland, College Park, MD 20742, USA

*To whom correspondence should be addressed. E-mail: daiwei@terpmail.umd.edu.

†These authors contributed equally.

(Dated: 23 June 2022)

Interaction is a powerful resource for quantum computation. It can be utilized in applications ranging from the verification of quantum algorithms all the way to verifying quantum mechanics itself. Here, we present the first implementation of interactive protocols for proofs of quantumness — which when suitably scaled promise the efficient verification of quantum computational advantage. The key feature that distinguishes these protocols from existing demonstrations of quantum advantage is that the classical verification is efficient and scales polynomially rather than exponentially with the number of qubits. The experimental implementation of such interactive protocols requires the ability to independently measure subsets of qubits in the middle of a quantum circuit and to continue coherent evolution afterwards. This is achieved by spatially isolating target qubits via shuttling, and opens the door to a range of quantum interactive protocols as well as new information processing architectures.

To date, the field of experimental quantum computation has largely operated in a non-interactive paradigm, where classical data is extracted from the computation only at the very last step. While this has led to many exciting advances, it has also become clear that in practice, interactivity—made possible by mid-circuit measurements performed on the quantum device—will be crucial to the operation of useful quantum computers. For example, within quantum error correction, mid-circuit measurements are used to project the state onto a single error syndrome, which can then be corrected^{1,2}. Certain quantum machine learning algorithms also leverage mid-circuit measurements to introduce essential non-linearities³. Recent work has shown that interaction can do much more: it has emerged as an indispensable tool for verifying the behavior of untrusted quantum devices⁴⁻⁶, and even for testing the fundamentals of quantum mechanics itself⁷.

Ultimately, any quantum computer is an interaction between a classical machine and a much more powerful quantum system with its exponentially large Hilbert space. Intuitively this asymmetry poses an insurmountable barrier for the classical machine’s ability to certify the behavior of an untrusted quantum device. This challenge shadows one pursued in the field of classical computing, which asks whether a skeptical, computationally-bounded “verifier,” who is not powerful enough to validate a given statement on their own, can be convinced of its veracity by a more powerful but untrusted “prover.” Several decades ago, the field of classical complexity theory began to pursue this idea through a novel tool called

an *interactive proof*. In these protocols, the verifier’s goal is to accept only valid statements, regardless of whether the prover behaves honestly or attempts to cheat. The crown jewel of computational complexity theory is a set of results showing that in certain scenarios multiple rounds of interaction allow the verifier to detect cheating by even *arbitrarily computationally powerful* provers^{8–10}. The key idea is that interaction can force the prover to *commit* to some data early in the protocol, upon which the verifier follows up with queries that can only be answered consistently if the prover is being truthful. In exciting recent developments, success has been found in the application of this idea to quantum computing: interactive proofs have been shown to allow the verification of a number of practical quantum tasks, including random number generation,⁵ remote quantum state preparation,⁶ and the delegation of computations to an untrusted quantum server.⁴ Connecting to seminal recent sampling experiments, perhaps the most direct application of an interactive protocol is for a “proof of quantumness”—the classically-verifiable demonstration of non-classical behavior from a single quantum device.

In practice, the experimental implementation of interactivity is extremely challenging. It requires the ability to independently measure subsets of qubits in the middle of a quantum circuit and to continue coherent evolution afterwards. Unfortunately, the measurement of a target qubit typically disturbs neighboring qubits, degrading the quality of computations following the mid-circuit measurement. One solution, which finds commonality among multiple quantum platforms is to spatially isolate target qubits via shuttling^{11–13}. While daunting from the perspective of quantum control, experimental progress toward coherent qubit shuttling opens the door not only to interactivity but also to distinct information processing architectures¹⁴.

In this work, we implement two complementary interactive proofs of quantumness, shown in the schematic of Fig. 1, on an ion trap quantum computer with up to 11 qubits and 145 gates. The interactions between verifier and prover are enabled by the experimental realization of mid-circuit measurements on a portion of the qubits (Fig. 2)^{2,13,15}. The first protocol involves two rounds of interaction and is based upon the learning with errors (LWE) problem^{16,17}. The LWE construction is unique because it exhibits a property known as the “adaptive hardcore bit”⁵, which enables a particularly simple measurement scheme. The second protocol circumvents the need for this special property and thus applies to a more general class of cryptographic functions; here we use a function from the Rabin cryptosystem^{18,19}.

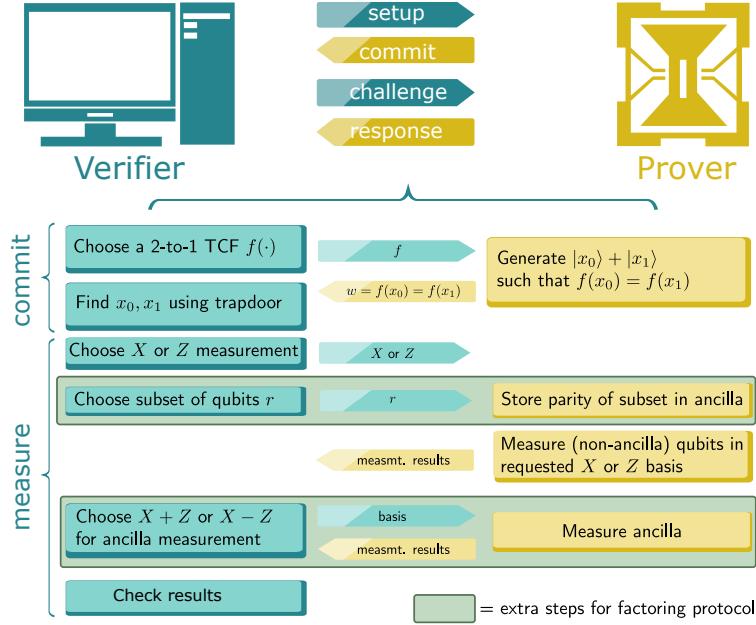


FIG. 1. Schematic of an interactive quantum verification protocol. The verifier’s goal is to test the “quantumness” of the prover through an exchange of classical information. The protocol begins with the verifier sending the prover an instance of a trapdoor claw-free function. By applying this function to a superposition of all possible inputs and projectively measuring the result, the prover commits to a particular quantum state $|x_0\rangle + |x_1\rangle$. Subsequent challenges issued by the verifier specify how to measure this state and enable the efficient validation the prover’s commitment. The LWE-based protocol requires two rounds of interaction, while the factoring-based protocol requires an additional round (green box).

By using an additional round of interaction, the cryptographic information is condensed onto the state of a single qubit. This makes it possible to implement an interactive proof of quantumness whose hardness is equivalent to that of factoring, but whose associated circuits can exhibit an asymptotic scaling much simpler than Shor’s algorithm, with the number of required gates that is almost linear in the problem size.

I. TRAPDOOR CLAW-FREE FUNCTIONS

Both interactive protocols (Fig. 1) rely upon a cryptographic primitive called a *trapdoor claw-free function* (TCF)²⁰—a 2-to-1 function f for which it is cryptographically hard to find two inputs mapping to the same output (i.e. a “claw”). The function also has a “trapdoor,”

a secret key with which it is easy to compute the pre-images x_0 and x_1 from any image $w = f(x_0) = f(x_1)$. The key intuition behind the protocols is the following: Despite the claw-free property, a quantum computer can efficiently generate a *superposition* of two pre-images that form a claw; this is most simply realized by evaluating f on a superposition of the entire domain, and then collapsing to a single image, w , via measurement. In this way, a quantum prover can generate the state $|\psi\rangle = (|x_0\rangle + |x_1\rangle)|w\rangle$, where w is the measurement result. Note that the value of w changes each time the protocol is executed, making it impossible to learn any additional information by repeating the protocol. After receiving w from the prover, the verifier can use the trapdoor to compute x_0 and x_1 , thus giving the verifier full knowledge of the prover’s quantum state. The verifier then asks the prover to measure $|\psi\rangle$. In particular, they request either a standard basis measurement (yielding x_0 or x_1 in full), or a measurement that *interferes* the states $|x_0\rangle$ and $|x_1\rangle$. In both cases, the verifier checks the measurement result on a per-shot basis.

The learning with errors problem—It is believed to be classically intractable to recover an input vector from the result of certain noisy matrix-vector multiplications—this constitutes the LWE problem^{16,17}. In particular, a secret vector, $s \in \{0, 1\}^n$, can be encoded into an output vector, $y = As + e$, where $A \in \mathbb{Z}_q^{m \times n}$ is a matrix and e is an error vector corresponding to the noise. Using the LWE problem, a TCF can be constructed as $f(b, x) = \lfloor Ax + b \cdot y \rfloor$, where b is a single bit that controls whether y gets added to Ax and $\lfloor \cdot \rfloor$ denotes a rounding operation^{21,22} (see Supplementary Information for additional details). Here, s and e play the role of the trapdoor, and a claw corresponds to colliding inputs $\{(0, x_0), (1, x_1)\}$ with $f(0, x_0) = f(1, x_1)$ and $x_0 = x_1 + s$. By implementing the protocol described above and illustrated in Figure 1, the prover is able to generate the state $|\psi\rangle = (|0, x_0\rangle + |1, x_1\rangle)|w\rangle$. For the aforementioned “interference” measurement, the prover simply measures each qubit of the superposition in the X basis. Crucially, the result of this measurement is cryptographically protected by the adaptive hardcore bit property⁵.

Rabin’s function—The function, $f(x) = x^2 \bmod N$, with N being the product of two primes, was originally introduced in the context of digital signatures^{18,19}. This function has the property that finding two colliding pre-images in the range $[0, N/2]$ is as hard as factoring N . Moreover, the prime decomposition $N = pq$ can serve as a trapdoor, enabling one to invert the function for any output. Thus, $f(x)$ is a trapdoor claw-free function. However, $f(x)$ does not have the adaptive hardcore bit property, making the simple X -basis

“interference” measurement (described in the LWE context above) not provably secure. To get around this, we perform the “interference” measurement differently. First, the verifier chooses a random subset of the qubits of the superposition, and the prover stores the parity of that subset on an ancilla. Then, the prover measures everything except the ancilla in the X basis; the polarization of this remaining ancilla qubit is cryptographically protected, in the same way that the state of one half of a Bell pair is protected (i.e. by “no communication”) when the other half is measured. Following this intuition, the verifier requests a measurement of the ancilla qubit in the $Z + X$ or $Z - X$ basis, effectively completing the Bell test^{23,24}; the verifier accepts if the prover returns the more likely measurement outcome. Crucially, the dependence of the measurement result on the claw renders it infeasible to guess classically²⁵.

II. IMPLEMENTING AN INTERACTIVE PROOF OF QUANTUMNESS

In order to implement an interactive proof of quantumness, we design quantum circuits for both the LWE- and factoring-based protocols. The high-level circuit diagrams are shown in Figs. 3(a,b). In both cases, the circuits are composed of several sections. First, the prover creates a uniform superposition $|\psi\rangle = \sum_{x=0}^{2^n-1} |x\rangle$ via Hadamard gates, where n is the number of input qubits. Then, they compute the TCF on an output register using this superposition as input [Fig. 3(a,d)], thereby generating the state $|\psi\rangle = \sum_x |x\rangle |f(x)\rangle$. Next, the prover performs a mid-circuit measurement on the output register, collapsing the state to $|\psi\rangle = (|x_0\rangle + |x_1\rangle) |w\rangle$. Finally, based on the verifier’s choice of measurement scheme (i.e. standard vs. interference), the prover must perform additional coherent gates and measurements (see Methods for a full description of the quantum circuits used).

We implement both interactive protocols using an ion trap quantum computer, with a base chain length of 15 ions (Fig. 2); for each $^{171}\text{Yb}^+$ ion, a qubit is encoded in a pair of hyperfine levels²⁶. The quantum circuits are implemented via the consecutive application of native single and two-qubit gates using individual optical addressing [Fig. 2(a)]²⁷. In order to realize rapid successive two-qubit interactions, we position the ions in a single, closely-spaced linear chain [Fig. 2(d)].

This geometry makes it challenging to implement mid-circuit measurements: light scattered from nearby ions during a state-dependent fluorescence measurement can destroy the state of the other ions. To overcome this issue, we vary the voltages on the trap electrodes

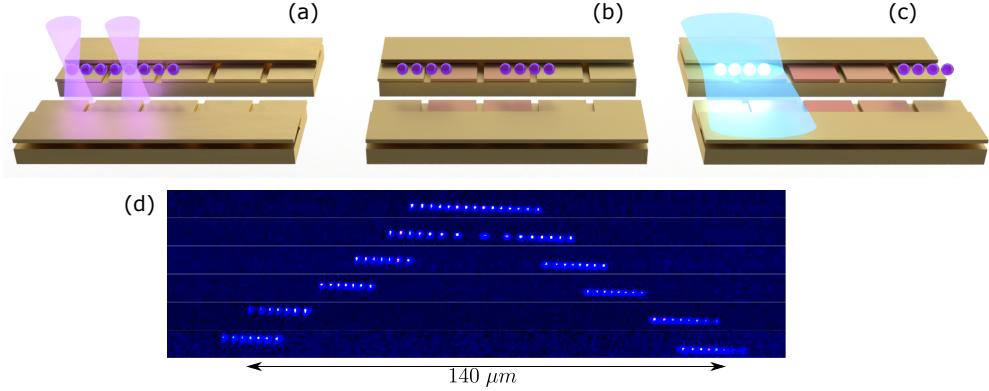


FIG. 2. (a-c) Schematic illustration of our mid-circuit measurement protocol. (a) To start, the ions are closely-spaced in a 1D chain above a surface trap. Coherent gates are implemented via a combination of individual addressing beams (purple) and global beams (not shown). (b) By tuning the electrodes of the surface trap, one can adjust the potential to deterministically split the ion chain. Depending on the protocol, we split the chain into either two or three individual segments. We optimize the rate of shuttling to minimize the perturbation of the motional state. (c) Once the segments are sufficiently far away from one another, it is possible to measure (blue beam) an individual segment without disturbing the coherence of the remaining ions. After the measurement, the shuttling is reversed and the ion chain is recombined. (d) Fluorescence image of an example shuttling protocol for a chain of $N = 15$ ions. At the start, the average spacing between ions is $\sim 4\mu\text{m}$. At the end of the splitting procedure, the distance between the two segments is $\sim 550\mu\text{m}$. Shown is the splitting up to a distance of $\sim 140\mu\text{m}$ until the two sub-chains reach the edge of the detection beam.

to split and shuttle the ion chain, thereby spatially isolating the ions not being measured (Fig. 2a-c). Depending on the protocol, the ion chain is split into either two or three segments. To measure the ions in a particular segment, we re-shape the electric potential to align the target segment with the detection system. In addition, we calibrate and correct for spatial drifts of the optical beams, variations of stray fields, and unwanted phase accumulation during shuttling (see supplementary information sections VB, VE for additional details).

In this demonstration, the qubits play the role of the prover and the classical control system plays the role of the verifier. This allows us to compile the decisions of the verifier

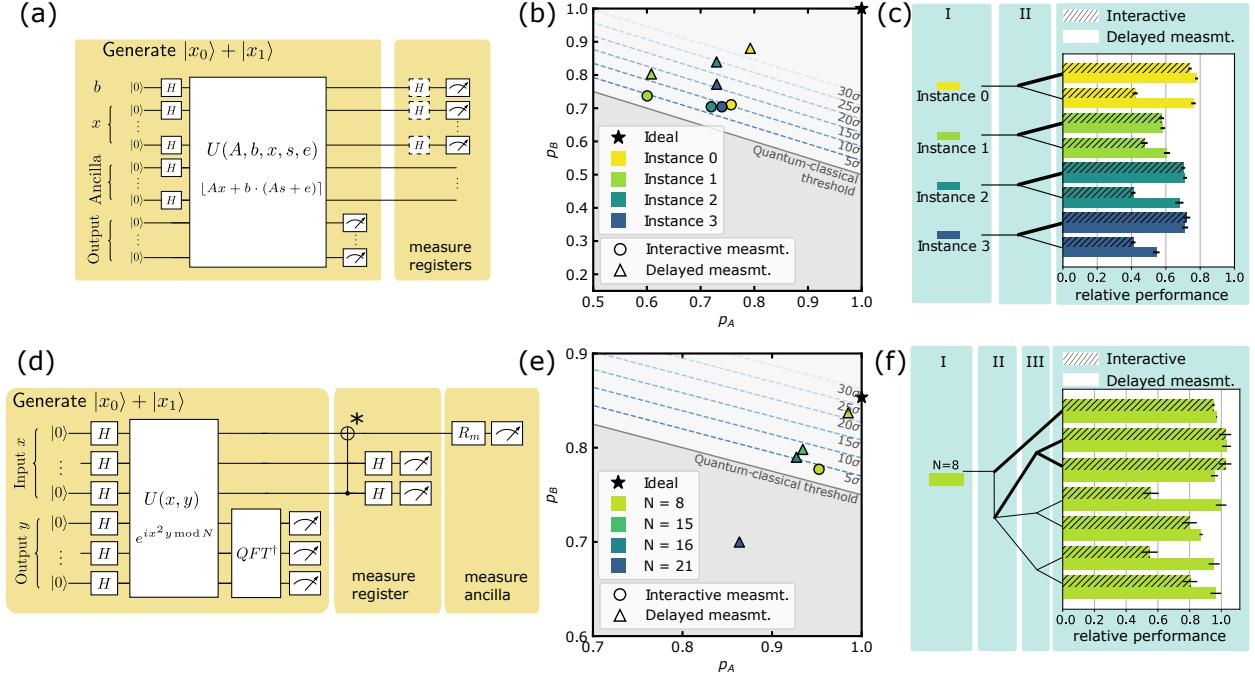


FIG. 3. (a),(d) depict the circuit diagrams for the LWE- and factoring-based protocols, respectively. Details about the implementation of $U(A, b, x, y)$ and $U(x, y)$ are provided in the supplementary information. In (d), the CNOT gate marked with an asterisk stands for the operations needed to store the parity of selected qubits into the ancilla. To reduce the impact of shuttling-induced gate fidelity degradation, we compute the parity for all of the verifier’s possible selections and then choose the relevant one once the prover receives the challenge. (b),(e): Experimentally measured probabilities of passing the standard-basis (p_A) and interference measurement (p_B) challenges for the LWE- and factoring-based protocols. These probabilities are compared against the asymptotic classical limits ($p_A + 2p_B \leq 2$ for LWE, derived in the supplementary information section VI, and $p_A + 4p_B \leq 4$ for factoring²⁵). Results for both interactive and delayed-measurement version of the protocols are presented. Numerical values of p_A and p_B for each experiment, and the corresponding values of statistical significance, are provided in the supplemental information. (c),(f): The relative performance, R , of the experiments for all possible branches. Certain branches (thick lines) are robust to phase errors and exhibit similar performance for both interactive and delayed-measurement protocols.

into the classical controller prior to execution of the quantum circuit.

III. BEATING THE CLASSICAL THRESHOLD

Much like a Bell test, even a classical prover can pass the verifier’s challenges with finite probability. Assuming that the trapdoor claw-free function is secure, this probability can be bounded by an asymptotic “classical threshold”—which a quantum prover must exceed to demonstrate advantage. For both protocols, this threshold is best expressed in terms of the probabilities of passing the verifier’s “standard basis” and “interference” checks, which we denote as p_A and p_B , respectively (see supplementary material Section V C for the definition of the verifier’s checks). For the LWE-based protocol, the classical threshold is given by $p_A + 2p_B - 2 \leq \epsilon$ (derivation in Supplementary Information); for the factoring-based protocol, it is given by $p_A + 4p_B - 4 \leq \epsilon$.²⁵ In both cases, ϵ is a function which goes to zero exponentially in the problem size. An intuition for the difference between the thresholds is that the factoring-based protocol requires an additional round of interaction during the “interference” test.

As depicted in Figure 3(b), we perform multiple instances of the LWE-based protocol for different matrices A and noise vectors e . For each of the verifier’s possible choices, we repeat the experiment $\sim 10^3$ times to collect statistics. This yields the experimental probabilities p_A and p_B , allowing us to confirm that the quantum prover exceeds the asymptotic classical threshold in all cases. The statistical significance by which the bound is exceeded (more than 6σ in all cases, see Table II in the supplementary information) is shown in Figure 3(b). Figure 3(e) depicts the analogous results for the factoring-based protocol, where the different instances correspond to different values of N . For all but $N = 21$, which requires the deepest circuit, the results exceed the asymptotic classical bound with more than 4σ statistical significance. We utilize an error-mitigation strategy based on excluding iterations where w is measured to be invalid, i.e. not in the range of f (see V D); effectively, this implements a post-selection which suppresses bit-flip errors²⁵.

To further analyze the performance of each branch of the interactive protocol, corresponding to the verifier’s choices [Figs. 3(c,f)], we define the relative performance $R = (p_{\text{exp}} - p_{\text{guess}})/(p_{\text{ideal}} - p_{\text{guess}})$ for each branch, where p_{ideal} is the probability that an error-free quantum prover would pass, p_{guess} is the probability that a random guesser would pass, and p_{exp} is the passing rate measured in the experiment. For a perfect quantum prover $R = 1$, and for a device with zero fidelity, $R = 0$.

For the LWE-based protocol, there are two rounds of interaction, corresponding to the two branches, I and II shown in Fig. 3(c), while for the factoring-based protocol there are three rounds of interaction [Fig. 3(f)]. By comparing the relative performance between the interactive and delayed-measurement versions of our experiment, we are able to probe a subtle feature of the protocols—namely, that certain branches are robust to additional decoherence induced by the mid-circuit measurements. Microscopically, this robustness arises because these branches (thick lines, Figs. 3c,f) do not depend on the phase coherence between $|x_0\rangle$ and $|x_1\rangle$. In particular, this is true for the standard-basis measurement branches in both protocols, and also for the branches of the factoring-based protocol where the ancilla is polarized in the Z basis (see supplementary information section VF). Noting that mid-circuit measurements are expected to induce mainly phase errors, one would predict that those branches insensitive to phase errors should yield similar performance in both the interactive and delayed-measurement cases. This is indeed borne out by the data.

Discussion and Outlook—The most direct application of our protocols is to validate an untrusted quantum device. Expanding the notion of an “untrusted prover,” one can use interactive protocols to test quantum mechanics itself, demonstrating quantum computational advantage. The “statement” to be proven in this case is that quantum mechanics operates as expected, even when scaled to system sizes that are impossible to classically simulate. As aforementioned, this idea is connected to recent sampling experiments, which have demonstrated the system sizes and fidelities necessary to make classical simulation extremely hard or impossible^{28–37}; however, these approaches do not operate as proofs of quantumness since there is no method to efficiently verify the output. Furthermore, practical strategies for a classical impostor to replicate the sampling are still being explored.^{38–44} To this end, as proofs of quantumness, interactive protocols satisfy three important criteria: 1) efficient (polynomial-time) classical verification, 2) classical hardness based on extremely well-studied cryptographic assumptions, and 3) the potential for realization on near-term quantum devices.^{5,25}

Our work represents the first experimental demonstration of an interactive protocol for a proof of quantumness and opens the door to a number of intriguing directions. First, by scaling up our proposed circuits [Fig. 3(a,b)], one should be able to perform a verifiable test of quantum advantage using ~ 1600 qubits (see supplementary information section VL). At these scales, the challenge on near-term devices will almost certainly be the circuit depth;

interestingly, recent advances suggest that our interactive protocols can be performed in constant depth at the cost of a larger number of qubits^{45,46}. Second, a clear next step is to apply the power of quantum interactive protocols to achieve more than just quantum advantage—for example, pursuing such tasks as certifiable random number generation, remote state preparation and the verification of arbitrary quantum computations^{4–6}. In another path forward, one can imagine generalizing our experiment to include interactions with a remote verifier, for example over the internet, which could serve as a loophole-free way of remotely benchmarking quantum cloud services. Finally, the advent of mid-circuit measurement capabilities in a number of platforms^{13,15,47,48}, enables the exploration of new phenomena such as entanglement phase transitions^{49–51} as well as the demonstration of coherent feedback protocols including quantum error correction².

IV. ACKNOWLEDGEMENTS

The authors are grateful to Vivian Uhlir for the design of the verifier and prover figures. This work is supported by the ARO through the IARPA LogiQ program, the U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers, Quantum Systems Accelerator (QSA), the AFOSR MURI on Quantum Measurement/Verification and Quantum Interactive Protocols (FA9550-18-1-0161) and Dissipation Engineering in Open Quantum Systems, the NSF STAQ Program, the ARO MURI on Modular Quantum Circuits, the DoE ASCR Accelerated Research in Quantum Computing program (award No. DE-SC0020312), the AFOSR YIP award number FA9550-16-1-0495, the NSF QLCI program through grant number OMA-2016245, the IQIM, an NSF Physics Frontiers Center (NSF Grant PHY-1125565), the Gordon and Betty Moore Foundation (GBMF-12500028), the Dr. Max Rössler, the Walter Haefner Foundation and the ETH Zürich Foundation, the NSF award DMR-1747426, a Vannever Bush Faculty Fellowship, the Office of Advanced Scientific Computing Research, under the Accelerated Research in Quantum Computing (ARQC) program, the A. P. Sloan foundation and the David and Lucile Packard Foundation. **Competing interests:** C.M. is Chief Scientist for IonQ, Inc. and has a personal financial interest in the company.

V. SUPPLEMENTARY MATERIALS

A. Result data

In Tables I and II we present the numerical results for each configuration of the experiment, along with the number of samples obtained (N_A and N_B), the measure of quantumness q , and the statistical significance of the result (see supp. info. Section V J for a description of how the significance is calculated).

We note that for the computational Bell test protocol, the sample size N_B is less than the actual number of shots that passed postselection (ultimately leading to slightly less statistical significance than might otherwise be expected). This is because the sample size varied for different values of the verifier's string r , yet we are interested in the passing rate p_B averaged uniformly over all r (not weighted by number of shots). To account for this, we simply took the r -value with the fewest number of shots, and computed N_B as if every r value had had that sample size (even if some values of r had more).

We also note that in some cases the statistical significance denoted here may be higher than that visually displayed in Figure 3 of the main text; this is because the contour lines in that figure correspond to the configuration with the smallest sample size.

N	Measurement scheme	p_A	p_B	N_A	N_B	Quantumness q	Stat. significance
8	interactive	0.952	0.777	4096	15267	0.061	4.3σ
8	delayed	0.985	0.837	2736	17361	0.334	24.1σ
15	delayed	0.934	0.798	2361	31353	0.127	10.0σ
16	delayed	0.927	0.790	3874	53550	0.087	8.8σ
21	delayed	0.864	0.700	2066	27944	-0.338	—

TABLE I. Results for various configurations of the computational Bell test protocol. For this protocol $q = p_A + 4p_B - 4$.

B. Trapped Ion Quantum Computer

The trapped ion quantum computer used for this study was designed, built, and operated at the University of Maryland and is described elsewhere^{27,52}. The system consists

Instance	Measurement scheme	p_A	p_B	N_A	N_B	Quantumness q	Stat. significance
0	interactive	0.757	0.710	8000	13381	0.178	18.6σ
0	delayed	0.793	0.880	10000	9415	0.553	60.3σ
1	interactive	0.601	0.737	8000	7622	0.075	6.2σ
1	delayed	0.608	0.803	8000	7547	0.215	18.0σ
2	interactive	0.720	0.704	14000	15310	0.129	15.0σ
2	delayed	0.730	0.839	4000	3735	0.409	24.6σ
3	interactive	0.740	0.704	8000	15189	0.148	16.2σ
3	delayed	0.730	0.772	8000	7528	0.274	23.1σ

TABLE II. Results for various configurations of the LWE-based protocol. For this protocol $q = p_A + 2p_B - 2$.

of a chain of fifteen single $^{171}\text{Yb}^+$ ions confined in a Paul trap and laser cooled close to their motional ground state. Each ion provides one physical qubit in the form of a pair of states in the hyperfine-split $^2S_{1/2}$ ground level with an energy difference of 12.642821 GHz, which is insensitive to magnetic fields to first order. The qubits are collectively initialized through optical pumping, and state readout is accomplished by state-dependent fluorescence detection⁵³. Qubit operations are realized via pairs of Raman beams, derived from a single 355-nm mode-locked laser⁵⁴. These optical controllers consist of an array of individual addressing beams and a counter-propagating global beam that illuminates the entire chain. Single qubit gates are realized by driving resonant Rabi rotations of defined phase, amplitude, and duration. Single-qubit rotations about the z-axis, are performed classically with negligible error. Two-qubit gates are achieved by illuminating two selected ions with beat-note frequencies near motional sidebands and creating an effective Ising spin-spin interaction via transient entanglement between the two ion qubits and all modes of motion^{55–57}. To ensure that the motion is disentangled from the qubit states at the end of the interaction, we used a pulse shaping scheme by modulating the amplitude of the global beam⁵⁸.

C. Verifier's checks

In this section we explicitly state the checks performed by the verifier to decide whether to accept or reject the prover's responses for each run of the protocol. We emphasize that these checks are performed on a per-shot basis, and the empirical success rates p_A and p_B are defined as the fraction of runs (after postselection, see below) for which the verifier accepted the prover's responses.

For both protocols, the “A” or “standard basis” branch check is simple. The prover has already supplied the verifier with the image value w ; for this test the prover is expected to measure a value x such that $f(x) = w$. Thus in this case the verifier simply evaluates $f(x)$ for the prover's supplied preimage x and confirms that it is equal to w .

For the “B” or “interference” measurement, the measurement scheme and verification check is different for the two protocols. For the LWE-based protocol, the interference measurement is an X -basis measurement of all of the qubits holding the preimage superposition $|x_0\rangle + |x_1\rangle$. This measurement will return a bitstring d of the same length as the number of qubits in that superposition, where for each qubit, the corresponding bit of d is 0 if the measurement returned the $|+\rangle$ eigenstate and 1 if the measurement returned the $|-\rangle$ eigenstate. The verifier has previously received the value w from the prover and used the trapdoor to compute x_0 and x_1 ; the verifier accepts the string d if it satisfies the equation

$$d \cdot x_0 = d \cdot x_1 \tag{1}$$

where (\cdot) denotes the binary inner product, i.e. $a \cdot b = \sum_i a_i b_i \bmod 2$. It can be shown that a perfect (noise-free) measurement of the superposition $|x_0\rangle + |x_1\rangle$ will yield a string d satisfying Eq. 1 with probability 1.

The interference measurement for the computational Bell test involves a sequence of two measurements (in addition to the first measurement of the string w). The first measurement yields a bitstring d as above. After performing that measurement, the prover holds the single-qubit state $(-1)^{d \cdot x_0} |r \cdot x_0\rangle + (-1)^{d \cdot x_1} |r \cdot x_1\rangle$, where (\cdot) is the binary inner product as above and r is a random bitstring supplied by the verifier. This state is one of $\{|0\rangle, |1\rangle, |+\rangle, |-\rangle\}$, and is fully known to the verifier after receiving d (via use of the trapdoor to compute x_0 and x_1). The second measurement is of this single qubit, in an intermediate basis $Z + X$ or $Z - X$ chosen by the verifier. For any of the four possible states, one eigenstate of

the measurement basis will be measured with probability $\cos^2(\pi/8) \approx 85\%$ (with the other having probability $\sim 15\%$), just as in a Bell test. The verifier accepts the measurement result if it corresponds to this more-likely result; an ideal (noise-free) prover will be accepted with probability $\sim 85\%$ (see Figure 3 of the main text).

D. Post-selection

Both the factoring-based and LWE-based protocols involve post-selection on the measurement results throughout the experiment.

Instance	Delayed Measurement	Interactive Measurement
0	3753/4000	13381/14000
1	7547/8000	7622/8000
2	3735/4000	15310/16000
3	7528/8000	15144/16000

TABLE III. This table displays the fractions of runs kept during post-selection for the LWE-based protocol, in the “interference” measurement branch. All runs are kept for the standard basis measurement.

For the factoring-based protocol, this post-selection is performed on the measured value of the output register w . Due to quantum errors in the experiment, in practice it is possible to measure a value of w that does not correspond to any preimages of the TCF—that is, there do not exist x_0, x_1 for which $f(x_0) = f(x_1) = w$, due to noise. Because such a result would not be possible without errors, measuring such a value indicates that a quantum error has occurred²⁵. Thus, we perform post-selection by discarding all runs for which the measured value w does not have two pre-images.

On the other hand, for the LWE protocol, we post-select in order to satisfy the conditions for the adaptive hardcore bit property to hold, as without this property, the protocol could be susceptible to attacks. In particular, the adaptive hardcore bit property requires that the result obtained from measuring the x register using the “interference” measurement scheme be a nonzero bitstring⁵. Hence, we simply post-select on this condition for the LWE case. Tables III and IV explicitly show how many runs are kept using each post selection scheme.

N	Interactive	Branch	Runs kept/Total	N	Interactive	Branch	Runs kept/Total
8	Yes	A	4096/9000	16	No	A	3874/6000
8	Yes	B, r=01	5093/12000	16	No	B, r=001	7842/12000
8	Yes	B, r=10	5089/12000	16	No	B, r=010	7847/12000
8	Yes	B, r=11	5492/12000	16	No	B, r=011	7732/12000
8	No	A	2736/6000	16	No	B, r=100	7936/12000
8	No	B, r=01	5787/12000	16	No	B, r=101	7870/12000
8	No	B, r=10	5818/12000	16	No	B, r=110	7841/12000
8	No	B, r=11	5865/12000	16	No	B, r=111	7650/12000
15	No	A	2361/6000	21	No	A	2066/6000
15	No	B, r=001	4636/12000	21	No	B, r=001	3992/12000
15	No	B, r=010	4532/12000	21	No	B, r=010	4273/12000
15	No	B, r=011	4666/12000	21	No	B, r=011	4137/12000
15	No	B, r=100	4496/12000	21	No	B, r=100	4182/12000
15	No	B, r=101	4727/12000	21	No	B, r=101	4193/12000
15	No	B, r=110	4479/12000	21	No	B, r=110	4261/12000
15	No	B, r=111	4673/12000	21	No	B, r=111	4221/12000

TABLE IV. Fraction of runs kept during postselection for each branch of the factoring-based protocol.

We note that in both cases, post-selection does not affect the soundness of the protocols. We only require that a non-negligible fraction of runs pass post-selection (to give good statistical significance for the results). This is indeed the case for our experiment, as can be seen in Tables III, IV, as well as the statistical significance of the results in Tables I, II.

E. Shuttling and Mid-circuit measurements

We control the position of the ions and run the split and shuttling sequences by changing the electrostatic trapping potential in a microfabricated chip trap⁵⁹ maintained at room-temperature. We generate 40 time-dependent signals using a multi channel DAC voltage source, which controls the voltages of the 38 inner electrodes at the center of the chip and the

voltages of two additional outer electrodes. Owing to the strong radial confining potential used (with secular trapping frequencies near 3 MHz), the central electrodes' potential effect predominantly the axial trapping potential, and in turn, generate movement predominantly along the linear trap axis. To maintain the ions at a constant height above the trap surface, we simulate the electric field based on the model in Ref.⁵⁹, and compensate for the average variation of its perpendicular component by controlling the voltages of the outer two electrodes.

In the first sequence, we split the 15 ion chain into two sub-chains of 7 and 8 ions, and shuttle the 8-ion group to $x = 0.55$ mm away from the trap center at $x = 0$. We then align the 7-ion chain with the individual-addressing Raman beams for the first mid-circuit measurement. For the LWE-based protocol, we then reverse the shuttling process and re-merge the ions to a 15-ion chain, completing the circuit and performing a final measurement. For the factoring-based protocol, we shuttle instead the 8-ion sub-chain to the trap center and the 7-ion sub-chain to $x = -0.55$ mm. We then split this chain into 5- and 3-ion sub-chains, shuttle the latter to $x = 0.55$ mm, and align the 5 ions at the center with the Raman beams for additional gates and a second mid-circuit measurement. Finally, we move away the measured ions and align the 3-ion group to the center of the trap to complete the protocol. Reversing of the sequence then prepares the ions in their initial state. For each protocol, all branches use the same shuttling sequences yet differ in the qubit assignment and the realized gates. The mid-circuit measurement duration was experimentally determined prior the experiment by maximizing the average fidelity of a Ramsey experiment using single-qubit gates, approximately optimizing for the trade-off between efficient detection of each sub-chain and stray light decoherence.

To enable efficient performance of the split and shuttling sequences we numerically simulate the electrostatic potential and the motional modes of the ions that are realized in the sequences. We minimize heating of the axial motion from low-frequency electric-field noise by ensuring that the calculated lowest axial frequency does not go below > 100 KHz. We also minimize frequent ions loss due to collisions with background gas by maintaining a calculated trap depth of at least 20 meV for each of the sub-chains throughout the shuttling sequences. The simulations enable efficient alignment of the sub chains with the Raman beams, taking into account the variation of the potential induced by all electrodes.

We account and correct for various systematic effects and drifts which appear in the

experiment. To eliminate the effect of systematic variation of the optical phases between the individual beams on the ions, we align each ion with the same individual beam throughout the protocol. Prior the experiment, we run several calibration protocols which estimate the electrostatic potential at the center of the trap through a Taylor series representation up to a fourth order, estimating the dominant effect of stray electric-fields on the pre-calculated potential. We then cancel the effect of these fields using the central electrodes during the alignment and split sequences, as they are most sensitive to the exact shape of the actual electrostatic potential. Additionally, we routinely measure the common-mode drift of the individually addressing optical Raman beams along the linear axis of the trap and correct for them by automatic re-positioning of the ions through variation of the potential.

During shuttling, the ions traverse an inhomogeneous magnetic field and consequently, each ion spin acquires a shuttling-induced phase $\phi_s^{(i)}$ which depends on its realized trajectory. We calibrate this by performing a Ramsey sequence in which each qubit is put in a superposition of $(|0\rangle_i + |1\rangle_i)/\sqrt{2}$ before shuttling, and after the shuttling apply $R_x^{(i)}(\pi/2)R_z^{(i)}(\phi)$ gates where ϕ is scanned between 0 to 2π . Fitting the realized fringe for each ion enables estimation of the phases $\phi_s^{(i)}$, which are corrected in the protocols by application of the inverse operation $R_z^{(i)}(-\phi_s^{(i)})$ after shuttling.

F. Circuit construction of the factoring-based protocol

In this section, we describe the procedure for generating a superposition of the claw $|x_0\rangle + |x_1\rangle$ in the factoring-based protocol, as shown in Fig. 3(a) of the main text.

This is achieved by generating

$$\sum_{0 \leq x \leq N/2} \frac{1}{\sqrt{2^{N/2}}} |x\rangle |f(x) = x^2 \bmod N\rangle. \quad (2)$$

and then measuring the $y = |f(x)\rangle$ register. We calculate $f(x)$ using a unitary $U(x, y)$ to encode the function into the phase of the y register and applying a Inverse Quantum Fourier Transform (QFT †) to extract the result.

To start, we apply Hadamard gates to all qubits to prepare a uniform superposition of all the possible bit strings for the x - and y - registers:

$$\sum_{0 \leq x \leq N/2, 0 \leq y \leq N} \alpha |x\rangle |y\rangle, \quad (3)$$

where α is the normalization factor.

Next, we evolve the state with the unitary $U(x, y) = e^{2\pi i \frac{x^2 y}{N}}$. Since the phase has period 2π , the unitary is equivalent to $U(x, y) = e^{2\pi i \frac{x^2 y \bmod N}{N}}$. We now show how to efficiently implement $U(x, y) = e^{2\pi i \frac{x^2 y}{N}}$ on the ion trap quantum computer.

First, note the multiplication in the phase can be expressed as a sum of bit-wise multiplication

$$U(x, y) = \prod_{i,j,k} \exp\left(2\pi i \frac{2^{i+j+k}}{N} x_i x_j y_k\right). \quad (4)$$

This bit-wise multiplication can be expressed using Pauli operators:

$$\prod_{i,j,k} \exp\left(2\pi i \frac{2^{i+j+k-3}}{N} (1 - \sigma_z^{(i)})(1 - \sigma_z^{(j)})(1 - \sigma_z^{(k)})\right) \quad (5)$$

We then organize the operators into three terms:

$$U(x, y) = \prod_{i,j,k} \exp(\alpha_{i,j,k} \sigma_z^{(i)} \sigma_z^{(j)} \sigma_z^{(k)}) \prod_{i,j} \exp(\beta_{i,j} \sigma_z^{(i)} \sigma_z^{(j)}) \prod_i \exp(\gamma_i \sigma_z^{(i)}). \quad (6)$$

We use α 's, β 's, and γ 's to represent the phases generated by these terms, which can be calculated from Eq.5. The third term contains single-qubit z-rotations that are implemented efficiently as software-phase-advances. The zz-interactions in the second term, are implemented as XX-gates sandwiched between single qubit rotations. The first term includes three body zzz-interactions, which can be decomposed using zz-interactions using the following relation:

$$\exp(-\pi/4i\sigma_y^{(i)}\sigma_y^{(j)})\exp(i\theta\sigma_x^{(j)}\sigma_x^{(k)})\exp(i\pi/4\sigma_y^{(i)}\sigma_y^{(j)}) = \exp(-i\theta\sigma_y^{(i)}\sigma_z^{(j)}\sigma_x^{(k)}) \quad (7)$$

This decomposition enables efficient construction of the following cascade of zzz-interactions:

$$\exp(-i\theta_1\sigma_y^{(a)}\sigma_z^{(b)}\sigma_x^{(1)})\exp(-i\theta_2\sigma_y^{(a)}\sigma_z^{(b)}\sigma_x^{(2)})\dots\exp(-i\theta_n\sigma_y^{(a)}\sigma_z^{(b)}\sigma_x^{(n)}) = \quad (8)$$

$$\exp(-\pi/4i\sigma_y^{(a)}\sigma_y^{(b)})\exp(i\theta_1\sigma_x^{(b)}\sigma_x^{(1)})\exp(i\theta_2\sigma_x^{(b)}\sigma_x^{(2)})\dots\exp(i\theta_n\sigma_x^{(b)}\sigma_x^{(n)})\exp(i\pi/4\sigma_y^{(a)}\sigma_y^{(b)}) \quad (9)$$

, which are efficiently implemented from the native xx-interaction and single-qubit rotations.

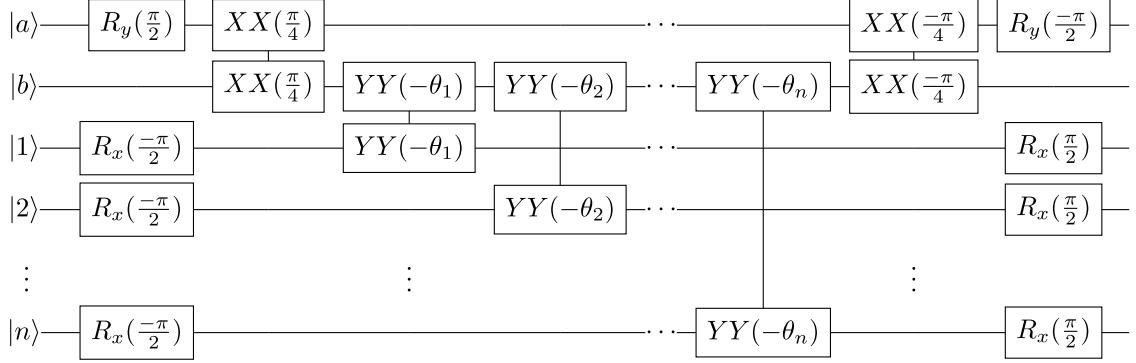


FIG. 4. Circuit implementing the first term in Eq. 6.

Using the above decomposition, we can implement the first term in Eq. 6 using the circuit shown Fig. 4.

With this term implemented, we complete the construction of the full unitary $U(x,y)$. After applying the unitary, we obtain the state

$$\alpha \sum_{0 \leq x \leq N/2} |x\rangle \sum_{0 \leq y \leq N} e^{2\pi i \frac{x^2 y \bmod N}{N}} |y\rangle. \quad (10)$$

We then apply the inverse Quantum Fourier Transforma QFT^\dagger to the y -register, which gives us:

$$\alpha \sum_{0 \leq x \leq N/2} |x\rangle |y = x^2 \bmod N\rangle \quad (11)$$

Next we measure the y -register to find an output w , and the x -register contains the superposition of a colliding input pair.

The number of qubits used to represent y in experiments are 3,4,4 and 5 for $N=8$, $N=15$, $N=16$ and $N=21$, respectively. The number of qubits used to represent x in experiments equals the length of the r string in table IV.

G. Circuit construction of the LWE based protocol

In this section we describe the procedure for implementing the circuit $U(A, b, x, y)$, displayed in Fig. 3(d). First, let us comment on the parameters on which this unitary depends. The matrix $A \in \mathbb{Z}_q^{m \times n}$ and vector $s \in \{0, 1\}^n$ are sampled uniformly at random by the verifier⁶⁰. The vector $e \in \mathbb{Z}_q^m$ is sampled from a discrete Gaussian distribution (see Braker-

ski et al.⁵ for more details on the parameter choices). The verifier constructs $y = As + e \in \mathbb{Z}_q^m$ and sends A and y to the prover.

Upon receiving A and y , the prover must evaluate the function $f(b, x) = \lfloor Ax + b \cdot y \rfloor$ in superposition, where $\lfloor \cdot \rfloor$ denotes a rounding operation corresponding to taking the most significant bit of each component in the vector $Ax + b \cdot y$. It should be noted that this specific function, which uses rounding, differs from the TCF used by Brakerski et al.⁵, but is nevertheless still a TCF⁴⁶.

To perform the coherent evaluation, the prover will use three registers (for the b and x inputs, as well as for the output of the TCF) to create the superposition state as well as a fourth ancilla register, which will be used to perform the unitary $U(A, b, x, y)$. The prover starts by applying a layer of Hadamard gates to all input qubits and the ancilla register (that were initialized as $|0\rangle$). The resulting state will be

$$\sum_{b \in \{0,1\}} \sum_{x \in \mathbb{Z}_q^n} \sum_{a \in \mathbb{Z}_q} \alpha |b\rangle |x\rangle |a\rangle |0\rangle \quad (12)$$

for some normalization constant α and where the third register is the ancilla register and the last register is the output register. In this output register, the prover must coherently add $\lfloor Ax + b \cdot y \rfloor$. As $Ax + b \cdot y$ is an m -component vector, we will explain the prover's operations, at a high level, for each component of the vector. For the i 'th component of this vector, the prover first computes the modulo q inner product between the i 'th row of A and x and places the result in the ancilla register. Since the prover has a classical description of A , this will involve a series of controlled operations between the x register and the ancilla register. Similar to the factoring case, this arithmetic operation is easiest to perform in the Fourier basis, hence why the ancilla register was Hadamarded. Once the inner product has been computed, the prover will perform a controlled operation between the b qubit and the ancilla register in order to add the i 'th component of y . Finally, the prover will “copy” the most significant bit of the result into the output register. This is done via another controlled operation. The prover then uncomputes the result in the ancilla, clearing that register. In this way, the i 'th component of $\lfloor Ax + b \cdot y \rfloor$ has been added into the output register. Repeating this procedure for all components will yield the desired state

$$\sum_{b \in \{0,1\}} \sum_{x \in \mathbb{Z}_q^n} \alpha' |b\rangle |x\rangle |0\rangle | \lfloor Ax + b \cdot y \rfloor \rangle \quad (13)$$

with normalization constant α' .

Having given the high level description, let us now discuss in more detail the specific circuits of the current implementation. From the above analysis, we can see that the total number of qubits is $N = 1 + n \log_2(q) + \log_2(q) + m$. In the instance for this experiment, we chose $m = 4, n = 2, q = 4$, resulting in $N = 11$ qubits. The first register contains $|b\rangle$ which requires only one qubit. In the second register, the vector $x = (x_0, x_1)$ consists of two components modulo 4, which is encoded in binary with four qubits as $|x\rangle = |x_{11}, x_{12}, x_{21}, x_{22}\rangle$. The third register, the ancilla, is one modulo 4 component and will thus consist of two qubits. Lastly, in the fourth register, we store the result of evaluating the function, which requires another four qubits. As mentioned, the matrix A and the vector y are specified classically. In the experiment, we considered four different input configurations, corresponding to four different choices for A, s and e . These choices are explicitly described later in the appendix.

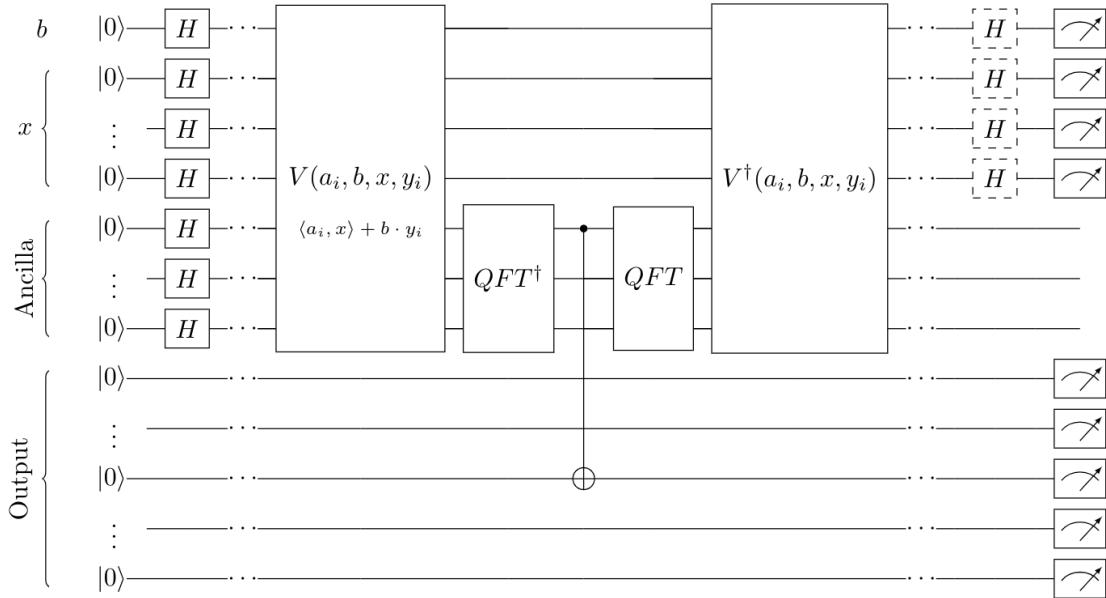


FIG. 5. Circuit used to “copy” the most significant bit of the result from the ancilla into the output register, adding the i ’th component of $[Ax + b \cdot y]$. Here, V represents the unitary used to compute $\langle a_i, x \rangle + b \cdot y_i$, in modular arithmetic, for the i ’th row a_i of the matrix A . Additionally, y_i denotes the i ’th entry of the vector $y = As + e$. Also, note that the target qubit of the CNOT in the diagram is the i ’th qubit. The step shown is repeated for each row of A , indexed by i .

To detail the operations implemented, as discussed previously, the prover first puts the ancilla register into the Fourier basis using the Quantum Fourier Transform (QFT). This allows them to more easily compute $\langle a_i, x \rangle + b \cdot y_i$ in the ancilla register, where a_i is the i ’th

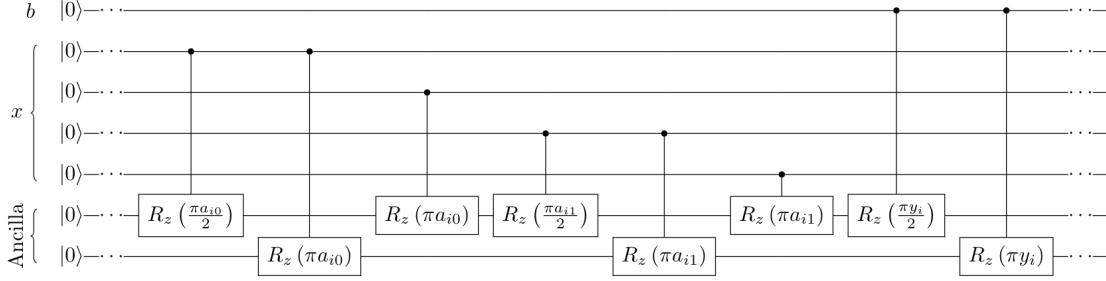


FIG. 6. The explicit rotation gates used to implement the unitary V from fig.5 for the case of $q = 4$. Here, a_{ij} denotes the entry in the i 'th row and j 'th column of the matrix A and y_i denotes the i 'th entry of the vector $y = As + e$. The output register is omitted as there are no operations performed on it in this section of the circuit. The step shown is repeated for each row of A , indexed by i .

row of the matrix A and $\langle \cdot, \cdot \rangle$ denotes the inner product modulo q . The explicit rotation gates to compute this in the Fourier basis are given in Fig. 6. After computing this for one row a_i , the prover converts the ancilla back into the computational basis and “copies” the most significant bit stored in the ancilla register into the output register, using a CNOT gate, to compute the rounding function. This completes the evaluation of the function for one bit. In order to reuse the qubits in the ancilla register, the prover then reverses this computation and repeats for each row of the matrix A . This process of evaluating the function and reversing that computation is depicted in Fig. 5.

Finally, after completing the evaluation of the TCF, the prover measures the output register to recover the rounded result of $\lfloor Ax + b \cdot y \rfloor$ for a certain value of x . The prover will then measure the b and x registers in either the Z basis or X basis, according to the challenge issued by the verifier. Should the verifier choose to measure in X basis, the prover applies Hadamard gates on all qubits in the b and x registers before measuring in the computational basis.

H. Instances of LWE Implemented

Here, we explicitly detail the LWE instances that were used in the experiment. Recall that such an instance is defined by A , s , and e , where $A \in \mathbb{Z}_q^{m \times n}$, $s \in \{0, 1\}^n$, and $e \in \mathbb{Z}_q^m$ for integers $m, n, q \in \mathbb{Z}$. In this experiment, we used $m = 4, n = 2, q = 4$ for all of the instances.

Instance	A^\top	e^\top	$(As + e)^\top$
0	$\begin{pmatrix} 0 & 2 & 0 & 1 \\ 2 & 0 & 1 & 2 \end{pmatrix}$	$\begin{pmatrix} 0 & 1 & 0 & 0 \end{pmatrix}$	$\begin{pmatrix} 0 & 3 & 0 & 1 \end{pmatrix}$
1	$\begin{pmatrix} 0 & 2 & 3 & 2 \\ 2 & 3 & 0 & 0 \end{pmatrix}$	$\begin{pmatrix} 0 & 0 & 0 & 1 \end{pmatrix}$	$\begin{pmatrix} 0 & 2 & 3 & 3 \end{pmatrix}$
2	$\begin{pmatrix} 2 & 0 & 0 & 1 \\ 0 & 3 & 2 & 1 \end{pmatrix}$	$\begin{pmatrix} 1 & 0 & 1 & 0 \end{pmatrix}$	$\begin{pmatrix} 3 & 0 & 1 & 1 \end{pmatrix}$
3	$\begin{pmatrix} 0 & 1 & 3 & 0 \\ 3 & 0 & 0 & 2 \end{pmatrix}$	$\begin{pmatrix} 1 & 0 & 1 & 0 \end{pmatrix}$	$\begin{pmatrix} 0 & 1 & 3 & 1 \end{pmatrix}$

TABLE V. Details of the LWE instances. Note that the entries are transposed and for all instances we use $s^\top = \begin{pmatrix} 0 & 1 \end{pmatrix}$.

Table V displays the explicit matrices and vectors used.

I. Derivation of Quantum-Classical Threshold for LWE based Protocol

Proposition 1. *For any classical prover, the probabilities that they pass branches A and B, p_A and p_B , must obey the relation*

$$p_A + 2p_B - 2 < \epsilon(\lambda) \quad (14)$$

where ϵ is a negligible function of the security parameter λ .

Proof. We first want to find the probability that the classical prover both responds correctly for Branch A and, for the same image w that they committed to the verifier, Branch B is also correct with probability greater than $1/2 + \mu(\lambda)$, where μ is a non-negligible function of the security parameter λ . Let this second probability be denoted as

$$p_{\text{good}} \equiv \Pr_w[p_{B,w} > 1/2 + \mu(\lambda)] \quad (15)$$

By a union bound, we arrive at a bound on the desired probability

$$\Pr[A \text{ correct and } p_{B,w} > 1/2 + \mu(\lambda)] > p_A + p_{\text{good}} - 1 \quad (16)$$

Now, we wish to write p_{good} in terms of p_B . Let S be the set of w values for which $p_{B,w} > 1/2 + \mu(\lambda)$. By definition, we know that with probability p_{good} , the prover samples a $w \in S$

so that they pass the verifier's Branch B test with probability at least $1/2 + \mu(\lambda)$ and at most 1. Similarly, we know that with probability $1 - p_{\text{good}}$, the prover samples a $w \notin S$ so that they pass the verifier's Branch B test with probability at most $1/2$. Hence, overall we see that the probability that the prover passes Branch B is at most the convex mixture of these two cases, i.e.

$$p_B < 1 \cdot p_{\text{good}} + 0.5 \cdot (1 - p_{\text{good}}) \quad (17)$$

Solving for p_{good} , we then obtain

$$p_{\text{good}} > 2p_B - 1 \quad (18)$$

Substituting this into Equation 16, we have

$$\Pr[A \text{ correct and } p_{B,w} > 1/2 + \mu(\lambda)] > p_A + 2p_B - 2 \quad (19)$$

However, notice that this probability on the left hand side is the probability of breaking the adaptive hardcore bit property, which we know⁵ must have

$$\Pr[A \text{ correct and } p_{B,w} > 1/2 + \mu(\lambda)] < \epsilon(\lambda) \quad (20)$$

where ϵ is a negligible function. Thus, combining this with Equation 19, we obtain the desired inequality

$$p_A + 2p_B - 2 < \epsilon(\lambda) \quad (21)$$

□

J. Computation of statistical significance contours

Here we describe the computation of the contour lines denoting various levels of statistical significance in Figure 3(b,e) of the main text. Recall the probabilities p_A and p_B introduced in Section III, which denote a prover's probability of passing the standard basis and interference test, respectively. Assuming the cryptographic soundness of the claw-free property of the TCF, and in the limit of large problem size, any classical cheating strategy must have true values of p_A^c and p_B^c that obey the bound $p_A^c + 2p_B^c - 2 < 0$ for the LWE protocol and $p_A^c + 4p_B^c - 4 < 0$ for the factoring-based protocol. To find the statistical significance of a pair of values p_A and p_B measured from an (ostensibly) quantum prover, we consider the null hypothesis that the data was generated by a classical cheater (which obeys the bounds

above), and compute the probability that the given data could be generated by that null hypothesis. In particular, since the bounds above exclude a region of a two-dimensional space, we consider an infinite “family” of null hypotheses which lie along the boundary, and define the overall statistical significance of measuring p_A and p_B to be the *minimum* of the statistical significances across the entire family of null hypotheses—that is, we define it as the significance with respect to the *least rejected* null hypothesis.

To compute the statistical significance of a result (p_A, p_B) with respect to a particular null hypothesis (p_A^c, p_B^c) , we define the “quantumness” q of an experiment as $q(p_A, p_B) = p_A + 4p_B - 4$ for the factoring-based protocol and $q(p_A, p_B) = p_A + 2p_B - 2$ for the LWE protocol. Letting N_A and N_B be the number of experimental runs performed for each branch respectively, we define the joint probability mass function (PMF) as the product of the PMFs of two binomial distributions $B(N_A, p_A^c)$ and $B(N_B, p_B^c)$. Mathematically the joint PMF is

$$f(k_A, k_B; p_A^c, p_B^c, N_A, N_B) = \binom{N_A}{k_A} \binom{N_B}{k_B} (p_A^c)^{k_A} (p_B^c)^{k_B} (1 - p_A^c)^{N_A - k_A} (1 - p_B^c)^{N_B - k_B} \quad (22)$$

where $k_A = p_A N_A$ and $k_B = p_B N_B$ are the “count” of passing runs for each branch respectively. Finally, we compute the statistical significance of a result (p_A, p_B) as the probability of achieving quantumness measure of at least $q' = q(p_A, p_B)$. Under a null hypothesis (p_A^c, p_B^c) , this is the sum of the PMF over all k_A, k_B for which $q(k_A/N_A, k_B/N_B) > q'$.

In practice, for the contour lines of Figure 3(b,e), we begin with a desired level of statistical significance (say, 5σ), and given the sample sizes N_A and N_B we compute the value of q' that would achieve at least that significance over all null hypotheses inside the classical bound.

K. Relative performance of additional instances of factoring-based protocols implemented with delayed-measurement.

In Fig. 3(f), we show the relative performance of the factoring-based protocol for $N = 8$, performed both interactively and with delayed measurement. In Fig. 7 we display the relative performance for $N \in \{15, 16, 21\}$ (for which experiments were run with delayed measurement only).

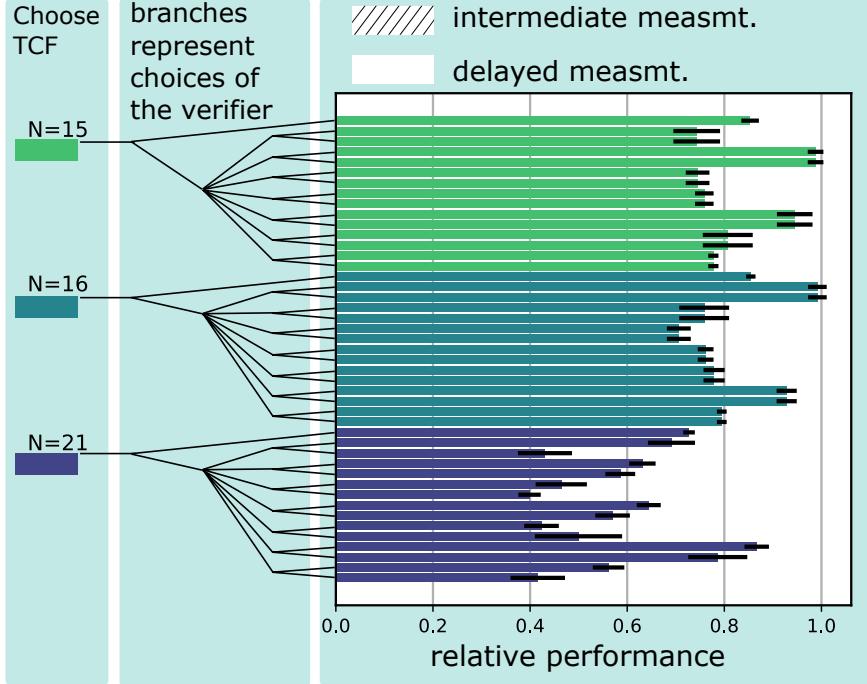


FIG. 7. Extra instances of factoring-based protocols implemented with delayed measurement.

L. Estimate of resources required to achieve quantum advantage

For a conclusive demonstration of quantum advantage, we desire the quantum machine to perform the protocol significantly faster than the amount of time a classical supercomputer would require to break the trapdoor claw-free function—ideally, orders of magnitude faster. To achieve this, we must set the parameters of the cryptographic problem to sufficiently large values. A major benefit of using protocols based off of established cryptographic assumptions (like factoring and LWE) is that the classical hardness of breaking these assumptions is extremely well studied, due to the implications for security. Thus the most straightforward way to choose parameters for our tests is to rely on publicly-available recommendations for *cryptographically secure* key sizes, which are used in practice. These parameter settings are designed to be not just slow for classical machines, but infeasible even for classical machines years from now—and thus certainly would constitute a definitive demonstration of quantum advantage. However, setting the parameters to these values may be considered overkill for our purposes, especially since we'd like the problem size to be as small as possible in order to make the protocols maximally feasible on near term quantum devices. With these considerations, in this section we provide two estimates for each protocol: we begin by

providing estimates for smaller problem sizes that still would demonstrate some level of quantum advantage, and then give estimates based on cryptographic parameters.

We conservatively estimate that a future quantum device running the protocols investigated in this work at scale would complete the protocols on a time scale of at most hours. Thus, to demonstrate quantum advantage by several orders of magnitude, we desire to set the parameters such that a classical supercomputer would require time on the order of thousands of hours to break the TCF. In 2020, Boudot et al. reported the record-breaking factorization of a 795-bit semiprime⁶¹. The cost of the computation was about 1000 core-years, meaning that a 1000-core cluster would complete it in a year. We consider this sufficient cost to demonstrate quantum advantage. We emphasize also that factoring is one of the most well-studied hard computational problems; the record of Boudot et al. is the product of decades of algorithm development and optimization and thus it is unlikely that any innovations will drastically affect the classical hardness of factoring in the near term. The computational Bell test protocol using a 795-bit prime could be performed using only about 800 qubits by computing and measuring the bits of the output value w one-by-one; however the gate count and circuit depth can be dramatically reduced by explicitly storing the full output value w , requiring roughly 1600 qubits total²⁵. Because it is so much more efficient in gate count, we use the 1600 qubit estimate as the space requirement to demonstrate quantum advantage with the computational Bell test protocol.

For LWE, estimating parameters for the same level of hardness (1000 core-years) is difficult to do exactly, because to our knowledge that amount of computational resources has never been applied to breaking an LWE instance. However, we may make a rough estimate. There is an online challenge (https://www.latticechallenge.org/lwe_challenge/challenge.php) intended to explore the practical classical hardness of LWE, in which users compete for who can break the largest possible instance. As of this writing, the largest instances which have been solved use LWE vectors of about 500-1000 bits (depending on the noise level of the error vector), but the computational cost of these calculations was only of order 0.5 core-years. To require 1000 core-years of computation time, we estimate that the LWE vectors would need to be perhaps 1000-2000 bits in length; by not explicitly storing the output vector w but computing it element-by-element (similar in principle to the scheme for evaluating $x^2 \bmod N$ using only $\log(N) + 1$ qubits²⁵) it may be possible to perform the LWE protocol using a comparable number of qubits to the bit length of one

LWE vector.

We now provide estimates for cryptographic parameters; that is, parameters for which it is expected to be completely infeasible for a classical machine to break the trapdoor claw-free function. For the factoring-based protocol, we may apply NIST’s recommended key sizes for the RSA cryptosystem, whose security relies on integer factorization. NIST recommends choosing a modulus N with length 2048 bits. By using circuits optimized to conserve qubits, it is possible to evaluate the function $x^2 \bmod N$ using only $\log(N) + 1$ qubits, yielding a total qubit requirement of 2049 qubits²⁵. However, the circuit depth can be improved significantly by including more qubits; a more efficient circuit can be achieved with roughly $2 \log(N) \sim 4100$ qubits. Because LWE is not yet broadly used in practice like RSA is, NIST does not provide recommendations for key sizes in its documentation. However, we can use the estimates of Lindner and Peikert⁶² to find parameters which are expected to be infeasible classically. In Fig. 3 of that work, the authors suggest using LWE vectors in \mathbb{Z}_q^n with $n = 256$ and $q = 4093$ for a “medium” level of security. Vectors with these parameters are $n \log(q) \sim 3072$ bits long. To store both an input and output vector would thus require roughly ~ 6200 qubits. By repeatedly reusing a set of qubits to compute the output vector element-by-element the computation could be performed using roughly 3100 qubits.

REFERENCES

¹M. A. Nielsen and I. L. Chuang, *Quantum computation and quantum information*, 10th ed. (Cambridge University Press, Cambridge ; New York, 2010).

²C. Ryan-Anderson, J. G. Bohnet, K. Lee, D. Gresh, A. Hankin, J. P. Gaebler, D. Francois, A. Chernoguzov, D. Lucchetti, N. C. Brown, T. M. Gatterman, S. K. Halit, K. Gilmore, J. Gerber, B. Neyenhuis, D. Hayes, and R. P. Stutz, “Realization of real-time fault-tolerant quantum error correction,” (2021), arXiv:2107.07505 [quant-ph].

³I. Cong, S. Choi, and M. D. Lukin, “Quantum convolutional neural networks,” *Nature Physics* **15**, 1273–1278 (2019), number: 12 Publisher: Nature Publishing Group.

⁴U. Mahadev, “Classical verification of quantum computations,” in *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)* (IEEE, 2018) pp. 259–267.

⁵Z. Brakerski, P. Christiano, U. Mahadev, U. Vazirani, and T. Vidick, “A cryptographic test of quantumness and certifiable randomness from a single quantum device,” in *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)* (IEEE, 2018) pp. 320–331.

⁶A. Gheorghiu and T. Vidick, “Computationally-secure and composable remote state preparation,” in *2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS)* (IEEE, 2019) pp. 1024–1033.

⁷D. Aharonov, M. Ben-Or, and E. Eban, “Interactive Proofs For Quantum Computations,” (2010) pp. 453–469.

⁸S. Goldwasser, S. Micali, and C. Rackoff, “The Knowledge Complexity of Interactive Proof Systems,” *SIAM Journal on Computing* **18**, 186–208 (1989), publisher: Society for Industrial and Applied Mathematics.

⁹C. Lund, L. Fortnow, H. Karloff, and N. Nisan, “Algebraic methods for interactive proof systems,” in *Proceedings [1990] 31st Annual Symposium on Foundations of Computer Science* (1990) pp. 2–10 vol.1.

¹⁰A. Shamir, “IP = PSPACE,” *Journal of the ACM* **39**, 869–877 (1992).

¹¹W. K. Hensinger, “Quantum computer based on shuttling trapped ions,” (2021).

¹²D. Bluvstein, H. Levine, G. Semeghini, T. T. Wang, S. Ebadi, M. Kalinowski, A. Keesling, N. Maskara, H. Pichler, M. Greiner, *et al.*, “A quantum processor based on coherent transport of entangled atom arrays,” *Nature* **604**, 451–456 (2022).

¹³J. M. Pino, J. M. Dreiling, C. Figgatt, J. P. Gaebler, S. A. Moses, M. Allman, C. Baldwin, M. Foss-Feig, D. Hayes, K. Mayer, *et al.*, “Demonstration of the trapped-ion quantum ccd computer architecture,” *Nature* **592**, 209–213 (2021).

¹⁴D. Kielpinski, C. Monroe, and D. J. Wineland, “Architecture for a large-scale ion-trap quantum computer,” *Nature* **417**, 709–711 (2002).

¹⁵Y. Wan, D. Kienzler, S. D. Erickson, K. H. Mayer, T. R. Tan, J. J. Wu, H. M. Vasconcelos, S. Glancy, E. Knill, D. J. Wineland, *et al.*, “Quantum gate teleportation between separated qubits in a trapped-ion processor,” *Science* **364**, 875–878 (2019).

¹⁶O. Regev, “On lattices, learning with errors, random linear codes, and cryptography,” *Journal of the ACM (JACM)* **56**, 1–40 (2009).

¹⁷O. Regev, “The learning with errors problem,” *Invited survey in CCC* **7**, 30 (2010).

¹⁸M. O. Rabin, “Digitalized signatures and public-key functions as intractable as factorization,” *Technical Report* (Massachusetts Institute of Technology, USA, 1979).

¹⁹S. Goldwasser, S. Micali, and R. L. Rivest, “A digital signature scheme secure against adaptive chosen-message attacks,” *SIAM Journal on computing* **17**, 281–308 (1988).

²⁰S. Goldwasser, S. Micali, and R. L. Rivest, “A “paradoxical” solution to the signature problem,” in *Advances in Cryptology, Proceedings of CRYPTO ’84, Santa Barbara, California, USA, August 19-22, 1984, Proceedings* (1984) p. 467.

²¹A. Banerjee, C. Peikert, and A. Rosen, “Pseudorandom functions and lattices,” in *Annual International Conference on the Theory and Applications of Cryptographic Techniques* (Springer, 2012) pp. 719–737.

²²J. Alwen, S. Krenn, K. Pietrzak, and D. Wichs, “Learning with rounding, revisited,” in *Annual Cryptology Conference* (Springer, 2013) pp. 57–74.

²³J. S. Bell, “On the Einstein Podolsky Rosen paradox,” *Physics* **1**, 195–200 (1964).

²⁴J. F. Clauser, M. A. Horne, A. Shimony, and R. A. Holt, “Proposed experiment to test local hidden-variable theories,” *Physical review letters* **23**, 880 (1969).

²⁵G. D. Kahanamoku-Meyer, S. Choi, U. V. Vazirani, and N. Y. Yao, “Classically-verifiable quantum advantage from a computational bell test,” *arXiv preprint arXiv:2104.00687* (2021).

²⁶C. Monroe, W. C. Campbell, L.-M. Duan, Z.-X. Gong, A. V. Gorshkov, P. Hess, R. Islam, K. Kim, N. M. Linke, G. Pagano, *et al.*, “Programmable quantum simulations of spin systems with trapped ions,” *Reviews of Modern Physics* **93**, 025001 (2021).

²⁷L. Egan, D. M. Debroy, C. Noel, A. Risinger, D. Zhu, D. Biswas, M. Newman, M. Li, K. R. Brown, M. Cetina, *et al.*, “Fault-tolerant operation of a quantum error-correction code,” *Nature* **598**, 281–286 (2021).

²⁸F. Arute *et al.*, “Quantum supremacy using a programmable superconducting processor,” *Nature* **574**, 505–510 (2019).

²⁹H.-S. Zhong *et al.*, “Quantum computational advantage using photons,” *Science* **370**, 1460–1463 (2020).

³⁰Y. Wu, W.-S. Bao, S. Cao, F. Chen, *et al.*, “Strong quantum computational advantage using a superconducting quantum processor,” arXiv preprint arXiv:2106.14734 (2021).

³¹Q. Zhu, S. Cao, F. Chen, M.-C. Chen, X. Chen, T.-H. Chung, H. Deng, Y. Du, D. Fan, M. Gong, *et al.*, “Quantum computational advantage via 60-qubit 24-cycle random circuit sampling,” *Science Bulletin* , 240–245 (2021).

³²S. Aaronson and A. Arkhipov, “The computational complexity of linear optics,” 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.

³³A. P. Lund, M. J. Bremner, and T. C. Ralph, “Quantum sampling problems, bosonsampling and quantum supremacy,” *npj Quantum Information* **3**, 1–8 (2017).

³⁴A. W. Harrow and A. Montanaro, “Quantum computational supremacy,” *Nature* **549**, 203–209 (2017).

³⁵S. Boixo, S. V. Isakov, V. N. Smelyanskiy, R. Babbush, N. Ding, Z. Jiang, M. J. Bremner, J. M. Martinis, and H. Neven, “Characterizing quantum supremacy in near-term devices,” *Nature Physics* **14**, 595–600 (2018).

³⁶A. Bouland, B. Fefferman, C. Nirkhe, and U. Vazirani, “On the complexity and verification of quantum random circuit sampling,” *Nature Physics* **15**, 159–163 (2019).

³⁷S. Aaronson and S. Gunn, “On the classical hardness of spoofing linear cross-entropy benchmarking,” arXiv preprint arXiv:1910.12085 (2019).

³⁸C. Huang, F. Zhang, M. Newman, J. Cai, X. Gao, Z. Tian, J. Wu, H. Xu, H. Yu, B. Yuan, M. Szegedy, Y. Shi, and J. Chen, “Classical Simulation of Quantum Supremacy Circuits,” arXiv:2005.06787 [quant-ph] (2020), arXiv: 2005.06787.

³⁹F. Pan and P. Zhang, “Simulating the Sycamore quantum supremacy circuits,” arXiv:2103.03074 [physics, physics:quant-ph] (2021), arXiv: 2103.03074.

⁴⁰J. Gray and S. Kourtis, “Hyper-optimized tensor network contraction,” *Quantum* **5**, 410

(2021), arXiv: 2002.01935.

⁴¹F. Pan, K. Chen, and P. Zhang, “Solving the sampling problem of the Sycamore quantum supremacy circuits,” arXiv:2111.03011 [physics, physics:quant-ph] (2021), arXiv: 2111.03011.

⁴²Yong, Liu, Xin, Liu, Fang, Li, H. Fu, Y. Yang, J. Song, P. Zhao, Z. Wang, D. Peng, H. Chen, C. Guo, H. Huang, W. Wu, and D. Chen, “Closing the ”Quantum Supremacy” Gap: Achieving Real-Time Simulation of a Random Quantum Circuit Using a New Sunway Supercomputer,” Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis , 1–12 (2021), arXiv: 2110.14502.

⁴³X. Liu, C. Guo, Y. Liu, Y. Yang, J. Song, J. Gao, Z. Wang, W. Wu, D. Peng, P. Zhao, F. Li, H.-L. Huang, H. Fu, and D. Chen, “Redefining the Quantum Supremacy Baseline With a New Generation Sunway Supercomputer,” arXiv:2111.01066 [quant-ph] (2021), arXiv: 2111.01066.

⁴⁴X. Gao, M. Kalinowski, C.-N. Chou, M. D. Lukin, B. Barak, and S. Choi, “Limitations of Linear Cross-Entropy as a Measure for Quantum Advantage,” arXiv:2112.01657 [cond-mat, physics:quant-ph] (2021), arXiv: 2112.01657.

⁴⁵S. Hirahara and F. L. Gall, “Test of quantumness with small-depth quantum circuits,” arXiv preprint arXiv:2105.05500 (2021).

⁴⁶Z. Liu and A. Gheorghiu, “Depth-efficient proofs of quantumness,” arXiv preprint arXiv:2107.02163 (2021).

⁴⁷A. D. Corcoles, M. Takita, K. Inoue, S. Lekuch, Z. K. Minev, J. M. Chow, and J. M. Gambetta, “Exploiting dynamic quantum circuits in a quantum algorithm with superconducting qubits,” arXiv preprint arXiv:2102.01682 (2021).

⁴⁸K. Rudinger, G. J. Ribeill, L. C. Goria, M. Ware, E. Nielsen, K. Young, T. A. Ohki, R. Blume-Kohout, and T. Proctor, “Characterizing mid-circuit measurements on a superconducting qubit using gate set tomography,” arXiv preprint arXiv:2103.03008 (2021).

⁴⁹B. Skinner, J. Ruhman, and A. Nahum, “Measurement-induced phase transitions in the dynamics of entanglement,” Physical Review X **9**, 031009 (2019).

⁵⁰Y. Li, X. Chen, and M. P. Fisher, “Quantum zeno effect and the many-body entanglement transition,” Physical Review B **98**, 205136 (2018).

⁵¹C. Noel, P. Niroula, A. Risinger, L. Egan, D. Biswas, M. Cetina, A. V. Gorshkov, M. Gulans, D. A. Huse, and C. Monroe, “Observation of measurement-induced quantum phases

in a trapped-ion quantum computer,” arXiv preprint arXiv:2106.05881 (2021).

⁵²M. Cetina, L. N. Egan, C. A. Noel, M. L. Goldman, A. R. Risinger, D. Zhu, D. Biswas, and C. Monroe, “Quantum gates on individually-addressed atomic qubits subject to noisy transverse motion,” arXiv preprint arXiv:2007.06768 (2021).

⁵³S. Olmschenk, K. C. Younge, D. L. Moehring, D. N. Matsukevich, P. Maunz, and C. Monroe, “Manipulation and detection of a trapped yb^+ hyperfine qubit,” Phys. Rev. A **76**, 052314 (2007).

⁵⁴S. Debnath, N. M. Linke, C. Figgatt, K. A. Landsman, K. Wright, and C. Monroe, “Demonstration of a small programmable quantum computer with atomic qubits,” Nature **536**, 63 (2016).

⁵⁵K. Mølmer and A. Sørensen, “Multiparticle entanglement of hot trapped ions,” Phys. Rev. Lett. **82**, 1835–1838 (1999).

⁵⁶E. Solano, R. L. de Matos Filho, and N. Zagury, “Deterministic bell states and measurement of the motional state of two trapped ions,” Phys. Rev. A **59**, R2539–R2543 (1999).

⁵⁷G. Milburn, S. Schneider, and D. James, “Ion trap quantum computing with warm ions,” Fortschritte der Physik **48**, 801–810 (2000).

⁵⁸T. Choi, S. Debnath, T. A. Manning, C. Figgatt, Z.-X. Gong, L.-M. Duan, and C. Monroe, “Optimal quantum control of multimode couplings between trapped ion qubits for scalable entanglement,” Phys. Rev. Lett. **112**, 190502 (2014).

⁵⁹P. L. W. Maunz, “High optical access trap 2.0.” (2016), 10.2172/1237003.

⁶⁰Technically, the matrix A is sampled together with the TCF trapdoor. However, as explained in⁵, the distribution from which the matrix is sampled is statistically close to a uniform distribution over $\mathbb{Z}_q^{m \times n}$.

⁶¹F. Boudot, P. Gaudry, A. Guillevic, N. Henger, E. Thomé, and P. Zimmermann, “Comparing the difficulty of factorization and discrete logarithm: a 240-digit experiment,” Tech. Rep. 697 (2020).

⁶²R. Lindner and C. Peikert, “Better Key Sizes (and Attacks) for LWE-Based Encryption,” in *Topics in Cryptology – CT-RSA 2011*, Lecture Notes in Computer Science, edited by A. Kiayias (Springer, Berlin, Heidelberg, 2011) pp. 319–339.