nature physics

Article

https://doi.org/10.1038/s41567-023-02162-9

Interactive cryptographic proofs of quantumness using mid-circuit measurements

Received: 20 October 2022

Accepted: 3 July 2023

Published online: 03 August 2023

Check for updates

Daiwei Zhu[1,2,3,4,16] ✉, Gregory D. Kahanamoku-Meyer[5,6,16], Laura Lewis[7,8], Crystal Noel[1,9,10], Or Katz[9,10], Bahaa Harraz[1,9,10], Qingfeng Wang[1,2,11], Andrew Risinger[1,2,3], Lei Feng[1,2], Debopriyo Biswas[1,2], Laird Egan[1,2], Alexandru Gheorghiu[7,12], Yunseong Nam[13], Thomas Vidick[7,14], Umesh Vazirani[5,6], Norman Y. Yao[5,6,15], Marko Cetina[1,9,10] & Christopher Monroe[1,2,3,9,10]

The ability to perform measurements in the middle of a quantum circuit is a powerful resource. It underlies a wide range of applications, from remote state preparation to quantum error correction. Here we apply mid-circuit measurements for a particular task: demonstrating quantum computational advantage. The goal of such a demonstration is for a quantum device to perform a computational task that is infeasible for a classical device with comparable resources. In contrast to existing demonstrations, the distinguishing feature of our approach is that the classical verification process is efficient, both in asymptotic complexity and in practice. Furthermore, the classical hardness of performing the task is based upon well-established cryptographic assumptions. Protocols with these features are known as cryptographic proofs of quantumness. Using a trapped-ion quantum computer, we perform mid-circuit measurements by spatially isolating portions of the ion chain via shuttling. This enables us to implement two interactive cryptographic proofs of quantumness, which when suitably scaled to larger systems, promise the efficient verification of quantum computational advantage. Our methods can be applied to a range of interactive quantum protocols.

To date, experimental quantum computation has largely operated in a non-interactive paradigm in which classical data are extracted from the computation only at the very last step. Although 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, for quantum error correction, projective mid-circuit measurements are used to convert a continuum of possible errors into a specific discrete set of errors that can be corrected, as has been demonstrated in a recent experiment[1,2]. Certain quantum machine learning algorithms also leverage mid-circuit measurements to introduce essential nonlinearities[3]. Recent work has shown that interaction can do much more: it has emerged as an indispensable tool for verifying the behaviour of untrusted quantum devices[4–6] and even for testing the fundamentals of quantum mechanics itself[7].

Consider a classical computer sending commands to an untrusted quantum device that it cannot feasibly simulate. This could consist of a lab computer testing a new, large quantum device but also, perhaps, a user connecting to a quantum cloud computing service over the internet. At first sight, the inability of the classical machine to simulate the quantum one seems to pose a difficulty for certifying the output. This

A full list of affiliations appears at the end of the paper. ✉e-mail: daiwei@terpmail.umd.edu

challenge mirrors one explored in classical computer science, which asks whether a sceptical, 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, this idea began to be pursued 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. One of the greatest achievements 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 essential idea is that interaction can force the prover to commit to some piece of information early in the protocol, upon which the verifier follows up with queries that can be answered consistently only if the prover is being truthful. In exciting recent developments, success has been achieved by applying 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 delegating computation to an untrusted quantum server[4]. Perhaps the most direct application of an interactive protocol is for a 'cryptographic proof of quantumness', a protocol that allows a quantum device to convincingly demonstrate its non-classical behaviour to a polynomial-time classical verifier by performing a task that is assumed to be computationally hard for a classical machine yet is efficient to check[5,11,12].

The simplest proof of quantumness, in general, is a Bell test (which does not rely on a computational hardness assumption)[13]. It uses entanglement to generate correlations that would be impossible to reproduce classically without communication. While the Bell test's simplicity is attractive, avoiding the communication loophole requires the use of multiple quantum devices that are separated by a considerable distance[14–16]. To prove the quantumness of a single 'black-box' quantum device whose inner workings are hidden from the verifier, one can, instead, rely on differences in classical and quantum computational power, in other words, asking the device to demonstrate its quantum computational advantage. In contrast to recent sampling-based tests of quantum computational advantage[17–26], in a cryptographic proof of quantumness, the verification step must also be efficient. Although in principle any algorithm that exhibits a quantum speedup and has an efficiently verifiable output could be used for this purpose, most such experiments are infeasible today because the necessary circuits are far too large to run successfully on current quantum computers. Remarkably, it has been shown that interactive proofs provide a way to reduce the experimental cost (in qubits and gate depth) of this type of test, while maintaining efficient verification and classical hardness.

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 neighbouring qubits, degrading the quality of computations following the mid-circuit measurement. One solution, which has some commonality among atomic quantum computing platforms, is to spatially isolate target qubits via shuttling[27–29]. Although 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 architectures[30].

In this work, we implement two complementary interactive cryptographic proof of quantumness protocols, shown in the schematic of Fig. 1, on a trapped-ion quantum computer with up to 11 qubits using circuits with up to 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)[2,29,31]. The first protocol involves two rounds of interaction and is based upon the learning with errors (LWE) problem[32,33]. The LWE construction is unique because it has a property known as the 'adaptive hardcore bit'[5] (described in more detail in the next section), 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[34,35]. By using an additional interaction round, the cryptographic information is condensed onto the state of a single qubit. This makes it possible to implement a cryptographic proof of quantumness that is as hard to spoof classically as factoring but whose associated circuits can exhibit an asymptotic scaling much simpler than Shor's algorithm ($\mathcal{O}(n\log n)$ instead of $\mathcal{O}(n^2\log n)$, in terms of gate counts)[12].

Trapdoor claw-free functions

Both interactive protocols (Fig. 1) rely upon a cryptographic primitive called a trapdoor claw-free function (TCF)[36], which is a 2-to-1 function f for which it is cryptographically hard to find two inputs mapping to the same output. Such pairs of colliding inputs are called 'claws', and the term 'claw-free' refers to the hardness of finding them. The function also has a 'trapdoor', a secret key with which it is easy to compute the inputs x_0 and x_1 from any output $w = f(x_0) = f(x_1)$. The intuition behind the protocols is the following. Despite the claw-free property, a quantum computer can efficiently generate a superposition of two inputs 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 output 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. The prover now sends w to the verifier, who then uses 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 $|x_1\rangle$. (Note that the value of w, and by association x_0 and x_1, changes each time the protocol is executed, so it is not possible to find a collision (x_0, x_1) by simply repeating this process with a standard-basis measurement multiple times.) The verifier checks the measurement result on a per-shot basis. Crucially, consistently producing correct values for these measurements results is impossible for a classical prover (assuming they cannot find a claw of the TCF), so reliably returning correct results constitutes a proof of quantumness.

The LWE problem

It is believed to be classically intractable to recover an input vector from the result of certain noisy matrix-vector multiplications, which constitutes the LWE problem[32,33]. In particular, a secret vector, $\mathbf{s} \in \{0,1\}^n$, can be encoded into an output vector, $\mathbf{y} = A\mathbf{s} + \mathbf{e}$, where $A \in \mathbb{Z}_q^{m\times n}$ is a matrix and $\mathbf{e}$ is an error vector corresponding to the noise. Using the LWE problem, a TCF can be constructed as $f(b, x) = \lfloor Ax + b\cdot\mathbf{y}\rceil$, where b is a single bit that controls whether $\mathbf{y}$ gets added to Ax and $\lfloor\cdot\rceil$ denotes a rounding operation[37,38] (see Circuit construction of the LWE-based protocol in Supplementary Information for additional details). Here, $\mathbf{s}$ and $\mathbf{e}$ play the role of then trapdoor, and a claw corresponds to colliding inputs $\{(0, x_0), (1, x_1)\}$ with $f(0, x_0) = f(1, x_1)$ and $x_0 = x_1 + \mathbf{s}$. By implementing the protocol described above and illustrated in Fig. 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 property, which is a strengthening of the claw-free property[5]. Informally, it says that for any input x_0 (of the prover's choosing), it is cryptographically hard to determine even a single bit of information about x_1 (as opposed to the entire value, which is the guarantee of the claw-free assumption).

Rabin's function

The function $f(x) = x^2 \bmod N$, with N being the product of two primes p and q, was originally introduced in the context of digital signatures[34,35].

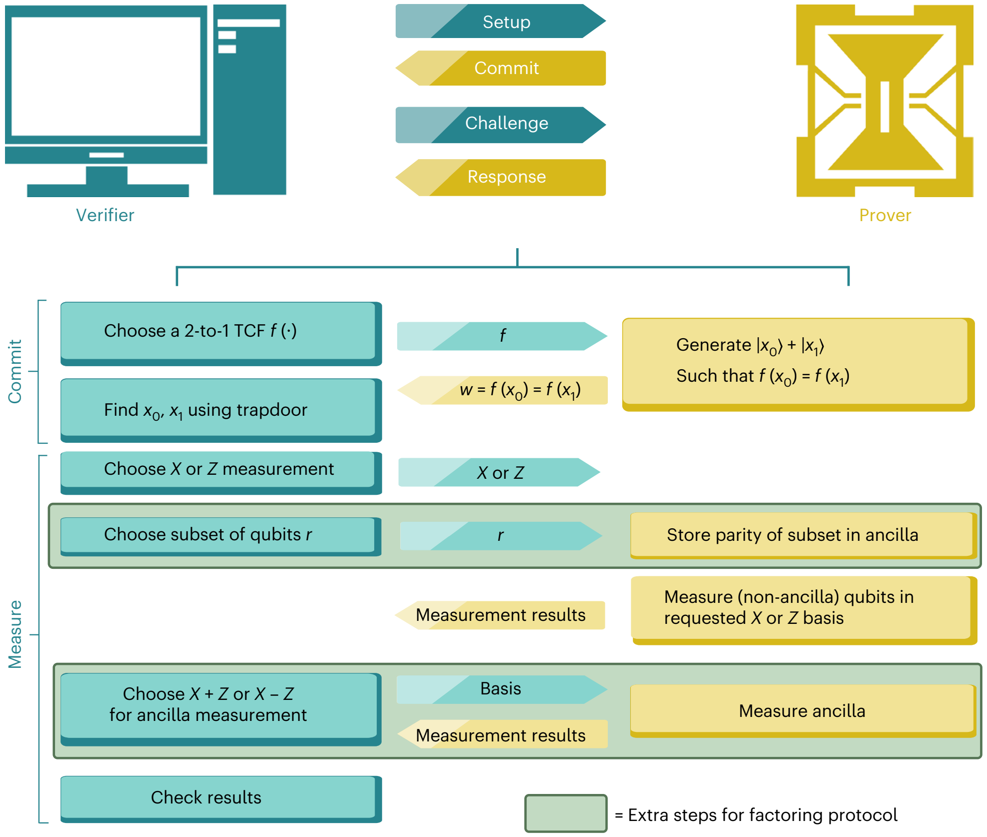
Setup
Commit
Challenge
Response
Verifier
Prover
Commit
Measure
Choose a 2-to-1 TCF $f(\cdot)$
f
Generate $|x_0\rangle + |x_1\rangle$
Such that $f(x_0) = f(x_1)$
Find x_0, x_1 using trapdoor
$w = f(x_0) = f(x_1)$
Choose X or Z measurement
X or Z
Choose subset of qubits r
r
Store parity of subset in ancilla
Measurement results
Measure (non-ancilla) qubits in requested X or Z basis
Choose $X+Z$ or $X-Z$ for ancilla measurement
Basis
Measurement results
Measure ancilla
Check results
🟩 = Extra steps for factoring protocol

Fig. 1 | Schematic diagram 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 TCF. 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 of the prover's commitment. The LWE protocol requires two rounds of interaction, and the factoring protocol requires an additional round (green box).

This function has the property that finding two colliding inputs (a claw) 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 TCF. 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. Given our cryptographic assumption that the prover cannot find a claw, the prover cannot guess the polarization of the remaining ancillary qubit. This is directly analogous to how, in Bell experiments, the assumption of no signalling faster than light implies that if Alice measures one half of an Bell pair, a space-like separated Bob who holds the other half is unable to immediately guess its polarization. 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[13,39]. The verifier accepts this if the prover returns the more likely measurement outcome. Crucially, the dependence of the measurement result on the claw renders it infeasible to guess classically[12].

Implementing an interactive cryptographic proof

To implement an interactive cryptographic proof of quantumness, we design quantum circuits for both the LWE and factoring protocols. The high-level circuit diagrams are shown in Fig. 3a,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. 3a,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 (that is standard versus 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 a trapped-ion 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 levels[40]. The quantum circuits are implemented via the consecutive application of native single- and two-qubit gates using individual optical addressing (Fig. 2a)[41]. To realize rapid, successive, two-qubit interactions, we position the ions in a single, closely spaced, linear chain (Fig. 2d).

This geometry makes it challenging to implement mid-circuit measurements, because 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 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 reshape 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 'Trapped-ion quantum computer' and 'Shuttling and mid-circuit measurements' sections of Methods for additional details).

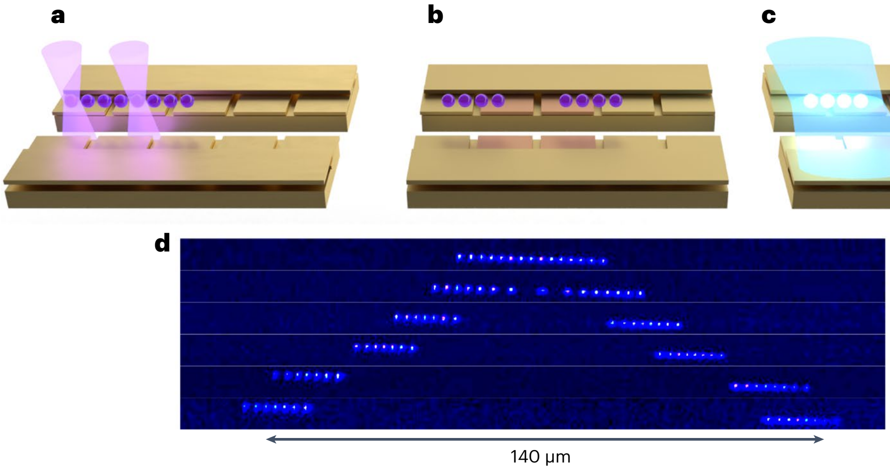

Fig. 2 | Mid-circuit measurements with shuttling. a–**c**, Schematic illustration of our mid-circuit measurement protocol. **a**, Initially, the ions are closely spaced in a one-dimensional chain above a surface trap. Coherent gates are implemented via a combination of individual-addressing beams (purple) and global beams (not shown). Both the coherent addressing beams and the detection optics are aligned to ions in the same section of the trap. **b**, By tuning the electrodes of the surface trap, we 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 images of an example shuttling protocol for a chain of 15 ions. Initially, the average spacing between ions is approximately 4 μm. At the end of the splitting procedure, the distance between the two segments is approximately 550 μm. The images show the splitting up to a distance of approximately 140 μm, at which point the two subchains reach the edge of the detection beam.

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 into the classical controller before execution of the quantum circuit.

Beating the classical threshold

As in a Bell test, even a classical prover can pass the verifier's challenges with finite probability. If the classical prover cannot find a claw in the TCF (which is assumed to be the case for a sufficiently large problem), this probability can be bounded by an asymptotic 'classical threshold', which a quantum prover must exceed to demonstrate advantage. (For a discussion of what it means for this threshold to be 'asymptotic' rather than absolute, see 'Discussion of the asymptotic classical threshold' section of Methods) 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 'Verifier's check' section of Methods for a definition of the verifier's checks). For the LWE protocol, the classical threshold is given by $p_A + 2p_B - 2 \le \epsilon$ (derivation in 'Quantum-classical threshold for the LWE protocol' section of Methods). For the factoring protocol, it is given by $p_A + 4p_B - 4 \le \epsilon$ (ref. 12). In both cases, ϵ is a function that goes to zero exponentially in the problem size. An intuition for the difference between the thresholds is that the factoring protocol requires an additional round of interaction during the interference test.

As depicted in Fig. 3b, we perform multiple instances of the LWE protocol for different matrices A and noise vectors $\mathbf{e}$ (see Instances of LWE implemented in Supplementary Information). For each of the verifier's possible choices, we repeat the experiment approximately 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 Result data in Supplementary Information) is shown in Fig. 3b. Figure 3e depicts the analogous results for the factoring 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, that is not in the range of f (see Post-selection in Supplementary Information). Effectively, this implements a post-selection that suppresses bit-flip errors[12].

To further analyse the performance of each branch of the interactive protocol, corresponding to the verifier's choices (Fig. 3c,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. This criterion is a way of isolating and evaluating the effect of noise on the success probabilities of each branch, as it removes effects such as if an error-free run if rejected by the verifier, which is inherent to the protocol. In particular, for a perfect (noise-free) quantum prover, $R = 1$ always. For a device so noisy that its outputs are uniformly random, $R = 0$. To probe the noise effects of the mid-circuit measurements, we implement two versions of the protocol, one interactive (the normal protocol) and the other with all measurements delayed until the end of the circuit. We compare the relative performance of the two cases. We emphasize that the delayed-measurement version is only a tool to probe our experimental system, and it may be vulnerable to classical spoofing even if it were run with a large problem where the other cryptographic assumptions hold, as the interaction due to the mid-circuit measurements is crucial.

For the LWE protocol, there are two rounds of interaction, corresponding to the two branches I and II in Fig. 3c. For the factoring protocol, there are three rounds of interaction (Fig. 3f). 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 in Fig. 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 protocol where the ancilla is polarized in the Z basis (see Circuit construction of the factoring-based protocol in Supplementary

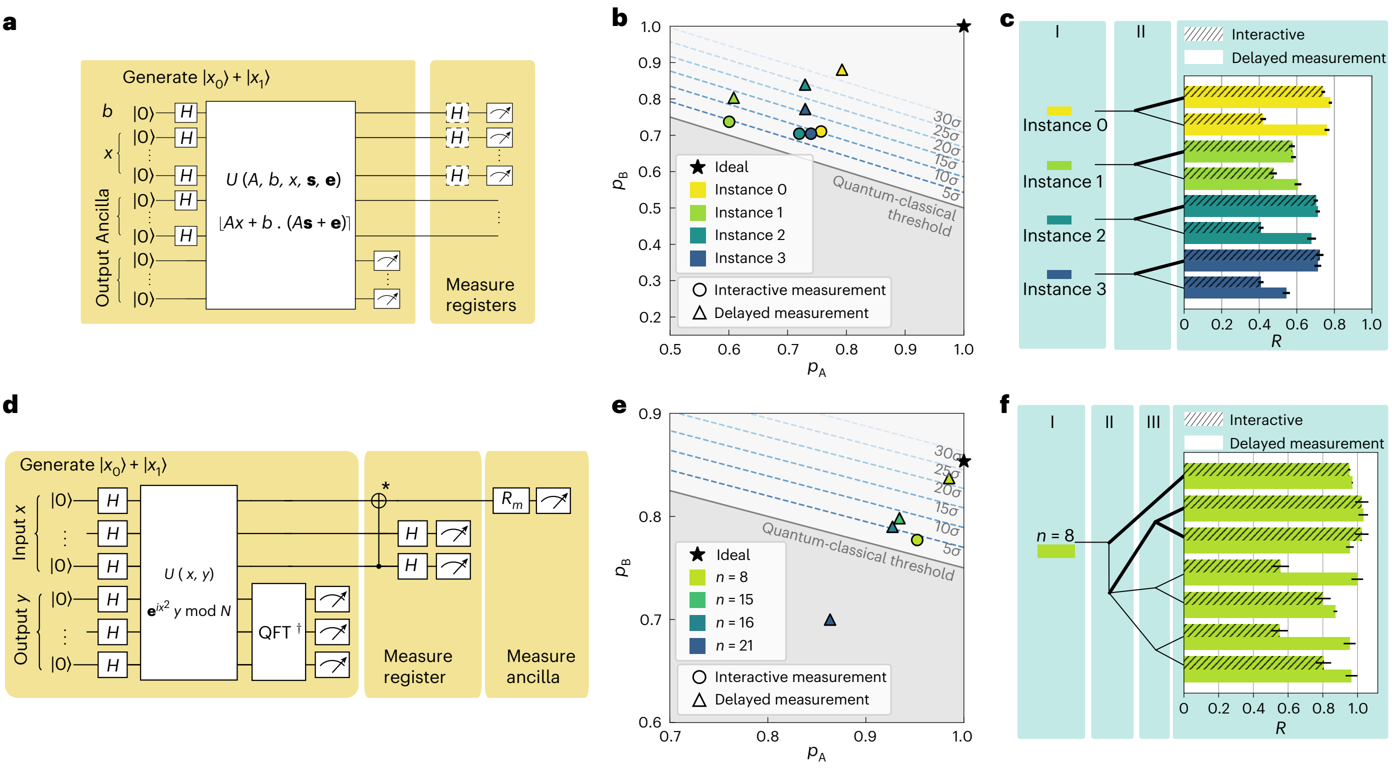

Fig. 3 | Circuit and experimental results. a, Circuit diagrams for the LWE protocol. **d**, Circuit diagrams for the factoring protocols. Details of the implementation of $U(A, b, x, y)$ and $U(x, y)$ are provided in Circuit construction of the LWE-based protocol of Supplementary Information. The CNOT gate marked with an asterisk represents the operations needed to store the parity of selected qubits in 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. In the circuit diagrams, QFT-inverse in the diagram stands for quantum Fourier transformation, H stands for Hadamard gate and R_m stands for single qubit rotations used to perform measurement in different basis. **b**, Experimentally measured probabilities of passing the standard-basis (p_A) and interference-measurement (p_B) challenges for the LWE protocols. **e**, Experimentally measured probabilities of passing the standard-basis (p_A) and interference-measurement (p_B) challenges for the factoring protocols. These probabilities are compared against the asymptotic classical limits ($p_A + 2p_B \leq 2$ for LWE, as derived in 'Discussion of the asymptotic classical threshold' section of Methods and $p_A + 4p_B \leq 4$ for factoring[12]). Results for both interactive and delayed-measurement versions of the protocols are presented. The numerical values of p_A and p_B for each experiment and the corresponding values of statistical significance are provided in Result data of Supplementary Information. **c**, The relative performance R of the experiments for all possible branches of the LWE protocols. **f**, The relative performance R of the experiments for all possible branches of the factoring protocols. Certain branches (thick lines) are robust to phase errors and exhibit similar performance for both interactive and delayed-measurement protocols. The number of shots (sample size) n for each bar is provided in Supplementary Information. Error bars are 95% confidence computed as a Wald interval.

Information). 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

There are two main experimental challenges to demonstrating quantum computational advantage via interactive protocols: (1) integrating mid-circuit measurements into arbitrary quantum circuits with sufficiently high overall fidelity to pass the verifier's tests and (2) scaling the protocols to large enough problems that it is classically infeasible to break the cryptographic assumptions. In this work, we have overcome the first obstacle, successfully implementing two interactive cryptographic proofs of quantumness with high enough fidelity to pass the verifier's challenges. We leave the second challenge, of scaling these protocols up, to future work. We estimate that one should be able to perform a cryptographic proof of quantum computational advantage using approximately 1,600 qubits (see 'Estimate of resources required to achieve a quantum advantage' section of Methods. Note that although this qubit count is comparable to some implementations of Shor's algorithm, the circuits are orders of magnitude smaller in gate count ($\mathcal{O}(n \log n)$ versus $\mathcal{O}(n^2 \log n)$) and depth[12]. Even with those smaller circuits, the challenge for near-term devices will almost certainly remain the circuit depth. Interestingly, recent advances suggest that our interactive protocols can be performed for constant depth at the cost of a larger number of qubits[42,43]. Once this scaling is achieved in an experiment, it will demonstrate a directly verifiable quantum computational advantage. This would mark a new step forward from recent sampling experiments, which have demonstrated the system sizes and fidelities necessary to make classical simulation extremely hard or impossible[17–26] but have no method to directly and efficiently verify the output (moreover, practical strategies for a classical impostor to replicate the sampling are still being explored[44–50]).

Our work may also lead to a number of other intriguing directions. A clear next step is to apply the power of quantum interactive protocols to achieve more than just quantum advantage, for example, with tasks such as certifiable random number generation, remote state preparation and verifying arbitrary quantum computations[4–6]. We emphasize that, unlike, for example, Bell-test protocols for random number generation, interactive proofs allow us to perform these cryptographic tasks with a single 'black-box' prover with which the verifier can interact only classically. This has the potential to allow these types of protocols (including our cryptographic proofs of quantumness) to be performed on a remote prover, such as a quantum cloud service on the internet, enabling a wide variety of practical applications. Finally, the advent of mid-circuit measurement capabilities in a number of platforms[29,31,51,52] enables the exploration of new phenomena, such as entanglement phase transitions[53–55] as well as the demonstration of coherent feedback protocols, including quantum error correction[2].

Online content

Any methods, additional references, Nature Portfolio reporting summaries, source data, extended data, supplementary information, acknowledgements, peer review information; details of author contributions and competing interests; and statements of data and code availability are available at https://doi.org/10.1038/s41567-023-02162-9.

References

1. Nielsen, M. A. & Chuang, I. L. *Quantum Computation and Quantum Information* 10th edn (Cambridge Univ. Press, 2010).
2. Ryan-Anderson, C. et al. Realization of real-time fault-tolerant quantum error correction. *Phys. Rev. X* **11**, 041058 (2021).
3. Cong, I., Choi, S. & Lukin, M. D. Quantum convolutional neural networks. *Nat. Phys.* **15**, 1273–1278 (2019).
4. Mahadev, U. Classical verification of quantum computations. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)* (ed. Rabani, Y.) 259–267 (IEEE, 2018).
5. Brakerski, Z., Christiano, P., Mahadev, U., Vazirani, U. & Vidick, T. 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)* (ed. Rabani, Y.) 320–331 (IEEE, 2018).
6. Gheorghiu, A. & Vidick, T. Computationally-secure and composable remote state preparation. In *2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS)* (ed. Rabani, Y.) 1024–1033 (IEEE, 2019).
7. Aharonov, D., Ben-Or, M. & Eban, E. Interactive proofs for quantum computations 453–469. Preprint at https://arxiv.org/abs/1704.04487 (2017).
8. Goldwasser, S., Micali, S. & Rackoff, C. The knowledge complexity of interactive proof systems. *SIAM J. Comput.* **18**, 186–208 (1989).
9. Lund, C., Fortnow, L., Karloff, H. & Nisan, N. Algebraic methods for interactive proof systems. In *Proc. 31st Annual Symposium on Foundations of Computer Science* Vol. 1 (ed. Yannakakis, M.) 2–10 (IEEE, 1990); https://doi.org/10.1109/FSCS.1990.89518
10. Shamir, A. IP = PSPACE. *J. ACM* **39**, 869–877 (1992).
11. Brakerski, Z., Koppula, V., Vazirani, U. & Vidick, T. Simpler proofs of quantumness. In *15th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2020)* (ed. Flammia, S. T.) 8:1–8:14 (Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2020).
12. Kahanamoku-Meyer, G. D., Choi, S., Vazirani, U. V. & Yao, N. Y. Classically verifiable quantum advantage from a computational Bell test. *Nat. Phys.* **18**, 918–924 (2022).
13. Bell, J. S. On the Einstein Podolsky Rosen paradox. *Physics* **1**, 195–200 (1964).
14. Hensen, B. et al. Loophole-free Bell inequality violation using electron spins separated by 1.3 kilometres. *Nature* **526**, 682–686 (2015).
15. Shalm, L. K. et al. Strong loophole-free test of local realism. *Phys. Rev. Lett.* **115**, 250402 (2015).
16. Giustina, M. et al. Significant-loophole-free test of Bell's theorem with entangled photons. *Phys. Rev. Lett.* **115**, 250401 (2015).
17. Arute, F. et al. Quantum supremacy using a programmable superconducting processor. *Nature* **574**, 505–510 (2019).
18. Zhong, H.-S. et al. Quantum computational advantage using photons. *Science* **370**, 1460–1463 (2020).
19. Wu, Y. et al. Strong quantum computational advantage using a superconducting quantum processor. *Phys. Rev. Lett.* **127**, 180501 (2021).
20. Zhu, Q. et al. Quantum computational advantage via 60-qubit 24-cycle random circuit sampling. *Sci. Bull.* **67**, 240–245 (2021).
21. Aaronson, S. & Arkhipov, A. The computational complexity of linear optics. In *Proc. 43rd Annual ACM Symposium on Theory of Computing* (ed. Fortnow, L.) 333–342 (Association for Computing Machinery, 2011).
22. Lund, A. P., Bremner, M. J. & Ralph, T. C. Quantum sampling problems, BosonSampling and quantum supremacy. *Npj Quantum Inf.* **3**, 15 (2017).
23. Harrow, A. W. & Montanaro, A. Quantum computational supremacy. *Nature* **549**, 203–209 (2017).
24. Boixo, S. et al. Characterizing quantum supremacy in near-term devices. *Nat. Phys.* **14**, 595–600 (2018).
25. Bouland, A., Fefferman, B., Nirkhe, C. & Vazirani, U. On the complexity and verification of quantum random circuit sampling. *Nat. Phys.* **15**, 159–163 (2019).
26. Aaronson, S. & Gunn, S. On the classical hardness of spoofing linear cross-entropy benchmarking. *Theory Comput. Syst.* **16**, 1–8 (2020).
27. Hensinger, W. K. Quantum computer based on shuttling trapped ions. *Nature* **592**, 190–191 (2021).
28. Bluvstein, D. et al. A quantum processor based on coherent transport of entangled atom arrays. *Nature* **604**, 451–456 (2022).
29. Pino, J. M. et al. Demonstration of the trapped-ion quantum ccd computer architecture. *Nature* **592**, 209–213 (2021).
30. Kielpinski, D., Monroe, C. & Wineland, D. J. Architecture for a large-scale ion-trap quantum computer. *Nature* **417**, 709–711 (2002).
31. Wan, Y. et al. Quantum gate teleportation between separated qubits in a trapped-ion processor. *Science* **364**, 875–878 (2019).
32. Regev, O. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM* **56**, 1–40 (2009).
33. Regev, O. The learning with errors problem. In *2010 IEEE 25th Annual Conference on Computational Complexity* (ed. van Melkebeek, D.) 191–204 (IEEE, 2010).
34. Rabin, M. O. *Digitalized Signatures and Public-key Functions as Intractable as Factorization*. Technical Report (Massachusetts Institute of Technology, 1979).
35. Goldwasser, S., Micali, S. & Rivest, R. L. A digital signature scheme secure against adaptive chosen-message attacks. *SIAM J. Comput.* **17**, 281–308 (1988).
36. Goldwasser, S., Micali, S. & Rivest, R. L. in *Advances in Cryptology, Proceedings of CRYPTO '84* (eds Blakley, R. R. & Chaum, D.) 467 (Springer, 1985).
37. Banerjee, A., Peikert, C. & Rosen, A. in *Advances in Cryptology—EUROCRYPT 2012* (eds Pointcheval, D. & Johansson, T.) 719–737 (Springer, 2012).
38. Alwen, J., Krenn, S., Pietrzak, K. & Wichs, D. in *Advances in Cryptology—CRYPTO 2013* (eds Canetti, R. & Garay, J. A.) 57–74 (Springer, 2013).
39. Clauser, J. F., Horne, M. A., Shimony, A. & Holt, R. A. Proposed experiment to test local hidden-variable theories. *Phys. Rev. Lett.* **23**, 880 (1969).
40. Monroe, C. et al. Programmable quantum simulations of spin systems with trapped ions. *Rev. Mod. Phys.* **93**, 025001 (2021).
41. Egan, L. et al. Fault-tolerant operation of a quantum error-correction code. *Nature* **598**, 281–286 (2021).
42. Hirahara, S. & Gall, F. L. Test of quantumness with small-depth quantum circuits. Preprint at https://arxiv.org/abs/2105.05500 (2021).
43. Liu, Z. & Gheorghiu, A. Depth-efficient proofs of quantumness. *Quantum* **6**, 807 (2022).

44. Huang, C. et al. Classical simulation of quantum supremacy circuits. Preprint at https://arxiv.org/abs/2005.06787 (2020).
45. Pan, F. & Zhang, P. Simulating the Sycamore quantum supremacy circuits. Preprint at https://arxiv.org/abs/2103.03074 (2021).
46. Gray, J. & Kourtis, S. Hyper-optimized tensor network contraction. *Quantum* **5**, 410 (2021).
47. Pan, F., Chen, K. & Zhang, P. Solving the sampling problem of the Sycamore quantum supremacy circuits. Preprint at https://arxiv.org/abs/2111.03011 (2021).
48. Yong et al. Closing the 'quantum supremacy' gap: achieving real-time simulation of a random quantum circuit using a new Sunway supercomputer. In *Proc. International Conference for High Performance Computing, Networking, Storage and Analysis* (ed. de Supinski, B. R.) 1–12 (Association for Computing Machinery, 2021).
49. Liu, X. et al. Redefining the quantum supremacy baseline with a new generation Sunway supercomputer. Preprint at https://arxiv.org/abs/2111.01066 (2021).
50. Gao, X. et al. Limitations of linear cross-entropy as a measure for quantum advantage. Preprint at https://arxiv.org/abs/2112.01657 (2021).
51. Córcoles, A. D. et al. Exploiting dynamic quantum circuits in a quantum algorithm with superconducting qubits. *Phys. Rev. Lett.* **127**, 100501 (2021).
52. Rudinger, K. et al. Characterizing midcircuit measurements on a superconducting qubit using gate set tomography. *Phys. Rev. Appl.* **17**, 014014 (2022).
53. Skinner, B., Ruhman, J. & Nahum, A. Measurement-induced phase transitions in the dynamics of entanglement. *Phys. Rev. X* **9**, 031009 (2019).
54. Li, Y., Chen, X. & Fisher, M. P. Quantum Zeno effect and the many-body entanglement transition. *Phys. Rev. B* **98**, 205136 (2018).
55. Noel, C. et al. Measurement-induced quantum phases realized in a trapped-ion quantum computer. *Nat. Phys.* **18**, 760–764 (2022).

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

[1]Joint Quantum Institute, and Department of Physics, NIST/University of Maryland, College Park, MD, USA. [2]Joint Center for Quantum Information and Computer Science, NIST/University of Maryland, College Park, MD, USA. [3]Departments of Electrical and Computer Engineering, University of Maryland, College Park, MD, USA. [4]IonQ, Inc., College Park, MD, USA. [5]Department of Physics, University of California, Berkeley, CA, USA. [6]Materials Sciences Division, Lawrence Berkeley National Laboratory, Berkeley, CA, USA. [7]Department of Computing and Mathematical Sciences, California Institute of Technology, Pasadena, CA, USA. [8]Division of Physics, Mathematics, and Astronomy, California Institute of Technology, Pasadena, CA, USA. [9]Duke Quantum Center and Department of Physics, Duke University, Durham, NC, USA. [10]Department of Electrical and Computer Engineering, Duke University, Durham, NC, USA. [11]Chemical Physics Program and Institute for Physical Science and Technology, University of Maryland, College Park, MD, USA. [12]Institute for Theoretical Studies, ETH Zürich, Zürich, Switzerland. [13]Department of Physics, University of Maryland, College Park, MD, USA. [14]Department of Computer Science and Applied Mathematics, Weizmann Institute of Science, Rehovot, Israel. [15]Department of Physics, Harvard University, Cambridge, MA, USA. [16]These authors contributed equally: Daiwei Zhu, Gregory D. Kahanamoku-Meyer. ✉e-mail: daiwei@terpmail.umd.edu

Methods

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[41,56]. The system consists of a chain of 15 single $^{171}\mathrm{Yb}^+$ 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 detecting state-dependent fluorescence[57]. Qubit operations are realized via pairs of Raman beams, derived from a single 355 nm mode-locked laser[58]. 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[59–61]. 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 beam[62].

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 post-selection, see below) for which the verifier accepted the prover's responses.

For both protocols, the check for the A or 'standard-basis' branch is simple. The prover has already supplied the verifier with the output 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 input x and confirms that it is equal to w.

For the B or interference measurement, the measurement scheme and verification check are different for the two protocols. For the LWE protocol, the interference measurement is an X-basis measurement of all of the qubits holding the input superposition $|x_0\rangle + |x_1\rangle$. This measurement will return a bit string d of the same length as the number of qubits in that superposition. 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, that is $a \cdot b = \sum_i a_i b_i \bmod 2$. 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 equation (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 bit string d as above. After performing that measurement, the prover holds the single-qubit state $(-1)^{d\cdot x_0}\,|r \cdot x_0\rangle + (-1)^{d\cdot x_0}\,|r \cdot x_0\rangle$, where $(\cdot)$ is the binary inner product as above and r is a random bit string 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 approximately 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 approximately 85% (Fig. 3).

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 trap[63], maintained at room temperature. (Technically, the matrix A is sampled together with the TCF trapdoor. However, as explained in ref. 5, the distribution from which the matrix is sampled is statistically close to a uniform distribution over $\mathbb{Z}_q^{m\times n}$.) We generate 40 time-dependent signals using a multi-channel digital-to-analogue converter voltage source, which controls the voltages of the 38 inner electrodes at the centre 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 potential of the central electrodes affects predominantly the axial trapping potential and, in turn, generates 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. 63 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 subchains of seven and eight ions and shuttle the eight-ion group to $x = 0.55$ mm away from the trap centre at $x = 0$. We then align the seven-ion chain with the individual-addressing Raman beams for the first mid-circuit measurement. For the LWE protocol, we then reverse the shuttling process and remerge the ions to a 15-ion chain, completing the circuit and performing a final measurement. For the factoring protocol, we shuttle the eight-ion subchain to the trap centre and the seven-ion subchain to $x = -0.55$ mm. We then split this chain into five- and three-ion subchains, shuttle the three-ion subchain to $x = 0.55$ mm and align the five ions at the trap centre with the Raman beams to perform additional gates and a second mid-circuit measurement. Finally, we move away the measured ions and align the three-ion group to the centre of the trap to complete the protocol. Reversing the sequence then prepares the ions in their initial state. For each protocol, all branches use the same shuttling sequences but differ in the qubit assignment and the realized gates. The duration of the mid-circuit measurement was experimentally determined before 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 subchain 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 the frequency of ion loss due to collisions with background gas by maintaining a calculated trap depth of at least 20 meV for each of the subchains throughout the shuttling sequences. The simulations enable efficient alignment of the subchains 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 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. Before the experiment, we run several calibration protocols that estimate the electrostatic potential at the centre of the trap through a Taylor series representation up to the fourth order, thus estimating the dominant effect of stray electric fields on the precalculated potential. We then cancel the effect of these fields using the central electrodes during the alignment and split sequences, as these sequences are most sensitive to the exact shape of the actual electrostatic potential. Additionally, we routinely measure the common-mode drift of the individual-addressing optical Raman beams along the linear axis of the trap and correct for them by automatically repositioning the ions by varying the potential.

During shuttling, the ions traverse an inhomogeneous magnetic field and, consequently, each ion spin acquires a shuttling-induced phase $\phi_{\mathrm{s}}^{(i)}$ that 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. After shuttling, $R_{\mathrm{x}}^{(i)}(\pi/2)R_{\mathrm{z}}^{(i)}(\phi)$ gates are applied, with ϕ scanned from 0 to 2π. Fitting the observed fringe for each ion enables us to estimate the phases $\phi_s^{(i)}$, which are corrected in the protocols by applying the inverse operation $R_{\mathrm{z}}^{(i)}(-\phi_{\mathrm{s}}^{(i)})$ after shuttling.

Discussion of the asymptotic classical threshold

In cryptography, showing that a new protocol is secure for practical use (meaning, in our case, that the proof cannot be spoofed by a classical prover) follows two broad steps: (1) proving that it is secure asymptotically (showing that the computational cost of cheating is at least superpolynomial in the problem size) and (2) picking a finite set of parameters such that cheating is not possible under certain classical resources (computational power and time, usually). What particular limitations are made for the resources available to the classical cheater are, ultimately, up to the user. In this section, we attempt to make precise exactly which statements are asymptotic (step 1), and how these statements make the jump in step 2 to finite, real parameters.

The first asymptotic statement, which is, perhaps, the most obvious, is that finding claws of the TCF is hard. In the theoretical papers upon which this work is based, this is shown by reducing the problem of finding claws to related problems for which there are standard cryptographic assumptions[5,12]. In particular, the assumptions are that the factoring and LWE problems have superpolynomial classical complexity. As discussed above, when using the test in practice, we would pick finite parameters in a way that finding a claw is infeasible for the set of classical resources that our quantum computer needs to outcompete (for a rigorous demonstration of quantum advantage, that would probably be a large supercomputer with ample runtime). Importantly, the reduction between the hardness of finding claws and breaking the cryptographic assumption is not in any sense asymptotic. For both TCFs, if a machine can find claws for a specific, finite set of parameters, these claws can be directly used to break the cryptographic assumption in practice. Therefore if the cryptographic assumption holds for a finite set of parameters, we can be sure that the claw-freeness does as well.

The second asymptotic statement used in the analysis of these protocols refers to the probability that a classical cheater passes a single iteration of the test. In 'Beating the classical threshold' section, we discuss the 'classical thresholds' that must be exceeded to demonstrate quantum capability. To be very precise about what we mean by this, we reproduce exactly what the theorems underlying these protocols state. If a classical prover's true success probabilities (not the empirically determined ones, which are subject to statistical fluctuation) exceed the given bound by a non-negligible amount, that prover could be used as a subroutine in a larger program that finds a claw in the TCF in polynomial time. Thus, if it is not possible for a classical prover with certain resources to find a claw (in a TCF with some specific parameters), it is provably also not possible for a classical prover with similar resources to non-negligibly exceed the threshold. There are two asymptotic portions of this statement: the polynomial time in which the larger program extracts a claw using the prover as a subroutine (which is the reason for the word 'similar' in the previous sentence) and the word 'negligible'. Negligible has a technical definition in cryptography, which is the sense in which we use it here. It means that a value (in this case, the amount by which the threshold can be exceeded) is bounded by a function that goes exponentially to zero in the problem size. The precise form of this exponential is not intended to be determined. Instead, the exponential decay is used to argue that the negligible function is 'essentially' zero for any reasonable problem size that would be used in practice.

Note that for the small problems we implement in this work, there is one instance in which this negligible function would meaningfully affect the classical success threshold, so we modify the protocol slightly to account for this. In the $x^2 \bmod N$ (Rabin's function) protocol, the value r sent by the verifier is supposed to be a uniformly random bit string. If r happens to be all zero, the product $r \cdot x$, whose value is supposed to be cryptographically hard to guess, is simply zero. This is not an issue for problem sizes that would be used for a full-scale test in practice, because an all-zero r is extremely unlikely to occur if r is of length several hundred bits. But for our smaller experiments with r of only a few bits, the all-zero string represents a sizeable fraction of possible r's. To prevent this from affecting the results, we simply choose our r from the set of non-zero bit strings rather than all bit strings. We note that excluding the all-zero string helps us better resolve the performance, too. When $r = 0^n$, the qubit measured in the last step of the protocol never interacts with any of the other qubits throughout the whole circuit, so the measurement result has nothing to do with the fidelity of the TCF circuit!

To close this discussion, it is worth taking a broader perspective and considering how the field of cryptography functions in general. Asymptotic proofs in cryptography are used to show that for any cheating machine with finite resources, a problem can always be made large enough to be hard in practice. The hardness grows quickly enough that this approach is, hopefully, not an unreasonable pursuit. But, ultimately, the question of how large the problem needs to be is an empirical one. Experts build the best possible algorithms and hardware they can and attempt to break the assumption. The parameters are then set to be larger than the largest problem size that can be broken this way (usually with an extra buffer added to secure against improvements in the attacks). In our case, the costs of breaking both factoring and LWE have been extensively explored, and the practical parameters needed for their security against current classical computing power are well understood. As described above, because there are no asymptotic statements in the reduction from the TCF to the underlying cryptographic assumptions, these parameters can be directly used to ensure that finding claws is hard in practice. As described above, the precise relation between the hardness of exceeding the thresholds and finding claws does rely on asymptotics, but the asymptotic function in the threshold has been shown to decay exponentially, which suggests strongly that this should not be an issue in practice.

Quantum-classical threshold for the LWE protocol

In this section, we state and prove the classical threshold for the LWE protocol. The corresponding proof for the factoring protocol is in the theoretical manuscript that first presented that protocol[12].

Below, the security parameter λ is used in the standard cryptographic sense, as a measure of the 'problem size'. It can be made larger to increase security or smaller to improve efficiency. The specifics of how each parameter of the LWE problem is defined as a function of λ can be found in the definition of the LWE TCF within the theoretical work that originally proposed it[5].

Proposition 1. For any classical prover, the probabilities that they pass branches A and B, namely p_{A} and p_{B}, must obey the relation

$$p_{\mathrm{A}} + 2p_{\mathrm{B}} - 2 < \epsilon(\lambda), \tag{2}$$

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

Proof. We first want to find the probability that the classical prover not only responds correctly for branch A but also (for the same output w that they committed to the verifier) correctly responds for branch B 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

$$p_{\mathrm{good}} \equiv \Pr_{w}[p_{\mathrm{B},w} > 1/2 + \mu(\lambda)]. \tag{3}$$

Pr[E] denotes the probability of event E. By a union bound, we arrive at a bound on the desired probability:

$$\Pr[A \text{ correct and } p_{\mathrm{B},w} > 1/2 + \mu(\lambda)] > p_{\mathrm{A}} + p_{\mathrm{good}} - 1. \tag{4}$$

Here, 'A correct' denotes the event where the prover correctly passes the verifier's challenge for branch A. Now, we wish to write p_{good} in terms of p_{B}. Let S be the set of w values for which $p_{\mathrm{B},w} > 1/2 + \mu(\lambda)$. 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 $1 - p_{\mathrm{good}}$, 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, that is

$$p_{\mathrm{B}} < 1 \cdot p_{\mathrm{good}} + 0.5 \cdot (1 - p_{\mathrm{good}}). \tag{5}$$

Solving for p_{good}, we then obtain

$$p_{\mathrm{good}} > 2p_{\mathrm{B}} - 1. \tag{6}$$

Substituting this into equation (4), we have

$$\Pr[A \text{ correct and } p_{\mathrm{B},w} > 1/2 + \mu(\lambda)] > p_{\mathrm{A}} + 2p_{\mathrm{B}} - 2. \tag{7}$$

However, notice that the probability on the left-hand side is the probability of breaking the adaptive hardcore bit property, which we know[5] must have

$$\Pr[A \text{ correct and } p_{\mathrm{B},w} > 1/2 + \mu(\lambda)] < \epsilon(\lambda), \tag{8}$$

where ϵ is a negligible function. Thus, combining this with equation (7), we obtain the desired inequality:

$$p_{\mathrm{A}} + 2p_{\mathrm{B}} - 2 < \epsilon(\lambda). \tag{9}$$

Computation of statistical significance contours

Here we describe the computation of the contour lines denoting various levels of statistical significance in Fig. 3b,e. Recall the probabilities p_{A} and p_{B} introduced in 'Beating the classical threshold' section, which denote the probabilities that the prover will pass the standard-basis and interference tests, 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 $p^{\mathrm{c}}_{\mathrm{A}}$ and $p^{\mathrm{c}}_{\mathrm{B}}$ that obey the bound $p^{\mathrm{c}}_{\mathrm{A}} + 2p^{\mathrm{c}}_{\mathrm{B}} - 2 < 0$ for the LWE protocol and $p^{\mathrm{c}}_{\mathrm{A}} + 4p^{\mathrm{c}}_{\mathrm{B}} - 4 < 0$ for the factoring 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 were 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 that 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 $(p_{\mathrm{A}}, p_{\mathrm{B}})$ with respect to a particular null hypothesis $\left(p^{\mathrm{c}}_{\mathrm{A}}, p^{\mathrm{c}}_{\mathrm{B}}\right)$, we define the 'quantumness' q of an experiment as $q(p_{\mathrm{A}}, p_{\mathrm{B}}) = p_{\mathrm{A}} + 4p_{\mathrm{B}} - 4$ for the factoring protocol and $q(p_{\mathrm{A}}, p_{\mathrm{B}}) = p_{\mathrm{A}} + 2p_{\mathrm{B}} - 2$ 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, $\mathrm{B}\left(N_{\mathrm{A}}, p^{\mathrm{c}}_{\mathrm{A}}\right)$ and $\mathrm{B}(N_{\mathrm{B}}, p^{\mathrm{c}}_{\mathrm{B}})$. Mathematically, the joint PMF is

$$f\left(k_{\mathrm{A}}, k_{\mathrm{B}}; p^{\mathrm{c}}_{\mathrm{A}}, p^{\mathrm{c}}_{\mathrm{B}}, N_{\mathrm{A}}, N_{\mathrm{B}}\right) = \binom{N_{\mathrm{A}}}{k_{\mathrm{A}}}\binom{N_{\mathrm{B}}}{k_{\mathrm{B}}}\left(p^{\mathrm{c}}_{\mathrm{A}}\right)^{k_{\mathrm{A}}}\left(p^{\mathrm{c}}_{\mathrm{B}}\right)^{k_{\mathrm{B}}}\left(1 - p^{\mathrm{c}}_{\mathrm{A}}\right)^{N_{\mathrm{A}} - k_{\mathrm{A}}}\left(1 - p^{\mathrm{c}}_{\mathrm{B}}\right)^{N_{\mathrm{B}} - k_{\mathrm{B}}}, \tag{10}$$

where $k_{\mathrm{A}} = p_{\mathrm{A}}N_{\mathrm{A}}$ and $k_{\mathrm{B}} = p_{\mathrm{B}}N_{\mathrm{B}}$ are 'counts' of the passing runs for each branch, respectively. Finally, we compute the statistical significance of a result $(p_{\mathrm{A}}, p_{\mathrm{B}})$ as the probability of achieving a quantumness measure of at least $q' = q(p_{\mathrm{A}}, p_{\mathrm{B}})$. Under a null hypothesis $\left(p^{\mathrm{c}}_{\mathrm{A}}, p^{\mathrm{c}}_{\mathrm{B}}\right)$, this is the sum of the PMFs over all k_{A} and k_{B} for which $q(k_{\mathrm{A}}/N_{\mathrm{A}}, k_{\mathrm{B}}/N_{\mathrm{B}}) > q'$.

In practice, for the contour lines of Fig. 3b,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.

Estimate of resources required to achieve a quantum advantage

For a conclusive demonstration of quantum advantage, the quantum machine must perform the protocol significantly faster than the amount of time a classical supercomputer would require to break the TCF, 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 on established cryptographic assumptions (like factoring and LWE) is that the classical hardness of breaking these assumptions has been extremely well studied, due to the implications for security[64]. 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. Thus, this would certainly constitute a definitive demonstration of quantum advantage. However, setting the parameters to these values may be considered overkill for our purposes, especially since we would like the problem size to be as small as possible 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 problems that would still 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 timescale of at most hours. Thus, to demonstrate quantum advantage by several orders of magnitude, we must set the parameters such that a classical supercomputer would require of the order of thousands of hours to break the TCF. In 2020, Boudot et al. reported the record-breaking factorization of a 795 bit semiprime[65]. The cost of the computation was about 1,000 core-years, meaning that a 1,000-core cluster would complete it in a year. We consider this a sufficient cost for demonstrating quantum advantage. We emphasize also that factoring is one of the most well-studied hard computational problems. The record set by Boudot et al. was 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. By computing and measuring the bits of the output value w one by one, the computational Bell-test protocol could be performed using only about 800 qubits with a 795 bit prime. However, the gate count and circuit depth can be dramatically reduced by explicitly storing the full output value w, thus requiring roughly 1,600 qubits in total[12]. Because it needs a much lower gate count, we use the 1,600 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 (1,000 core-years) is difficult to do exactly, because, to the best of 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 to see who can break the largest possible instance. As of this writing, the largest instances that have been solved use LWE vectors of about 500–1,000 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 1,000 core-years of computation time, we estimate that the LWE vectors would need to be perhaps 1,000–2,000 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 $\log(N)+1$ qubits[12]), it may be possible to perform the LWE protocol using a number of qubits comparable to the bit length of one LWE vector.

We now provide estimates for the cryptographic parameters. These parameter values should be such that it would be expected to be completely infeasible for a classical machine to break the TCF. For the factoring protocol, we apply the key sizes recommended by the National Institute of Standards and Technology (NIST) for the RSA cryptosystem, whose security relies on integer factorization. NIST recommends choosing a modulus N with a length of 2,048 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 requirement of 2,049 qubits[12]. However, the circuit depth can be improved significantly by including more qubits, so that a more efficient circuit can be realized with roughly $2\log(N) \approx 4{,}100$ qubits. Because LWE is not yet broadly used in practice, unlike RSA, NIST does not provide recommendations for key sizes in its documentation. However, we can use the estimates of Lindner and Peikert[66] to find parameter values that 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 $q = 4{,}093$ for a 'medium' level of security. Vectors with these parameters are $n\log(q) \approx 3{,}072$ bits long. To store both an input and output vector would, thus, require roughly 200 qubits. By repeatedly reusing a set of qubits to compute the output vector element by element, the computation could be performed using roughly 3,100 qubits.

Data availability

All data supporting the findings of this study are available in the paper or Methods. The raw experimental data are available from the corresponding author upon reasonable request.

References

56. Cetina, M. et al. Control of transverse motion for quantum gates on individually addressed atomic qubits. *PRX Quantum* **3**, 010334 (2022).
57. Olmschenk, S. et al. Manipulation and detection of a trapped Yb$^+$ hyperfine qubit. *Phys. Rev. A* **76**, 052314 (2007).
58. Debnath, S. et al. Demonstration of a small programmable quantum computer with atomic qubits. *Nature* **536**, 63 (2016).
59. Mølmer, K. & Sørensen, A. Multiparticle entanglement of hot trapped ions. *Phys. Rev. Lett.* **82**, 1835–1838 (1999).
60. Solano, E., de Matos Filho, R. L. & Zagury, N. Deterministic Bell states and measurement of the motional state of two trapped ions. *Phys. Rev. A* **59**, R2539–R2543 (1999).
61. Milburn, G., Schneider, S. & James, D. Ion trap quantum computing with warm ions. *Fortschr. Phys.* **48**, 801–810 (2000).
62. Choi, T. et al. Optimal quantum control of multimode couplings between trapped ion qubits for scalable entanglement. *Phys. Rev. Lett.* **112**, 190502 (2014).
63. Maunz, P. L. W. *High Optical Access Trap 2.0* Technical Report (Office of Scientific and Technical Information, 2016); https://doi.org/10.2172/1237003
64. Barker, E. *Recommendation for Key Management Part 1: General*. Technical report NIST SP 800-57pt1r4 (National Institute of Standards and Technology, 2016); https://doi.org/10.6028/NIST.SP.800-57pt1r4
65. Boudot, F. et al. in *Advances in Cryptology—CRYPTO 2020* (eds Micciancio, D. & Ristenpart, T.) 62–91 (Springer, 2020).
66. Lindner, R. & Peikert, C. in *Topics in Cryptology—CT-RSA 2011* (ed. Kiayias, A.) 319–339 (Springer, 2011).

Acknowledgements

We are grateful to V. Uhlir for the design of the verifier and prover figures. This work is supported by the Army Research Office (ARO) through the LogiQ programme of the Intelligence Advanced Research Projects Activity; the National Quantum Information Science Research Centers, Office of Science, US Department of Energy (DoE); the Quantum Systems Accelerator; the Air Force Office of Scientific Research (AFOSR) Multidisciplinary University Research Initiatives (MURIs) on Quantum Measurement/Verification and Quantum Interactive Protocols (Grant No. FA9550-18-1-0161) and Dissipation Engineering in Open Quantum Systems; the Software-Tailored Architecture for Quantum Co-design Program of the National Science Foundation (NSF); the ARO MURI on Modular Quantum Circuits; the Accelerated Research in Quantum Computing Program, Advanced Scientific Computing Research, DoE (Award No. DE-SC0020312); the Young Investigator Program of the AFOSR (Award No. FA9550-16-1-0495); the Quantum Leap Challenge Institutes, NSF (Grant No. OMA-2016245); the Institute for Quantum Information and Matter, an NSF Physics Frontiers Center (NSF Grant PHY-1125565); the Gordon and Betty Moore Foundation (Grant No. GBMF-12500028); Dr Max Rössler and the Walter Haefner Foundation through the ETH Zürich Foundation; the NSF (Award No. DMR-1747426); a Vannevar Bush Faculty Fellowship; the Office of Advanced Scientific Computing Research under the Accelerated Research in Quantum Computing Program; the A. P. Sloan Foundation; and the David and Lucile Packard Foundation.

Author contributions

D.Z., G.K.M., L.L., C.N., A.G., T.V., U.V., N.Y., M.C. and C.M. designed the research. D.Z., C.N., O.K., B.H., Q.W., A.R., L.F., D.B. and L.E. performed the experiments and collected the data. D.Z., G.K.M., L.L., C.N., L.E., A.G. and Y.N. compiled and optimized the circuit. D.Z., G.K.M. and L.L. analysed the data. All authors contributed to the preperation of the paper.

Competing interests

C.M. is the chief scientist for IonQ, Inc. and has a personal financial interest in the company. The other authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at https://doi.org/10.1038/s41567-023-02162-9.

Correspondence and requests for materials should be addressed to Daiwei Zhu.

Peer review information *Nature Physics* thanks Robin Blume-Kohout and the other, anonymous, reviewer(s) for their contribution to the peer review of this work.

Reprints and permissions information is available at www.nature.com/reprints.