Experimental Verifiable Multiclient Blind Quantum Computing on a Qline Architecture

Beatrice Polacchi,[1] Dominik Leichtle,[2] Gonzalo Carvacho,[1] Giorgio Milani,[1] Nicolò Spagnolo,[1] Marc Kaplan,[3] Elham Kashefi,[4,2,*] and Fabio Sciarrino[1,†]

[1]*Dipartimento di Fisica—Sapienza Università di Roma, Piazzale Aldo Moro 5, I-00185 Roma, Italy*
[2]*School of Informatics, University of Edinburgh, 10 Crichton Street, EH8 9AB Edinburgh, United Kingdom*
[3]*VeriQloud, 13 rue Victor Hugo, 92 120 Montrouge, France*
[4]*Laboratoire d'Informatique de Paris 6, CNRS, Sorbonne Université, 75005 Paris, France*

(Received 26 July 2024; revised 20 December 2024; accepted 13 February 2025; published 21 May 2025)

The exploitation of certification tools by end users represents a fundamental aspect of the development of quantum technologies as the hardware scales up beyond the regime of classical simulability. Certifying quantum networks becomes even more crucial when the privacy of their users is exposed to malicious quantum nodes or servers as in the case of multiclient distributed blind quantum computing, where several clients delegate a joint private computation to remote quantum servers, such as federated quantum machine learning. In such protocols, security must be provided not only by keeping data hidden but also by verifying that the server is correctly performing the requested computation while minimizing the hardware assumptions on the employed devices. Notably, standard verification techniques fail in scenarios where the clients receive quantum states from untrusted sources such as in a recently demonstrated linear quantum network performing multiclient blind quantum computation. However, recent theoretical results provide techniques to verify blind quantum computations even in the case of untrusted state preparation. Equipped with such theoretical tools, in this Letter we provide the first experimental implementation of a multiclient verifiable blind quantum computing protocol in a distributed architecture. Our results represent novel perspectives for the verification of multitenant distributed quantum computation in large-scale networks.

DOI: 10.1103/PhysRevLett.134.200603

Introduction—Quantum technologies promise to outperform classical devices in several tasks ranging from cryptography [1–4] and computation [5–9] to randomness generation [10–12]. However, analogously to the development of classical computers, near-term quantum devices are expected to be available to end users on the cloud [13–15]. This opens several privacy issues on the users' side, who must be equipped with tools to keep their data private, *blindness* of the computation, but also to check whether the quantum server they employ correctly performs the desired computations, *verification* of the computation. Certifying whether a quantum device is performing the correct operation is a challenging and widely investigated task [10,16–22]. On one hand, the challenges mainly come from the required resource overhead, e.g., for quantum tomography procedures, or from the need for partial assumptions about the inner details of the quantum hardware under investigation. This is often seen as a major drawback, since, once quantum technologies are available to general users, the latter may completely ignore the technical details of the hardware they are using and may need certification protocols that rely on minimal assumptions about it, i.e., the so-called device-independent framework [21]. On the other hand, such certification procedures represent a key ingredient for realizing a secure quantum internet [13–15]. Their importance becomes even more evident when a user needs to perform delegated tasks, such as delegated blind quantum computing (BQC) [23–30], a class of protocols that allow users with minimal quantum resources to delegate hard computations to powerful remote quantum servers while keeping hidden input and outcome data, as well as the calculation itself. Because of the flourishing of federated protocols [31,32], several multiclient versions of BQC have been theoretically proposed [33–37] and recently experimentally demonstrated in a two-client setting [38]. The theoretical proposals included the possibility of verifying the computation, which was demonstrated in one-client settings [39–43], with the assumption that the clients owned trusted single-qubit sources. The protocol presented in [38], instead, is based on a linear quantum network structure, known as Qline [44], featuring an untrusted quantum source that distributes qubits along linear quantum channels. However, the use of the Qline

*Contact author: elham.kashefi@lip6.fr
†Contact author: fabio.sciarrino@uniroma1.it

architecture triggered a fundamental question regarding the verifiability of the computation in this scenario, since the single-qubit sources are not owned by the clients and are, therefore, untrusted. Notably, the issue of verifying BQC with untrusted sources of quantum states was recently addressed in [45], where the authors demonstrate that, if the clients can also perform bit-flip operations together with z rotations, they can verify that the server is correctly performing the computation, even without trusting the qubit source.

In this Letter, we present the first experimental implementation of a verifiable multiclient BQC protocol where the clients receive single qubits from an untrusted source. In detail, we exploit the versatile adaptive photonic platform introduced in [38], where the clients are connected through fiber links and perform single-qubit transformations sequentially over a Qline. This platform exploits a trusted third party (TTP) that secures the classical communication between the clients and the server, therefore reducing the time latency on the server's side. We tailor this platform to the verifiable BQC protocol introduced in [45]. Our protocol can be extended to more parties and arbitrary distances. In particular, we experimentally demonstrate the protocol for two and three clients, thus providing a step forward to realizing secure and scalable distributed quantum cloud infrastructures.

Multiclient BQC—We suppose that a group of untrusted clients own only single-qubit transformation devices, and want to delegate the joint computation shown in Fig. 1(a) to a remote quantum server while keeping their private data hidden. In our protocol, we suppose that the clients desire to perform the classical function described by the two-qubit measurement pattern $\{\phi_1, \phi_2\} = \{(\pi/2), (\pi/2)\}$. The clients respectively insert two classical input bits, x_1 and x_2. To mask the computation data, the clients carry out the multiclient BQC protocol introduced in [38] and depicted in Fig. 1(b). The other parties involved in the protocol are (i) an untrusted source of qubits called server S_1, (ii) a trusted third party (TTP) to secure classical communications between the clients and the server, and (iii) a remote quantum server that performs the computation, indicated as S_2. When the source behaves honestly, it distributes maximally entangled bipartite states along two Qlines of the form

$$|\psi\rangle = \frac{1}{2}(|00\rangle + |01\rangle + |10\rangle - |11\rangle). \tag{1}$$

The clients sequentially apply the following one-time pad transformation to the two qubits: $R_z(\theta_1^{C_j}) \otimes R_z(\theta_2^{C_j})$, where the angles $\theta_i^{C_j}$ are randomly drawn in the set $\mathcal{A} = \{0, \pi/4, \ldots, 7\pi/4\}$. We use $i = 1, 2$ as the qubit index and $j = 1, 2, 3, \ldots, N$ as the client index. Moreover, each client chooses two random bits $r_i^{C_j}$ to hide the outcome of the computation. The clients then communicate their secret

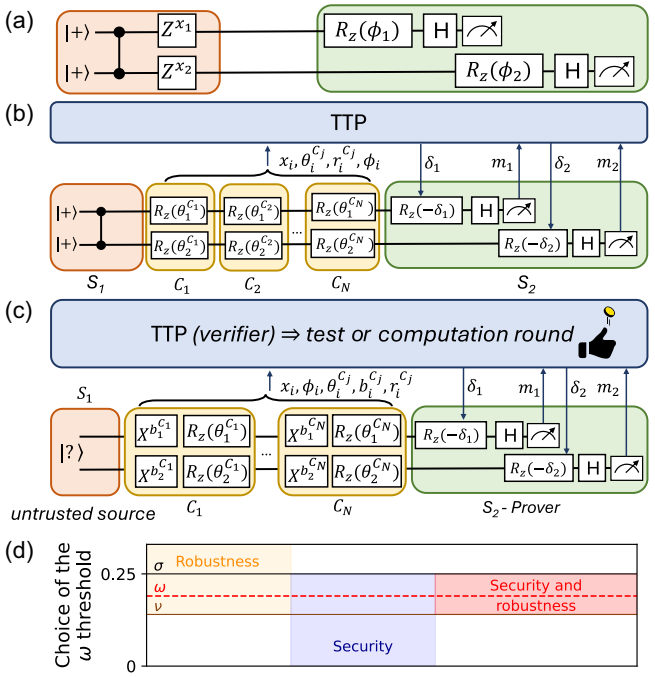

FIG. 1. Conceptual scheme of verifiable multiclient BQC with N clients and two Qlines. (a) Target circuit without any masking. (b) Nonverifiable BQC with N clients. (c) Modification of the scheme in (b) to perform verifiable BQC with N clients. (d) Choice of the threshold of tolerated failed test rounds ω. This threshold must be chosen in the interval between the maximum theoretical threshold for security, σ, and the maximum expected level of hardware noise, ν.

parameters $\theta_i^{C_j}$, $r_i^{C_j}$ to the TTP. We define $\theta_i = \sum_j \theta_i^{C_j}$ and $r_i = \bigoplus_j r_i^{C_j}$, for the ith qubit. The resulting quantum state is then sent to server S_2. From now on, the clients and S_2 only communicate classically, through the TTP. The protocol requires one round for each qubit, and, at the ith round, the TTP computes the blind measurement angle δ_i according to the formula

$$\delta_i = \theta_i + x_i\pi + (-1)^{m_{i-1}^{\text{true}}}\phi_i + r_i\pi, \tag{2}$$

where $m_0^{\text{true}} = 0$ and $m_1^{\text{true}} = m_1 \oplus r_1$. Both outcomes are corrected by the TTP according to $m_i^{\text{true}} = m_i \oplus r_i$ at the end of the protocol. The protocol is repeated many times and the outcome of the computation is chosen through majority vote among all rounds.

Verifiable multiclient BQC—To allow for verification of the computation, two main modifications are needed. A conceptual illustration of the modified protocol is depicted in Fig. 1(c). First, we need to repeat the experiment n times, m of which are used as test rounds, while $n - m$ as computational rounds [45,46]. We notice that test rounds introduce only a polynomial overhead in terms of the round complexity of the protocol, thus preserving the efficiency and scalability of the original protocol. The results from the

test rounds are used to decide whether the computational results can be accepted as correct or not. The choice between test or computation is performed run by run from the TTP.

A second modification is needed in the state preparation procedure. In detail, the clients need to randomly prepare the received qubits according to the transformation $R_z(\theta_1^{C_j})X^{b_1^{C_j}} \otimes R_z(\theta_2^{C_j})X^{b_2^{C_j}}$, as shown in Fig. 1(c). The bits $b_i^{C_j}$ are randomly chosen by the clients to perform the verification protocol [45] and secretly communicated to the TTP. For easiness of notation, we also introduce the quantity $b_i = \bigoplus_j b_i^{C_j}$. In computation rounds, the blind measurement angle δ_i at the ith round must be computed in an adaptive way by the TTP according to the formula given in Eq. (2), where the angles θ_i are substituted with the angles $\theta'_i = \theta_i^{C_N} + \sum_{j=1}^{N-1}(-1)^{\oplus_{k=j+1}^{N} b_i^{C_k}}\theta_i^{C_j} + b_{(i\%2+1)}\pi$, where N is the number of clients involved.

In test rounds, the target algorithm of the clients is the stabilizer measurement $Y \otimes Y$. Indeed, since the outcome of such measurements is known, the clients can detect eventual malicious deviation on the server side [45,46]. In practice, this means that the second measurement is fixed to $\delta_2 = \theta'_2 + x_2\pi + \phi_2 + r_2\pi$.

Test rounds can fail because of deviations (malicious or not) from the ideal behavior of the server or because of overly noisy hardware. While the *security* of the protocol ensures blindness and verifiability in the case of malicious deviations, even if the server behaves honestly, real-life hardware is affected by noise. Hence, the *robustness* property ensures that in the realistic, noisy case the protocol accepts with high probability. As a core parameter of the protocol, the clients need to choose a threshold for the number of tolerated failed test rounds [46,47]. The need for security and robustness yields two different bounds on the tolerable test failure rate ω in the protocol, as we depict in Fig. 1(d). On one hand, the security threshold σ, an upper bound on ω, depends on both the inherent, algorithmic error probability of the target bounded-error quantum polynomial computation p and the number of different types of tests k. It amounts to $\sigma = 0.25$ in our protocol (see Theorem 1 of [46]), and it does not depend on the number of clients involved. Further detail about the computation of this threshold can be found in Supplemental Material [48]. Only for a choice of $\omega < \sigma$, the security of the protocol is guaranteed. On the other hand, the expected magnitude of hardware noise implies a natural lower bound ν on ω to obtain robustness, where ν is the expected maximum test failure rate on the used hardware (see Theorem 2 of [46]). Such a parameter depends on the features of the untrusted qubit source and can be a public property of a given service provider. However, we stress that the protocol security does not depend on the choice of ν. Indeed, it is only provided as an indication to the clients about the lower bound on the choice of the ω threshold that must be taken into account in order to achieve sufficient robustness to the noise present in the setup, i.e., a reasonable rate of protocol round acceptance, and could not lead to a breach of security.

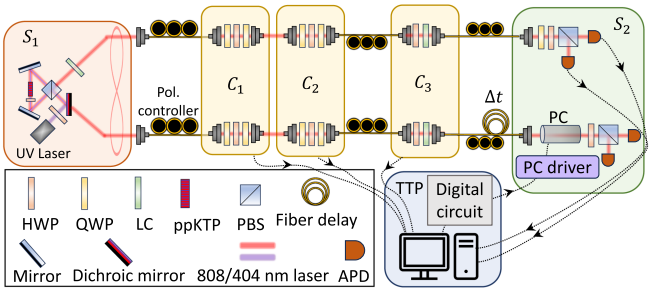

FIG. 2. Experimental setup. A Sagnac-based source [49] generates polarization-entangled photon pairs that are sent via a fiber link to the clients C_1, C_2, and C_3, who perform their transformation through wave plates and liquid crystals. Server S_2 performs measurements on the two qubits. The second qubit measurement is chosen according to the first measurement outcome via an active feed-forward system acting through an electro-optical modulator, i.e., a Pockels cell (PC). The computation is orchestrated by the TTP.

Once the clients have chosen the threshold of tolerated failed test rounds ω, if the recorded error fraction during test rounds, ε, is such that $\varepsilon \leq \omega$, they perform a majority vote on the computation rounds. Otherwise, the protocol aborts. These parameters are tunable depending on the desired level of privacy and robustness from the clients' perspective, analytically providing the required number of protocol rounds.

Results—We now present the results of our implementation in a two- and a three-client setting. The experimental setup employed in the three-client protocol is described in Fig. 2. Further details also on the two-client setup one can be found in the Supplemental Material. The experimental results for the two-client setting are shown in Figs. 3(a), 3(b), and 3(e). We set the threshold for the tolerated fraction of failed test rounds to $\omega_2 = 0.18$, well in between the lower robustness bound $\nu_2 \approx 0.14$, the error rate expected when considering a noisy model based on the characteristics of the devices we employ (see Supplemental Material), and the upper security bound $\sigma = 0.25$. The subscript 2 indicate the two-client case. This choice of ω_2 is the result of equally minimizing the security error and the robustness error of the protocol, for which concrete expressions can be found in the proofs of Theorems 3 and 4 from [46], respectively. In Figs. 3(a) and 3(b), we summarize the results for the classical computation of the algorithm $(\phi_1, \phi_2, x_1, x_2) = [(\pi/2), (\pi/2), 0, 0]$. Test rounds [Fig. 3(a)] are used by the clients to draw conclusions on the server behavior. In particular, since in test rounds the clients measure stabilizers of the two-qubit cluster state employed, the outcomes can only be equal for the two qubits, i.e., one can only have $(m_1, m_2) = (0, 0)$ or $(m_1, m_2) = (1, 1)$. We record an experimental error rate

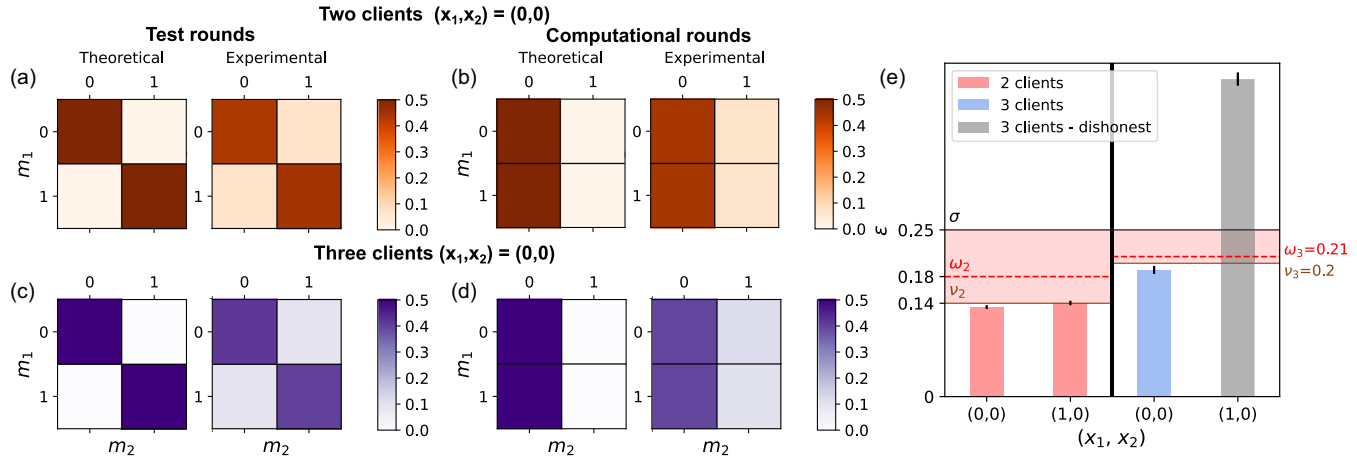

FIG. 3. Experimental results. (a),(b) Theoretical and experimental outcome probabilities of the test (a) and computation (b) rounds in the two-client protocol, for the algorithm $(\phi_1, \phi_2, x_1, x_2) = [(\pi/2), (\pi/2), 0, 0]$. (c),(d) Theoretical and experimental outcome probabilities of the test (c) and computation (d) rounds in the three-client setting for input $(x_1, x_2) = (0, 0)$. (e) Robustness and security of the protocol. We find that, for the first three algorithms, the experimental test error fraction ε is such that $\varepsilon < \omega$ and therefore the protocol is robust and secure. In the last case (gray bar), the source acts dishonestly by sending separable states. Consequently, the test failure rate exceeds the security threshold.

equal to $\varepsilon_2 = (13.4 \pm 0.3)\%$, as we show in Fig. 3(e). Since $\varepsilon_2 < \omega_2 < \sigma$, the clients can apply the majority vote to get the correct outcome. The outcomes for the computation rounds related to this algorithm are shown in Fig. 3(b). The figure shows that in the majority of rounds ($\approx$86.6%) the clients obtained the results $(m_1, m_2) = (0, 0)$ or $(m_1, m_2) = (1, 0)$, which corresponds to the theoretical expectations for this algorithm. The total number of performed rounds is equal to 27441, which suppresses the probability that the protocol rejects because of the observed noise below the order of 10^{-26}. In all other cases, the clients can conclude that the computation is correct with a soundness error below 10^{-22}. These numbers can be obtained from the concrete formulas for the security and robustness errors in the proofs of Theorems 3 and 4 from [46]. We repeat the experiment for the algorithm $(\phi_1, \phi_2, x_1, x_2) = [(\pi/2), (\pi/2), 1, 0]$. In this case, the experimental error rate averaged over all test rounds amounts to $\varepsilon_2 = (14.0 \pm 0.3)\%$, which is still compatible with security and robustness, as also visible in Fig. 3(e). Further detail about this algorithm can be found in Supplemental Material.

The experimental frequencies for the algorithm with input $(x_1, x_2) = (0, 0)$ in the three-client setting are summarized in Figs. 3(c) and 3(d) and compared with the theoretical expectations in test and computation rounds. In a honest protocol execution, the failure rate averaged over all test rounds amounts to $\varepsilon_3 = (19.0 \pm 0.6)\%$. The total number of rounds amounts to 8969. With three clients, the expected hardware noise is higher than in the two-client protocol and is estimated to be $\nu_3 \approx 0.2$. For $\omega_3 = 0.21$, we obtain a security bound of 0.012 with a maximum robustness error of 0.041. Since the recorded failure rate is such that $\varepsilon_3 < \omega_3$, the clients accept the results of the computation, shown in Fig. 3(d).

In Fig. 3(e), we summarize our experimental test failure rates (colored bars) by comparing them with the ω, ν, and σ thresholds for both multiclient settings. In the first three algorithms, the error rate is below the ω threshold and, therefore, the clients accept the protocol outcome. We repeated the experiment by deliberately considering a scenario where the source S_1 acts dishonestly, by sending separable states to the clients, with input data $(x_1, x_2) = (1, 0)$. In this case, we expect that the failure rate exceeds the security threshold ω_3, which is indeed what we observe in the gray bar of Fig. 3(e), thereby prompting the clients to reject the outcome of the delegated computation. We conclude that our protocol satisfies both the robustness and the security properties to a satisfyingly high degree, for both configurations with two and three clients.

To demonstrate the blindness of the initial quantum state, we measured the density matrix averaged over all 64 combinations obtained by varying the parameters $\theta_1^{C_1}$ and $\theta_2^{C_2}$. For the two-client protocol, the fidelity with the completely mixed state amounts to $F_2 = 0.9901 \pm 0.0002$. For the three-client protocol, it is $F_3 = 0.9969 \pm 0.0002$ (see also Supplemental Material). These results indicate that the server can get almost no information about the received quantum state.

Discussions—In this Letter, we addressed a fundamental open question about the verification of BQC in multiclient scenarios featuring untrusted sources of quantum states. The main advantage of our platform, compared with other platforms and protocols, is the presence of the TTP that secures the communication between the clients and the server, thus reducing the latency on the server side. We

stress that the introduced extra test rounds cause only a polynomial overhead in terms of the round complexity of the protocol. Hence, the overall protocol remains efficient and scalable similarly to the original protocol [38], while being equipped with verifiability. We showed that the protocol remains blind even after such modifications while being able to provide information about the server's honesty with an arbitrarily high confidence level, depending on the number of runs performed. Overall, our findings support the first demonstration of verifiable multiclient BQC, while further implementations [39–43] are restricted to the one-client setting.

Our results provide evidence of the scalability of this approach toward large networks. Our protocol can be scaled up by adding clients or qubits. Adding clients would require more rotation stations along each Qline and increasing the distances. This would affect photon losses proportionally to the number of stations and to the physical distance among clients, as well as affecting the time needed to store and manipulate the classical information. However, notably, in our platform, this does not translate into the need to increase the optical time delay between the two measurements in the server's station. Consequently, the losses due to this part of the apparatus do not scale with the number of clients. Analogously, the feedback mechanism would only need minimal adjustment.

Furthermore, the security of the protocols realized in this work rely on the fact that each participating client is present on each Qline. From the perspective of any honest client, it is therefore trivially guaranteed that at least one honest client has randomized the qubits along each of the Qlines. Alternatively, we could choose to adopt a distinct perspective at the security of the protocol, considering every rotation on every Qline as being performed by a separate client, and making the direct assumption that at least one client per Qline is behaving honestly during the execution of the protocol. This effectively bans any coalition of colluding malicious parties to take over the control of all operations along any one Qline. From this perspective, the experiments presented previously can be considered involving four or even six clients, amending the necessary assumptions for security.

On the other hand, increasing the size of the computation would require the generation and storage of larger cluster states for the computation [23], a task that is currently under development according to the present technology. However, our architecture is versatile and can be adapted to scenarios involving larger states and independent quantum state sources.

Concerning the verification protocol, scaling up the number of qubits and the amount of entanglement in the resource state would potentially increase the number of stabilizer measurements that the clients would need to perform to be able to detect all possible types of errors [33], thereby increasing the number of types of tests and possibly lowering the detection rate. While this number directly depends on the connectivity between the qubits, it is known that some families of graph states, such as brickwork states, achieve universality for quantum computing while retaining a constant detection rate independent of the number of qubits. Therefore, a small, constant number of types of tests are sufficient to achieve verifiability for universal quantum computing of arbitrary size.

In conclusion, our work motivates the theoretical investigation and the experimental demonstration of secure and verifiable BQC protocols in scenarios where the clients do not have control over the quantum sources of states as well as the remote quantum computational server node. Therefore, we believe that our findings represent a significant step toward implementing safe and densely connected quantum networks.

Acknowledgments—All authors acknowledge the support of the European Union's Horizon 2020 research and innovation program through the FET project PHOQUSING ("PHOtonic QUantum SamplING machine"—Grant Agreement No. 899544) and by ICSC—Centro Nazionale di Ricerca in High Performance Computing, Big Data and Quantum Computing, funded by European Union—NextGenerationEU. B. P. acknowledges support from the Sapienza research grant Avvio alla Ricerca 2022 (No. AR1221816B79122A). D. L. and E. K. acknowledge the support of the ANR research grant ANR-21-CE47-0014 (SecNISQ) and the EPSRC grant EP/X026167/1. D. L. acknowledges support by the Engineering and Physical Sciences Research Council [grant reference EP/T001062/1], and the UK National Quantum Computer Centre [NQCC200921], which is a UKRI Centre and part of the UK National Quantum Technologies Programme (NQTP).

Data availability—The data are available upon reasonable request from the corresponding authors.

[1] C. H. Bennett and G. Brassard, Quantum cryptography: Public key distribution and coin tossing, Theor. Comput. Sci. **560**, 7 (2014).
[2] A. K. Ekert, Quantum cryptography based on Bell's theorem, Phys. Rev. Lett. **67**, 661 (1991).
[3] A. K. Ekert, Quantum cryptography and Bell's theorem, in *Quantum Measurements in Optics* (Springer, New York, 1992), pp. 413–418.
[4] S. Pirandola, U. L. Andersen, L. Banchi, M. Berta, D. Bunandar, R. Colbeck, D. Englund, T. Gehring, C. Lupo, C. Ottaviani *et al.*, Advances in quantum cryptography, Adv. Opt. Photonics **12**, 1012 (2020).
[5] D. Deutsch and R. Jozsa, Rapid solution of problems by quantum computation, Proc. R. Soc. A **439**, 553 (1992).
[6] P. W. Shor, Algorithms for quantum computation: Discrete logarithms and factoring, in *Proceedings 35th Annual*

Symposium on Foundations of Computer Science (IEEE, New York, 1994), pp. 124–134.
[7] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. Brandao, D. A. Buell *et al.*, Quantum supremacy using a programmable superconducting processor, Nature (London) **574**, 505 (2019).
[8] H.-S. Zhong, H. Wang, Y.-H. Deng, M.-C. Chen, L.-C. Peng, Y.-H. Luo, J. Qin, D. Wu, X. Ding, Y. Hu *et al.*, Quantum computational advantage using photons, Science **370**, 1460 (2020).
[9] S. Aaronson, A. Cojocaru, A. Gheorghiu, and E. Kashefi, Complexity-theoretic limitations on blind delegated quantum computation, in 46th International Colloquium on Automata, Languages, and Programming (Schloss Dagstuhl—Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 2019).
[10] A. Acín and L. Masanes, Certified randomness in quantum physics, Nature (London) **540**, 213 (2016).
[11] Y. Liu, Q. Zhao, M.-H. Li, J.-Y. Guan, Y. Zhang, B. Bai, W. Zhang, W.-Z. Liu, C. Wu, X. Yuan *et al.*, Device-independent quantum random-number generation, Nature (London) **562**, 548 (2018).
[12] M. Herrero-Collantes and J. C. Garcia-Escartin, Quantum random number generators, Rev. Mod. Phys. **89**, 015004 (2017).
[13] A. S. Cacciapuoti, M. Caleffi, F. Tafuri, F. S. Cataliotti, S. Gherardini, and G. Bianchi, Quantum internet: Networking challenges in distributed quantum computing, IEEE Netw. **34**, 137 (2019).
[14] S. Wehner, D. Elkouss, and R. Hanson, Quantum internet: A vision for the road ahead, Science **362**, eaam9288 (2018).
[15] H. J. Kimble, The quantum internet, Nature (London) **453**, 1023 (2008).
[16] I. Agresti, B. Polacchi, D. Poderini, E. Polino, A. Suprano, I. Šupić, J. Bowles, G. Carvacho, D. Cavalcanti, and F. Sciarrino, Experimental robust self-testing of the state generated by a quantum network, PRX Quantum **2**, 020346 (2021).
[17] I. Šupić and J. Bowles, Self-testing of quantum systems: A review, Quantum **4**, 337 (2020).
[18] T. H. Yang, T. Vértesi, J.-D. Bancal, V. Scarani, and M. Navascués, Robust and versatile black-box certification of quantum devices, Phys. Rev. Lett. **113**, 040401 (2014).
[19] A. Gheorghiu, T. Kapourniotis, and E. Kashefi, Verification of quantum computation: An overview of existing approaches, Theory Comput. Syst. **63**, 715 (2019).
[20] N. D'Alessandro, B. Polacchi, G. Moreno, E. Polino, R. Chaves, I. Agresti, and F. Sciarrino, Machine-learning-based device-independent certification of quantum networks, Phys. Rev. Res. **5**, 023016 (2023).
[21] S. Pironio, V. Scarani, and T. Vidick, Focus on device independent quantum information, New J. Phys. **18**, 100202 (2016).
[22] J. Eisert, D. Hangleiter, N. Walk, I. Roth, D. Markham, R. Parekh, U. Chabaud, and E. Kashefi, Quantum certification and benchmarking, Nat. Rev. Phys. **2**, 382 (2020).
[23] A. Broadbent, J. Fitzsimons, and E. Kashefi, Universal blind quantum computation, in *2009 50th Annual IEEE Symposium on Foundations of Computer Science* (IEEE, New York, 2009), pp. 517–526.
[24] A. Mantri, C. A. Pérez-Delgado, and J. F. Fitzsimons, Optimal blind quantum computation, Phys. Rev. Lett. **111**, 230502 (2013).
[25] T. Morimae and K. Fujii, Blind quantum computation protocol in which Alice only makes measurements, Phys. Rev. A **87**, 050301(R) (2013).
[26] T. Morimae, Verification for measurement-only blind quantum computing, Phys. Rev. A **89**, 060302(R) (2014).
[27] Y.-B. Sheng and L. Zhou, Deterministic entanglement distillation for secure double-server blind quantum computation, Sci. Rep. **5**, 7815 (2015).
[28] M. Hayashi and T. Morimae, Verifiable measurement-only blind quantum computing with stabilizer testing, Phys. Rev. Lett. **115**, 220502 (2015).
[29] J. F. Fitzsimons and E. Kashefi, Unconditionally verifiable blind quantum computation, Phys. Rev. A **96**, 012303 (2017).
[30] Y.-B. Sheng and L. Zhou, Blind quantum computation with a noise channel, Phys. Rev. A **98**, 052343 (2018).
[31] J. Konečný, H. B. McMahan, F. X. Yu, P. Richtárik, A. T. Suresh, and D. Bacon, Federated learning: Strategies for improving communication efficiency, arXiv:1610.05492.
[32] Q. Yang, Y. Liu, T. Chen, and Y. Tong, Federated machine learning: Concept and applications, ACM Trans. Intell. Syst. Technol. (TIST) **10**, 1 (2019).
[33] T. Kapourniotis, E. Kashefi, D. Leichtle, L. Music, and H. Ollivier, Asymmetric quantum secure multi-party computation with weak clients against dishonest majority, Quantum Sci. Technol. **10**, 025015 (2025).
[34] R.-T. Shan, X. Chen, and K.-G. Yuan, Multi-party blind quantum computation protocol with mutual authentication in network, Sci. China Inf. Sci. **64**, 162302 (2021).
[35] G.-J. Qu and M.-M. Wang, Secure multi-party quantum computation based on blind quantum computation, Int. J. Theor. Phys. **60**, 3003 (2021).
[36] M. Ciampi, A. Cojocaru, E. Kashefi, and A. Mantri, Secure two-party quantum computation over classical channels, arXiv:2010.07925.
[37] E. Kashefi and A. Pappa, Multiparty delegated quantum computing, Cryptography **1**, 12 (2017).
[38] B. Polacchi, D. Leichtle, L. Limongi, G. Carvacho, G. Milani, N. Spagnolo, M. Kaplan, F. Sciarrino, and E. Kashefi, Multi-client distributed blind quantum computation with the Qline architecture, Nat. Commun. **14**, 7743 (2023).
[39] P. Drmota, D. P. Nadlinger, D. Main, B. C. Nichol, E. M. Ainley, D. Leichtle, A. Mantri, E. Kashefi, R. Srinivas, G. Araneda, C. J. Ballance, and D. M. Lucas, Verifiable blind quantum computing with trapped ions and single photons, Phys. Rev. Lett. **132**, 150604 (2024).
[40] H.-L. Huang, Q. Zhao, X. Ma, C. Liu, Z.-E. Su, X.-L. Wang, L. Li, N.-L. Liu, B. C. Sanders, C.-Y. Lu *et al.*, Experimental blind quantum computing for a classical client, Phys. Rev. Lett. **119**, 050503 (2017).
[41] C. Greganti, M.-C. Roehsner, S. Barz, T. Morimae, and P. Walther, Demonstration of measurement-only blind quantum computing, New J. Phys. **18**, 013020 (2016).
[42] S. Barz, J. F. Fitzsimons, E. Kashefi, and P. Walther, Experimental verification of quantum computation, Nat. Phys. **9**, 727 (2013).

[43] S. Barz, E. Kashefi, A. Broadbent, J. F. Fitzsimons, A. Zeilinger, and P. Walther, Demonstration of blind quantum computing, Science **335**, 303 (2012).

[44] M. Doosti, L. Hanouz, A. Marin, E. Kashefi, and M. Kaplan, Establishing shared secret keys on quantum line networks: Protocol and security, arXiv:2304.01881.

[45] E. Kashefi, D. Leichtle, L. Music, and H. Ollivier, Verification of quantum computations without trusted preparations or measurements, arXiv:2403.10464.

[46] D. Leichtle, L. Music, E. Kashefi, and H. Ollivier, Verifying BQP computations on noisy devices with minimal overhead, PRX Quantum **2**, 040302 (2021).

[47] T. Kapourniotis, E. Kashefi, D. Leichtle, L. Music, and H. Ollivier, Unifying quantum verification and error-detection: Theory and tools for optimisations, Quantum Sci. Technol. **9**, 035036 (2024).

[48] See Supplemental Material at http://link.aps.org/supplemental/10.1103/PhysRevLett.134.200603 for further technical details about the experimental setup, the computation of the protocol security threshold, and the experimental outcomes.

[49] A. Fedrizzi, T. Herbst, A. Poppe, T. Jennewein, and A. Zeilinger, A wavelength-tunable fiber-coupled source of narrowband entangled photons, Opt. Express **15**, 15377 (2007).