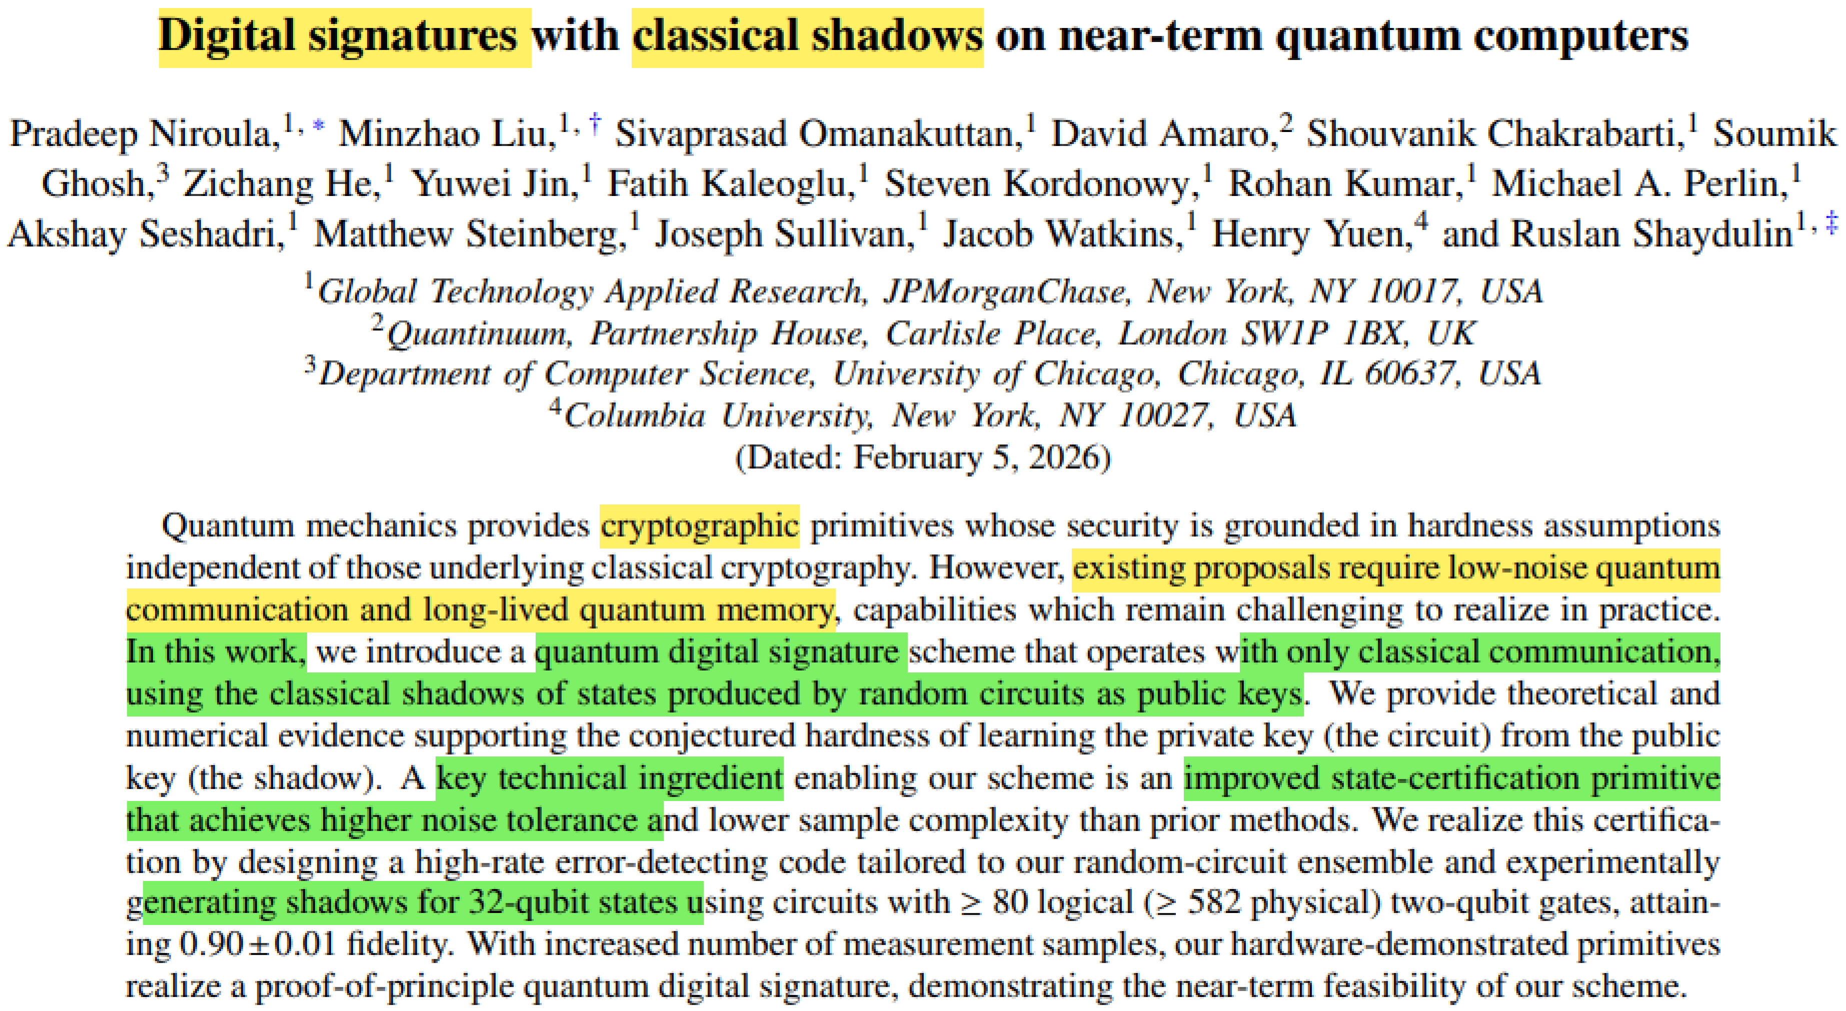

Digital signatures with classical shadows on near-term quantum computers

Pradeep Niroula,[1,*] Minzhao Liu,[1,†] Sivaprasad Omanakuttan,[1] David Amaro,[2] Shouvanik Chakrabarti,[1] Soumik Ghosh,[3] Zichang He,[1] Yuwei Jin,[1] Fatih Kaleoglu,[1] Steven Kordonowy,[1] Rohan Kumar,[1] Michael A. Perlin,[1] Akshay Seshadri,[1] Matthew Steinberg,[1] Joseph Sullivan,[1] Jacob Watkins,[1] Henry Yuen,[4] and Ruslan Shaydulin[1,‡]

[1]*Global Technology Applied Research, JPMorganChase, New York, NY 10017, USA*
[2]*Quantinuum, Partnership House, Carlisle Place, London SW1P 1BX, UK*
[3]*Department of Computer Science, University of Chicago, Chicago, IL 60637, USA*
[4]*Columbia University, New York, NY 10027, USA*

(Dated: February 5, 2026)

Quantum mechanics provides cryptographic primitives whose security is grounded in hardness assumptions independent of those underlying classical cryptography. However, existing proposals require low-noise quantum communication and long-lived quantum memory, capabilities which remain challenging to realize in practice. In this work, we introduce a quantum digital signature scheme that operates with only classical communication, using the classical shadows of states produced by random circuits as public keys. We provide theoretical and numerical evidence supporting the conjectured hardness of learning the private key (the circuit) from the public key (the shadow). A key technical ingredient enabling our scheme is an improved state-certification primitive that achieves higher noise tolerance and lower sample complexity than prior methods. We realize this certification by designing a high-rate error-detecting code tailored to our random-circuit ensemble and experimentally generating shadows for 32-qubit states using circuits with ≥ 80 logical (≥ 582 physical) two-qubit gates, attaining 0.90 ± 0.01 fidelity. With increased number of measurement samples, our hardware-demonstrated primitives realize a proof-of-principle quantum digital signature, demonstrating the near-term feasibility of our scheme.

https://arxiv.org/pdf/2602.04859

Digital signatures

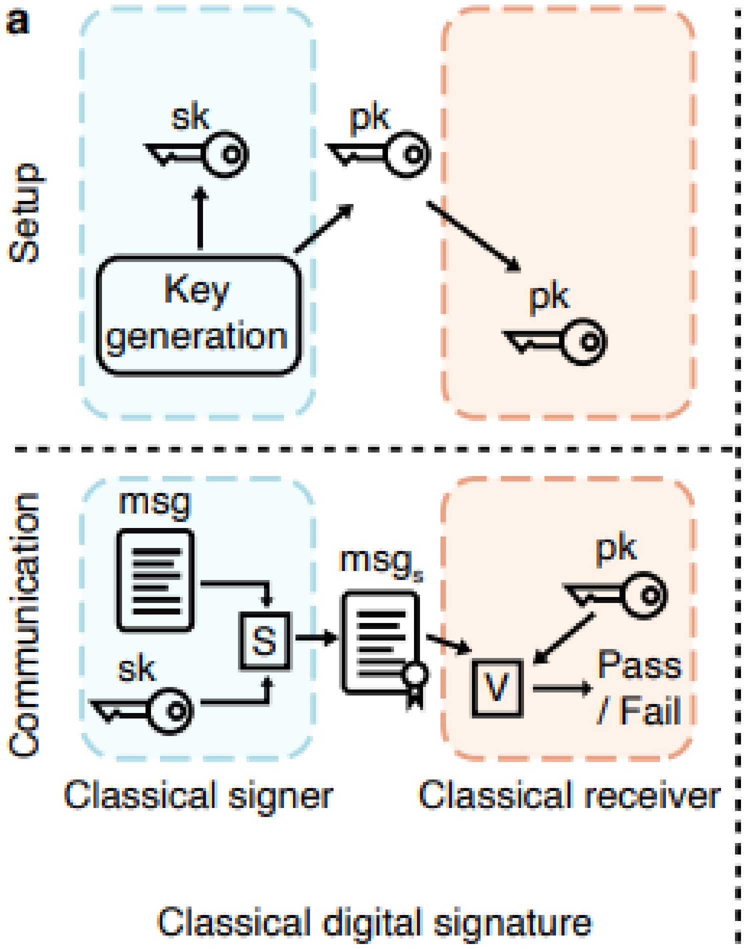

Classically, based on 1-way functions
E.g. N = p*q, pk = N, sk = (p,q)

For pk = N = 15,
can find p=3, q=5

For pk = N = 1234567890123456
45678901234567890123456789
Challenge : find sk = (p,q)

Quantum mechanics offers independent routes to secure cryptographic schemes. While it has been long known that quantum *communication* can enable unconditionally-secure key distribution [7], recently there have been significant advances in using quantum *computation* to propose cryptographic primitives that either rest on alternative assumptions than classical schemes or have no classical counterpart. Notably, some of these primitives are compatible with noisy devices [8, 9], unlike many traditional algorithms that typically require fault tolerance, opening opportunities to strengthen

In this work, we present a digital signatures scheme that does not require OWFs or quantum communication. Digital signatures enable a recipient to verify the message's true author using a signature generated by the sender. Our signature scheme modifies earlier proposals from [13] to use classical shadows, instead of quantum states, as public keys, making the scheme friendly to near-term experiments. The security of our scheme is based on the conjectured computational hardness of learning quantum states from their classical shadows ("computational no-learning" conjecture). This conjecture is independent of the existence of one-way functions as illustrated in Fig. 1.

More precisely, this protocol implements the quantum cryptographic primitive known as "one-way puzzles", distinct from OWFs. One-way puzzles have been used [15] to construct secure quantum multiparty computation [16] and non-interactive quantum bit commitments [16–18], providing OWF-free building blocks for quantum cryptography.

At present, no algorithm exists to efficiently learn a quantum circuit given its classical shadows. To provide evidence for the "computational no-learning" conjecture, we instead consider algorithms requiring a more powerful access model. First, we show that existing algorithms for learning shallow

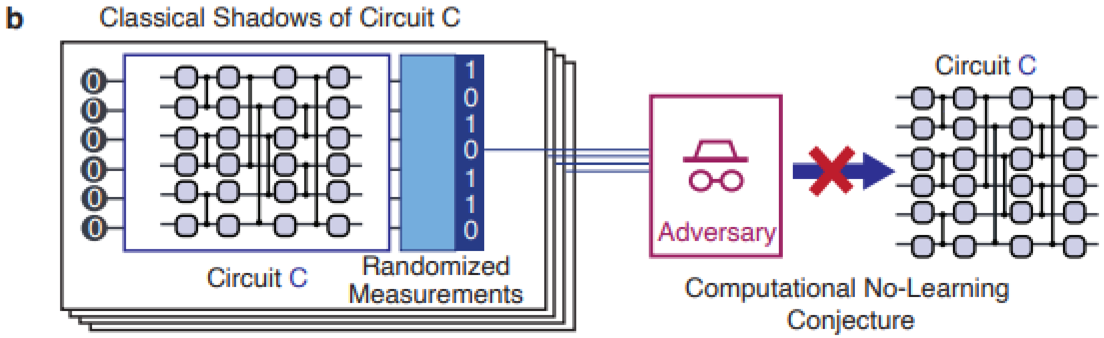
b
Classical Shadows of Circuit C
0
0
0
0
0
0
Circuit C
Randomized Measurements
1
0
1
0
1
1
0
Adversary
Computational No-Learning Conjecture
Circuit C

circuits fail in our setting [13, 19–22]. Second, we derive a learning algorithm specialized to the class of all-to-all circuits considered in this work, extending the approach from Ref. [22], and show security against it. In addition to support for CNL, we also provide evidence that "spoofing" the protocol is exponentially harder than verifying the signature. Specifically, we prove the separation for two general classes of potential attacks: low-degree polynomial learning algorithms [23] and variational circuit learning [24–26].

We demonstrate the near-term feasibility of the proposed scheme by generating shadows on a trapped-ion processor, successfully certifying the fidelity of the prepared quantum state. To enable the experiment, we introduce improved classical shadow schemes with better tolerance to experimental imperfections and a variant of Iceberg error-detecting code [27–29] tailored to our random circuits. Using multi-block Iceberg encoding, we achieve beyond-break-even performance on circuits with 40 logical qubits and 680 logical gates (560 single-qubit and 120 two-qubit), where the encoded fidelity exceeds the fidelity without error detection, a result of independent interest. For hard-to-learn 32-qubit states, we achieve a fidelity of 0.90 ± 0.01 with 17,316 samples. We show that increasing the number of samples by 19.5 times is sufficient for a proof-of-principle demonstration of a signature scheme, demonstrating that our proposal can be realized in the near term.

Such a one-way state generator may also be leveraged to devise a quantum analog of a digital signature [13]. In the quantum digital signature protocol proposed by Ref. [13], to sign a single bit, the sender begins with classical descriptions of two quantum circuits C_0 and C_1, which serve as the secret key. For the public key, the sender distributes copies of the corresponding quantum states, $|C_0\rangle = C_0\,|0\rangle$ and $|C_1\rangle = C_1\,|0\rangle$, for example via a trusted cloud repository. To transmit the bit b, the sender reveals the signed message (b, C_b) to the receiver. The receiver then downloads copies of the quantum state $|C_b\rangle$ from the cloud and verifies that the state matches the circuit C_b—for instance, by checking that $C_b^\dagger\,|C_b\rangle$ yields the all-zero state. This protocol is illustrated in the left half of Fig. 2**b**.

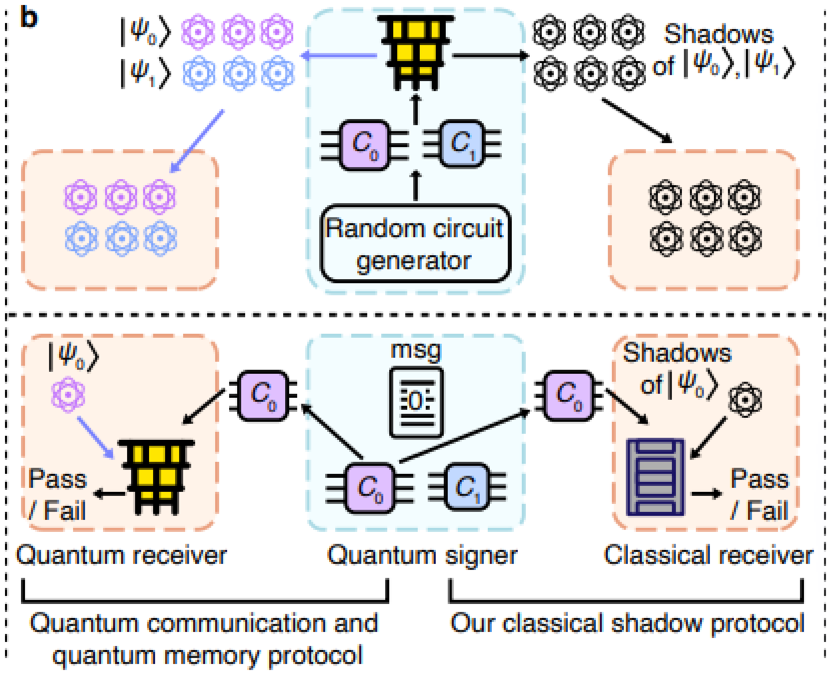

classical.

We propose a protocol where the public keys are not quantum states but classical shadows. In particular, for secret keys (C_0, C_1), the sender publishes polynomial-length classical shadows $\mathcal{S}(|C_0\rangle), \mathcal{S}(|C_1\rangle)$ as public keys. Here, $\mathcal{S}$ denotes the quantum procedure that, across a polynomial number of samples, selects random local measurement bases on the qubits, measures the state accordingly, and for each sample records the chosen bases together with the corresponding measurement outcomes; the collection of these recorded basis choices and outcomes is the classical shadow. To communicate a bit b, the signer releases (b, C_b) as before. The receiver verifies the message by certifying that $\mathcal{S}(C|0\rangle) \approx \mathcal{S}(C_b|0\rangle)$

falsify CNL and compromise the security of our protocol.

The arguments above provide theoretical and empirical evidence against the efficient learning of private circuits from public shadows. Since, for shadow-based signatures, the verification step – wherein a verifier validates a signature – is also expected to take exponential time, it would be desirable to choose circuit families such that verification can be done in a feasible wall-time while spoofing remains prohibitively expensive. We expect such a "gap" between verification time and learning or spoofing time to arise since circuit learning involves a search from the space of all possible circuits in the ensemble, which is as large as $\exp(O(nd))$ where d is the depth of the ensemble. In this work, we call a protocol secure if the time of verification T_V and the time of learning the circuits T_L satisfy $T_L/T_V \geq \Omega(\exp(n))$, where n, the circuit size, serves as the security parameter. In other words, there is an *exponential* gap between the cost to cheat and the cost to verify. In SI Sec. III, we formalize this conjecture and provide evidence towards it.