

Experimental Single-Setting Quantum State Tomography

Roman Stricker^{1,*}, Michael Meth,¹ Lukas Postler,¹ Claire Edmunds¹, Chris Ferrie,² Rainer Blatt^{1,3,4}, Philipp Schindler,¹ Thomas Monz^{1,4}, Richard Kueng,⁵ and Martin Ringbauer¹


¹*Institut für Experimentalphysik, Universität Innsbruck, Innsbruck 6020, Austria*

²*Centre for Quantum Software and Information, University of Technology Sydney, Ultimo, NSW 2007, Australia*

³*Institut für Quantenoptik und Quanteninformatik, Österreichische Akademie der Wissenschaften, Innsbruck 6020, Austria*

⁴*Alpine Quantum Technologies GmbH, Innsbruck 6020, Austria*

⁵*Institute for Integrated Circuits, Johannes Kepler University Linz, Linz 4040, Austria*

 (Received 30 May 2022; revised 27 July 2022; accepted 22 August 2022; published 21 October 2022)

Quantum computers solve ever more complex tasks using steadily growing system sizes. Characterizing these quantum systems is vital, yet becoming increasingly challenging. The gold-standard method for this task is *quantum state tomography* (QST), capable of fully reconstructing a quantum state without prior knowledge. The measurement and classical computing costs, however, increase exponentially with the number of constituents (e.g., qubits)—a daunting bottleneck given the scale of existing and near-term quantum devices. Here, we demonstrate a scalable and practical QST approach that only uses a single measurement setting, namely symmetric informationally complete (SIC) positive operator-valued measures (POVMs). We implement these nonorthogonal measurements on an ion trap quantum processor by utilizing additional energy levels within each ion—without requiring ancillary ions to assist in measurements. More precisely, we locally map the SIC POVM to orthogonal states embedded in a higher-dimensional system, which we read out using repeated in-sequence detections, thereby providing full tomographic information in every shot. Combining this *SIC tomography* with the recently developed *randomized measurement toolbox* (“classical shadows”) proves to be a powerful combination. SIC tomography alleviates the need for choosing measurement settings at random (“derandomization”), while classical shadows enable the estimation of arbitrary polynomial functions of the density matrix orders of magnitudes faster than standard methods. The latter enables in-depth entanglement characterization, which we experimentally showcase on a five-qubit absolutely maximally entangled state. Moreover, the fact that the full tomography information is available in every shot enables online QST in real time (i.e., while the experiment is running). We demonstrate this on an eight-qubit entangled state (which has $2^8 \cdot 2^8 - 1 = 65\,535$ degrees of freedom), as well as for fast state identification. All in all, these features single out SIC-based classical shadow estimation as a highly scalable and convenient tool for quantum state characterization.

DOI: [10.1103/PRXQuantum.3.040310](https://doi.org/10.1103/PRXQuantum.3.040310)

Quantum systems are prepared in laboratories and in engineered devices such that their state delicately encodes quantum information essential for achieving goals in both science and technology. Any small adjustments, changes in the environment, or active control all change the state. Yet, an accurate mathematical description of the state

is a necessary component for most higher-level tasks. A crucial requirement for ensuring the performance of quantum devices is thus having methods for accurately determining the quantum states that have been prepared. The gold-standard approach for this fundamental task is quantum state tomography (QST) or simply tomography [1]. QST enables the full reconstruction of the system’s quantum state from an exponential number of measurements. Often, however, we are not even interested in the full quantum state, but rather certain features, like entanglement across a particular bipartition. However, it is not clear *a priori* how to access such nonlinear properties without resorting to full QST.

Formally, QST methods use an *informationally complete* set of measurements to reconstruct the complete

*roman.stricker@uibk.ac.at

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article’s title, journal citation, and DOI.

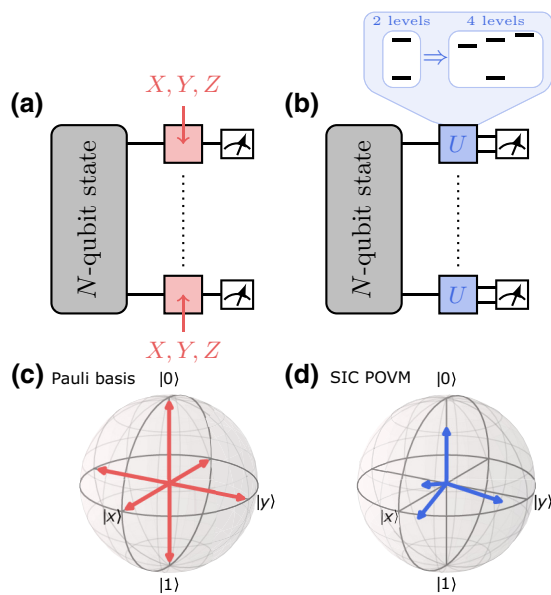


FIG. 1. Schematic illustration of Pauli and SIC tomography. *Pauli tomography* (a) uses three basis measurements per qubit to obtain tomographic information about an unknown N -qubit system; see (c). Each basis measurement is contingent on one of three possible unitary rotations—red boxes in (a). This produces a total of 3^N different measurement settings that need to be accessed. *SIC tomography* (b), on the other hand, uses the same measurement setting for each qubit; see (d). This nonorthogonal measurement is achieved by isometrically embedding each two-level system (qubit) into a larger four-level system (ququart)—blue boxes in (b)—and subsequently measuring this larger system. The experimental realization of this embedding within each ion is shown in Fig. 2.

description of the quantum state. The optimal measurement for collecting the necessary tomography data has long been known to be the so-called symmetric informationally complete positive operator-valued measures (SIC POVMs) [2,3]. SIC POVMs are constructed from the minimal number of d^2 measurements for a d -dimensional system, which are arranged in a way that maximizes the pairwise distance in Hilbert space. SIC POVMs are known to exist for several low-dimensional systems [4,5], and for qubits, take the form of four nonorthogonal vectors arranged as a tetrahedron in the Bloch sphere; see Fig. 1(d). While SIC POVMs uniquely offer access to complete tomographic information in every single experimental run (shot), implementing these measurements in practice is very challenging, requiring purpose-built setups [6,7], sequential measurement [8,9], or ancilla-assisted schemes [10,11]. Hence, tomography remains almost exclusively performed using the simpler, but overcomplete Pauli basis, requiring 3^N orthogonal measurement *settings*, each with 2^N outcomes for an N -qubit system. The resulting overhead effectively limits full tomography to system sizes of only a few qubits.

From a conceptual point of view, the qubit SIC POVM is favorable, as a single experimental shot already contains complete tomographic information. This distinct advantage has also been recognized in recent theoretical work on adaptive tomography for linear cost functions [11], as well as neural network quantum state tomography [12,13]. Experimentally, it is also much cheaper to repeat the same measurement setting many times than to switch settings an exponential number of times as the system size grows with Pauli tomography. So, this feature can have a significant impact in practice. Moreover, since full tomographic information is contained in every shot, the experimenter is free to stop the tomography at any point, e.g., when certain quantities of interest have converged. In contrast, other QST approaches would require at least one shot for each measurement setting to collect sufficient information in the first place. This discrepancy is particularly relevant when we are not interested in the full density matrix, but only in certain (nonlinear) properties, which often require far fewer shots than the 3^N minimum in Pauli tomography [14,15]. Finally, for randomized measurement schemes [16], where ideally a different measurement setting is required in each shot, the SIC approach obviates this requirement completely (“derandomization”), making these schemes even more practical. Hence, in most situations, SIC tomography has the potential to substantially outperform standard methods for tomography or for the direct estimation of state properties.

Here, we describe our realization of SIC POVMs in a trapped-ion quantum processor and their use for characterizing unknown quantum states. We put an emphasis on demonstrating the speed and robustness obtained from reducing the number of measurement settings in conjunction with new data processing techniques that come with rigorous accuracy guarantees. With our approach, we are able to comfortably reconstruct the full eight-qubit quantum state encoded in the electronic energy levels of calcium ions *in real time* using a standard laboratory computer. Moreover, we demonstrate the simultaneous real-time estimation of Rényi entropies across all bipartitions using a sampling-free classical shadow method [14]. This enables full entanglement characterization of arbitrary (but close to pure) quantum states with orders of magnitude fewer experimental shots than standard QST methods.

I. SIC TOMOGRAPHY

QST aims at reconstructing an unknown quantum state ρ from an informationally complete set of measurements, which spans the entire Hilbert space of the quantum system. The minimal number of measurement outcomes to reconstruct an arbitrary d -dimensional quantum state then is d^2 . An experimenter performs these measurements on many copies of ρ referred to as experimental shots, and attempts to reconstruct ρ from the observed measurement

counts. The standard approach to QST of N -qubit systems combines tomographic measurements of each individual qubit. The overcomplete Pauli basis is a particularly prominent choice; see Fig. 1(c). Three distinct measurement settings are required to evenly cover the Bloch sphere and obtain tomographic single-qubit information. Extending this to N -qubit systems produces 3^N distinct measurement settings that need to be explored; see Fig. 1(a).

In contrast to (single-qubit) Pauli basis measurements, (single-qubit) SIC POVMs provide access to a complete set of tomographic data from a single experimental shot. No change in measurement settings is required. This desirable feature extends to N -qubit measurements: a single measurement setting per qubit suffices to obtain tomographically complete data; see Fig. 1(b). SIC tomography utilizes (tensor products of) the single-qubit SIC POVM depicted in Fig. 1(d) (or a local rotation thereof). Following Naimark's dilation theorem [18], every POVM can be realized as a projective measurement on a higher-dimensional Hilbert space. Using this result, together with a qudit quantum processor [17], we realize the qubit SIC POVM as a projective measurement on a four-level system (ququart) employing two more states within each calcium ion. For this purpose, we map each qubit locally to a ququart using the unitary

$$\hat{M} = \begin{pmatrix} 1/\sqrt{2} & 0 & 0 & 1/\sqrt{2} \\ 1/\sqrt{6} & 1/\sqrt{3} & 1/\sqrt{3} & -1/\sqrt{6} \\ 1/\sqrt{6} & (-1)^{2/3}/\sqrt{3} & (-1)^{4/3}/\sqrt{3} & -1/\sqrt{6} \\ 1/\sqrt{6} & (-1)^{4/3}/\sqrt{3} & (-1)^{2/3}/\sqrt{3} & -1/\sqrt{6} \end{pmatrix}. \quad (1)$$

Here, the four two-dimensional vectors contained in the first two columns represent the measurement vectors of the qubit basis given in Fig. 1(d). The optimized gate sequence for locally mapping the measurement states from qubit to ququart is shown in Fig. 2(b). It consists of five single-qubit rotations, sequentially applied to each qubit, which we optimize such that local phase shifts are absorbed into the rotation angles of the single-qubit operations; see Appendix A 1. Therefore, our single-setting SIC tomography implementation remains the very same, independent of the qubit number. Our approach is particularly well suited to other quantum architectures, since many of today's information carriers manifest multilevel systems. Higher-dimensional systems have been demonstrated, frequently featuring in Rydberg ions [19], atomic and molecular systems [20,21], photonic systems [22], solid states [23], and superconducting platforms [24].

II. RECONSTRUCTION METHODS AND CLASSICAL SHADOWS

Even with a complete set of measurements, reconstructing ρ is computationally demanding, especially if one

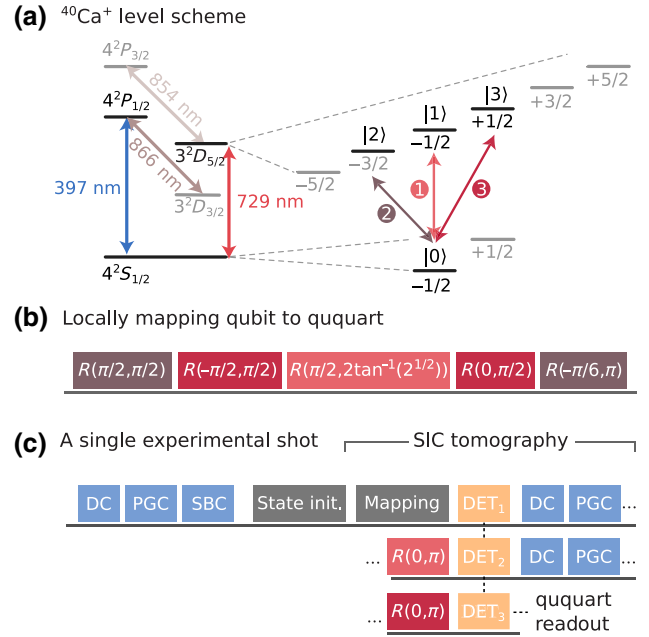


FIG. 2. Experimental implementation of SIC tomography from Figs. 1(b) and 1(d). (a) Level scheme of a $^{40}\text{Ca}^+$ ion representing a qubit or ququart with important transitions marked: (blue) dipole transition for cooling and detection, (red) metastable quadrupole transition for encoding qubits and ququarts within the Zeeman submanifold, and (brown) additional transitions for repumping. (b) Gate sequence for locally mapping the SIC POVM from Fig. 1(d) from qubit level to four orthogonal basis states of a ququart denoted in (a). This enables full readout of the SIC POVM in a single experimental shot by means of a four-outcome projective measurement. (c) Experimental realization of SIC tomography composed of cooling (DC, doppler cooling; PGC, polarization gradient cooling; SBC, side-band cooling), preparation of the state to be analyzed, mapping from qubit to ququart according to the SIC POVM, and finally the four-outcome projective measurement. For the latter, three sequential fluorescence detections (DETs) are required [17]; see Appendix A 1 for details.

insists on enforcing physicality constraints. Two standard QST methods in the field are linear inversion (LI) and maximum likelihood estimation (MLE) [25], which can be readily applied to both Pauli and SIC data. LI provides an analytical approach to estimating ρ from a complete set of projectors Π_j that span the entire Hilbert space. Access to (approximations of) the associated probabilities $\hat{p}_j \approx \text{tr}(\Pi_j \rho)$ allows us to reconstruct the underlying state:

$$|\hat{\rho}\rangle\rangle = \sum_j \hat{p}_j \cdot S_p^{-1} |\Pi_j\rangle\rangle \quad (2)$$

with $|\Pi_j\rangle\rangle$ the vectorized projector, obtained by stacking the columns of Π_j . Furthermore, S_p^{-1} denotes the Moore-Penrose pseudoinverse of the measurement superoperator $S_p = \sum_j |\Pi_j\rangle\rangle \langle\langle \Pi_j|$ [26], and $\hat{p}_j = n_j/N_j$ is the observed

frequency of outcome j after averaging over N_j experimental shots. As an unconstrained method, the LI version bears the risk of producing unphysical estimators for ρ featuring negative eigenvalues. This is particularly pronounced when few experimental shots are used and is very problematic for estimating nonlinear observables. Physical constraints are thus typically introduced through MLE, which, following Ref. [26], can be approximated by a convex optimization problem

$$\begin{aligned} & \text{minimize} && \|\mathcal{W}(S|\hat{\rho}) - |f\rangle\|_2 \\ & \text{subject to} && \hat{\rho} \geq 0, \quad \text{tr}(\hat{\rho}) = 1. \end{aligned} \quad (3)$$

Here, $S = \sum_j |j\rangle\langle\Pi_j|$ denotes a change of basis operator, $|f\rangle = \sum_j n_j/N_j |j\rangle$ a column vector of the observed frequencies, and \mathcal{W} a diagonal matrix of statistical weights. Optimization is performed under the constraints that the estimator for ρ is positive semidefinite ($\rho \geq 0$) with unit trace [$\text{tr}(\rho) = 1$], i.e., it must be a valid quantum state. The convex optimization in Eq. (3) is computationally more efficient than full MLE and recovers the latter in the limit of large sample sizes. Nonetheless, the computational complexity remains intractable for anything but very small systems. LI is much more efficient by comparison. Nonetheless, inverting the superoperator S_p also becomes more challenging as the system size increases. Viewed as a matrix, every tomographically complete N -qubit superoperator S_p must have (at least) 4^N rows and (at least) 4^N columns. Performing the inversion row by row can offer some relief in terms of memory load, but the exponential number of multiplications remains challenging. Finally, physical constraints [$\rho \geq 0$ and $\text{tr}(\rho) = 1$] can be incorporated into LI by truncating negative eigenvalues to obtain the closest quantum state under the Frobenius norm [27–29], referred to as *projected least squares* (PLS). It should be noted that more principled, yet ever more computationally challenging, approaches exist [30,31].

So far, we have considered full QST, i.e., experimentally extracting a complete description of ρ , which is traditionally required for predicting certain properties of complex quantum systems, especially nonlinear functions, most prominently purity or entanglement. In large-scale systems, however, predicting such properties becomes very costly independent of the data acquisition (SIC, Pauli) and reconstruction (LI, MLE) method, both in regards to the number of required shots and in regards to the computational power required to analyze the data.

A promising alternative comes in the form of classical shadows [14,15] as a general-purpose method to construct classical descriptions of quantum states using very few experimental shots. Consequently, the classical shadow framework allows for the prediction of L different functions of the state with high accuracy, using order $\log(L)$ experimental shots. Importantly, the number of shots is

independent of the system size and saturates information-theoretic lower bounds. Moreover, target properties can be selected after the measurements are completed. A big drawback of existing classical shadow methods, however, is that they require a different measurement to be sampled randomly for each shot [16], which is demanding and slows down data acquisition. We show in the following that SIC POVMs naturally alleviate this sampling requirement (“derandomization”). SIC POVMs are thus an ideal choice for unlocking the full potential of the classical shadow framework. This has, in parts, been already pointed out in Ref. [11], which explores adaptive SIC tomography for linear cost functions inspired by variational quantum eigensolver (VQE). Instead, we are here interested in a general framework for efficiently predicting general linear and nonlinear properties of the quantum state.

Formally, classical shadows provide an alternative approach for a linear-inversion estimator deduced from SIC measurements on an N -qubit state ρ . Each experimental shot m , containing complete tomographic information, can be assigned to a size- N string $\hat{i}_{m,1}, \dots, \hat{i}_{m,N} \in \{1, 2, 3, 4\}^{\times N}$, where each quartic value keeps track of the SIC POVM outcome observed. For each shot m , an N -qubit estimator for the density matrix $\hat{\sigma}_m = \bigotimes_{n=1}^N (3|\psi_{i_{m,n}}\rangle\langle\psi_{i_{m,n}}| - \mathbb{I})$ is obtained, referred to as a *classical shadow*. A total of M such estimators can be experimentally inferred and accumulated to approximate ρ as

$$\hat{\rho} = \frac{1}{M} \sum_{m=1}^M \bigotimes_{n=1}^N (3|\psi_{i_{m,n}}\rangle\langle\psi_{i_{m,n}}| - \mathbb{I}) \xrightarrow{M \rightarrow \infty} \rho. \quad (4)$$

It is worthwhile to emphasize that each term in this average is highly unphysical, because each contains many negative eigenvalues that also become exponentially large. The crucial insight from Ref. [14] was that these unphysicalities quickly average out. In fact, empirical averages of such unbiased single-shot estimators converge much quicker than physical density matrix estimators that necessarily contain a bias. In the following, we present rigorous convergence guarantees that underscore this claim and refer to Appendix A 8 for additional context and exposition. Another crucial observation is that, compared to standard linear inversion in Eq. (2), the processing of classical shadows is performed in the dimension of the quantum state $2^N \cdot 2^N$, which is the minimal possible dimension for full tomography; see Appendix A 9 b. Moreover, predicting linear observables using classical shadows is even more efficient as it suffices to reconstruct a subset of ρ solely where operators act on. In Appendix A 10, we show how we can formalize these considerations to derive a measurement budget for estimating linear observables. Suppose that we are interested in estimating a total of $L \gg 1$ subsystem observables $\text{tr}(O_l \rho)$, where each O_l only

acts nontrivially on (at most) $K \leq N$ qubits. Then,

$$M \geq \frac{8}{3} 6^K \log(2L/\delta)/\epsilon^2 \quad (5)$$

measurements suffice to jointly ϵ -approximate all observables with probability (at least) $1 - \delta$. We emphasize that this is a novel, rigorous *a priori* bound based on minimal assumptions. In practice, convergence sets in (much) earlier. A full derivation and additional context is provided in the Appendix. For now, we merely point out that improvements of order 2^K are possible for the exponential scaling in case the observables in question have small Hilbert-Schmidt norm, as is the case for fidelities. Apart from linear observables, classical shadows also promise to allow for efficient estimation of nonlinear functions; see Appendix A 11. Whereas the full scope of nonlinear functions is covered in the Appendix, here we focus on a quadratic estimator in form of the (subsystem) purity

$$\hat{p}_{(M)} = \binom{M}{2}^{-1} \sum_{m < m'} \text{tr}(\hat{\sigma}_m \hat{\sigma}_{m'}), \quad (6)$$

as purity generally obeys hard convergence properties and also provides a means to measuring entanglement via Rényi entropies (see below). The latter are given by the negative logarithm of subsystem purities, corresponding to certain bipartitions of the state. Similar to before we can derive a measurement budget, where

$$M \geq 6L3^K/(\epsilon^2\delta) \quad (7)$$

measurements allow for ϵ -approximating L subsystem purities of size (at most) K with probability (at least) $1 - \delta$. We emphasize that this is again a novel, rigorous *a priori* bound based on minimal assumptions; actual convergence sets in much quicker. However, this bound still marks an improvement over the best available results for purity estimation with randomized single-qubit measurements [32,33]. The improvement follows from exploiting the geometric structure of SIC POVMs.

At this point it is also worthwhile to note that the proof techniques behind Eqs. (5) and (7) are novel. The original proof techniques from Ref. [14] do not extend to SIC POVM measurements, because they rely on a powerful structural property—a so-called 3-design [34–37]—that Pauli measurements have, but SIC POVMs do not. We refer the reader to Appendix A 11 for additional context and complete proofs. These results demonstrate that classical shadows in combination with SIC measurements offer a powerful tool set for measuring entanglement in a scalable fashion [32,33]. Whereas our experimental studies primarily focus on quadratic estimators of Rényi entropy [Eq. (9)], classical shadows can be extended to higher-order estimators following the same principles: (i) rewrite a degree- d polynomial as $\text{tr}(O\rho^{\otimes d})$, (ii) replace each ρ with

an independent classical shadow $\hat{\sigma}_m$, and (iii) average over all different subselections of distinct classical shadows. We refer the reader to Refs. [14,32,33] for details. Finally, we remark that classical shadow estimators from Pauli basis measurements can in principle be obtained by randomly sampling over the measurement settings from shot to shot. Although experimentally feasible, this is highly impractical. Remarkably, Pauli basis measurements also appear to lead to slower convergence than SIC measurements.

III. EXPERIMENTAL SETUP AND SIC IMPLEMENTATION

Experimental results in this work are obtained on a trapped-ion quantum processor based on a linear string of $^{40}\text{Ca}^+$ ions, each encoding a single qubit in the (meta)stable electronic states $\{S_{1/2}(m = -1/2) = |0\rangle, D_{5/2}(m = -1/2) = |1\rangle\}$ [38]. A universal set of quantum gate operations is realized upon coherent laser-ion interaction and comprises arbitrary local single-qubit rotations together with two-qubit entangling operations, enabling all-to-all connectivity. A binary qubit measurement is implemented by scattering on the dipole transition, where fluorescence is only observed if the ion is in the $|0\rangle$ state, thereby separating the computational basis states $\{S_{1/2}(m = -1/2) = |0\rangle, D_{5/2}(m = -1/2) = |1\rangle\}$; see Fig. 2(a). Equivalent control over the entire S - and D -state Zeeman manifolds allows for encoding a higher-dimensional quantum decimal digit (qudit) with up to eight levels in each ion, combined with full fluorescence readout of the whole qudit space [17]; see Appendix A 1.

The present work builds on this capability by utilizing up to four levels per ion to implement SIC POVMs. To this extent, two additional levels $D_{5/2}(m = -3/2) = |2\rangle$ and $D_{5/2}(m = +1/2) = |3\rangle$ are taken into account; see Fig. 2(a). Upon applying the mapping sequence depicted in Fig. 2(b), each qubit is locally extended to a ququart where each basis state encodes one SIC vector. A four-outcome projective measurement is implemented by three sequential fluorescence detections, where before the second detection the population between states $|0\rangle$ and $|1\rangle$ is flipped and likewise before the final detection the states $|0\rangle$ and $|2\rangle$ are flipped. This enables us to evaluate the full ququart state probabilities from three binary outcomes. The entire experimental sequence comprising cooling, state preparation, mapping the SIC POVM to the ququart, and four-outcome readout is shown in Fig. 2(c), which we refer to as a single experimental shot. We remark that this SIC tomography procedure works independently on each qubit and that such a single experimental shot delivers the complete tomographic information of the N -qubit system.

IV. RECONSTRUCTION TIME

While SIC POVMs can significantly speed up data acquisition, the classical resources needed for

reconstructing and storing the quantum state ρ is typically an additional bottleneck in QST [see Eqs. (2) and (3)]. In the following, we compare the computational time for reconstructing ρ following various tomography approaches. For the moment, we solely focus on the classical reconstruction time, which is dominated by the size of the involved matrices, and discuss the convergence properties of the various methods later and in Figs. 7, 12, and 14 in the Appendix. For a system of N qubits, we consider tomography data composed of $M = 100 \cdot 3^N$ shots. This corresponds to 100 shots for each measurement setting used in Pauli tomography, which, on the trapped-ion platform, has proven to be a good trade-off accounting for statistics, systematic drifts, and measurement time.

Figure 3 illustrates the classical reconstruction time versus the number of qubits N for experimental data used throughout this manuscript. Whereas absolute time reflects a laboratory desktop computer, relative scaling between methods remains generally valid. Note that the plot is double logarithmic in the number of shots $M = 100 \cdot 3^N$. While MLE methods always obey physical constraints, solving the convex optimization problems is costly and only feasible for small system sizes. We find MLE with SIC measurements to be more efficient, due to handling matrices of maximum size $4^N \cdot 4^N$, in contrast to $6^N \cdot 4^N$ for the overcomplete Pauli basis. However, MLE quickly becomes infeasible as the number of qubits increases. SIC LI suffers an initial offset to Pauli LI due to computing ququart instead of qubit state probabilities, which, for the MLE approaches, is masked in the overhead of convex optimization from Eq. (3). As the number of qubits increases, SIC makes up for this, as the computations are performed in a smaller dimension. Although computationally much cheaper than MLE, even LI becomes increasingly costly due to the memory requirements of processing the inverse superoperator S_p from Eq. (2). Already for six qubits this requires 268.4 and 3100 MB for SIC and Pauli measurements, respectively. Scaled up further, this will rapidly exceed the memory of today's computers. Alternatively, inversion of S_p could be done row by row to reduce memory load, but this would be more time consuming than precalculating the inverse S_p^{-1} as we have done here. While linear observables under LI are proven to quickly converge (see Appendix A 10), nonlinear functions suffer from non-physical properties in the form of negative eigenvalues; see Fig. 4(b). Furthermore, PLS adds negligible computational overhead over LI and is thus neglected in this comparison. PLS does, however, affect the convergence and accuracy of the estimators, as shown and discussed further below.

Finally, we find the best scaling for the SIC-based classical shadows from Eq. (4), where data are processed at the dimension of the density matrix, $2^N \cdot 2^N$, avoiding matrix inversion or optimization altogether. Instead, individual experimental shots are accumulated, offering convenient




FIG. 3. Classical run time comparison for tomographic reconstruction methods. Comparison of the computational run time for state reconstruction using LI [Eq. (2)] or MLE [Eq. (3)] for both Pauli [Fig. 1(c)] and SIC tomography [Fig. 1(d)], as well as SIC-based classical shadows [Eq. (4)] as a function of the system size. For each qubit number N , the reconstruction considers $M = 100 \cdot 3^N$ experimental shots (i.e., 100 shots per Pauli basis; see the text). The analysis is conducted on a standard desktop computer and plotted double logarithmically in the number of experimental shots. We find that both MLE methods require the highest computational resources, with SIC MLE significantly faster as the processed dimension is lower with $4^N \cdot 4^N$ compared to $6^N \cdot 4^N$ for Pauli tomography. LI with SIC measurements shows an initial time offset to Pauli tomography arising from handling ququart ($\text{dim} = 4^N$) instead of qubit data ($\text{dim} = 2^N$), but grows much more slowly with system size. Among the LI methods, we find the classical shadow reconstruction to be the fastest for an increasing number of qubits, as it accumulates each individual shot in dimension $2^N \cdot 2^N$ [see Eq. (4)], avoiding the costly matrix inversion in dimension $4^N \cdot 4^N$ or $6^N \cdot 4^N$ as required for SIC and Pauli tomography, respectively.

updates of ρ for every new set of data. As a consequence of this individual accounting for every shot, the computational complexity of this method grows linearly with the total number of shots. While this linear overhead leads to slightly worse performance for very small systems, it is more than compensated by the improved exponential scaling with qubit number ($2^N \cdot 2^N$ versus at least $4^N \cdot 4^N$) for large systems. Hence, the SIC-based classical shadows clearly outperform all other methods for six or more qubits. Note that, for large-scale systems, the gap between classical shadows and standard LI becomes even larger as row by row LI becomes requisite.

V. ESTIMATING PROPERTIES OF THE STATE

Here we shift our attention towards convergence of the different tomography estimators. In particular, we showcase the classical shadows' unique feature of efficiently predicting nonlinear properties of even large-scale quantum systems. To this end, we experimentally perform




FIG. 4. SIC-based classical shadow tomography of a five-qubit absolutely maximally entangled (AME) state. We compare the convergence of LI [Eq. (2)] and classical shadow tomography using SIC measurements [Eqs. (4)–(6)]. We additionally use PLS following Refs. [27–29], explained in the text. (a) In terms of convergence, we find both SIC LI from Eq. (2) and SIC shadows LI from Eq. (4) to overlap, which is expected due to their similarity. We remark that in terms of reconstruction SIC shadows LI is computationally more efficient; see Fig. 3. PLS incorporates physical constraints [$\text{tr}(\rho) = 1$ and $\rho > 0$], but underestimates the fidelity for a small number of shots and converges very slowly, requiring at least an order of magnitude more shots (indeed slower than MLE, additionally confirmed by experiments in Appendix A2 and numerical simulations in Appendix A6). Here, connected data points represent an individual experimental run, while shaded regions represent 1 standard deviation around the mean value of multiple batches of experimental data at identical shot numbers. (b) When estimating purity, LI shows highly unphysical behavior under insufficient statistics and does not converge in the plot. Physical constraints can be corrected by PLS, although at the cost of much slower convergence. On the other hand, classical shadow purity estimators [Eq. (6)] display very quick convergence. The dash-dot line indicates the ideal value. (c) Estimated Rényi entropies [Eq. (9)] from subset purities across all bipartitions converge very quickly for all methods. The lack of difference between the methods is likely due to the small system sizes, and for increasing bipartition sizes, we expect similar behavior as in (b). Note that bipartitions are denoted by tuples (2, 3) qubits and (1, 4) qubits, referring to the number of qubits in each part.

tomography on a five-qubit AME state [39]

$$\begin{aligned}
 2\sqrt{2} |\Omega_{5,2}\rangle = & |00000\rangle + |00011\rangle + |01100\rangle \\
 & - |01111\rangle + |11010\rangle + |11001\rangle \\
 & + |10110\rangle - |10101\rangle. \quad (8)
 \end{aligned}$$

AME states are the most entangled states in the sense that they are maximally entangled in all bipartitions [40]. This makes them interesting for applications in quantum error correction [41], quantum teleportation, quantum secret sharing, and superdense coding [42]. Alas, their general existence remains unknown for all but the smallest systems.

We characterize the five-qubit AME state from Eq. (8) using SIC tomography data. The results in Fig. 4(a) indicate that both LI methods converge very quickly in fidelity, as is likewise expected for all linear observables; see Appendix A10. In terms of convergence SIC LI from Eq. (2) and SIC shadows LI from Eq. (4) expectantly overlap due to their similarity, whereas the latter is found to be more efficient in terms of reconstruction; see Fig. 3. In contrast, incorporating physical constraints with PLS drastically slows convergence [29], because truncation of negative eigenvalues produces a bias [43]; see also Fig. 12 in Appendix A6. We repeatedly confirm this bias by experiments in Appendix A2 as well as numerical simulations in Appendix A6. We further estimate the state’s

purity, as an example of an archetypal nonlinear function of the full state. Here LI shows highly unphysical results under insufficient statistics that do not converge in the plot, while the classical shadow purity estimators from Eq. (6) converge rapidly after only about 3000 shots; see Fig. 4(b). Finally, most inspired by applications, we probe the state’s entanglement by estimating all second-order Rényi entropies

$$S^{(2)}(\rho_A) = -\log_2 \text{tr}(\rho_A^2) \quad (9)$$

with the reduced density matrix ρ_A for part A of a bipartition (A, \bar{A}) together forming ρ [44]. In Fig. 4(c), we present results on Rényi entropies following Eq. (6) and particularly cover all bipartitions denoted by tuples (1, 4) qubits and (2, 3) qubits, referring to the number of qubits in each part. Note that, for classical shadow prediction, only the subset qubits in the smaller partition need to be taken into account. This leads to a drastic speedup of the analysis for larger scale systems. Moreover, all predictions can be analyzed after the data have been acquired.

For comparison, we also analyze the AME state with Pauli tomography; see Fig. 7 in Appendix A2. Slowest convergence is consistently found for PLS, which is also notably slower than regular MLE. We additionally confirm this by numerical simulations, considering only shot noise, i.e., statistical noise; see Appendix A6). Given only a few experimental shots, SIC tomography outperforms

Pauli tomography in the case of MLE reconstruction, likely because the SIC POVM provides the optimal information gain per shot. Curiously, however, for large shot numbers, Pauli MLE starts to outperform SIC MLE in terms of convergence. We suspect this to be a result of the over-completeness of the Pauli basis, where very large shot numbers may lead to improved accuracy for each orthogonal direction. Note that all methods converge to the same point, as verified in numerical simulations; see Fig. 12 in Appendix A 6.

In conclusion, we find SIC tomography to be preferable over Pauli tomography in regards to both classical computation time and convergence speed. At the same time, the underlying classical shadow formalism provides the potential for scaling to large quantum systems. We emphasize that the moderately lower quantitative performance of SIC tomography observed in our data is not inherent to the method, but due to experimental imperfections, i.e., the additional overhead in mapping SIC POVMs to ququarts, and the four-outcome readout; see Appendix A 1. These technical imperfections can be overcome in future devices.

VI. LIVE-UPDATE TOMOGRAPHY

Since the SIC POVM contains full tomographic information within each shot, it provides a unique way to speedup QST. Combined with classical shadows, which work by accumulating estimates in each shot according to Eq. (4), SIC tomography can be performed in real time (or “online”), i.e., a live update is performed for every new set of experimental shots. Apart from reducing time overheads by performing experiment and analysis at the same time, this approach has the big advantage that the experiment can be stopped as soon as all properties of interest are (believed to be) accurately estimated. Based on these ideas, we demonstrate a live reconstruction of a maximally entangled eight-qubit state. Specifically, we use a Greenberger-Horne-Zeilinger (GHZ) state with an additional local $\pi/4$ rotation of all qubits as a proxy for a generic maximally entangled state that is not aligned with any of the tomography bases to provide a fair comparison.

Figure 5(a) presents results on the estimation of fidelity, purity, and Rényi entropies [Eq. (9)], the latter for all possible (2, 6)-qubit bipartitions. Purity of the full eight-qubit state is found to converge the slowest, after around 1000 s. This is still a drastic speedup compared to LI, which, after an order of magnitude longer averaging time still produces unphysical results. We qualitatively conclude that an estimated property has converged by the time the observed values no longer change significantly with additional data, while neglecting slow experimental drifts. Fidelity converges after less than 500 s and Rényi entropies saturate almost immediately on the presented timescale. We remark

that both curves for SIC shadows LI and SIC shadows quadratic overlap as a consequence of the fast converging two-qubit subsets. Figure 8 in Appendix A 3 additionally characterizes data on the eight-qubit state in postprocessing to deliver both uncertainty estimates as well as results on all bipartitions. In Fig. 5(a), live updates are tracked up to 2500 s, which is the limit beyond where the analysis starts to take longer than the data acquisition [due to the quadratic scaling in the number of shots for purity estimation; see Eq. (6)]. We extend the discussion about the time relation between data acquisition and data processing in Fig. 5(b), where we acquire 100 experimental shots in about 2.4 s and show the entire 12 500 s of data taking. Over this time, we observe that the computational time nicely follows the expected quadratic growth with shot number, relating to the number of approximated ρ to compare in Eq. (6). When focusing on entanglement properties, the full-system purity can be excluded from the analysis. We then find that all bipartition Rényi entropies can be estimated in real time throughout the entire time of data acquisition. On top, simulations suggest that even on an 18-qubit state all bipartition Rényi entropies can be estimated live for around 1000 s.

VII. DISCUSSION AND OUTLOOK

We have demonstrated that real-time SIC tomography enables the prediction of Rényi entropies in less experimental shots than required for a minimal Pauli tomography implementation (see Fig. 5). Depending on the state and estimator, SIC tomography has the potential to significantly speed up the prediction of many more properties of quantum states. Beyond Rényi entropies, we challenge SIC tomography in a state identification game partly inspired by Ref. [45]. We show that it excels over other state-of-the-art methods (see Fig. 10 in Appendix A 5). For the challenge in question, SIC tomography required less than 20 shots to correctly identify a randomly chosen four-qubit linear cluster state among 16 orthogonal state possibilities, clearly outperforming any Pauli QST method. Note that such a speedup in favor of SIC tomography will become even more pronounced on up-scaled systems.

In practice, we accumulate 100 shots for each classical shadow approximating ρ following Eq. (6) that we refer to as *batching*, which in this particular case allows us to analyze the data in real time until all properties of interest converge. This postprocessing trick enables a trade-off between computational time and convergence speed, which is studied thoroughly through numerical simulations in Fig. 9 in Appendix A 4. A suitable batch size must be decided on a case-by-case basis. From a practical point of view, the experimental noise might also fall into consideration as it affects the targeted accuracy. Thus, batching experimental shots for the analysis comes as a handy tool for reducing analysis time with a limited effect




FIG. 5. Live-update SIC tomography of a maximally entangled eight-qubit state. We prepare a locally rotated GHZ state and probe the state via live-update SIC tomography. The local $\pi/4$ rotation ensures that the state is not aligned with any tomography basis and can serve as a proxy for an arbitrary entangled state. Data analysis is performed in parallel to data acquisition to get the quickest possible feedback. Shot-by-shot reconstruction is realized via classical shadows [Eqs. (4)–(A20)]. (a) Experimental results on purity (left), fidelity (middle top), and all (2,6)-qubit bipartition Rényi entropies (middle bottom) obtained live in a time regime where the data acquisition time dominates. We find the purity from classical shadows to have converged after less than 1000 s, with the other measures converging significantly faster. Fidelity converges very fast below 500 s and Rényi entropies saturate almost immediately. Note that, for the latter, the curves for SIC shadows LI and SIC shadows quadratic overlap due to the fast convergence for small subsystem sizes. The dash-dot line indicates the ideal value. (b) Comparison of the time for data acquisition (left-pointing triangles) following the expected linear curve along times curves for data analysis utilizing different methods and metrics covering the entire 12 500 s of data taking. When analyzing fidelity, purity, and all (2,6)-qubit bipartition Rényi entropies (circles) as shown in the experimental results of (a), live update is possible for around 2500 s. This is expected, since the shot-by-shot data analysis scales quadratically with the total shot number; see Eq. (6). Estimating only Rényi entropies for all bipartitions (diamonds), on the other hand, remains feasible for the entire time as it overlaps with data acquisition (left-pointing triangles) following the linear curve. Additionally, we simulate the estimation of all Rényi entropies of an 18-qubit (18q) state (crosses) by considering the data acquisition overhead (left-pointing triangles), demonstrating that this in principle also remains feasible to perform live.

on convergence and thereby extends the window for doing real-time analysis.

Though SIC POVMs are optimal for tomography [4], they are not known to exist in every dimension. Indeed, just resolving this issue would have far-reaching consequences for the foundations of mathematics [46]. For a given Hilbert space of dimension $d = 2^N$, a SIC constructed in that space is referred to as a *global* SIC, whereas a measurement composed of N two-dimensional SICs is referred to as a *local* SIC. If they indeed exist, global SIC tomography would be sample optimal in the sense that it saturates fundamental lower bounds from information theory [29,47] (joint measurements across many state copies could still yield further improvements [47,48]). Global SIC measurements would, however, be challenging to realize on quantum hardware. Quadratic circuit sizes (in the number of qubits) may be necessary, because the associated SIC states form a 2-design—a concept closely related to chaos and information scrambling [49]. And recent works provide lower bounds on the minimum circuit depth required to achieve information scrambling [50]. We conclude that, although local SICs may be less efficient than global SICs in terms of measurement complexity [29,51], they are much cheaper to implement. In addition, they are informationally complete and optimal amongst all possible local measurements [51].

To further improve system predictions, one might want to adapt the measurement basis depending on the state to analyze, potentially reducing quantum shot noise and offering a speedup in convergence. The concept of adaptive tomography follows those means, where the measurement setting is adjusted based on the outcomes of prior measurements [11,52–56]. In the most extreme cases, the measurement settings are changed after each shot. Although this has been demonstrated in some settings [57–61], it requires both additional classical computation and physical setting adjustments, rendering it potentially very time consuming.

In contrast, the SIC POVM representation from Fig. 1(d) turns out to be very efficient and practical. Moreover, we studied convergence properties of states with different overlap to either SIC POVM or Pauli basis; see Appendix A 7. Interestingly, purely local states analyzed by SIC significantly increase convergence when a component along one of the nonorthogonal SIC POVMs vanishes, obeying the concept of unambiguous state discrimination [62]. For randomly aligned, or correlated states, however, the effect vanishes, making the local rotation of the SIC POVM irrelevant. However, for estimating Pauli observables, the rotation does matter, with the optimal alignment given such that the overlap with the Pauli basis is symmetric; see Fig. 15(b) in Appendix A 7. This rotated SIC will be notably useful for VQE applications, as those rely on

efficient Pauli observable measurements. Moreover, investigating state dependencies more comprehensively holds the potential for interesting future research.

Finally, we emphasize that the combination of SIC-POVM measurements with the classical shadow formalism is well suited for directly estimating higher-order polynomials of an unknown density matrix ρ . As discussed in Appendix A 12 and following ideas from Refs. [32,33], this opens the door for mixed-state entanglement characterization of large-scale systems in real time, a likely requirement in the development of scalable quantum technology.

DATA AVAILABILITY

The data underlying the findings of this work is available at [63].

ACKNOWLEDGMENTS

R.S., L.P., M.M., C.E., M.R., P.S., T.M., and R.B. gratefully acknowledge funding by the U.S. ARO Grant No. W911NF-21-1-0007. We also acknowledge funding by the Austrian Science Fund (FWF), through the SFB BeyondC (FWF Project No. F7109), by the Austrian Research Promotion Agency (FFG) under Contract No. 872766, by the EU H2020-FETFLAG-2018-03 under Grant Agreement No. 820495, and by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via the U.S. ARO Grant No. W911NF-16-1-0070 and the US Air Force Office of Scientific Research (AFOSR) via IOE Grant No. FA9550-19-1-7044 LASCEM. C.F. is supported by the Australian Department of Industry, Innovation and Science (Grant No. AUSMURI000002). This project has received funding from the European Union's Horizon 2020 research

and innovation programme under the Marie Skłodowska-Curie Grant Agreement No. 840450. It reflects only the author's view, the EU Agency is not responsible for any use that may be made of the information it contains. T.M. and R.B. acknowledge support by the IQI GmbH.

M.R. conceived the project. C.F. and R.K. derived the theory results. R.S., L.P., M.M., C.E., P.S., T.M., and M.R. performed the experiments. R.S. analyzed the data. R.K. (theory), M.R. and R.B. (experiment) supervised the project. All authors contributed to writing the manuscript.

The authors declare no competing interests.

Note added.—In the final stages of this project we became aware of independent and complementary research using SIC POVMs on a superconducting quantum processor [64].

APPENDIX: EXPERIMENTAL SINGLE-SETTING QUANTUM STATE TOMOGRAPHY

1. Experimental toolbox

Experimental implementations here and in the main text are performed on a trapped-ion quantum computer, which is schematically shown in Fig. 6(a). The device operates on a string of $^{40}\text{Ca}^+$ ions stored in ultrahigh vacuum using a linear Paul trap. Each ion acts as a qubit encoded in the electronic levels $S_{1/2}(m = -1/2) = |0\rangle$ and $D_{5/2}(m = -1/2) = |1\rangle$ denoting the computational subspace [38].

Quantum state manipulation is realized upon coherent laser-ion interaction. A universal gate set comprises addressed single-qubit rotations with an angle θ around the rotation angle ϕ in the equatorial plane of the form $\mathbf{R}^j(\theta, \phi) = \exp(-i\theta(\cos\phi\sigma_j^x + \sin\phi\sigma_j^y)/2)$ with the Pauli operators $\sigma_j^x = X_j$ or $\sigma_j^y = Y_j$ acting on the j th

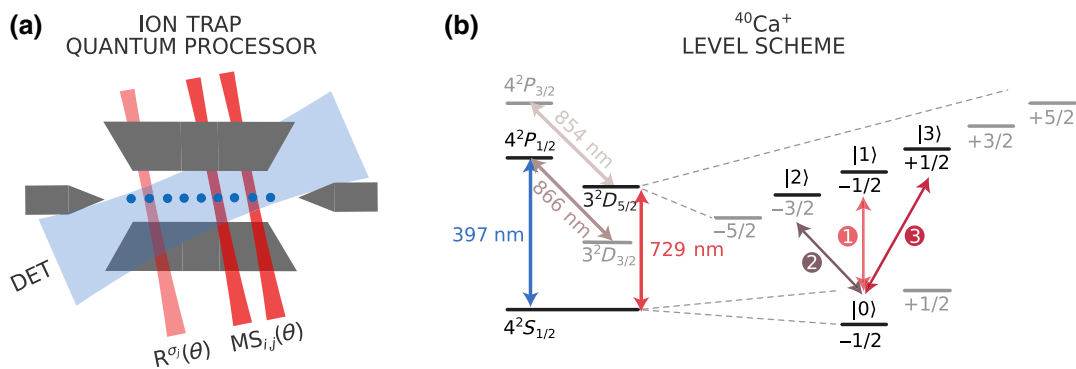


FIG. 6. Schematic of the trapped-ion quantum processor using $^{40}\text{Ca}^+$. (a) Each ion within the linear string encodes a qubit ($|0\rangle$, $|1\rangle$) or ququart ($|0\rangle$, $|1\rangle$, $|2\rangle$, $|3\rangle$). A universal gate set is realized upon coherent laser-ion interaction using tightly focused beams, addressing single ions for local gates (bright red) and pairwise ions for entangling gates (dark red). Alternatively, we can globally address all ions simultaneously. Readout is performed via collective fluorescence detection (DET); see the text for details. (b) $^{40}\text{Ca}^+$ level scheme with a dipole transition (397 nm) for cooling and detection, a metastable quadrupole transition (729 nm) for encoding qubits and ququarts within the Zeemann submanifold as well as transitions for repumping at 854 and 866 nm. The labeled transitions (1, 2, 3) provide coherent connection between all ququart states.

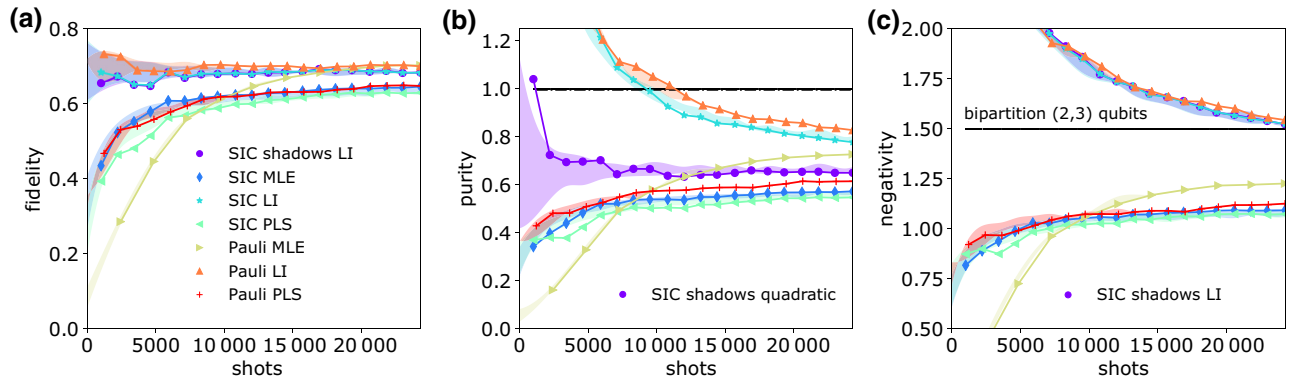


FIG. 7. Comparison between SIC and Pauli tomography for a standard set of reconstruction methods. Using the five-qubit AME state from Fig. 4 in the main text, we analyze fidelity, purity, and negativity from Eq. (A1). Because of the increased experimental complexity, we observe that SIC methods generally converge to slightly lower values than Pauli; however, this is not inherent to the method, but rather the implementation. (a) Fidelity converges quickest for LI approaches, as expected for general linear functions; see Appendix A 10. Convergence of PLS, however, is very slow compared to all other methods. SIC MLE at low shot numbers performs better than Pauli MLE as SIC measurements provide the maximum information gain. (b) SIC-based classical shadows (“SIC shadows quadratic”) demonstrate the best convergence for purity. MLE methods converge similarly to (a), but LI methods show very slow convergence with highly unphysical results, making their use problematic in practice. (c) As a commonly used entanglement measure, we evaluate the negativity of the reconstructed states; see the text for details. As expected, this property converges the slowest, since it is a global property of the density matrix (making this non-scalable), with LI again showing highly unphysical results. Bright shaded regions represent 1 standard deviation around the mean value when averaging multiple sets at the given shot number. Black horizontal lines denote ideal values.

qubit, together with two-qubit Mølmer-Sørensen entangling gate operations $\text{MS}_{i,j}(\theta) = \exp(-i\theta X_i X_j / 2)$ [65]. Multiple addressed laser beams, coherent among themselves, allow for arbitrary two-qubit connectivity across the entire ion string [17]. Optionally, all ions can be addressed simultaneously using a global beam to enable both collective local operations as well as collective entangling operations $\text{MS}(\theta) = \exp(-i\theta \sum_{j < \ell} X_j X_\ell / 2)$. We choose whatever is more efficient for the underlying experiment. Initial state preparation in $|0\rangle$ is reached after a series of Doppler cooling, polarization-gradient cooling, and sideband cooling. Readout is realized by exciting a dipole transition coupled to the lower qubit level $|0\rangle$ and collecting its scattered photons, from which the computational basis states $|0\rangle$ and $|1\rangle$ can be identified. Thereby, a qubit’s state is revealed by accumulating probabilities from multiple experimental runs. The dipole laser collectively covers the entire ion string, which enables a complete readout in a single measurement round. Additional pump lasers support efficient state preparation as well as cooling and prevent the occupation of unwanted metastable states outside the computational subspace $\{|0\rangle, |1\rangle\}$. Beyond the qubit level, we hold equivalent control over the entire S - and D -state Zeeman manifolds with up to eight levels in each ion, allowing us to encode a higher-dimensional quantum decimal digit (qudit); see Fig. 6(b). In this work we make use of up to four levels, denoting a ququart via additionally employing $D_{5/2}(m = -3/2) = |2\rangle$ and $D_{5/2}(m = +1/2) = |3\rangle$ alongside both qubit states. Ququart readout can be performed via three consecutive

fluorescence detections, where before the second detection the population between $|0\rangle$ and $|1\rangle$ is switched and before the third and final detection the population between $|0\rangle$ and $|2\rangle$ is switched; see Fig. 2(c) in the main text. Combining these three binary outcomes enables us to evaluate the ququart state probability within a single experimental run. Note that fluorescence detection in the case of measuring a bright state, even on only a single ion, heats up the entire ion string due to photon scattering as the qubits motion is coupled through the coulomb interaction. This is counteracted by a sequence of Doppler and polarization-gradient cooling after each individual detection, independent of its outcome, to keep the quality of postmeasurement bit-flip operations high and thereby suppress detection errors.

The last paragraph of this experimental setup section is dedicated to technical errors limiting our tomography experiments. Performance on *SIC tomography* is generally found to be moderately lower compared to Pauli tomography; see Fig. 4 in the main text. Evidently, this performance decrease is not inherent to the method, but rather owed to technical errors for two main reasons. (1) The SIC tomography implementation generates an overhead of five local pulses per qubit used for mapping qubit to ququart, depicted in Fig. 2(b), as well as two additional bit flips realizing the four-outcome readout; see Fig. 2(c). In contrast each Pauli *setting* requires just one local pulse per qubit, yet requires three orthogonal measurements per qubit to extract full tomography information. Our trapped-ion setup has a single-qubit gate fidelity of 0.9994(3) estimated from

randomized benchmarking as well as a two-qubit gate infidelity of roughly 0.98(1) estimated from a decay of fully entangling MS gates [17]. The latter two-qubit gate fidelity might slightly fluctuate from pair to pair. Moreover crosstalk to adjacent ions have an influence. (2) After each insequence detection the CCD camera demands for a 3-ms pause to process the data, before the upcoming sequence continues. However, we utilize this time for recooling the ion string via Doppler and polarization-gradient cooling. During this pause the ions are exposed to amplitude damping due to spontaneous decay from the upper D -state states, having a lifetime of about 1 s. Accounting for all this throughout measurement taking, we observe a loss of fidelity per qubit between Pauli and SIC tomography of less than 1%. Thus, extracting complete tomography information in a single experimental run, i.e., shot, comes at the expense of a more complex experiment, which is however of a technical nature and can be overcome in future devices with reduced single-qubit error rates as well as with faster processing CCD cameras, which nowadays already exist. More importantly, only SIC tomography offers the unique potential to predict nonlinear properties in large-scale systems, as pointed out here further below and in the main text.

2. Comparing tomography methods

We start off by presenting complementary experimental data covering the five-qubit AME state from Fig. 4 in the main text. Whereas in the main text the focus was on scalable approaches, especially SIC-based classical shadows, here we compare these results to Pauli tomography according to Figs. 1(a) and 1(c) using linear inversion and MLE reconstruction following Eq. (3). Generally, linearly reconstructed density matrices exhibit negative eigenvalues in the case of insufficient statistics, which manifest themselves particularly in unphysical values for nonlinear functions of the density matrix, as indicated by purity values above 1. Physical constraints can be imposed on LI, through truncation of negative eigenvalues following Refs. [27–29], which we refer to as PLS. Importantly, PLS adds only a negligible computational overhead to LI. Consequently, the covered set of tomography approaches is representative in the field of quantum computation and quantum information. Figure 7 depicts results on fidelity, purity, and negativity, with the latter being a common measure of entanglement, albeit one that is challenging to access with experiments.

Note that, for five-qubit Pauli tomography, $3^5 = 243$ settings are required, where for this particular case, each setting is repeated 100 times, leading to the stated maximum shot number of 24 300. A hundred shots per setting proves to be a good trade-off in the trapped-ion platform, accounting for both statistics and systematic drifts in the experiment. The moderately lower performance (in terms

of the numerical values) of SIC tomography is not inherent to the method, but comes from technical imperfections due to experimental overhead. In particular, the mapping of the SIC POVM to ququart states as well as the four-outcome readout, both essential for single-setting tomography, add experimental complexity; see Appendix A 1.

Figure 7(a) shows that the fidelity converges very quickly for LI approaches, as is expected for general linear operators; see Appendix A 10. Interestingly, Pauli MLE performs worse than all other methods at very few shots, while SIC-based methods perform very well, even for low shot numbers, since this measurement extracts the maximal amount of information for a generic state. PLS, on the other hand, despite the computational efficiency, converges much slower than the other methods, even including MLE [29]. For quadratic measures in Fig. 7(b), it becomes clear that LI methods produce highly unphysical results that take a long time to converge, limiting the usefulness of these methods in practice. PLS solves this problem, but again shows very slow convergence. Both problems are solved by the SIC-based classical shadow purity estimator from Eq. (A20), which demonstrates both fast convergence and accurate (physical) estimates. Finally, we study negativity as a commonly used measure of quantum entanglement [66]:

$$\mathcal{N}(\rho) = \frac{\|\rho^{\Gamma_A}\|_1 - 1}{2}. \quad (\text{A1})$$

Here ρ^{Γ_A} represents the partial transpose with respect to subsystem A of a bipartition (A, B) together forming ρ . The 1-norm in Eq. (A1) denotes the absolute sum of all negative eigenvalues given by ρ^{Γ_A} . By construction, the partial transpose of a separable state cannot have negative eigenvalues, such that the negativity vanishes. Quantum negativity is an entanglement monotone: if it is positive then the underlying state must be entangled. The converse, however, need not be true in general [67].

As in the main text, we consider the bipartition (2, 3) for the five-qubit system. We find a significantly slower convergence than for Rényi entropy (see Fig. 4). This is due to the requirement for processing the entire density matrix ρ for quantum negativity as opposed to the classical shadow subsystem purity estimator, which is only evaluated on the smaller partition. The same convergence behavior is confirmed by numerical simulations discussed in Fig. 12 below. Classical shadows, on the other hand, allow for tighter classifications of entanglement [32,33]. The key idea is to probe the presence of negative eigenvalues in the partial transpose by comparing degree- d polynomials in the underlying density matrices. As d increases, these tests become tighter and eventually recover the negativity condition for entanglement [$\mathcal{N}(\rho) > 0$]. Classical shadows allow the estimation of all polynomials involved, but

the classical postprocessing cost becomes less and less favorable as the polynomial degree d increases [33].

We emphasize that, from the given set of tomography schemes, SIC-based classical shadow estimators deliver the best results in terms of both convergence as well as practicability. MLE reconstruction typically fails due to a lack of computational power and LI neglects physical constraints—a shortcoming that becomes very pronounced for nonlinear observables. Incorporating physical constraints by projection (PLS) remains computationally efficient, but leads to considerably poorer convergence behavior. Finally, classical shadows are the only approach for efficiently predicting nonlinear functions, such as mixed-state entanglement [32,33] of large-scale systems.

3. Complementary results on the rotated eight-qubit GHZ state

In the live-update discussion of Fig. 5 in the main text all two-qubit bipartitions were evaluated on top of fidelity and purity. Here we analyze the same data in postprocessing to present results on all bipartitions ((1, 7) qubits, (2, 6) qubits, (3, 5) qubits, and (4, 4) qubits). This evaluation for all possible pairs is not possible in real time on a standard desktop computer. We average until 50 000 shots (roughly 1200 s of data taking), where the SIC-based classical shadow purity estimator [Eq. (A20)] of the eight-qubit state, representing the most demanding property, has converged. Data are taken for a total of 12 500 s. Note that at around 40 000 shots almost no change is visible in the classical shadow purity as well as the respective standard deviation. The latter is due to systematic experimental drifts over the course of the long time measurement. Here we batch 100 shots for each approximate

ρ following Eq. (4), which speeds up the analysis without significant loss in accuracy; see Appendix A 4 for a thorough study of batch sizes considering both analysis time and accuracy. We also observe that the convergence of bipartitions from LI become significantly slower than classical shadows as the subsets get larger. For PLS, individual bipartitions even visually separate [see Fig. 8(c)], indicating very slow convergence behavior, confirming similar observations throughout the main text and Appendix.

Pauli tomography is neglected here as all bipartition Rényi entropies from SIC-based classical shadows converge faster than the time it takes to obtain just a single measurement per Pauli setting. Moreover, there is no efficient way of adding physical constraints to the reconstructed state ρ as MLE for an eight-qubit state is unfeasible on standard desktop computers, and PLS delivers significantly worse convergence. The findings here agree with those of the five-qubit AME state, previously discussed. SIC-based classical shadow estimators demonstrated to be most suitable for predicting nonlinear functions towards larger system sizes. Apart from the exemplified quadratic measures utilized in purity and Rényi entropy, classical shadow estimators support higher polynomial functions following the same principles [32].

4. Classical shadow convergence and practicability

In the live-update studies from Fig. 5(b) in the main text, we demonstrated real-time analysis of ongoing SIC tomography experiments, until all properties of interest were accurately estimated. How long the analysis can be performed in real time, however, depends on the size of the subsystems to be analyzed, as well as the type and number of functions to be estimated in parallel. Whereas time




FIG. 8. Detailed postprocessing analysis of the rotated eight-qubit GHZ state from Fig. 5 in the main text. Results now cover all bipartitions, which is not feasible in real time on a desktop computer. Shaded regions represent 1 standard deviation around the mean value when averaging over multiple sets of a total of 12 500 s of data. (a) We find fidelity from LI to converge before 20 000 shots, i.e., 500 s of data taking. (b) Classical shadow purity as converges well by 40 000 shots, or about 1000 s of data taking, which in the postprocessing analysis here is also confirmed by the saturation of the standard deviation towards more shots. The remaining fluctuations arise from systematic drifts in the experiment over the course of the measurements. (c) SIC-based classical shadows converge even quicker than the time it takes to measure just one shot per Pauli basis. On the other hand, LI and PLS are significantly slower, with PLS even showing a distinct separation between the individual bipartitions [(1, 7) qubits, (2, 6) qubits, (3, 5) qubits, (4, 4) qubits]. Dash-dot horizontal lines denote ideal values.




FIG. 9. Numerical simulations of batching SIC-based classical shadow purities. We plot the purity’s standard deviation from 100 repetitions of each point in the 2D grid using a logarithmic color scheme. A three-qubit linear cluster state is utilized in these numerical simulations, although the particular state and qubit number does not effect the qualitative picture. Regions of similar accuracy are found to have a certain almost vertical extension, indicating that bunching does not significantly degrade accuracy, while greatly reducing computation complexity. The ideal batch size must be evaluated on a case-by-case basis, weighing convergence against analysis time.

consumption for linear observables by means of SIC classical shadows remains constant over the course of data acquisition—individual experimental shots are simply processed and accumulated according to Eq. (A15)—nonlinear functions generally require higher-order products of all combinations from the given set of shots; see Eq. (A20). The resulting scaling is governed by the maximum polynomial order of the function in question minus one. For quadratic functions, in particular, the computation time for every new shot grows linearly with the number of already accumulated shots M . Hence, over the course of the data acquisition, this can eventually become computationally demanding, especially for large problem sizes where a large number of shots must be accumulated. Thus, to perform nonlinear function analysis in real time, one either keeps subsystems (and by that shot requirements) relatively small or uses a so-called batching approach, where multiple shots are bundled to estimate a more accurate ρ and thereby reduce the number of costly higher-order product combinations. In this section, we discuss the potential, as well as limits, of increasing the batch size in contrast to comparing single shots. To develop a quantitative statement, we perform

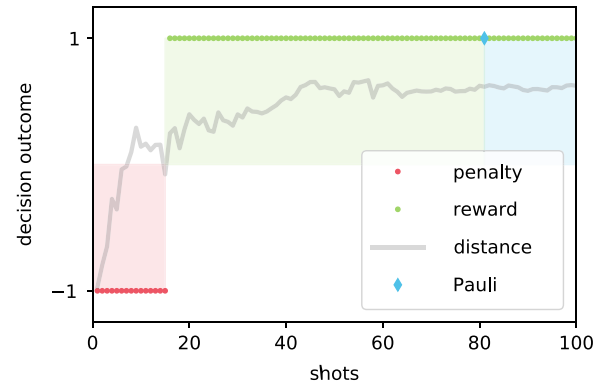


FIG. 10. Decision game. A quiz master constructs a quantum game where she provides us experimenters, having access to a quantum computer, with a circuit to prepare a state and perform measurements. After the measurement stage, we must report a certain property. Importantly, the property is only revealed after the experiment has been completed, preventing us from simply measuring the property in question. The aim is to win this challenge with as few experimental shots as possible. In the present case a single state from a fixed set is randomly chosen and implemented, which we then have to identify in as few shots as possible. SIC tomography delivers a reward in less than 20 shots, whereas the smallest Pauli tomography instance requires $3^4 = 81$ shots. Distance refers to the minimum fidelity difference between the target state and all others; see Fig. 11(b).

simulations on a three-qubit linear cluster state considering only statistical quantum projection noise. Along those lines, simulated tomography data are sampled from a multinomial distribution considering different sample sizes to mimic experimental shots. Each noisy set of tomography data is then reconstructed by means of SIC-based classical shadows, in particular, focusing on the quadratic purity estimator from Eq. (A20). Since we are only interested in changes in convergence behavior, the choice of state does not affect the qualitative statement of these numerical simulations. We compare 100 different shot numbers for 100 different batch sizes in a practical regime for three-qubit states. On the one extreme we compare all combinations of shadows obtained from a single shot each, which is known to be statistically optimal [see also the discussions around Eq. (A20)]. On the other extreme we compare just two shadows, each linearly more accurate, since they are obtained from averaging half the shots. In between we can trade-off the quality of the individual estimators versus the number of comparisons between estimators. By additionally accounting for analysis time, a sweet spot can be determined on a case-by-case basis.

We analyze the convergence behavior of these different strategies in Fig. 9 by means of the standard deviation of the classical shadow purity estimator, when repeating every point in the two-dimensional (2D) grid 100 times. We plot “shots” against “batch size” using a logarithmic color coding. Darker regions refer to a more accurate purity

estimate for the underlying state. A region at a specific color always shows a certain almost vertical extension, indicating that batching has a very small effect on accuracy. At the same time, however, it offers to speedup the data analysis significantly. We note that the observed pattern qualitatively remains the same for larger system sizes and accordingly more shots. Especially for larger subsystems, where a lot of statistics is required, batching has the potential to significantly speedup analysis. We made use of this method in Fig. 5 in the main text, where 100 experimental shots were used for each classical shadow. This turned out to be sufficient to estimate all relevant target properties in real time.

5. Fast state identification with classical shadows

We now consider a quantum game, where a quiz master targets us experimenters, having access to a quantum computer with a state to prepare and a question about a certain property. Importantly, the question is only revealed after performing the experiment. Such a setup is partly inspired by recent works on quantum-enhanced learning [45]. Here, we follow a game where the goal is to prepare a random target state from a fixed set of 16 states and then pinpoint this target state in as few experimental shots as possible. Figure 10 depicts the results, where SIC tomography allows us to receive a reward in less than 20 shots. In stark contrast, the minimum number of shots for the same task using Pauli tomography is $3^4 = 81$. The figure of merit for this game is the estimated distance between the states as the minimum difference in fidelity between the target state and all others. A reward is obtained as soon as that minimum distance remains positive.

Specifically, in the decision game of Fig. 10, we randomly prepare 1 out of 16 orthogonal four-qubit linear cluster states. These states correspond to all combinations of input states $|\pm\rangle$ on the four qubits. The target state is then identified by comparing the linear-inversion fidelities between the prepared state and all 16 possible targets. Note that fidelity is a good option for state identifications as linear observables generally converge quickest under LI, as we showed experimentally (see Figs. 4 and 5) and also formally derived in Appendix A 10. Figure 11(a) depicts fidelities with respect to all 16 states for SIC and Pauli tomography against the number of experimental shots. As expected, only one of the curves is close to fidelity 1, indicating the target state, whereas all others approach 0 within experimental uncertainties. We remark that, for the minimal Pauli implementation given at 81 shots, the state fidelities bunch at two values distinguishable in the figure due to the limited data. Going to the next higher settings at 162 and 243 shots, errors start to distribute indicated by more traces popping up. The distinguishing performance is even better visualized by plotting the difference between the target state

fidelity and all others in Fig. 11(b). The minimum of these distances (i.e., the worst case) is also shown in the background of Fig. 10 and used as our state distinguishability criterion. Surprisingly, less than 20 experimental shots are required to pinpoint the state, clearly undercutting the minimum for Pauli tomography. This argument can in principle be extended to larger systems, where we have seen that properties such as linear observables or Rényi entropies (see Fig. 5) converge much faster than the 3^N experimental shots required for the smallest instance of Pauli tomography. Fast state identification represents another fruitful example where SIC-based classical shadows not only appear more practical but also impart quicker convergence for certain tasks than Pauli tomography. The difference in performance between the tomography approaches originates from SIC's bigger experimental overhead (Appendix A 1) as well as the fact that the analyzed states are stabilizer states in the Pauli basis. Hence, they favor the analysis using Pauli measurements in terms of convergence; see Appendix. A 7.

6. Numerical simulations comparing tomography approaches

This section aims to support previous experimental findings through numerical simulations of convergence under statistical quantum shot noise. For this purpose, simulated tomography data are generated by sampling from a multinomial distribution, where the sample size is given by the number of shots. For the sake of comparability, we perform these numerical convergence simulations using the five-qubit AME state from Fig. 4 in the main text and Eq. (8). We study both SIC and Pauli tomography and cover all reconstruction methods as experimentally studied in Appendix A 2, i.e., LI, PLS, MLE, and, for quadratic functions, also SIC-based classical shadows. Numerical results are presented in Fig. 12 on a double-logarithmic scale, covering infidelity (to better illustrate convergence), and trace distance as another commonly used property for state distinguishability:

$$T(\rho, \sigma) = \frac{1}{2} \text{tr} \left[\sqrt{[(\rho - \sigma)^\dagger (\rho - \sigma)]} \right]. \quad (\text{A2})$$

LI approaches are again found to converge significantly faster than all other methods in terms of infidelity, while trace distance shows the opposite behavior due to being much more complicated to estimate. This is indicative of the unphysical nature of LI estimates. We further note that linear-inversion infidelities more often produce negative eigenvalues in the case of SIC tomography, since the five-qubit AME state is aligned with the Pauli basis and thus favors this basis for the estimation of linear observables; see Appendix A 7. In contrast, MLE approaches respect those physical boundaries, which result in valid values on all estimators at the cost of a slower

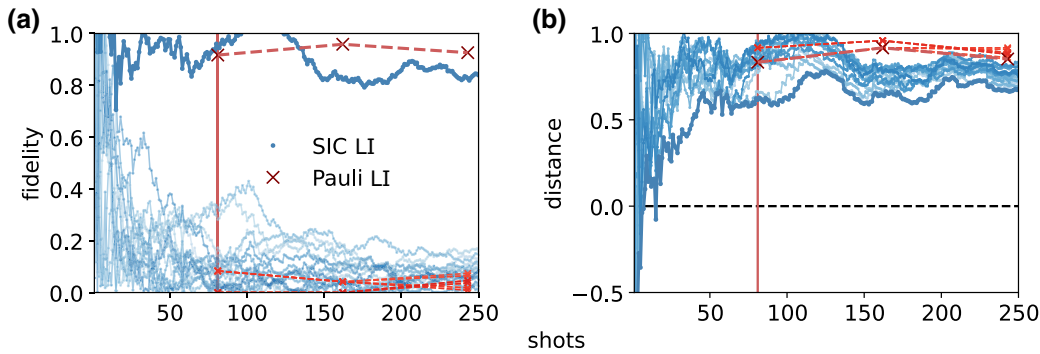


FIG. 11. Fast state identification from SIC tomography. Complementary analysis of the quantum game from Fig. 10. (a) Estimated fidelity with each of the 16 possible states. Via SIC tomography, we manage to identify the state with less than 20 shots, clearly undercutting the smallest Pauli implementation using $3^4 = 81$ shots. (b) For illustration purposes, we plot the fidelity difference between the generated state and all possible states as individual curves, which provides a distance measure. The minimum of these differences is used as a state distinguishability criterion to gain a reward; see the text for details.

convergence. Here we also confirm previous experimental observations (see Fig. 7) that SIC MLE performs better for small shot numbers than Pauli MLE. This is likely due to the fact that SIC POVMs provide the optimal information gain in each shot. Curiously, however, for larger numbers of shots, Pauli MLE eventually converges faster. This might be due to the overcompleteness of the

Pauli basis, but remains to be fully understood. While not specifically presented in this manuscript, we find the very same behavior for other states having different overlaps with the Pauli basis and SIC POVM; hence, this effect does not seem to be due to the choice of state. PLS again produces the slowest converges as we have already seen across our experimental studies.

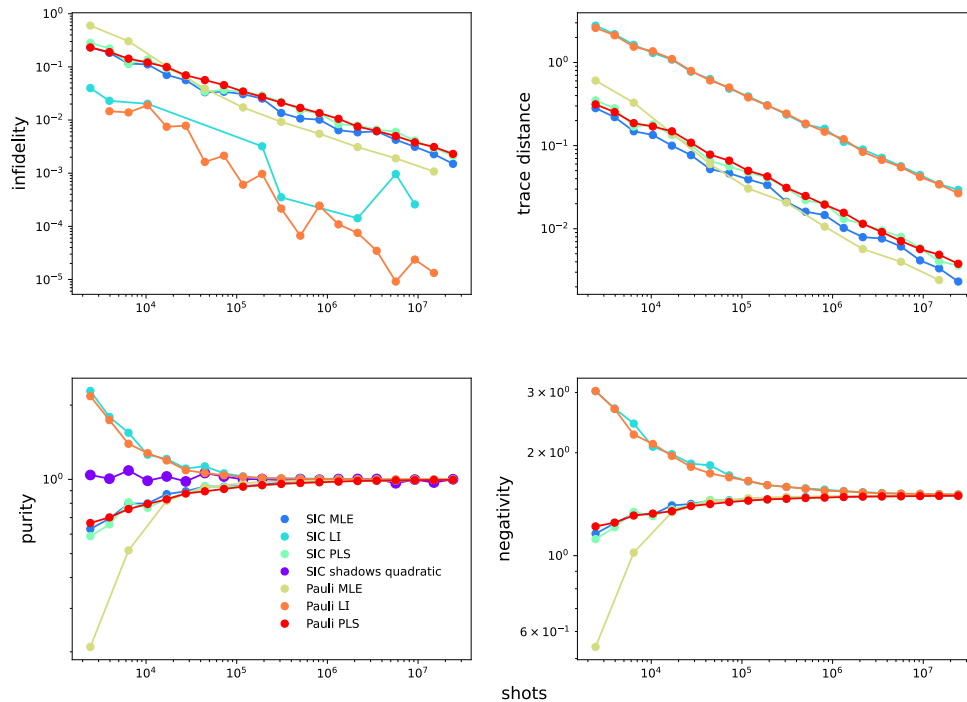


FIG. 12. Numerical simulations incorporating statistical noise covering all tomography approaches. To this extent, quantum shot noise is incorporated via sampling ideal tomography data from a multinomial distribution, with sample sizes corresponding to the number of shots. We additionally present infidelity (1-fidelity) and complementary trace distance following Eq. (A2) and plot in double-logarithmic scale to more clearly visualize several orders of magnitudes. Note that infidelity estimation with LI sometimes delivers unphysical results, indicated by values out of range of the plot. Note that the smaller number of data points for SIC LI comes from ignoring unphysical results, which are more likely when the target state is misaligned with the measurement basis; see Appendix A 7 for details.

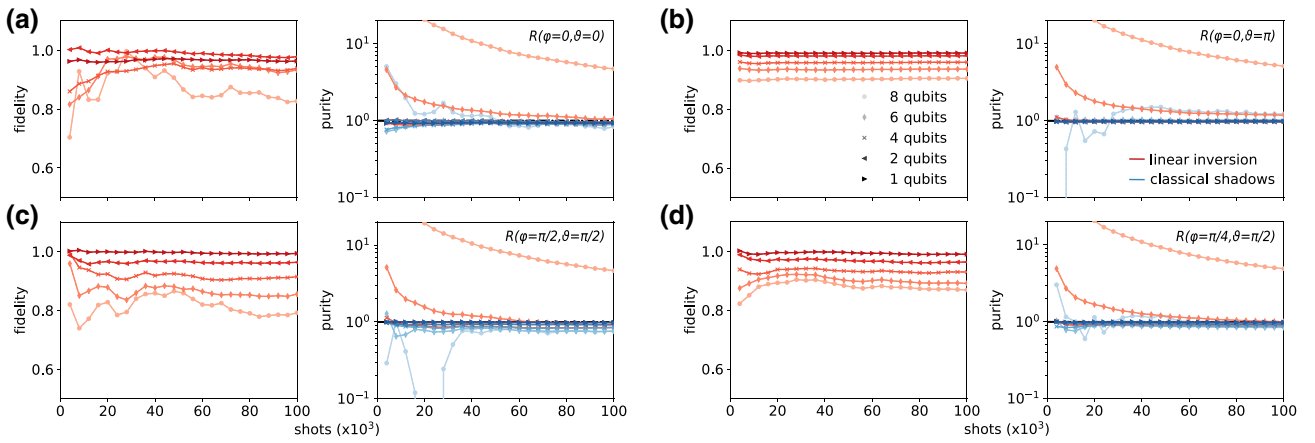


FIG. 13. Experimental convergence investigation of local states with different orientation, i.e., basis overlap. The given states are collectively prepared on eight qubits and partially traced to study multiple subsets. (a) The $|0\rangle$ state perfectly aligned with the first SIC vector. (b) The $|1\rangle$ state orthogonal to the first SIC vector, whose component completely vanishes. Convergence in (b) is significantly faster than in (a) as only three of the nonorthogonal SIC vectors take part. Boosted convergence is related to unambiguous state discrimination [62], limiting nonorthogonal measurements. (c) Superposition state with maximized overlap with one of the SIC vectors. (d) Superposition state with minimum overlap with the SIC vectors. Here, convergence of the state in (d) is better than (a) as the state's information is more regularly distributed over the SIC vectors, not favoring one, resulting in a higher information gain. Also, note that the achievable fidelity generally drops with higher qubit number due to experimental imperfections.

Overall, quantum negativity exhibits the slowest convergence, which confirms what we experimentally observed in Fig. 8. We note that the negativity calculation always requires the full density matrix independent of the bipartition. This is in stark contrast to Rényi-entropy-based measures and general nonlinear functions supported on a subset of the system, which can be estimated efficiently using SIC-based classical shadows. The latter again show the fastest convergence for purity according to Eq. (A20), which will generally be true for classical shadows by construction. To keep simulations efficient, the batch size for the purity estimator is chosen as a constant fraction (0.01) of the total number of shots, which has a negligible effect on convergence; see Appendix A 4. Finally, these numerical simulations could reproduce all features and findings from the experimental studies, thus indicating that there are no principal limitations to our experimental implementation of both SIC and Pauli tomography. Moreover, all reconstruction methods converge to the same values, indicating no principle drawback across the various methods.

7. Overlap with tomography basis in experiment and simulation

Our experimental studies covered by Figs. 4 and 5 focus on maximally entangled states that are differently aligned with respect to Pauli basis and SIC POVM in order to not particularly favor either of them. We resume this discussion in more detail by investigating states of varying orientation, with respect to the measurement

states, to study its effect on both experiments and numerical simulations. To this extent, convergence behavior of first purely local states followed by entangled states is demonstrated.

We start off by presenting experimental results on SIC tomography of local states up to eight qubits having different overlap with the Pauli basis and SIC POVM, depicted in Fig. 13. Analyzed metrics include fidelity and purity using LI and classical shadows as those scale more favorably in the number of qubits than MLE and PLS. While MLE becomes computationally too demanding for moderate systems sizes, PLS shows extremely slow convergence behavior, notably demanding more experimental shots. The given states are collectively prepared and subsequently partial traced to study multiple qubit numbers; see Appendix A 1. Brighter colors denote higher qubit numbers. Among the Pauli basis states $|0\rangle$ and $|1\rangle$, the latter performs significantly better under the SIC POVM. This reflects a general theme, where in using nonorthogonal bases it is preferable if one component vanishes, which is related to the concept of unambiguous state discrimination [62]. In the example of $|1\rangle$ this is indeed true for the first SIC vector aligned with $|0\rangle$ [see Fig. 15(a)]. For superposition states, we find that states that maximize the overlap with one SIC vector [Fig. 13(c)] perform worse than those that feature more even overlap with all SIC vectors [Fig. 13(d)]. Generally, higher qubit number states result in lower fidelities, due to experimental imperfections. We again emphasize the particularly fast convergence of the SIC-based classical shadow purity estimator, being here only moderately slower than fidelity.

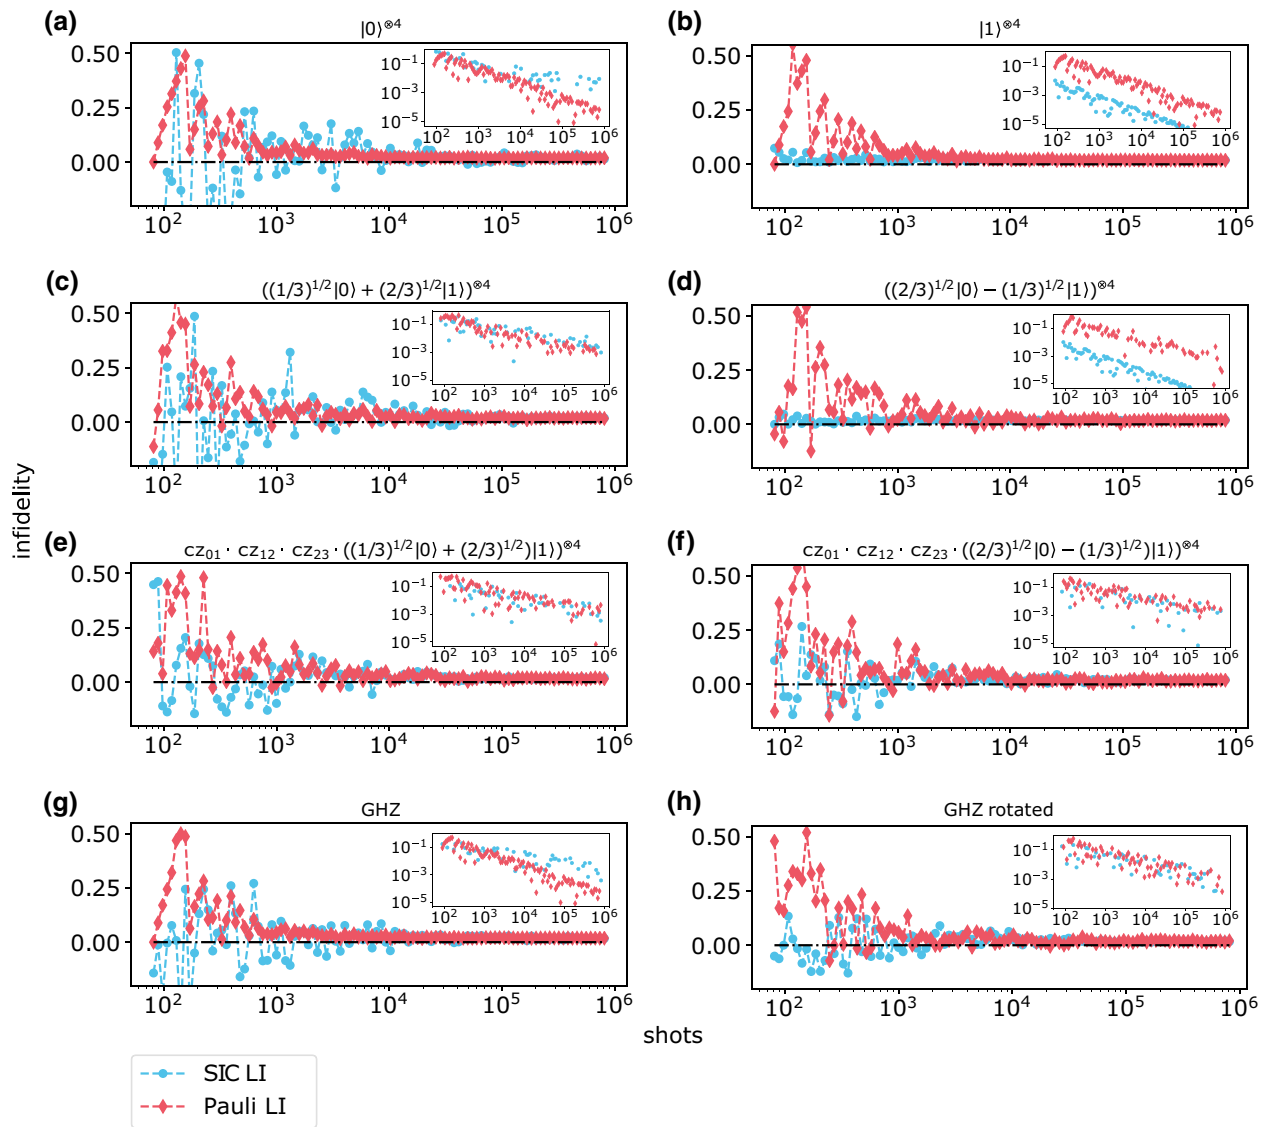


FIG. 14. Convergence of SIC and Pauli tomography in numerical simulations for noisy states with different basis overlap. To incorporate noise, a depolarizing channel is applied to the ideal density matrices; see the text for details. Complementary, insets represent numerical results on projection noise only to more distinctively illustrate convergence behavior. (a)–(d) Local states that align with a SIC vector (a),(c) show slower convergence for SIC tomography than those orthogonal to a SIC vector (b),(d), confirming the experimental observations of Fig. 13. Pauli tomography, on the other hand, generally performs best for states aligned with the basis (a),(b). (e)–(h) The differences in convergence performance vanish for SIC tomography when considering entangled states. In the case of the Pauli basis, however, some improved convergence can still be observed for specific states, such as the GHZ state (g), which is aligned with the Pauli basis and thus performs better. Overall, the fast converging situations (more pronounced in the insets) for SIC (b),(d) and Pauli (a),(b),(g) tomography do beyond a certain shot number that depends on the noise, no longer carry unphysical components in the resulting infidelities.

To confirm these results for purely local states and to extend the discussion to entangled states of different orientation, we perform numerical simulations on four qubit states under quantum shot noise as previously explained within Appendix A 6. Figure 14 contains results for SIC tomography as well as now for Pauli tomography for predicting infidelity via LI, which scales more efficiently than MLE, that however would not change the essence

of the statements. On top, we numerically simulate noisy states by applying a depolarizing channel $(1 - p_{\text{depol}})\rho + p_{\text{depol}}\mathbb{1}/d$ with $p_{\text{depol}} = 0.02$ to the ideal density matrix ρ to perform the study under more realistic conditions. This enables us to observe the influence of unphysical properties, especially in regards of LI. Figures 14(a)–14(d) cover local states both along a SIC vector (a),(c) as well as orthogonal to it (b),(d), i.e., in the opposite direction of the

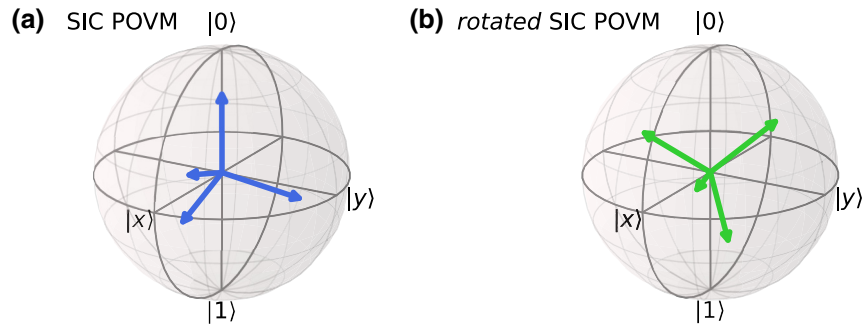


FIG. 15. Various representations of SIC POVM for maximizing information gain. (a) Standard representation of SIC POVM as used throughout this manuscript in both experiments and numerical simulations. (b) Rotated SIC optimized for predicting Pauli observables.

Bloch sphere. These simulations reproduce the effects seen experimentally in Fig. 13, where states aligned with a SIC vector (a),(c) converge more slowly than those orthogonal to it (b),(d). In contrast for Pauli tomography, representing an orthogonal basis, best results are obtained for states that are aligned with the basis, as seen in Figs. 14(a) and 14(b).

In the case of local states, measurements are uncorrelated and the above state-dependent convergence is to be expected. The situation might change when moving to entangled states. In Figs. 14(e) and 14(f) we apply controlled-phase gate states along the SIC vectors from Figs. 14(c) and 14(d) to generate an entangled state that is still (in a sense) aligned with the SIC basis. The resulting convergence is similar for both tomography methods. Curiously, SIC tomography performs equally well as for special states like a GHZ state [Fig. 14(g)] or a more generic rotated GHZ state [Fig. 14(h)]. In contrast, Pauli tomography, which also shows little dependence on the state once entanglement is involved, does outperform for the GHZ state [Fig. 14(g)], which is perfectly aligned with the Pauli basis. Importantly, the rotated GHZ state, which is somewhat randomly aligned with both the SIC POVM and Pauli basis [Fig. 14(h)] shows no difference between the tomography methods. The same rotated state, yet on eight qubits, is utilized for the real-time analysis in Fig. 5 in the main text to make sure that the comparison does not favor any approach. In the case of fast converging situations (see the insets of Fig. 14 without depolarizing noise) negative infidelities, influenced by unphysical properties of LI, disappear beyond a certain shot number depending on the amount of noise. This effect is illustrated by both tomography methods, particularly in Figs. 14(b) and 14(d) for SIC and in Figs. 14(a), 14(b), and 14(g) for Pauli tomography.

Inspired by these findings, it can be beneficial to rotate the SIC POVM used throughout this manuscript from Fig. 15(a) to favor the particular state or application that it is used for. In particular, the alignment depicted in Fig. 15(b), which has equal overlap with every Pauli basis vector, leads to improved prediction of Pauli observables.

Intuitively, this alignment ensures that each shot contains equal information about every Pauli observable. Geometrically, imagine a triangle spanned by the three Pauli basis vectors in the positive direction. Then, the orthogonal state to the first SIC vector ($\{-1/\sqrt{3}, -1/\sqrt{3}, -1/\sqrt{3}\}$) perpendicularly intersects this triangle area through its center. The front area of the SIC tetrahedron orientates parallel with the triangle, thus favoring Pauli eigenstates or stabilizer states, for which we find plenty of applications across the entire field of quantum computation and quantum information. In particular, VQE applications, which rely on the efficient estimation of many Pauli observables, could benefit from this choice of SIC POVM.

Note that experiments on rotated SICs can be straightforwardly realized by changing the local mapping sequence from Fig. 2(b), and do not add further complexity to the implementation.

8. SIC-POVM-based classical shadow framework

Renes *et al.* [2] discovered SIC POVMs, because of their exceptional tomographic capabilities. Ever since, both (tensor products of N) single-qubit SIC POVMs and global 2^N -dimensional SIC POVMs have served as idealized measurements for state reconstruction tasks. Single-qubit SIC POVMs, also known as tetrahedral POVMs, have also been used to acquire training data for neural network quantum state tomography [12,13].

Geometrically speaking, SIC POVMs [2] and the overcomplete Pauli basis [68] both form complex projective 2-designs (the single-qubit Pauli basis is actually a 3-design [35–37]). Roughly speaking, this means that the first two (three) moments exactly reproduce the moments of uniformly (Haar) random states. As detailed below, closed-form expressions for Haar-random moments can then be used to compute measurement operators and estimators analytically. For Pauli basis measurements, this observation culminated in efficient PLS estimators for full state tomography [29], as well as the classical shadow formalism for directly predicting (non)linear properties of the

underlying state [14,32]. Subsequently, some of these ideas have been extended to (single-qubit) SIC POVM measurements. García-Pérez *et al.* [11], in particular, highlighted that SIC POVM measurements can outperform Pauli basis measurements in VQE-type energy measurements and, more generally, for predicting linear state properties.

Here, we build on all these ideas and provide a self-contained derivation of classical shadows from (single-qubit) SIC POVM measurements, as well as rigorous sample complexity bounds for general linear and quadratic property estimation. The actual definition of SIC POVM shadows is virtually identical to existing (Pauli basis) classical shadows, because both form complex projective 2-designs. Sample complexity bounds, on the other hand, require novel proof techniques. They require computing variances that correspond to third-order polynomials in the measurement ensemble. This is comparatively easy for Pauli basis measurements, which form a 3-design. But for SIC POVMs—which only form a 2-design—these existing techniques do not apply. We overcome this drawback with new proof methods that directly use the symmetries within a SIC POVM rather than an abstract 3-design property. To our knowledge, these theoretical arguments are novel and may be of independent interest.

9. Classical shadows from SIC POVM measurements

a. The single-qubit case

Single-qubit density matrices live in the (real-valued) vector space of Hermitian 2×2 matrices that we denote by \mathbb{H}_2 . On this space, single-qubit SIC POVMs are known to form so-called projective 2-designs [2]. Mathematically, this means that discrete averages over (outer products of) SIC vectors reproduce uniform averages over all possible pure states up to second moments [34]. This is captured by the two averaging formulas

$$\frac{1}{4} \sum_{i=1}^4 |\psi_i\rangle\langle\psi_i| = \int |v\rangle\langle v| dv = \frac{1}{2} \mathbb{I} \in \mathbb{H}_2, \quad (\text{A3})$$

$$\frac{1}{4} \sum_{i=1}^4 (|\psi_i\rangle\langle\psi_i|)^{\otimes 2} = \int (|v\rangle\langle v|)^{\otimes 2} dv = \frac{1}{6} (\mathbb{I} \otimes \mathbb{I} + \mathbb{F}). \quad (\text{A4})$$

Here, dv denotes the unique pure state measure (normalized to $\int dv = 1$) that assigns the same infinitesimal weight to each state ($dv = dw$). The two-qubit Slashed Down operator $\mathbb{F}|v\rangle \otimes |w\rangle = |w\rangle \otimes |v\rangle$ acts by permuting tensor factors.

The first averaging formula (A3) confirms that the collection $\{\frac{1}{2}|\psi_i\rangle\langle\psi_i| : i = 1, 2, 3, 4\} \subset \mathbb{H}_2$ forms a valid quantum measurement for single-qubit systems (POVM). Let

$\rho \in \mathbb{H}_2$ be a density matrix, i.e., a Hermitian matrix with unit trace whose eigenvalues are non-negative. Then,

$$\Pr[i|\rho] = \text{tr}\left(\frac{1}{2}|\psi_i\rangle\langle\psi_i|\rho\right) = \frac{1}{2}\langle\psi_i|\rho|\psi_i\rangle \quad \text{for } i = 1, 2, 3, 4$$

is the probability of obtaining POVM outcome i upon measuring state ρ , obeying $\Pr[i|\rho] \geq 0$, because density matrices do not have negative eigenvalues ($\langle\psi_i|\rho|\psi_i\rangle \geq 0$). Moreover, Eq. (A3) ensures proper normalization:

$$\begin{aligned} \sum_{i=1}^4 \Pr[i|\rho] &= \text{tr}\left(2\left(\frac{1}{4} \sum_{i=1}^4 |\psi_i\rangle\langle\psi_i|\right)\rho\right) \\ &= \text{tr}(\mathbb{I}\rho) = \text{tr}(\rho) = 1. \end{aligned}$$

The second averaging property (A4) is more interesting. It ensures that SIC POVM measurements are informationally complete, i.e., we can reconstruct every density matrix ρ based on outcome probabilities. This may not directly be obvious from the display itself, but it follows from the following chain of arguments. If we weigh SIC projectors $|\psi_i\rangle\langle\psi_i|$ with the probability $\Pr[i|\rho]$ of observing this outcome, Eq. (A4) allows us to compute the resulting average. Let $\text{tr}_1(\cdot)$ denote the partial trace over the first of two qubits [$\text{tr}_1(A \otimes B) = \text{tr}(A)B$ and linearly extended to all of $\mathbb{H}_2^{\otimes 2}$]. Then,

$$\begin{aligned} \sum_{i=1}^4 \Pr[i|\rho] |\psi_i\rangle\langle\psi_i| &= \sum_{i=1}^4 \frac{1}{2} \langle\psi_i|\rho|\psi_i\rangle |\psi_i\rangle\langle\psi_i| \\ &= 2\text{tr}_1\left(\left(\frac{1}{4} \sum_{i=1}^4 (|\psi_i\rangle\langle\psi_i|)^{\otimes 2}\right)\rho \otimes \mathbb{I}\right) \\ &= \frac{1}{3} \text{tr}_1((\mathbb{I} \otimes \mathbb{I} + \mathbb{F})\mathbb{I} \otimes \rho) \\ &= \frac{1}{3} (\text{tr}(\rho)\mathbb{I} + \rho), \end{aligned} \quad (\text{A5})$$

where the last equation follows from the interplay between the partial trace and SWAP operator. The final expression is equivalent to applying a depolarizing channel with parameter $p = 1/3$ to the quantum state in question:

$$\mathcal{D}_{1/3}(\rho) = \frac{1}{3}\rho + \left(1 - \frac{1}{3}\right)\frac{\text{tr}(\rho)}{2}\mathbb{I} \in \mathbb{H}_2.$$

Viewed as a linear map on \mathbb{H}_2 , this channel has a uniquely defined inverse:

$$\mathcal{D}_{1/3}^{-1}(A) = 3A - \text{tr}(A)\mathbb{I} \quad \text{for all } A \in \mathbb{H}_2. \quad (\text{A6})$$

Although a linear map, this is not a physical operation. We can, however, use it in the classical postprocessing stage to

counterbalance the effect of averaging over SIC elements. Indeed, linearity and Eq. (A5) ensure that

$$\begin{aligned}
\sum_{i=1}^4 \Pr[i|\rho](3|\psi_i\rangle\langle\psi_i| - \mathbb{I}) &= \sum_{i=1}^4 \Pr[i|\rho]\mathcal{D}_{1/3}^{-1}(|\psi_i\rangle\langle\psi_i|) \\
&= \mathcal{D}_{1/3}^{-1}\left(\sum_{i=1}^4 \Pr[i|\rho]|\psi_i\rangle\langle\psi_i|\right) \\
&= \mathcal{D}_{1/3}^{-1}(\mathcal{D}_{1/3}(\rho)) \\
&= \rho. \tag{A7}
\end{aligned}$$

The left-hand side of this display features a linear combination involving SIC outcome probabilities $\Pr[i|\rho]$, while the right-hand side exactly reproduces the underlying state ρ . This equips us with a concrete state reconstruction formula—the so-called linear inversion estimator. But, at least at first sight, this formula is only useful if we have precise knowledge of the SIC outcome probabilities $\Pr[i|\rho]$. And, with current quantum technology, these probabilities must be estimated from repeatedly performing SIC POVM measurements on independent copies of ρ and approximating these probabilities by frequencies.

The classical shadow formalism provides an alternative perspective on this estimation process. Suppose that we perform a single SIC POVM measurement of an unknown quantum state ρ (single shot). Then, we obtain a random measurement outcome $\hat{i} \in \{1, 2, 3, 4\}$ with probability $\Pr[\hat{i}|\rho]$ each. Inspired by the left-hand side of Eq. (A7), we can use this outcome \hat{i} to construct a Monte Carlo estimator of ρ :

$$\hat{i} \mapsto \hat{\sigma} = (3|\psi_{\hat{i}}\rangle\langle\psi_{\hat{i}}| - \mathbb{I}) = \mathcal{D}_{1/3}^{-1}(|\psi_{\hat{i}}\rangle\langle\psi_{\hat{i}}|) \in \mathbb{H}_2.$$

This is a *random* 2×2 matrix that can assume four different forms—one for each possible outcome $\hat{i} \in \{1, 2, 3, 4\}$. It does exactly reproduce the underlying quantum state ρ in expectation over the observed single-shot outcome:

$$\mathbb{E}[\hat{\sigma}] = \sum_{i=1}^4 \Pr[i|\rho](3|\psi_i\rangle\langle\psi_i| - \mathbb{I}) = \rho. \tag{A8}$$

Here, we have used the definition of the expectation value (sum of outcomes weighed by their probability), as well as Eq. (A7). It is worthwhile to emphasize that each $\hat{\sigma}$ has the same eigenvalue structure: $\lambda_+ = 2$ and $\lambda_- = -1$. In turn, these random matrices all have unit trace [$\text{tr}(\hat{\sigma}) = 2 - 1 = 1$], but are unphysical in the sense that one eigenvalue is always negative. Equation (A8) represents a physical density matrix ρ as the expectation of four unphysical estimators. This desired expectation value can be approximated by empirically averaging M independently generated Monte Carlo estimators. Suppose that $\hat{\sigma}_1, \dots, \hat{\sigma}_M$ are

M independently and identically distributed (IID) Monte Carlo estimators. Then, their empirical average obeys

$$\hat{\rho} = \frac{1}{M} \sum_{m=1}^M \hat{\sigma}_m \xrightarrow{M \rightarrow \infty} \mathbb{E}[\hat{\sigma}] = \rho,$$

and the rate of convergence can be controlled with arguments from probability theory. This will be the content of the next two subsections.

b. Extension to multiqubit systems

The formalism and ideas presented above readily extend to quantum systems composed of multiple qubits. Let $\rho \in \mathbb{H}_2^{\otimes N} \simeq \mathbb{H}^{2^N}$ be an N -qubit density matrix. We can perform a single-qubit SIC POVM measurement on each of the N qubits. As in the single-qubit case, each such measurement yields one out of four possible outcomes. In total, a single-shot measurement produces a string (i_1, \dots, i_N) of outcomes. There are 4^N such outcomes and the probability of obtaining any one of them is given by

$$\begin{aligned}
\Pr[i_1, \dots, i_N|\rho] &= \text{tr}\left(\bigotimes_{n=1}^N \left(\frac{1}{2}|\psi_{i_n}\rangle\langle\psi_{i_n}|\right)\rho\right) \\
&= \frac{1}{2^N} \langle\psi_{i_1}, \dots, \psi_{i_N}|\rho|\psi_{i_1}, \dots, \psi_{i_N}\rangle \\
&\quad \text{for each } i_1, \dots, i_N \in \{1, 2, 3, 4\}. \tag{A9}
\end{aligned}$$

Here, we have introduced the shorthand notation $|\psi_{i_1}, \dots, \psi_{i_N}\rangle = \bigotimes_{n=1}^N |\psi_{i_n}\rangle \in (\mathbb{C}^2)^{\otimes N} \simeq \mathbb{C}^{2^N}$. These expressions are non-negative, because the N -qubit density matrix does not have negative eigenvalues. Equation (A3), applied to each qubit separately, moreover ensures proper normalization:

$$\begin{aligned}
\sum_{i_1, \dots, i_N=1}^4 \Pr[i_1, \dots, i_N|\rho] &= \text{tr}\left(\bigotimes_{n=1}^N \left(\sum_{i_n=1}^4 \frac{1}{2}|\psi_{i_n}\rangle\langle\psi_{i_n}|\right)\rho\right) \\
&= \text{tr}(\mathbb{I}^{\otimes N} \rho) = \text{tr}(\rho) = 1.
\end{aligned}$$

Again, the second averaging property is more interesting: for single-qubit density matrices $\tilde{\rho} \in \mathbb{H}_2$, we already know that Eq. (A4) implies that $\sum_{i=1}^4 \frac{1}{2} \langle\psi_i|\rho|\psi_i\rangle = \mathcal{D}_{1/3}(\rho)$, where $\mathcal{D}_{1/3} : \mathbb{H}_2 \rightarrow \mathbb{H}_2$ is the depolarizing channel from Eq. (A6). It is now easy to check that this equation extends to general Hermitian 2×2 matrices:

$$\sum_{i=1}^4 \frac{1}{4} \langle\psi_i|A|\psi_i\rangle|\psi_i\rangle\langle\psi_i| = \mathcal{D}_{1/3}(A) \quad \text{for all } A \in \mathbb{H}_2. \tag{A10}$$

We can use this observation to show that N single-qubit SIC POVM measurements are tomographically complete.

To achieve this, it is helpful to first decompose $\rho \in \mathbb{H}_2^{\otimes N}$ into a sum of elementary tensor products:

$$\rho = \sum_{W_1, \dots, W_N} r(W_1, \dots, W_N) \bigotimes_{n=1}^N W_n, \quad \text{where}$$

$$r(W_1, \dots, W_N) = \text{tr} \left(\bigotimes_{n=1}^N W_n \rho \right) \in [-1, 1]$$

and the summation goes over all four single-qubit Pauli matrices $W_n = \mathbb{I}, X, Y, Z$. Combine this with Eq. (A10) to compute

$$\begin{aligned} & \sum_{i_1, \dots, i_N=1}^4 \text{Pr}[i_1, \dots, i_N | \rho] \bigotimes_{n=1}^N |\psi_{i_n}\rangle\langle\psi_{i_n}| \\ &= \sum_{W_1, \dots, W_N} r(W_1, \dots, W_N) \bigotimes_{n=1}^N \\ & \quad \times \left(\frac{1}{2} \sum_{i_n=1}^4 \langle\psi_{i_n} | W_n | \psi_{i_n}\rangle |\psi_{i_n}\rangle\langle\psi_{i_n}| \right) \\ &= \sum_{W_1, \dots, W_N} r(W_1, \dots, W_N) \bigotimes_{n=1}^N \mathcal{D}_{1/3}(W_n) \\ &= \mathcal{D}_{1/3}^{\otimes N}(\rho), \end{aligned} \quad (\text{A11})$$

where the last equality follows from linearity of depolarizing channels. The final expression is equivalent to applying N -independent, single-qubit depolarizing channels to the N -qubit quantum state ρ . Viewed as a linear map on $\mathbb{H}_2^{\otimes N}$, this tensor product channel has a uniquely defined inverse $\mathcal{D}_{1/3}^{-\otimes N} : \mathbb{H}_2^{\otimes N} \rightarrow \mathbb{H}_2^{\otimes N}$. For elementary tensor products, this tensor product of inverse depolarizing channels factorizes nicely into tensor products. In particular,

$$\begin{aligned} \mathcal{D}_{1/3}^{-\otimes N} \left(\bigotimes_{n=1}^N |\psi_{i_n}\rangle\langle\psi_{i_n}| \right) &= \bigotimes_{n=1}^N \mathcal{D}_{1/3}^{-1}(|\psi_{i_n}\rangle\langle\psi_{i_n}|) \\ &= \bigotimes_{n=1}^N (3|\psi_{i_n}\rangle\langle\psi_{i_n}| - \mathbb{I}). \end{aligned}$$

Again, this is not a physical operation, because it produces matrices with negative eigenvalues. However, we can nonetheless use it in the classical postprocessing stage to counterbalance the N -qubit averaging effect encountered

in Eq. (A11):

$$\begin{aligned} & \sum_{i_1, \dots, i_N=1}^4 \text{Pr}[i_1, \dots, i_N | \rho] \bigotimes_{n=1}^N (3|\psi_{i_n}\rangle\langle\psi_{i_n}| - \mathbb{I}) \\ &= \mathcal{D}_{1/3}^{-\otimes N} \left(\sum_{i_1, \dots, i_N=1}^4 \text{Pr}[i_1, \dots, i_N | \rho] \bigotimes_{n=1}^N |\psi_{i_n}\rangle\langle\psi_{i_n}| \right) \\ &= \mathcal{D}_{1/3}^{-\otimes N} (\mathcal{D}_{1/3}^{\otimes N}(\rho)) \\ &= \rho. \end{aligned} \quad (\text{A12})$$

The left-hand side of this display features a linear combination involving single-qubit SIC outcome probabilities $\text{Pr}[i_1, \dots, i_N | \rho]$, while the right-hand side exactly reproduces the underlying N -qubit density matrix. This provides us with a concrete reconstruction formula for arbitrary N -qubit states. In fact, it is a natural and relatively straightforward extension of the single-qubit linear inversion estimator (A7) to N qubits.

As was the case for single qubits, the classical shadow formalism provides an alternative perspective on such a linear-inversion estimation process. Suppose that we perform N single-qubit SIC POVM measurements of an unknown N -qubit state ρ (single shot). Then, we obtain a random outcome string $(\hat{i}_1, \dots, \hat{i}_N) \in \{1, 2, 3, 4\}^{\times N}$ with probability $\text{Pr}[\hat{i}_1, \dots, \hat{i}_N | \rho]$ each. Inspired by the left-hand side of Eq. (A12), we can use this random outcome string to construct a Monte Carlo estimator of ρ :

$$(\hat{i}_1, \dots, \hat{i}_N) \mapsto \hat{\sigma} = \bigotimes_{n=1}^N (3|\psi_{\hat{i}_n}\rangle\langle\psi_{\hat{i}_n}| - \mathbb{I}) \in \mathbb{H}_2^{\otimes N}. \quad (\text{A13})$$

This is a random $2^N \times 2^N$ matrix that decomposes nicely into tensor products of single-qubit contributions. Each tensor factor contributes a matrix with eigenvalues $\lambda_{n,+} = +2$ and $\lambda_{n,-} = -1$. The eigenvalues of the tensor product $\hat{\sigma}$ then correspond to N -fold products of these two possible numbers. The largest eigenvalue is $+2^N$, the smallest is -2^{N-1} , so $\hat{\sigma}$ is a very unphysical random matrix. The randomness stems from an actual quantum measurement and depends on the underlying quantum state. This ensures that $\hat{\sigma}$ reproduces ρ in expectation, i.e.,

$$\mathbb{E}[\hat{\sigma}] = \sum_{i_1, \dots, i_N=1}^4 \text{Pr}[i_1, \dots, i_N | \rho] \bigotimes_{n=1}^N (3|\psi_{i_n}\rangle\langle\psi_{i_n}| - \mathbb{I}) = \rho, \quad (\text{A14})$$

courtesy of Eq. (A12). This expectation value can now be approximated by empirically averaging M independently generated Monte Carlo estimators, so-called *classical shadows*. Let $\hat{\sigma}_1, \dots, \hat{\sigma}_M \in \mathbb{H}_2^{\otimes N}$ be estimators generated from repeatedly preparing ρ and performing single-qubit

SIC POVM measurements (IID). Then, their empirical average obeys

$$\hat{\rho} := \frac{1}{M} \sum_{m=1}^M \hat{\sigma}_m \xrightarrow{M \rightarrow \infty} \rho, \quad (\text{A15})$$

in full analogy to the single-qubit case. As detailed in the next section, the rate of convergence will depend on the number of qubits N . The larger the space, the longer it takes for convergence to kick in, and this scaling can become unfavorable. It is therefore worthwhile to emphasize another distinct advantage of the tensor product structure of estimators (A13): marginalization to subsystem density operators is straightforward. Let ρ be an N -qubit state, but suppose that we are only interested in a subsystem $K \subset [N] = \{1, \dots, N\}$ composed of only $|K| \leq N$ qubits. Such a subsystem is fully described by the reduced $|K|$ -qubit density matrix $\rho_K = \text{tr}_{-K}(\rho) \in \mathbb{H}_2^{\otimes |K|}$ that results from tracing out all qubits not in K . This partial trace is a linear operation that plays nicely with the tensor product structure in Eq. (A13). Each tensor product factor has unit trace, which ensures that

$$\hat{\sigma}_K = \text{tr}_{-K}(\hat{\sigma}) = \bigotimes_{k \in K} (3|\psi_{i_k}\rangle\langle\psi_{i_k}| - \mathbb{I}) \in \mathbb{H}_2^{\otimes |K|} \quad \text{obeys}$$

$$\mathbb{E}[\hat{\sigma}_K] = \text{tr}_{-K}(\mathbb{E}[\hat{\sigma}]) = \text{tr}_{-K}(\rho) = \rho_K.$$

The object at the very left is a random $2^{|K|} \times 2^{|K|}$ matrix that can be generated from performing a complete N -qubit SIC POVM measurement to obtain outcomes $(\hat{i}_1, \dots, \hat{i}_N) \in \{1, \dots, 4\}^{\times N}$. Subsequently, we only use outcomes that correspond to qubits in K to directly construct an estimator for the subsystem density matrix ρ_K in question. This trick reduces the question of convergence to a problem that only involves $|K|$ qubits, not N . This can be highly advantageous if $|K| \ll N$. What is more, we can use the same N -qubit measurement outcome $(\hat{i}_1, \dots, \hat{i}_N)^{\times N}$ to construct estimators for multiple subsystems $K_1, \dots, K_L \subset [N]$ at once. This allows us to use the same N -qubit measurement statistics to estimate many subsystem properties in parallel.

10. Convergence for predicting linear observables

Suppose that we have access to M -independent Monte Carlo approximations $\hat{\sigma}_1, \dots, \hat{\sigma}_M$ of an unknown N -qubit state ρ . Each of them arises from measuring N single-qubit SIC POVMs on an independent copy of ρ . We can then use these approximations to estimate observable expectation values $\text{tr}(O\rho)$ with $O \in \mathbb{H}_2^{\otimes N}$:

$$\hat{\sigma} = \text{tr}(O\hat{\rho}) = \frac{1}{M} \sum_{m=1}^M \text{tr}(O\hat{\sigma}_m).$$

This is an empirical average of M IID random numbers $X_1, \dots, X_M \stackrel{\text{IID}}{\sim} X = \text{tr}(O\hat{\sigma})$ that converges to the true expectation $\mathbb{E}[X] = \text{tr}(O\mathbb{E}[\hat{\sigma}_m]) = \text{tr}(O\rho)$ as M increases. The rate of convergence is controlled by the variance $\text{Var}[X]$.

Lemma 1. Fix an N -qubit observable O and let $\hat{\sigma} \in \mathbb{H}_2^{\otimes N}$ be a (SIC POVM) classical shadow as defined in Eq. (A13). Then,

$$\text{Var}[\text{tr}(O\hat{\sigma})] \leq 3^N \text{tr}(O^2) \quad \text{for any underlying } N\text{-qubit state } \rho.$$

This bound is reminiscent of existing variance bounds for randomized single-qubit Pauli measurements [14], but slightly weaker [random Pauli basis measurements achieve $\text{Var}[\text{tr}(O\hat{\sigma})] \leq 2^N \text{tr}(O^2)$]. The proof exploits the 2-design property of SIC POVM measurements and will be supplied in Appendix A 14 below. For now, we use this variance bound to conclude strong convergence guarantees for observable estimation with SIC POVM measurements. The key ingredient is a strong tail bound for sums of IID random variables with known variance and bounded magnitude. The *Bernstein inequality* is a stronger version of the better known Hoeffding inequality; see, e.g., Ref. [69, Corollary 7.31] or [70, Theorem 2.8.4].

Let $X_1, \dots, X_M \stackrel{\text{IID}}{\sim} X \in \mathbb{R}$ be IID random variables with expectation $\mu = \mathbb{E}[X]$ and variance $\sigma^2 = \mathbb{E}[(X - \mu)^2]$ that also obey $|X_m| \leq R$ (with probability 1). Then, for $\epsilon > 0$,

$$\begin{aligned} \Pr\left[\left|\frac{1}{M} \sum_{m=1}^M X_m - \mu\right| \geq \epsilon\right] &\leq 2 \exp\left(-\frac{M\epsilon^2/2}{\sigma^2 + R\epsilon/3}\right) \\ &\leq \begin{cases} 2 \exp\left(-\frac{3}{8}M\epsilon^2/\sigma^2\right) & \text{if } \epsilon \leq \sigma^2/R, \\ 2 \exp\left(-\frac{3}{8}M\epsilon/R\right) & \text{if } \epsilon \geq \sigma^2/R. \end{cases} \end{aligned} \quad (\text{A16})$$

For the task at hand, we write $\hat{\sigma} = (1/M) \sum_{m=1}^M \text{tr}(O\hat{\sigma}_m) = (1/M) \sum_{m=1}^M X_m$, where $X_m = \text{tr}(O\hat{\sigma}_m) \stackrel{\text{IID}}{\sim} X = \text{tr}(O\hat{\sigma})$. For technical reasons, we also assume that $\text{tr}(O^2) \geq (5/9)^N > 2^{-N}$. Note that this is achieved by (i) physical observables [$\text{tr}(O^2) \geq \|O\|_\infty^2 = 1$], as well as quantum states [$O = \rho$ obeys $\text{tr}(\rho^2) \geq \text{tr}[(\mathbb{I}/2^N)^2] = 2^{-N}$]. The random variable X obeys

$$\begin{aligned} \mu &= \text{tr}(O\mathbb{E}[\hat{\sigma}]) = \text{tr}(O\rho), \quad \text{Var}[\text{tr}(O\hat{\sigma})] \leq 3^N \text{tr}(O^2) \\ &=: \sigma^2, \text{ and } |\text{tr}(O\hat{\sigma})| \leq 5^{N/2} \sqrt{\text{tr}(O^2)} \leq 3^N \text{tr}(O^2) =: R. \end{aligned}$$

The first equality is Eq. (A14), the second bound is Lemma 1, and the last bound is a consequence of the

Cauchy-Schwarz inequality: $|\text{tr}(O\hat{\sigma})| \leq \sqrt{\text{tr}(O^2)}\sqrt{\text{tr}(\hat{\sigma}^2)}$. Since classical shadows are tensor products of single-qubit blocks with eigenvalues $\lambda_+ = 2, \lambda_- = -1$, we can readily conclude that $\text{tr}(\hat{\sigma}^2) = (\lambda_+^2 + \lambda_-^2)^N = 5^N$. The final bound is rather loose and follows from the assumption that $\text{tr}(O^2) \geq (5/9)^N$. Inserting these bounds into Eq. (A16) now ensures that

$$\Pr[|\hat{\sigma} - \text{tr}(O\rho)| \geq \epsilon] = \Pr[|\hat{\sigma} - \text{tr}(O\rho)| \geq \epsilon] \leq 2 \exp \times \left(-\frac{3M\epsilon^2}{8 \times 3^N \text{tr}(O^2)} \right) \quad \text{for all } 0 < \epsilon \leq 1. \quad (\text{A17})$$

This is a bound on the probability of an ϵ -deviation (or more) that diminishes exponentially in the number of Monte Carlo samples (measurements) M . For a fixed confidence $\delta \in (0, 1)$, setting

$$M \geq \frac{8}{3} 3^N \text{tr}(O^2) \log(1/\delta)/\epsilon^2$$

ensures that

$$|\hat{\sigma} - \text{tr}(O\rho)| \leq \epsilon \quad \text{with probability (at least) } 1 - \delta.$$

Note that the required measurement budget M scales exponentially in the number N of involved qubits. At this point it is helpful to remember the marginalization property of classical shadows. Suppose that O is localized in the sense that it only affects a subsystem $\mathbf{K} \subset [N]$ composed of $|\mathbf{K}| \leq N$ qubits. Then, $\text{tr}(O\rho) = \text{tr}(O_{\mathbf{K}}\rho_{\mathbf{K}})$, where $\rho_{\mathbf{K}} = \text{tr}_{-\mathbf{K}}(\rho)$ is the reduced $|\mathbf{K}|$ -qubit density matrix and $O_{\mathbf{K}}$ is the nontrivial part of O . In turn, we can use appropriately marginalized classical shadows $\hat{\sigma}_{m,\mathbf{K}} = \text{tr}_{-\mathbf{K}}(\hat{\sigma}_m)$ to directly approximate this subsystem property:

$$\hat{\sigma} = \frac{1}{M} \sum_{m=1}^M \text{tr}(O\hat{\sigma}_m) = \frac{1}{M} \sum_{m=1}^M \text{tr}(O_{\mathbf{K}}\hat{\sigma}_{m,\mathbf{K}}).$$

The reformulation on the right-hand side now involves only the $|\mathbf{K}| < N$ relevant qubits. We can now redo the argument from above to obtain a measurement budget that only scales exponentially in $|\mathbf{K}|$:

$$\Pr[|\hat{\sigma} - \text{tr}(O\rho)| \geq \epsilon] \leq 2 \exp \left(-\frac{3M\epsilon^2}{8 \times 3^{|\mathbf{K}|} \text{tr}(O_{\mathbf{K}}^2)} \right) \quad \text{for } \epsilon \in (0, 1). \quad (\text{A18})$$

This refinement asserts that the probability of an ϵ -deviation for a single observable estimation diminishes exponentially in the number of measurements. We can use this exponential concentration to bound the probability of a single deviation among many. This allows us to use the same measurement data to predict many observables $\text{tr}(O_1\rho), \dots, \text{tr}(O_L\rho)$ in parallel.

Theorem 1. *Let O_1, \dots, O_L be N -qubit observables that are all localized to (at most) K qubits and fix $\epsilon, \delta \in (0, 1)$. Then,*

$$M \geq \frac{8}{3} 3^K \max_{1 \leq l \leq L} \text{tr}(O_{l,\mathbf{K}_l}^2) \log(2L/\delta)/\epsilon^2 \quad (\text{A19})$$

N -qubit SIC POVM measurements of an unknown state ρ are very likely to ϵ -approximate all observables simultaneously. More precisely, the resulting classical shadows $\hat{\sigma}_1, \dots, \hat{\sigma}_M$ obey

$$\max_{1 \leq l \leq L} \left| \frac{1}{M} \sum_{m=1}^M \text{tr}(O_l \hat{\sigma}_m) - \text{tr}(O_l \rho) \right| \leq \epsilon$$

with probability (at least) $1 - \delta$.

The convergence bound given in Eq. (5) of the main text is a simplified consequence of this result. Note that, by and large, physical observables are normalized in the operator norm: $\|O_l\|_\infty = \|O_{l,\mathbf{K}}\|_\infty = 1$. Equation (A19) features squared Hilbert-Schmidt norms $\|O_{l,\mathbf{K}}\|_2^2 = \text{tr}(O_{l,\mathbf{K}}^2)$ on the $|\mathbf{K}|$ -qubit subsystems in question. These Hilbert-Schmidt norms can be related to the operator norm, which is bounded:

$$\text{tr}(O_{l,\mathbf{K}}^2) = \|O_{l,\mathbf{K}}\|_2^2 \leq 2^{|\mathbf{K}|} \|O_{l,\mathbf{K}}\|_\infty^2 = 2^{|\mathbf{K}|} \quad \text{for all } 1 \leq l \leq L.$$

Here, we have used the fact that each $O_{l,\mathbf{K}}$ is a matrix of size (at most) $2^{|\mathbf{K}|} \cdot 2^{|\mathbf{K}|}$. This implies the bound $\max_{1 \leq l \leq L} \text{tr}(O_{l,\mathbf{K}}^2) \leq 2^{|\mathbf{K}|}$, which can be very pessimistic. Inserting it into Eq. (A19) yields Eq. (5) in the main text.

Proof of Theorem 1. A maximum deviation larger than ϵ occurs if at least one individual prediction $\hat{\sigma}_l$ is further than ϵ off from the actual target $\text{tr}(O_l\rho)$. The union bound, also known as Boole's inequality, tells us that the probability of such a maximum deviation is upper bounded by the sum of individual deviation probabilities. These, in turn, can be controlled via the tail bound from Eq. (A18):

$$\begin{aligned} \Pr \left[\max_{1 \leq l \leq L} \left| \frac{1}{M} \sum_{m=1}^M \text{tr}(O_l \hat{\sigma}_m) - \text{tr}(O_l \rho) \right| \geq \epsilon \right] \\ \leq \sum_{l=1}^L \Pr \left[\left| \frac{1}{M} \sum_{m=1}^M \text{tr}(O_l \hat{\sigma}_m) - \text{tr}(O_l \rho) \right| \geq \epsilon \right] \\ \leq \sum_{l=1}^L 2 \exp \left(-\frac{3M\epsilon^2}{8 \times 3^{|\mathbf{K}_l|} \text{tr}(O_{l,\mathbf{K}_l}^2)} \right). \end{aligned}$$

We see that each of these summands diminishes exponentially in the measurement budget M . The right-hand side of Eq. (A19) ensures that each term contributes at most δ/L to this sum. Since there are L summands in total, we conclude that $\Pr[\max_{1 \leq l \leq L} |(1/M) \sum_{m=1}^M \text{tr}(O_l \hat{\sigma}_m) -$

$\text{tr}(O\rho) \geq \epsilon] \leq \delta$. This is equivalent to the advertised display. ■

11. Convergence for predicting (subsystem) purities

Classical shadows can also be used to predict nonlinear quantum state properties; see, e.g., Refs. [14,32]. A prototypical example is the purity of an N -qubit density matrix ρ :

$$p(\rho) = \text{tr}(\rho^2) = \text{tr}(\rho\rho) \in (0, 1].$$

The purity equals one if and only if ρ describes a pure quantum state $|\phi\rangle\langle\phi|$. Conversely, it achieves its minimum value for the maximally mixed state: $\rho = (\frac{1}{2}\mathbb{I})^{\otimes N}$ achieves $p(\rho) = 1/2^N \ll 1$. We now describe how to obtain a purity estimator based on classical shadows $\hat{\sigma}_1, \dots, \hat{\sigma}_M$ that arise from measuring N single-qubit SIC POVMs on (independent copies of) ρ . By construction, each $\hat{\sigma}_m$ is a Monte Carlo estimator of ρ . Indeed, Eq. (A14) asserts that $\mathbb{E}[\hat{\sigma}_m] = \rho$ for all $1 \leq m \leq M$. What is more, distinct Monte Carlo estimators $\hat{\sigma}_m$ and $\hat{\sigma}_{m'}$ with $m \neq m'$ are statistically independent. The expectation over statistically independent random matrices factorizes. This ensures that the trace of the product of any two distinct classical shadows reproduces the purity in expectation:

$$\begin{aligned} \text{tr}(\hat{\sigma}_m \hat{\sigma}_{m'}) &\text{ obeys } \mathbb{E}[\text{tr}(\hat{\sigma}_m \hat{\sigma}_{m'})] = \text{tr}(\mathbb{E}[\hat{\sigma}_m] \mathbb{E}[\hat{\sigma}_{m'}]) \\ &= \text{tr}(\rho\rho) = p(\rho) \end{aligned}$$

whenever $m \neq m'$. To boost convergence to this desired expectation, we can form the empirical average over all distinct pairs:

$$\hat{p} = \frac{1}{M(M-1)} \sum_{m \neq m'} \text{tr}(\hat{\sigma}_m \hat{\sigma}_{m'}) = \binom{M}{2}^{-1} \sum_{m < m'} \text{tr}(\hat{\sigma}_m \hat{\sigma}_{m'}). \quad (\text{A20})$$

This formula describes an empirical average of $\binom{M}{2}$ random variables with the correct expectation value $p(\rho)$. This, in turn, ensures that $\mathbb{E}[\hat{p}] = p_2(\rho)$. However, in contrast to before, the individual random variables are not necessarily statistically independent. The first two terms $\text{tr}(\hat{\sigma}_1 \hat{\sigma}_2)$ and $\text{tr}(\hat{\sigma}_1 \hat{\sigma}_3)$, for instance, both depend on $\hat{\sigma}_1$. This prevents us from reusing exponential concentration inequalities, like the Bernstein inequality, to establish rapid convergence to this desired expectation value. More general, albeit weaker, concentration arguments still apply. Chebyshev's inequality, for instance, implies that

$$\begin{aligned} \Pr[|\hat{p} - p(\rho)| \geq \epsilon] &= \Pr[|\hat{p} - \mathbb{E}[\hat{p}]| \geq \epsilon] \leq \frac{\text{Var}[\hat{p}]}{\epsilon^2} \\ &\text{for any } \epsilon > 0. \end{aligned} \quad (\text{A21})$$

In words, the probability of an ϵ -deviation (or larger) is bounded by the variance $\text{Var}[\hat{p}]$ of our estimator divided

by ϵ^2 . This variance can be decomposed into individual contributions:

$$\begin{aligned} \text{Var}[\hat{p}] &= \mathbb{E}[\hat{p}^2] - \mathbb{E}[\hat{p}]^2 \\ &= \mathbb{E}[\hat{p}^2] - \text{tr}(\rho^2)^2 \\ &= \binom{M}{2}^{-2} \sum_{m_1 < m'_1} \sum_{m_2 < m'_2} (\text{tr}(\hat{\sigma}_{m_1} \hat{\sigma}_{m'_1}) \text{tr}(\hat{\sigma}_{m_2} \hat{\sigma}_{m'_2}) \\ &\quad - \text{tr}(\rho^2) \text{tr}(\rho^2)). \end{aligned}$$

Now, note that $\mathbb{E}[\hat{\sigma}_{m_1}] = \mathbb{E}[\hat{\sigma}_{m'_1}] = \mathbb{E}[\hat{\sigma}_{m_2}] = \mathbb{E}[\hat{\sigma}_{m'_2}] = \rho$ implies that these summands vanish unless either two or all four summation indices coincide. A careful case-by-case analysis yields

$$\begin{aligned} \text{Var}[\hat{p}] &= \binom{M}{2}^{-1} 2(M-2) \text{Var}[\text{tr}(\rho \hat{\sigma})] \\ &\quad + \binom{M}{2}^{-1} \text{Var}[\text{tr}(\hat{\sigma} \hat{\sigma}')] \\ &= \frac{4(M-2)}{M(M-1)} \text{Var}[\text{tr}(\rho \hat{\sigma})] \\ &\quad + \frac{2}{M(M-1)} \text{Var}[\text{tr}(\hat{\sigma} \hat{\sigma}')], \end{aligned} \quad (\text{A22})$$

and we refer the reader to [32, Supplemental Material] for details. Here, ρ is the underlying state and $\hat{\sigma}, \hat{\sigma}'$ denote independent instances of a classical shadow approximation. This reformulation contains two variance terms that depend on one (linear contribution) and two independent classical shadows (quadratic contribution), respectively. We can use Lemma 1 to control the first term. Set $O = \rho$ to conclude that

$$\text{Var}[\text{tr}(\rho \hat{\rho})] \leq 3^N \text{tr}(\rho^2) \leq 3^N, \quad (\text{A23})$$

because $\text{tr}(\rho^2) \leq 1$ for any underlying quantum state. Bounding the quadratic variance term requires more work. The following statement is a consequence of the geometric structure of SIC POVM measurements and substitutes existing arguments that rely on 3-design properties that do not apply here.

Lemma 2. *Let $\hat{\sigma}, \hat{\sigma}'$ be independent classical shadows of an underlying N -qubit state ρ . Then,*

$$\text{Var}[\text{tr}(\hat{\sigma} \hat{\sigma}')] = \mathbb{E}[\text{tr}(\hat{\sigma} \hat{\sigma}')^2] - \mathbb{E}[\text{tr}(\hat{\sigma} \hat{\sigma}')]^2 \leq 9^N.$$

We provide a detailed argument in Appendix A15 below. For now, we insert both bounds into Eq. (A22) to

obtain

$$\text{Var}[\hat{p}] \leq \frac{4(M-2)}{M(M-1)}3^N + \frac{2}{M(M-1)}9^N \leq \frac{4 \times 3^N}{M-1} + \left(\frac{\sqrt{2} \times 3^N}{M-1}\right)^2.$$

This variance bound diminishes as M increases. For fixed $\epsilon \in (0, 1)$ and $\delta \in (0, 1)$, a measurement budget of $M \geq (5 \times 3^N + 1)/(\delta\epsilon^2)$ ensures that $\text{Var}[\hat{p}] \leq \frac{4}{5}\epsilon^2\delta + \frac{2}{25}\epsilon^4\delta^2 < \epsilon^2\delta$. We can insert this implication into the Chebyshev bound (A21) to obtain a rigorous convergence guarantee for purity estimation:

$$M \geq (5 \times 3^N + 1)/(\delta\epsilon^2) \text{ ensures that } \Pr[|\hat{p} - \text{tr}(\rho^2)| \geq \epsilon] \leq \delta.$$

In words, with probability (at least) $1 - \delta$, the purity estimator \hat{p} is ϵ -close to the true purity. Again, the required measurement budget M scales exponentially in the number of qubits involved. For global purities, this exponentially increasing measurement demand quickly becomes prohibitively expensive—a situation that cannot be avoided due to recent fundamental lower bounds [71]. However, once more, the situation changes if we consider *subsystem purities* instead. Let $\mathbf{K} \subseteq [N]$ be a subsystem composed of $|\mathbf{K}|$ qubits. The associated density matrix is $\rho_{\mathbf{K}} = \text{tr}_{-\mathbf{K}}(\rho)$ and we can estimate it by averaging appropriately marginalized classical shadows:

$$\hat{p}_{\mathbf{K}} = \binom{M}{2} \sum_{m \neq m'} \text{tr}(\text{tr}_{-\mathbf{K}}(\hat{\sigma}_m)\text{tr}_{-\mathbf{K}}(\hat{\sigma}'_{m'})) \text{ obeys } \mathbb{E}[\hat{p}_{\mathbf{K}}] = \text{tr}(\rho_{\mathbf{K}}^2) = p(\rho_{\mathbf{K}}). \tag{A24}$$

Importantly, this estimation process now depends only on the $|\mathbf{K}| < N$ qubits involved, such that

$$M \geq (5 \times 3^{|\mathbf{K}|} + 1)/(\delta\epsilon^2) \text{ ensures that } \Pr[|\hat{p}_{\mathbf{K}} - p(\rho_{\mathbf{K}})| \geq \epsilon] \leq \delta.$$

This scaling is much more favorable, especially for small subsystems ($|\mathbf{K}| \ll N$). Similar to linear observable estimation, we can use this assertion to predict many subsystem purities based on the same classical shadows. A union bound argument, similar to the proof of Theorem 1 above, readily implies the following statement.

Theorem 2. *Suppose that we are interested in predicting L subsystem purities $p(\rho_{\mathbf{K}_l})$ of an unknown N -qubit state ρ . Let $K = \max_{1 \leq l \leq L} |\mathbf{K}_l|$ be the largest subsystem*

size involved and set $\epsilon, \delta \in (0, 1)$. Then,

$$M \geq 6L3^K/(\epsilon^2\delta) \tag{A25}$$

N -qubit SIC POVM measurements on (independent copies of) ρ are likely to ϵ -approximate all subsystem purities simultaneously. More precisely, the resulting subsystem purity estimators $\hat{p}_{\mathbf{K}}$ defined in Eq. (A24) obey

$$\max_{1 \leq l \leq L} |\hat{p}_{\mathbf{K}_l} - \text{tr}(\rho_{\mathbf{K}_l}^2)| \leq \epsilon \text{ with probability (at least) } 1 - \delta.$$

The dependence on subsystem size K and accuracy ϵ is virtually identical to convergence guarantees for linear observable prediction; see Theorem 1. However, the dependence on the number of subsystem purities L and the inverse confidence $1/\delta$ now enter linearly, not logarithmically. This is a consequence of the fact that the individual contributions to $\hat{p}_{\mathbf{K}_l}$ are not statistically independent. In turn, we have to resort to Chebyshev’s inequality instead of stronger exponential tail bounds like the Bernstein inequality.

It is possible to obtain a scaling proportional to $\log(2L/\delta)$ by using a more sophisticated estimation procedure known as *median of means estimation*; see, e.g., Ref. [14] for details. Practical tests with real data do, however, suggest that median of means estimation actually reduces the approximation quality overall [32]. This is not a contradiction, because statements like Theorem 2 are conservative mathematical statements about the worst-case rate of convergence. In practical applications, convergence can—and usually does—set in much earlier.

12. Convergence for higher-order polynomials and entanglement detection (outlook)

Quadratic estimation with classical shadows readily extends to higher-order polynomials. Such higher-order polynomials can be used, for instance, to probe entanglement in mixed states [32,33]. This is important because quadratic entanglement conditions—like subsystem Rényi entropies (purities)—only apply to global states that are reasonably pure [$\text{tr}(\rho^2) \approx 1$]. To see this, consider the maximally mixed state $\tau = (\frac{1}{2}\mathbb{I})^{\otimes N}$ on N qubits. This state is certainly not entangled, but nonetheless

$$\begin{aligned} R_2(\tau) &= -\log_2[\text{tr}(\rho_{\mathbf{K}}^2)] \\ &= -\log_2[\text{tr}((\frac{1}{2}\mathbb{I})^{\otimes |\mathbf{K}|})] \\ &= -\log_2(2^{-|\mathbf{K}|}) \\ &= |\mathbf{K}| \text{ for all subsystems } \mathbf{K} \subseteq N. \end{aligned}$$

In words, second Rényi entropy is maximal for all subsystems simultaneously. This, however, is not a consequence of entanglement, but a trivial consequence of the fact that the state is very (maximally) mixed.

Fortunately, there exist entanglement criteria that extend to (very) mixed states. Chief among them is the positive partial transpose (PPT) criterion [72–74]. Let ρ be an N -qubit quantum state, and let $(\mathbf{A}, \bar{\mathbf{A}})$ be a bipartition of the qubits into two disjoint sets. Then, ρ is entangled (across the bipartition) if the partial transpose density matrix is not positive semidefinite (i.e., it has negative eigenvalues):

$$\rho^{T\mathbf{A}} \not\geq 0 \quad \text{implies that } \rho \text{ is entangled} \\ \text{across the bipartition.} \quad (\text{A26})$$

The partial transpose is defined by transposing tensor factors belonging to subsystem \mathbf{A} , i.e., $(\rho_1 \otimes \cdots \otimes \rho_N)^{T\mathbf{A}} = \bigotimes_{a \in \mathbf{A}} \rho_a^T \bigotimes_{\bar{a} \in \bar{\mathbf{A}}} \rho_{\bar{a}}$ and linearly extended to all N -qubit density matrices. A quick sanity check confirms that the maximally mixed state does not pass the PPT condition ($\mathbb{I}^T = \mathbb{I}$):

$$\tau^{T\mathbf{A}} = \bigotimes_{a \in \mathbf{A}} (1/2\mathbb{I})^T \bigotimes_{\bar{a} \in \bar{\mathbf{A}}} (1/2\mathbb{I})^T = \bigotimes_{n \in [N]} (1/2\mathbb{I}) = \tau \geq 0.$$

Very entangled states, like the two-qubit Bell state $|\Omega\rangle = 1/\sqrt{2}(|00\rangle + |11\rangle)$ do, in contrast, have partial transposes with negative eigenvalues:

$$(|\Omega\rangle\langle\Omega|)^{T_1} = (|\Omega\rangle\langle\Omega|)^{T_2} = 1/\sqrt{2}\mathbb{F} \in \mathbb{H}_2^{\otimes 2},$$

where $\mathbb{F}|x\rangle \otimes |y\rangle = |y\rangle \otimes |x\rangle$ denotes the SWAP operator that has one negative eigenvalue [$\lambda_{\min}(\mathbb{F}) = -1$]. We call a state ρ with $\rho^{T\mathbf{A}} \not\geq 0$ a *PPT-entangled state* [with respect to the bipartition $(\mathbf{A}, \bar{\mathbf{A}})$]. The PPT condition is a sufficient, but not necessary, condition for entanglement. It is known that there exist states that are entangled, but nonetheless obey $\rho^{T\mathbf{A}} \geq 0$ [67]. So, it is fruitful to view the PPT criterion as a one-sided test for entanglement: if $\rho^{T\mathbf{A}} \not\geq 0$, we can be sure that the state is entangled. But, $\rho^{T\mathbf{A}} \geq 0$ does not necessarily imply that the state is not entangled (i.e., separable).

The PPT criterion (A26) is conceptually appealing, but it does require full and accurate knowledge of the density matrix ρ . This, in turn, typically requires full state tomography that quickly becomes prohibitively expensive. It is, however, possible to test consequences of $\rho^{T\mathbf{A}} \geq 0$ by comparing moments of the partially transposed density matrix. The simplest consistency check is the so-called p_3 criterion [32]:

$$\rho^{T\mathbf{A}} \geq 0 \quad \implies \quad \text{tr}((\rho^{T\mathbf{A}})^3) \geq \text{tr}((\rho^{T\mathbf{A}})^2)^2 = \text{tr}(\rho^2)^2. \quad (\text{A27})$$

The final simplification follows from the fact that partial transposition preserves the purity. The contrapositive of this implication serves as a (one-sided) test for entangle-

ment: if $\text{tr}((\rho^{T\mathbf{A}})^3) < \text{tr}(\rho^2)^2$ [for some bipartition $(\mathbf{A}, \bar{\mathbf{A}})$] then the underlying state must be PPT entangled (across this bipartition).

Classical shadows can be used to directly estimate the trace moments involved in this test. Indeed, $\text{tr}(\rho^2)$ is just the purity, while $\text{tr}((\rho^{T\mathbf{A}})^3)$ can be rewritten as a linear function on three copies of the underlying state:

$$\text{tr}((\rho^{T\mathbf{A}})^3) = \text{tr}(O\rho \otimes \rho \otimes \rho).$$

We refer the reader to Ref. [32, Eq. (4)] for a precise reformulation. Subsequently, we can approximate this function by averaging over triples of distinct (and therefore independent) classical shadows:

$$\text{tr}(O\rho \otimes \rho \otimes \rho) \approx \frac{1}{6} \binom{M}{3}^{-1} \sum_{m \neq m' \neq m''} \text{tr}(O\hat{\sigma}_m \otimes \hat{\sigma}_{m'} \otimes \hat{\sigma}_{m''}).$$

The convergence analysis from above can, in principle, be extended to this form of cubic approximation. For randomized Pauli basis measurements, this has been done in the supplemental material of Ref. [32]. We leave a parallel treatment of cubic estimation with SIC POVM shadows for future work. The experimental and numerical results from the present work indicate that a SIC-POVM-based approach is expected to be both cheaper and easier than existing approaches based on Pauli basis measurements.

Finally, we point out that Ref. [33] extended the intuition behind the p_3 criterion (A27) to a complete family of polynomial consistency checks that compare polynomials of degree d with polynomials of degree $(d-1)$ and lower. This produces a hierarchy of in total $d_{\max} = 2^N$ consistency checks that is complete in the sense that a state ρ passes all of them if and only if $\rho^{T\mathbf{A}} \geq 0$. Although polynomial estimation with classical shadows becomes more and more challenging as the degree d increases, the lower levels of this hierarchy may still be attainable with (comparatively) modest experimental and postprocessing effort.

13. Technical auxiliary results

14. Variance bounds for observable estimation

Here we prove Lemma 1, which we recall for convenience.

Lemma 3. *Fix an N -qubit observable O and let $\hat{\sigma} \in \mathbb{H}_2^{\otimes N}$ be a (SIC POVM) classical shadow as defined in Eq. (A13). Then,*

$$\text{Var}[\text{tr}(O\hat{\sigma})] \leq 3^N \text{tr}(O^2) \\ \text{for any underlying } N\text{-qubit state } \rho.$$

Proof. The classical shadow $\hat{\sigma}$ is constructed from performing single-qubit SIC POVM measurements on an underlying N -qubit quantum state ρ . Recall from Eq. (A9) that each of the 4^N possible outcome strings $i_1, \dots, i_N \in \{1, 2, 3, 4\}$ occurs with probability

$$\Pr[i_1 \dots i_N | \rho] = 2^{-N} \langle \psi_{i_1}, \dots, \psi_{i_N} | \rho | \psi_{i_1}, \dots, \psi_{i_N} \rangle \leq 2^{-N}.$$

In words, the probability of any particular outcome occurring is bounded by 2^{-N} . This allows us to bound the dominating part of the variance by

$$\begin{aligned} \mathbb{E}[\text{tr}(O\hat{\sigma})^2] &= \sum_{i_1, \dots, i_N=1}^4 \Pr[i_1 \dots i_N | \rho] \text{tr}(O(3|\psi_{i_1}\rangle\langle\psi_{i_1}| - \mathbb{I}) \otimes \dots \otimes (3|\psi_{i_N}\rangle\langle\psi_{i_N}| - \mathbb{I}))^2 \\ &\leq 2^{-N} \sum_{i_1, \dots, i_N=1}^4 \text{tr}(O(3|\psi_{i_1}\rangle\langle\psi_{i_1}| - \mathbb{I}) \otimes \dots \otimes (3|\psi_{i_N}\rangle\langle\psi_{i_N}| - \mathbb{I}))^2. \end{aligned}$$

Next, we expand the N -qubit observable O in terms of tensor products of single-qubit Pauli matrices $W_1, \dots, W_N \in \{\mathbb{I}, X, Y, Z\}$:

$$O = \sum_{W_1, \dots, W_N} o(W_1, \dots, W_N) W_1 \otimes \dots \otimes W_N$$

with $o(W_1, \dots, W_N) = 2^{-N} \text{tr}(W_1 \otimes \dots \otimes W_N O) \in \mathbb{R}$.

Such a decomposition into tensor products allows us to factorize the above bound into a product of single-qubit contributions:

$$\begin{aligned} \mathbb{E}[\text{tr}(O\hat{\sigma})^2] &\leq 2^{-N} \sum_{i_1, \dots, i_N=1}^4 \left(\sum_{W_1, \dots, W_N} o(W_1, \dots, W_N) \text{tr}((3|\psi_{i_1}\rangle\langle\psi_{i_1}| - \mathbb{I})W_1) \times \dots \times \text{tr}((3|\psi_{i_N}\rangle\langle\psi_{i_N}| - \mathbb{I})W_N) \right)^2 \\ &= \sum_{V_1, \dots, V_N} \sum_{W_1, \dots, W_N} o(V_1, \dots, V_N) o(W_1, \dots, W_N) \\ &\quad \times \prod_{n=1}^N \underbrace{\left(\frac{1}{2} \sum_{i_n=1}^4 \text{tr}((3|\psi_{i_n}\rangle\langle\psi_{i_n}| - \mathbb{I})V_n) \text{tr}((3|\psi_{i_n}\rangle\langle\psi_{i_n}| - \mathbb{I})W_n) \right)}_{f(V_n, W_n)}. \end{aligned}$$

These single-qubit averages can be computed individually. Using the 1-design property (A3) [i.e., $\frac{1}{2} \sum_{i_n=1}^4 \frac{1}{2} \langle \psi_{i_n} | A | \psi_{i_n} \rangle = \text{tr}(A)$], we obtain

$$\begin{aligned} f(V_n, W_n) &= \frac{1}{2} \sum_{i_n=1}^4 \text{tr}((3|\psi_{i_n}\rangle\langle\psi_{i_n}| - \mathbb{I})V_n) \text{tr}((3|\psi_{i_n}\rangle\langle\psi_{i_n}| - \mathbb{I})W_n) \\ &= \frac{9}{2} \sum_{i_n=1}^4 \langle \psi_{i_n} | V_n | \psi_{i_n} \rangle \langle \psi_{i_n} | W_n | \psi_{i_n} \rangle - 3 \text{tr}(V_n) \frac{1}{2} \sum_{i_n} \langle \psi_{i_n} | W_n | \psi_{i_n} \rangle \\ &\quad - 3 \text{tr}(W_n) \frac{1}{2} \sum_{i_n=1}^4 \langle \psi_{i_n} | V_n | \psi_{i_n} \rangle + \frac{1}{2} \sum_{i_n=1}^4 \text{tr}(V_n) \text{tr}(W_n) \\ &= \frac{9}{2} \sum_{i_n=1}^4 \langle \psi_{i_n} | V_k | \psi_{i_n} \rangle \langle \psi_{i_n} | W_n | \psi_{i_n} \rangle - 4 \text{tr}(V_n) \text{tr}(W_n). \end{aligned}$$

Next, we use the 2-design property (A4) of single-qubit SIC POVMs to obtain

$$\begin{aligned} f(V_n, W_n) &= \frac{9}{2} \sum_{i_n=1}^4 \langle \psi_{i_n} | V_n | \psi_{i_n} \rangle \langle \psi_{i_n} | W_n | \psi_{i_n} \rangle - 4 \text{tr}(V_n) \text{tr}(W_n) \\ &= 3(\text{tr}(V_n W_n) + \text{tr}(V_n) \text{tr}(W_n)) - 4 \text{tr}(V_n) \text{tr}(W_n) \\ &= 3 \text{tr}(V_n W_n) - \text{tr}(V_n) \text{tr}(W_n). \end{aligned}$$

For Pauli matrices V_n, W_n , this expression vanishes whenever $V_n \neq W_n$. It equals 2 if $V_n = W_n = \mathbb{I}$ and 6 if $V_n = W_n \neq \mathbb{I}$. That is,

$$f(V_n, W_n) = 2\delta(V_n, W_n)3^{1-\delta(V_n, \mathbb{I})}.$$

Inserting this closed-form expression into the original expression yields

$$\begin{aligned} \mathbb{E}[\text{tr}(O\hat{\sigma})^2] &\leq \sum_{V_1, \dots, V_N} \sum_{W_1, \dots, W_N} o(V_1, \dots, V_N) o(W_1, \dots, W_N) \prod_{n=1}^N 2\delta(V_n, W_n) 3^{1-\delta(V_n, \mathbb{I})} \\ &\leq 3^N 2^N \sum_{W_1, \dots, W_N} o(W_1, \dots, W_N)^2 \\ &= 3^N \sum_{W_1, \dots, W_N} 2^{-N} \text{tr}(W_1 \otimes \dots \otimes W_N O)^2 \\ &= 3^N \text{tr}(O^2), \end{aligned}$$

where the last equation follows from the fact that normalized N -qubit Pauli matrices form an orthonormal basis of $\mathbb{H}_2^{\otimes N}$ with respect to the Hilbert-Schmidt inner product (Parseval's identity). This completes the proof. \blacksquare

15. Variance bounds for purity estimation

Proposition 1 (Purity variance bound). *Let $\hat{\sigma}, \hat{\sigma}' \in \mathbb{H}_2^{\otimes N}$ be two independent classical shadows that arise from performing single-qubit SIC POVM measurements on a K -qubit state ρ . Then,*

$$\text{Var}[\text{tr}(\hat{\sigma}\hat{\sigma}')] = \mathbb{E}[\text{tr}(\hat{\sigma}\hat{\sigma}')^2] - \mathbb{E}[\text{tr}(\hat{\sigma}\hat{\sigma}')]^2 \leq 9^N.$$

The proof strategy behind this statement differs from existing arguments in the literature, most notably Refs. [14,32,33]. These use the fact that Pauli basis measurements form a 3-design, a structural property that does not apply to SIC POVMs. The key idea behind this new proof technique is to note that trace inner products of SIC POVM shadows can only assume very discrete values. Recall that

$$\hat{\sigma} = \bigotimes_{n=1}^N (3|\psi_{i_n}\rangle\langle\psi_{i_n}| - \mathbb{I}) \quad \text{and} \quad \hat{\sigma}' = \bigotimes_{n=1}^N (3|\psi_{j_n}\rangle\langle\psi_{j_n}| - \mathbb{I}),$$

where $i_1, \dots, i_N \in \{1, 2, 3, 4\}$ and $j_1, \dots, j_N \in \{1, 2, 3, 4\}$ record the outcomes of each single-qubit SIC POVM measurement. This tensor product structure then implies that

$$\text{tr}(\hat{\sigma}\hat{\sigma}') = \prod_{n=1}^N \text{tr}((3|\psi_{i_n}\rangle\langle\psi_{i_n}| - \mathbb{I})(3|\psi_{j_n}\rangle\langle\psi_{j_n}| - \mathbb{I})) = \prod_{n=1}^N (9|\langle\psi_{i_n}|\psi_{j_n}\rangle|^2 - 4),$$

and, because $|\psi_{i_n}\rangle, |\psi_{j_n}\rangle \in \mathbb{C}^2$ are SIC vectors, each contribution can only assume one of two discrete values:

$$9|\langle\psi_{i_n}|\psi_{j_n}\rangle|^2 - 4 = \begin{cases} +5 & \text{if } i_n = j_n, \\ -1 & \text{else if } i_n \neq j_n. \end{cases} \quad (\text{A28})$$

So, the magnitude of $\text{tr}(\hat{\sigma}\hat{\sigma}')$ scales exponentially in the number of coincidental measurement outcomes ($i_n = j_n$). This observation can be used to control the variance of this trace inner product. We first illustrate this for $N = 2$ qubits,

which is enough to convey the main gist. The proof of Proposition 1 is then a straightforward, yet somewhat technical, generalization to an arbitrary number of qubits.

In the two-qubit case, $\text{tr}(\hat{\sigma}\hat{\sigma}')^2$ can only assume three values: 25^2 if all single-qubit outcomes coincide, 25 if exactly one single-qubit outcome coincides, and 1 if no outcomes coincide. That is,

$$\begin{aligned} \text{tr}(\hat{\sigma}\hat{\sigma}')^2 &= 25^2\mathbf{1}\{i_1 = j_1 \wedge i_2 = j_2\} + 25\mathbf{1}\{i_1 = j_1 \wedge i_2 \neq j_2\} \\ &\quad + 25\mathbf{1}\{i_1 \neq j_1 \wedge i_2 = j_2\} + \mathbf{1}\{i_1 \neq j_1 \wedge i_2 \neq j_2\}, \end{aligned} \quad (\text{A29})$$

where $\mathbf{1}\{E\}$ denotes the indicator function of the event E . Next, we reexpress these indicator functions in terms of simpler ones:

$$\begin{aligned} \mathbf{1}\{i_1 = j_1 \wedge i_2 \neq j_2\} &= \mathbf{1}\{i_1 = j_1\} - \mathbf{1}\{i_1 = j_1 \wedge i_2 = j_2\}, \\ \mathbf{1}\{i_1 \neq j_1 \wedge i_2 = j_2\} &= \mathbf{1}\{i_2 = j_2\} - \mathbf{1}\{i_1 = j_1 \wedge i_2 = j_2\}, \\ \mathbf{1}\{i_1 \neq j_1 \wedge i_2 \neq j_2\} &= 1 - \mathbf{1}\{i_1 = j_1\} - \mathbf{1}\{i_2 = j_2\} + \mathbf{1}\{i_1 = j_1 \wedge i_2 = j_2\}. \end{aligned}$$

Inserting these reformulations into Eq. (A29) and rearranging terms yields

$$\begin{aligned} \text{tr}(\hat{\sigma}\hat{\sigma}')^2 &= (25^2 - 2 \times 25 + 1)\mathbf{1}\{i_1 = j_1 \wedge i_2 = j_2\} + (25 - 1)\mathbf{1}\{i_1 = j_1\} + (25 - 1)\mathbf{1}\{i_2 = j_2\} + 1 \\ &= (25 - 1)^2\mathbf{1}\{i_1 = j_1 \wedge i_2 = j_2\} + (25 - 1)\mathbf{1}\{i_1 = j_1\} + (25 - 1)\mathbf{1}\{i_2 = j_2\} + 1 \\ &= 8^2 \times 3^2\mathbf{1}\{i_1 = j_1 \wedge i_2 = j_2\} + 8 \times 3\mathbf{1}\{i_1 = j_1\} + 8 \times 3\mathbf{1}\{i_2 = j_2\} + 1, \end{aligned}$$

where we have used the fact that $(25 - 1) = 24 = 8 \times 3$. Now, we are ready to take expectation values. Recall that taking the expectation of an indicator function produces the probability of the associated event:

$$\mathbb{E}[\text{tr}(\hat{\sigma}\hat{\sigma}')^2] = 8^2 \times 3^2\text{Pr}[i_1 = j_1 \wedge i_2 = j_2] + 8 \times 3\text{Pr}[i_1 = j_1] + 8 \times 3\text{Pr}[i_2 = j_2] + 1. \quad (\text{A30})$$

These probabilities for coincidental measurement outcomes can be computed explicitly. This is the content of the following auxiliary result.

Lemma 4. *Suppose that we perform two N -qubit SIC POVM measurements on (distinct copies of) a quantum state ρ and let $\mathbf{K} \subseteq [N] = \{1, \dots, N\}$ be a subset of $K = |\mathbf{K}|$ qubits. Then, the probability that the obtained measurement outcomes are equal ($i_k = j_k$) for all $k \in \mathbf{K}$ obeys*

$$\text{Pr}\left[\bigwedge_{k \in \mathbf{K}} \{i_k = j_k\}\right] = \text{tr}(\rho_{\mathbf{K}} \mathcal{D}_{1/3}^{\otimes K}(\rho_{\mathbf{K}})) \leq 3^{-K},$$

where $\rho_{\mathbf{K}} = \text{tr}_{-\mathbf{K}}(\rho)$ is the reduced density matrix supported on the relevant qubit subset and each $\mathcal{D}_{1/3}$ is a single-qubit depolarizing channel.

The proof follows from exploiting the fact that the two SIC POVM measurements are statistically independent, as well as the 2-design property of SIC POVMs. We defer it to the end of this section. For the task at hand, Lemma 3 bounds all remaining probabilities in Eq. (A30). Doing so, conveniently cancels the existing powers of 3 and produces the bound given in Proposition 1 for $K = 2$ qubits:

$$\begin{aligned} \mathbb{E}[\text{tr}(\hat{\sigma}\hat{\sigma}')^2] &= 8^2 \times 3^2\text{Pr}[i_1 = j_1 \wedge i_2 = j_2] + 8 \times 3\text{Pr}[i_1 = j_1] + 8 \times 3\text{Pr}[i_2 = j_2] + 1 \\ &\leq 8^2 + 2 \times 8 + 1 \\ &= (8 + 1)^2 \\ &= 9^2. \end{aligned}$$

This argument can be readily extended to an arbitrary number of qubits.

Proof of Proposition 1. The trace inner product between two N -qubit SIC POVM shadows can only assume discrete values. Indeed, Eq. (A28) states that $\text{tr}(\hat{\sigma}\hat{\sigma}') = \pm 5^c$, where c is the number of coincidental measurement outcomes. We can use indicator functions to single out all possibilities for coincidences and multiplying them with the correct scaling factor provides a closed-form expression of the trace inner product in terms of measurement outcomes alone. In the following, we use $\mathbf{K} \subseteq [N]$ to denote a subset of coincidental indices. The complementary set (where indices must not coincide) will be denoted by $\bar{\mathbf{K}} = [N] \setminus \mathbf{K}$. For the squared trace inner product, we then obtain

$$\begin{aligned} \text{tr}(\hat{\sigma}\hat{\sigma}')^2 &= \sum_{\mathbf{K} \subseteq [N]} 25^{|\mathbf{K}|} \mathbf{1} \left(\bigwedge_{k \in \mathbf{K}} \{i_k = j_k\} \bigwedge_{\bar{k} \in \bar{\mathbf{K}}} \{i_{\bar{k}} \neq j_{\bar{k}}\} \right) \\ &= \sum_{\mathbf{K} \subseteq [N]} 25^{|\mathbf{K}|} \left(\sum_{\mathbf{T} \subseteq \bar{\mathbf{K}}} (-1)^{|\mathbf{T}|} \mathbf{1} \left(\bigwedge_{u \in \mathbf{K} \cup \mathbf{T}} \{i_u = j_u\} \right) \right) \\ &= \sum_{\mathbf{U} \subseteq [N]} \left(\sum_{\mathbf{T} \subseteq \mathbf{U}} (-1)^{|\mathbf{T}|} 25^{|\mathbf{U}| - |\mathbf{T}|} \right) \mathbf{1} \left(\bigwedge_{u \in \mathbf{U}} \{i_u = j_u\} \right). \end{aligned}$$

In the second line, we have reexpressed conditions for noncoincidence ($\{i_{\bar{k}} \neq j_{\bar{k}}\}$) as linear combinations of coincidences on larger subsystems ($\mathbf{K} \cup \mathbf{T}$ with $\mathbf{T} \subseteq \bar{\mathbf{K}}$). The last line follows from introducing the union $\mathbf{U} = \mathbf{K} \cup \mathbf{T}$ and rewriting \mathbf{K} as $\mathbf{K} \setminus \mathbf{T}$. The inner sum over subsets $\mathbf{T} \subseteq \mathbf{U}$ has no effect on the indicator function. The size of such sets ranges from $T = |\mathbf{T}| = 0$ up to $T = |\mathbf{T}| = |\mathbf{U}|$ and, for each T , there are $\binom{|\mathbf{U}|}{T}$ subsets of that size. For a fixed set \mathbf{U} , we therefore obtain

$$\begin{aligned} \sum_{\mathbf{T} \subseteq \mathbf{U}} (-1)^{|\mathbf{T}|} 25^{|\mathbf{U}| - |\mathbf{T}|} &= \sum_{T=0}^{|\mathbf{U}|} \binom{|\mathbf{U}|}{T} (-1)^T 25^{|\mathbf{U}| - T} \\ &= (25 - 1)^{|\mathbf{U}|} = 8^{|\mathbf{U}|} \times 3^{|\mathbf{U}|}, \end{aligned}$$

which considerably simplifies the entire function. Taking the expectation now produces

$$\begin{aligned} \mathbb{E}[\text{tr}(\hat{\sigma}\hat{\sigma}')] &= \sum_{\mathbf{U} \subseteq [N]} 8^{|\mathbf{U}|} \times 3^{|\mathbf{U}|} \mathbb{E} \left[\mathbf{1} \left(\bigwedge_{u \in \mathbf{U}} \{i_u = j_u\} \right) \right] \\ &= \sum_{\mathbf{U} \subseteq [N]} 8^{|\mathbf{U}|} \times 3^{|\mathbf{U}|} \Pr \left[\bigwedge_{u \in \mathbf{U}} \{i_u = j_u\} \right], \end{aligned}$$

and we can use Lemma 3 to complete the argument. Indeed, $\Pr[\bigwedge_{u \in \mathbf{U}} \{i_u = j_u\}] \leq 3^{-|\mathbf{U}|}$ ensures that

$$\mathbb{E}[\text{tr}(\hat{\sigma}\hat{\sigma}')] \leq \sum_{\mathbf{U} \subseteq [N]} 8^{|\mathbf{U}|} = \sum_{U=0}^N \binom{N}{U} 8^U = (8 + 1)^N = 9^N.$$

This is the advertised bound for the variance of N -qubit purity estimators. \blacksquare

Finally, we provide the proof for the bound on coincidental SIC POVM measurement outcomes.

Proof of Lemma 3. For simplicity, we assume that the subset $\mathbf{K} = \{1, \dots, K\} \subseteq [N]$ encompasses the first $K = |\mathbf{K}|$ qubits. The general case works analogously, but notation becomes somewhat cumbersome. We perform two independent single-qubit SIC POVM measurements on (two copies of) an N -qubit quantum state ρ . The probability of getting $K = |\mathbf{K}|$ particular outcomes depends only on the reduced density matrix $\rho_{\mathbf{K}} = \text{tr}_{-\mathbf{K}}(\rho)$ of the relevant qubit subset:

$$\Pr[i_1 \cdots i_K | \rho_{\mathbf{K}}] = 2^{-K} \langle \psi_{i_1}, \dots, \psi_{i_K} | \rho_{\mathbf{K}} | \psi_{i_1}, \dots, \psi_{i_K} \rangle$$

for $i_1, \dots, i_K \in \{1, 2, 3, 4\}$.

This observation allows us to rewrite the probability for K coincidental measurement outcomes as

$$\begin{aligned} \Pr \left[\bigwedge_{k=1}^K \{i_k = j_k\} \right] &= \sum_{i_1=1}^4 \cdots \sum_{i_K=1}^4 \Pr[i_1 \cdots i_K | \rho_{\mathbf{K}}]^2 \\ &= \frac{1}{4^K} \sum_{i_1=1}^4 \cdots \sum_{i_K=1}^4 \langle \psi_{i_1}, \dots, \psi_{i_K} | \rho_{\mathbf{K}} | \\ &\quad \times \psi_{i_1}, \dots, \psi_{i_K} \rangle^2. \end{aligned}$$

At this point it is helpful to decompose (one) $\rho_{\mathbf{K}}$ into a linear combination of tensor products, e.g., $\rho_{\mathbf{K}} = \sum_{W_1, \dots, W_K} r(W_1, \dots, W_K) W_1 \otimes \cdots \otimes W_K$. Doing so allows us to rewrite

$$\begin{aligned} \Pr \left[\bigwedge_{k=1}^K \{i_k = j_k\} \right] &= 2^{-K} \text{tr} \left(\rho_{\mathbf{K}} \sum_{W_1, \dots, W_K} r(W_1, \dots, W_K) \right. \\ &\quad \times \left. \left(\frac{1}{2} \sum_{i_1=1}^4 |\psi_{i_1}\rangle \langle \psi_{i_1}| \langle \psi_{i_1} | W_1 | \psi_{i_1}\rangle \right) \otimes \cdots \right. \\ &\quad \left. \otimes \left(\frac{1}{2} \sum_{i_K=1}^4 |\psi_{i_K}\rangle \langle \psi_{i_K}| \langle \psi_{i_K} | W_K | \psi_{i_K}\rangle \right) \right) \\ &= 2^{-K} \text{tr} \left(\rho_{\mathbf{K}} \sum_{W_1, \dots, W_K} r(W_1, \dots, W_K) \mathcal{D}_{1/3}(W_1) \right. \\ &\quad \left. \otimes \cdots \otimes \mathcal{D}_{1/3}(W_K) \right), \\ &= 2^{-K} \text{tr}(\rho_{\mathbf{K}} \mathcal{D}_{1/3}^{\otimes K}(\rho_{\mathbf{K}})), \end{aligned}$$

as advertised. Here, we have used the 2-design property (A4) of SIC POVMs, more precisely, $\sum_{i=1}^4 \frac{1}{2} |\psi_i\rangle \langle \psi_i| \langle \psi_i | A | \psi_i\rangle = \mathcal{D}_{1/3}(A)$. To get the state-independent upper bound, we note that each depolarizing channel is a linear combination between the identity channel

$[\mathcal{I}(A) = A]$ and the projection onto the identity matrix $[\mathcal{T}(A) = \text{tr}(A)\mathbb{I}]$: $\mathcal{D}_{1/3} = \frac{1}{3}(\mathcal{I} + \mathcal{T})$. We also drop the subscript K in ρ to declutter notation somewhat: $\rho_K \mapsto \rho$. Then,

$$\begin{aligned} 2^{-K} \text{tr}(\rho \mathcal{D}_{1/3}^{\otimes K}(\rho)) &= 2^{-K} \text{tr}(\rho 3^{-K} (\mathcal{I} + \mathcal{T})^{\otimes K}(\rho)) \\ &= \frac{1}{3^K} 2^{-K} \sum_{T \subseteq \{1, \dots, s\}} \text{tr}(\rho \mathcal{T}_T \otimes \mathcal{I}_{\bar{T}}(\rho)) \\ &= 3^{-K} \left(2^{-K} \sum_{T \subseteq \{1, \dots, K\}} \text{tr}(\rho \rho_T \otimes \mathbb{I}_{\bar{T}}) \right) \\ &= 3^{-K} \left(2^{-K} \sum_{T \subseteq \{1, \dots, s\}} \text{tr}(\rho_T^2) \right). \end{aligned}$$

The remaining bracket averages over all subsystem purities $\text{tr}(\rho_T^2)$. Each of them obeys $\text{tr}(\rho_T^2) \leq 1$ and there are a total of 2^K of them (a finite set of size K has 2^K subsets). Upper bounding each of them by 1 produces

$$2^{-K} \text{tr}(\rho \mathcal{D}_{1/3}^{\otimes K}(\rho)) = 3^{-K} \left(2^{-K} \sum_{T \subseteq \{1, \dots, s\}} \text{tr}(\rho_T^2) \right) \leq 3^{-K}.$$

This completes the proof. \blacksquare

-
- [1] M. Paris and J. Řeháček, *Quantum State Estimation*, edited by M. Paris and J. Řeháček, Lecture Notes in Physics, Vol. 649 (Springer Berlin, Heidelberg, 2004).
- [2] J. M. Renes, R. Blume-Kohout, A. J. Scott, and C. M. Caves, Symmetric informationally complete quantum measurements, *J. Math. Phys.* **45**, 2171 (2004).
- [3] A. J. Scott, Tight informationally complete quantum measurements, *J. Phys. A* **39**, 13507 (2006).
- [4] A. J. Scott and M. Grassl, Symmetric informationally complete positive-operator-valued measures: A new computer study, *J. Math. Phys.* **51**, 042203 (2010).
- [5] C. A. Fuchs, M. C. Hoang, and B. C. Stacey, The SIC question: History and state of play, *Axioms* **6**, 00 (2017).
- [6] G. N. M. Tabia, Experimental scheme for qubit and qutrit symmetric informationally complete positive operator-valued measurements using multipoint devices, *Phys. Rev. A* **86**, 062107 (2012).
- [7] Z. Bian, J. Li, H. Qin, X. Zhan, R. Zhang, B. C. Sanders, and P. Xue, Realization of Single-Qubit Positive-Operator-Valued Measurement via a One-Dimensional Photonic Quantum Walk, *Phys. Rev. Lett.* **114**, 203602 (2015).
- [8] M. Proietti, M. Ringbauer, F. Graffitti, P. Barrow, A. Pickston, D. Kundys, D. Cavalcanti, L. Aolita, R. Chaves, and A. Fedrizzi, Enhanced Multiqubit Phase Estimation in Noisy Environments by Local Encoding, *Phys. Rev. Lett.* **123**, 180503 (2019).
- [9] N. Bent, H. Qassim, A. A. Tahir, D. Sych, G. Leuchs, L. L. Sánchez-Soto, E. Karimi, and R. W. Boyd, Experimental Realization of Quantum Tomography of Photonic Qudits via Symmetric Informationally Complete Positive Operator-Valued Measures, *Phys. Rev. X* **5**, 041006 (2015).
- [10] P.-X. Chen, J. A. Bergou, S.-Y. Zhu, and G.-C. Guo, Ancilla dimensions needed to carry out positive-operator-valued measurement, *Phys. Rev. A* **76**, 060303 (2007).
- [11] G. García-Pérez, M. A. Rossi, B. Sokolov, F. Tacchino, P. K. Barkoutsos, G. Mazzola, I. Tavernelli, and S. Maniscalco, Learning to Measure: Adaptive Informationally Complete Generalized Measurements for Quantum Algorithms, *PRX Quantum* **2**, 040342 (2021).
- [12] G. Torlai, G. Mazzola, J. Carrasquilla, M. Troyer, R. Melko, and G. Carleo, Neural-network quantum state tomography, *Nat. Phys.* **14**, 447 (2018).
- [13] J. Carrasquilla, G. Torlai, R. G. Melko, and L. Aolita, Reconstructing quantum states with generative models, *Nat. Machine Intell.* **1**, 155 (2019).
- [14] H.-Y. Huang, R. Kueng, and J. Preskill, Predicting many properties of a quantum system from very few measurements, *Nat. Phys.* **16**, 1050 (2020).
- [15] M. Pains and A. Kalev, An approximate description of quantum states, (2019), Preprint at [ArXiv:1910.10543](https://arxiv.org/abs/1910.10543).
- [16] A. Elben, S. T. Flammia, H.-Y. Huang, R. Kueng, J. Preskill, B. Vermersch, and P. Zoller, The randomized measurement toolbox, arXiv preprint (2022), [ArXiv:2203.11374](https://arxiv.org/abs/2203.11374).
- [17] M. Ringbauer, M. Meth, L. Postler, R. Stricker, R. Blatt, P. Schindler, and T. Monz, A universal qudit quantum processor with trapped ions, (2021), Preprint at [ArXiv:2109.06903](https://arxiv.org/abs/2109.06903).
- [18] M. Naimark, Spectral functions of a symmetric operator, *Bull. Acad. Sci. URSS. Sér. Math. [Izvestia Akad. Nauk SSSR]* **4**, 277 (1940).
- [19] J. Ahn, T. C. Weinacht, and P. H. Bucksbaum, Information storage and retrieval through quantum phase, *Science* **287**, 463 (2000).
- [20] B. E. Anderson, H. Sosa-Martinez, C. A. Riofrío, I. H. Deutsch, and P. S. Jessen, Accurate and Robust Unitary Transformations of a High-Dimensional Quantum System, *Phys. Rev. Lett.* **114**, 240401 (2015).
- [21] C. Godfrin, A. Ferhat, R. Ballou, S. Klyatskaya, M. Ruben, W. Wernsdorfer, and F. Balestro, Operating Quantum States in Single Magnetic Molecules: Implementation of Grover's Quantum Algorithm, *Phys. Rev. Lett.* **119**, 187702 (2017).
- [22] X.-M. Hu, Y. Guo, B.-H. Liu, Y.-F. Huang, C.-F. Li, and G.-C. Guo, Beating the channel capacity limit for superdense coding with entangled ququarts, *Sci. Adv.* **4**, eaat9304 (2018).
- [23] Y. Chi, *et al.*, A programmable qudit-based quantum processor, *Nat. Commun.* **13**, 1166 (2022).
- [24] M. S. Blok, V. V. Ramasesh, T. Schuster, K. O'Brien, J. M. Kreikebaum, D. Dahlen, A. Morvan, B. Yoshida, N. Y. Yao, and I. Siddiqi, Quantum Information Scrambling on a Superconducting Qutrit Processor, *Phys. Rev. X* **11**, 021010 (2021).
- [25] Z. Hradil, Quantum-state estimation, *Phys. Rev. A* **55**, R1561 (1997).
- [26] C. J. Wood, Initialization and characterization of open quantum systems, Ph.D. thesis, University of Waterloo (2015).

- [27] J. A. Smolin, J. M. Gambetta, and G. Smith, Efficient Method for Computing the Maximum-Likelihood Quantum State from Measurements with Additive Gaussian Noise, *Phys. Rev. Lett.* **108**, 070502 (2012).
- [28] T. Sugiyama, P. S. Turner, and M. Muraio, Precision-Guaranteed Quantum Tomography, *Phys. Rev. Lett.* **111**, 160406 (2013).
- [29] M. Guță, J. Kahn, R. Kueng, and J. A. Tropp, Fast state tomography with optimal error bounds, *J. Phys. A: Math. Theor.* **53**, 204001 (2020).
- [30] R. Blume-Kohout, Optimal, reliable estimation of quantum states, *New J. Phys.* **12**, 043034 (2010).
- [31] J. M. Lukens, K. J. H. Law, A. Jasra, and P. Lougovski, A practical and efficient approach for Bayesian quantum state estimation, *New J. Phys.* **22**, 063038 (2020).
- [32] A. Elben, R. Kueng, H.-Y. R. Huang, R. van Bijnen, C. Kokail, M. Dalmonte, P. Calabrese, B. Kraus, J. Preskill, P. Zoller, and B. Vermersch, Mixed-State Entanglement from Local Randomized Measurements, *Phys. Rev. Lett.* **125**, 200501 (2020).
- [33] A. Neven, J. Carrasco, V. Vitale, C. Kokail, A. Elben, M. Dalmonte, P. Calabrese, P. Zoller, B. Vermersch, R. Kueng, and B. Kraus, Symmetry-resolved entanglement detection using partial transpose moments, *npj Quantum Inf.* **7**, 152 (2021).
- [34] A. Ambainis and J. Emerson, in *Twenty-Second Annual IEEE Conference on Computational Complexity (CCC'07)* (2007), p. 129.
- [35] R. Kueng and D. Gross, Qubit stabilizer states are complex projective 3-designs, arXiv preprint (2015), [ArXiv:1510.02767](https://arxiv.org/abs/1510.02767).
- [36] H. Zhu, Multiqubit Clifford groups are unitary 3-designs, *Phys. Rev. A* **96**, 062336 (2017).
- [37] Z. Webb, The Clifford group forms a unitary 3-design, *Quantum Info. Comput.* **16**, 1379 (2016).
- [38] P. Schindler, D. Nigg, T. Monz, J. T. Barreiro, E. Martinez, S. X. Wang, S. Quint, M. F. Brandl, V. Nebendahl, C. F. Roos, M. Chwalla, M. Hennrich, and R. Blatt, A quantum information processor with trapped ions, *New J. Phys.* **15**, 123012 (2013).
- [39] M. Enríquez, I. Wintrowicz, and K. Życzkowski, Maximally entangled multipartite states: A brief survey, *J. Phys.: Conf. Ser.* **698**, 012003 (2016).
- [40] W. Helwig and W. Cui, Absolutely maximally entangled states: Existence and applications, (2013), Preprint at [ArXiv:1306.2536](https://arxiv.org/abs/1306.2536).
- [41] M. Grassl, T. Beth, and M. Roetteler, On optimal quantum codes, *Int. J. Quantum Inf.* **02**, 55 (2003).
- [42] S. Muralidharan and P. K. Panigrahi, Perfect teleportation, quantum-state sharing, and superdense coding through a genuinely entangled five-qubit state, *Phys. Rev. A* **77**, 032321 (2008).
- [43] C. Ferrie and R. Blume-Kohout, *Maximum likelihood quantum state tomography is inadmissible*, (2018).
- [44] T. Brydges, A. Elben, P. Jurcevic, B. Vermersch, C. Maier, B. P. Lanyon, P. Zoller, R. Blatt, and C. F. Roos, Probing Rényi entanglement entropy via randomized measurements, *Science* **364**, 260 (2019).
- [45] H.-Y. Huang, M. Broughton, J. Cotler, S. Chen, J. Li, M. Mohseni, H. Neven, R. Babbush, R. Kueng, and J. Preskill, *et al.*, Quantum advantage in learning from experiments, (2021), arXiv preprint [ArXiv:2112.00778](https://arxiv.org/abs/2112.00778).
- [46] M. Appleby, S. Flammia, G. McConnell, and J. Yard, SICs and algebraic number theory, *Found. Phys.* **47**, 1042 (2017).
- [47] J. Haah, A. W. Harrow, Z. Ji, X. Wu, and N. Yu, Sample-optimal tomography of quantum states, *IEEE Trans. Inform. Theory* **63**, 5628 (2017).
- [48] R. O'Donnell and J. Wright, in *STOC'16—Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing* (ACM, New York, 2016), p. 899.
- [49] D. A. Roberts and B. Yoshida, *Chaos and complexity by design*, *J. High Energy Phys.*, 121, front matter+63 (2017).
- [50] F. G. Brandão, W. Chemissany, N. Hunter-Jones, R. Kueng, and J. Preskill, Models of Quantum Complexity Growth, *PRX Quantum* **2**, 030316 (2021).
- [51] Z. Huangjun, Quantum state estimation and symmetric informationally complete POMs, Ph.D. thesis, National University of Singapore (2012).
- [52] D. H. Mahler, L. A. Rozema, A. Darabi, C. Ferrie, R. Blume-Kohout, and A. Steinberg, Adaptive Quantum State Tomography Improves Accuracy Quadratically, *Phys. Rev. Lett.* **111**, 183601 (2013).
- [53] F. Huszár and N. M. Housley, Adaptive Bayesian quantum tomography, *Phys. Rev. A* **85**, 052120 (2012).
- [54] C. Ferrie, Self-Guided Quantum Tomography, *Phys. Rev. Lett.* **113**, 190404 (2014).
- [55] Z. Hou, H. Zhu, G.-Y. Xiang, C.-F. Li, and G.-C. Guo, Achieving quantum precision limit in adaptive qubit state tomography, *npj Quantum Inf.* **2**, 1 (2016).
- [56] S. M. Kazim, A. Farooq, J. ur Rehman, M. Sohail, M. Pasha, M. Nadeem, and A. Ali, Adaptive quantum state tomography with iterative particle filtering, *Quantum Inf. Process.* **20**, 348 (2021).
- [57] R. J. Chapman, C. Ferrie, and A. Peruzzo, Experimental Demonstration of Self-Guided Quantum Tomography, *Phys. Rev. Lett.* **117**, 040402 (2016).
- [58] Z. Hou, J.-F. Tang, C. Ferrie, G.-Y. Xiang, C.-F. Li, and G.-C. Guo, Experimental realization of self-guided quantum process tomography, *Phys. Rev. A* **101**, 022317 (2020).
- [59] M. Rambach, M. Qaryan, M. Kewming, C. Ferrie, A. G. White, and J. Romero, Robust and Efficient High-Dimensional Quantum State Tomography, *Phys. Rev. Lett.* **126**, 100402 (2021).
- [60] B. Qi, Z. Hou, Y. Wang, Q. Guo, Z.-Y. Fang, X.-H. Bao, X. Zhu, and R.-B. Liu, Adaptive quantum state tomography via linear regression estimation: Theory and two-qubit experiment, *npj Quantum Inf.* **3**, 19 (2017).
- [61] G. I. Struchalin, I. A. Pogorelov, S. S. Straupe, K. S. Kravtsov, I. V. Radchenko, and S. P. Kulik, Experimental adaptive quantum tomography of two-qubit states, *Phys. Rev. A* **93**, 012103 (2016).
- [62] D. Dieks, Overlap and distinguishability of quantum states, *Phys. Lett. A* **126**, 303 (1988).
- [63] R. Stricker, M. Meth, L. Postler, C. Edmunds, C. Ferrie, R. Blatt, P. Schindler, T. Monz, R. Kueng, and M. Ringbauer, Data for: "Experimental Single-Setting Quantum State Tomography", *PRX Quantum*, (2022), <https://doi.org/10.5281/zenodo.7054827>.
- [64] L. E. Fischer, D. Miller, F. Tacchino, P. K. Barkoutsos, D. J. Egger, and I. Tavernelli, *Ancilla-free implementation of*

- generalized measurements for qubits embedded in a qudit space*, (2022).
- [65] K. Mølmer and A. Sørensen, Multiparticle Entanglement of Hot Trapped Ions, *Phys. Rev. Lett.* **82**, 1835 (1999).
- [66] G. Vidal and R. F. Werner, Computable measure of entanglement, *Phys. Rev. A* **65**, 032314 (2002).
- [67] P. Horodecki, Separability criterion and inseparable mixed states with positive partial transposition, *Phys. Lett. A* **232**, 333 (1997).
- [68] A. Klappenecker and M. Rotteler, in *Proceedings. International Symposium on Information Theory, 2005. ISIT 2005.* (2005), p. 1740.
- [69] S. Foucart and H. Rauhut, *A Mathematical Introduction to Compressive Sensing*, Applied and Numerical Harmonic Analysis (Birkhäuser/Springer, New York, 2013), p. xviii+625.
- [70] R. Vershynin, *High-Dimensional Probability*, Cambridge Series in Statistical and Probabilistic Mathematics, Vol. 47 (Cambridge University Press, Cambridge, 2018), p. xiv+284, an introduction with applications in data science, With a foreword by Sara van de Geer.
- [71] S. Chen, J. Cotler, H.-Y. Huang, and J. Li, in *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science—FOCS 2021* (IEEE Computer Soc., Los Alamitos, CA, [2022] ©2022), p. 574.
- [72] A. Peres, Separability Criterion for Density Matrices, *Phys. Rev. Lett.* **77**, 1413 (1996).
- [73] R. Horodecki and M. Horodecki, Information-theoretic aspects of inseparability of mixed states, *Phys. Rev. A* **54**, 1838 (1996).
- [74] R. Horodecki, P. Horodecki, M. Horodecki, and K. Horodecki, Quantum entanglement, *Rev. Mod. Phys.* **81**, 865 (2009).