A Tailor-made Quantum State Tomography Approach

D. Binosi,[1] G. Garberoglio,[1] D. Maragnano,[2] M. Dapor,[1] and M. Liscidini[2]
[1)]*European Centre for Theoretical Studies in Nuclear Physics and Related Areas (ECT*, Fondazione Bruno Kessler); Villa Tambosi, Strada delle Tabarelle 286, I-38123 Villazzano (TN), Italy*
[2)]*Dipartimento di Fisica, Università di Pavia, via Bassi 6, 27100 Pavia, Italy*

(*Electronic mail: diego.maragnano01@universitadipavia.it)

(Dated: 4 July 2024)

Quantum state tomography (QST) aims at reconstructing the state of a quantum system. However in conventional QST the number of measurements scales exponentially with the number of qubits. Here we propose a QST protocol, in which the introduction of a threshold allows one to drastically reduce the number of measurements required for the reconstruction of the state density matrix without compromising the result accuracy. In addition, one can also use the same approach to reconstruct an approximated density matrix tailoring the number of measurements on the available resources. We experimentally demonstrate this protocol by performing the tomography of states up to 7 qubits. We show that our approach can lead to results in agreement with those obtained by QST even when the number of measurements is reduced by more than two orders of magnitude.

I. INTRODUCTION

In the last few years, the race for quantum technologies has favoured the realization of large and complex quantum states in various platforms, including superconducting,[1] atomic,[2,3] and photonic systems.[4–9] Such advances are an important resource for the implementation of useful quantum computers and an extraordinary opportunity for fundamental studies in quantum mechanics.

Ideally, one would like to know the complete state of a quantum system, for this is sufficient to compute the value of any other observable. However, as the dimension of the quantum system increases, this task is not trivial at all. For example, consider quantum state tomography (QST),[10] which is arguably the most famous approach for quantum state reconstruction. The idea is to obtain the density matrix of the system from the measurements of a set of observables on a large enough ensemble of identical copies of the state. This method is general and does not make any *a priori* assumption on the state. However, a system composed of n qudits requires determining the expectation value of at least d^{2n} observables, that is the number of independent entries of the density matrix. Such an exponential growth of the number of measurements with the number of qudits can make the experimental implementation of QST unfeasible even for just three or four qudits, depending on the quantum system under investigation.

Over the years, strategies have been developed to decrease the number of required measurements, usually by leveraging certain assumptions about the state or the anticipated outcome. For example, the efficiency of compressed sensing QST depends on the rank of the density matrix, with the state characterization done via a certain number of random Pauli measurements.[11] Further improvements of this approach have reduced the number of necessary measurement sets significantly. [12,13] In QST via reduced density matrices, one necessitates the state to be uniquely determined by local reduced density matrices or the assumption that the global state is pure.[14] In Bayesian QST, one defines conditional probabilities starting from measurement results and a suitable *a priori* distribution over the space of possible states.[15,16] In self-guided QST one addresses tomography through optimization rather than estimation, iteratively learning the quantum state by treating tomography as a projection measurement optimization problem.[17–19]

Other strategies avoid the reconstruction of the density matrix by focusing only on *some* relevant information about the state. For instance, in shadow tomography, one can estimate the value of a large number of observables with a few copies of the unknown state.[20–22] Finally, quantum witnesses are useful to verify some crucial properties of a quantum state, such as its degree of entanglement and the kind of quantum correlation.[23] While these last two approaches are extremely useful and applicable to states of large dimensions, it is somewhat disappointing that, after all the efforts to implement a state, one cannot look at it as a whole.

In this work, we describe and analyze *threshold* Quantum State Tomography (tQST),[24] a protocol that allows one to reconstruct the density matrix of a quantum state, even of large dimension, by allowing an efficient trade-off between the number of measurements and the accuracy of the state reconstruction through the presence of a *threshold* parameter. We show that tQST can drastically reduce the resources required for state reconstruction, but it can also be used to obtain an approximated density matrix by further reducing the number of measurements and the experimental efforts. The protocol can be applied to any quantum system and does not make any assumption about the state to be characterized. First, we outline the tQST procedure and discuss its implementations in the case of n-qubits. Second, we experimentally verify the protocol by performing tQST on systems up to 7 qubits with the IBMQ platform and compare the results with conventional QST. Finally, we draw our conclusions.

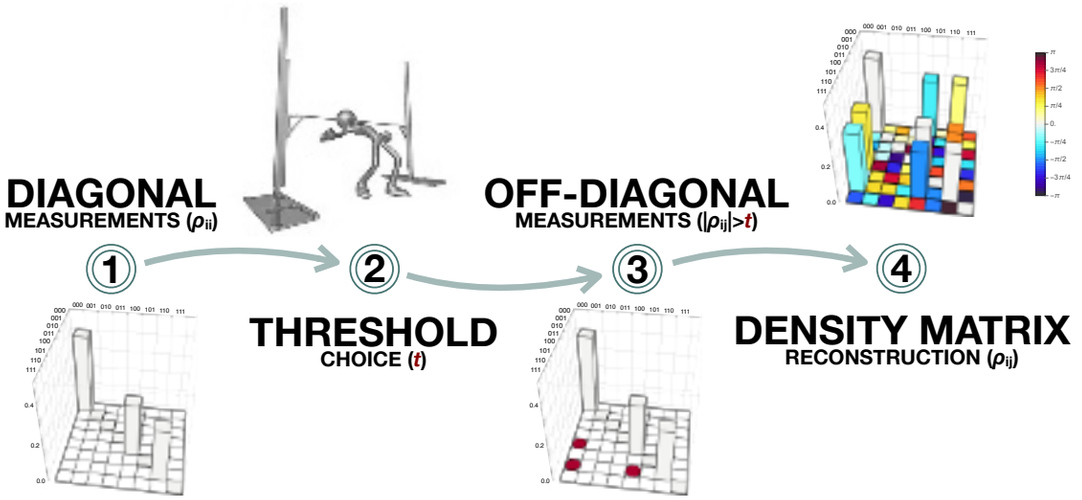

FIG. 1. Pictorial representation of the tQST process. (1) Measurement of the diagonal terms ρ_{ii}; (2) Choice of the threshold t. The choice of t resembles that of positioning the bar in a limbo competition. If the bar is very low ($t \ll 1$), a successful performance will certainly be spectacular. However, if the bar is too low, the dance may be prohibitive even for a most skilled athlete. This is the situation of conventional QST when the number of measurements to be performed is unreasonable. On the contrary, for a sufficiently high bar, even a beginner may be able to dance, but the result may not be particularly noteworthy. (3) Construction and measurement of the observable associated with the off-diagonal terms; (4) Reconstruction of the density matrix; we plot the absolute values of the matrix elements colored by their phase as indicated by the color bar.

II. RESULTS

A. The threshold quantum state tomography protocol

In conventional QST, the number of measurements required to reconstruct the density matrix is uniquely determined by its dimension.[10] Specifically, a system of n qubits requires at least 4^n measurements, that is the number of independent parameters of the state density matrix. These measurements can be performed in an arbitrary sequence, and the results are finally combined to obtain the state density matrix through maximum likelihood estimation or other approaches.[25–27]

Here, we begin by noting that a density matrix, denoted as ρ, is required to have a unit trace, be Hermitian, and positive semi-definite. Specifically, this latter property entails that any element ρ_{ij} of a physical density matrix ρ must fulfill the condition $|\rho_{ij}| \leq \sqrt{\rho_{ii}\rho_{jj}}$. Thus, measuring the diagonal elements of the density matrix (*i.e.*, projecting on the states of the computational basis) provides some information on the off-diagonal terms. Indeed, if ρ_{ii} is found to be zero, then one can immediately set to zero all the elements of the i-th row and column of ρ. Similarly, if ρ_{ii} and ρ_{jj} are different from zero but small compared to the other diagonal elements, one knows that the modulus of ρ_{ij} will be small too.

These considerations are at the basis of tQST, whose protocol is illustrated in Fig. 1. First, one measures the diagonal elements $\{\rho_{ii}\}$ of the density matrix, which are directly accessible by projecting on the chosen computational basis. Second, one chooses a threshold t and, by exploiting the information on $\{\rho_{ii}\}$, identifies those ρ_{ij} for which $\sqrt{\rho_{ii}\rho_{jj}} \geq t$. Third, one constructs a proper set of observables associated with these ρ_{ij} (see Methods) and performs only those measurements. Finally, one uses these results to reconstruct the density matrix, for example, using a maximum-likelihood estimation. We stress that maximum likelihood estimation is not the only possible choice, and one can consider other approaches, such as linear inversion, Bayesian mean estimation,[28] or linear regression estimation.[29]

In tQST, the information achieved by measuring the diagonal terms of ρ is immediately used to decide the subsequent measurements to be performed, by choosing to neglect the terms of ρ that in modulus are smaller than a certain threshold. Unlike conventional QST, the resources necessary to reconstruct the state are not uniquely determined by the dimension of the quantum system but can be controlled with the threshold t. For example, if one sets $t = 0$, all the elements of ρ are considered, and the protocol reduces to conventional QST. On the contrary, for $t > 0$, the protocol may require fewer measurements than those needed with conventional QST and, in any case, no more than them. Importantly, one does not make any *a priori* assumption on the state or the result of the characterization. It should be noted that unlike adaptive approaches,[12,13] in which the necessary measurement sets cannot be known *a priori*, for each one is chosen based on the outcomes of the previous measurement set, in tQST the projectors can all be determined once the system is measured in the chosen computational basis. In this respect, we hasten to emphasize that, once a threshold value is chosen, the protocol does not simply reduce to measuring a subspace of the Hilbert space of the whole system.

The threshold t determines the amount of information one is willing to trade in exchange for fewer measurements. This

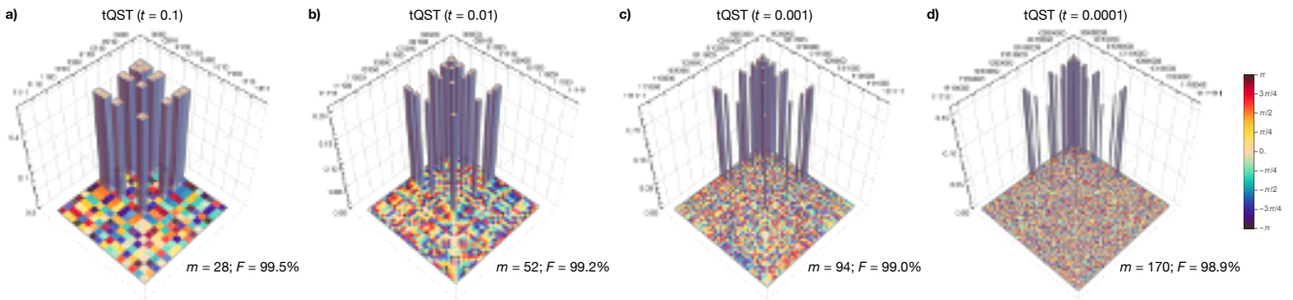

FIG. 2. Reconstruction of the density matrix of a W state via tQST. We consider different numbers of qubits and threshold values: a) 4 qubits, $t = 0.1$ b) 5 qubits, $t = 0.01$, c) 6 qubits, $t = 0.001$, and d) 7 qubits, $t = 0.0001$. For each case, we indicate the number of measurements used and the fidelity with the corresponding target state.

has several consequences. First, one may be able to reconstruct the entire density matrix of a large state by reducing the amount of measurements to a level compatible with the available experimental resources. Second, fewer measurements can give a significant advantage in terms of storing and handling experimental data. Third, one may be able to avoid useless measurements and increase the integration time for the remaining measurements, leading to an improvement of the signal-to-noise ratio. The amount of information obtainable in the characterization of a quantum state is always limited, for example by noise or the finite precision of the experimental setup. Such experimental constraints *de facto* bound the accuracy with which ρ can be determined, even for traditional QST. Thus, although one may naively expect that reducing the number of measurements will decrease the quality of the results, as we shall see in the following, a wise choice of t can still guarantee practically reaching the best achievable result while requiring fewer resources.

In Fig. 2 we show the simulated reconstruction of the density matrix via tQST with maximum likelihood estimation of several W states for different numbers of qubits ranging from 4 to 7.[30] Even for the case of 7 qubits, in which traditional QST would have required 16,384 measurements, one can reconstruct ρ with a fidelity of about 99% with the target state by implementing only 170 measurements. We stress that here the dramatic reduction of measurements is not simply given by the choice of a particular state, but rather by the employed computational basis, which affects the final state representation. As in compressed sensing tomography or similar approaches, sparse matrices are usually easier to reconstruct, because most of the information is contained in fewer elements of the density matrix. Yet, for sparse matrices the advantage of tQST can be significantly larger than with other approaches. Take for example the 7-qubit state reconstructed in Ref. 31 via compressed sensing QST, in which the characterization required 16,256 projective measurements. While this number is significantly smaller than 6^n, which is the number of observables of the typical tomographically overcomplete set, this is still 99% of the 4^n observables that are strictly necessary.[10] On the contrary, the very same state could be reconstructed via tQST by performing only 184 measurements, which is about 1% of those required by compressed sensing (see Supplementary Material). We stress that tQST can also be applied to matrices that are not sparse, where the threshold t will determine both the number of measurements and the reliability of the reconstruction.

As shown in Fig. 2, a notable feature of tQST is that a significant reduction in the number of measurements does not necessarily lead to large errors in the reconstruction of the density matrix. However, this feature is strongly dependent on the state representation. More in general, one may be interested in using tQST because the available resources are simply not enough to implement the traditional QST. In this case, it is important to have an estimate of the largest error that is associated with the threshold choice. Specifically, given the diagonal elements $\{\rho_{ii}\}$ and a corresponding threshold t, one can estimate a lower bound for the fidelity achievable through a tQST reconstruction (see Sec. IV C):

$$F_{\text{bound}}(\{\rho_{ii}\},t) = \left(1 - \sqrt{r(\bar{\rho}) \sum_{\{ij\}\in\mathscr{C}_t} \rho_{ii}\rho_{jj}}\right)^2, \qquad (1)$$

with $\mathscr{C}_t$ the ensemble of the elements ρ_{ij} that are below threshold (*i.e.*, $\mathscr{C}_t = \left\{\{ij\} \mid \sqrt{\rho_{ii}\rho_{jj}} < t\right\}$), and $r(\bar{\rho})$ is the rank of the ideal density matrix. We stress that in some cases this can be a quite conservative lower bound but still useful to verify whether a resulting low fidelity depends on the choice of an excessively high threshold, or rather a real dissimilarity between ρ and the target state.

Our approach can be extended to systems composed of n qudits (see Supplementary Material). In the case of qubits, we developed a Python package which implements tQST for an arbitrary number of qubits that is freely available to users on Github.[32]

B. Implementation on an IBMQ processor

We experimentally demonstrate tQST by using the IBMQ processor `lagos`,[33,34] which allows one to prepare states with up to 7 superconducting qubits. In our implementation, each

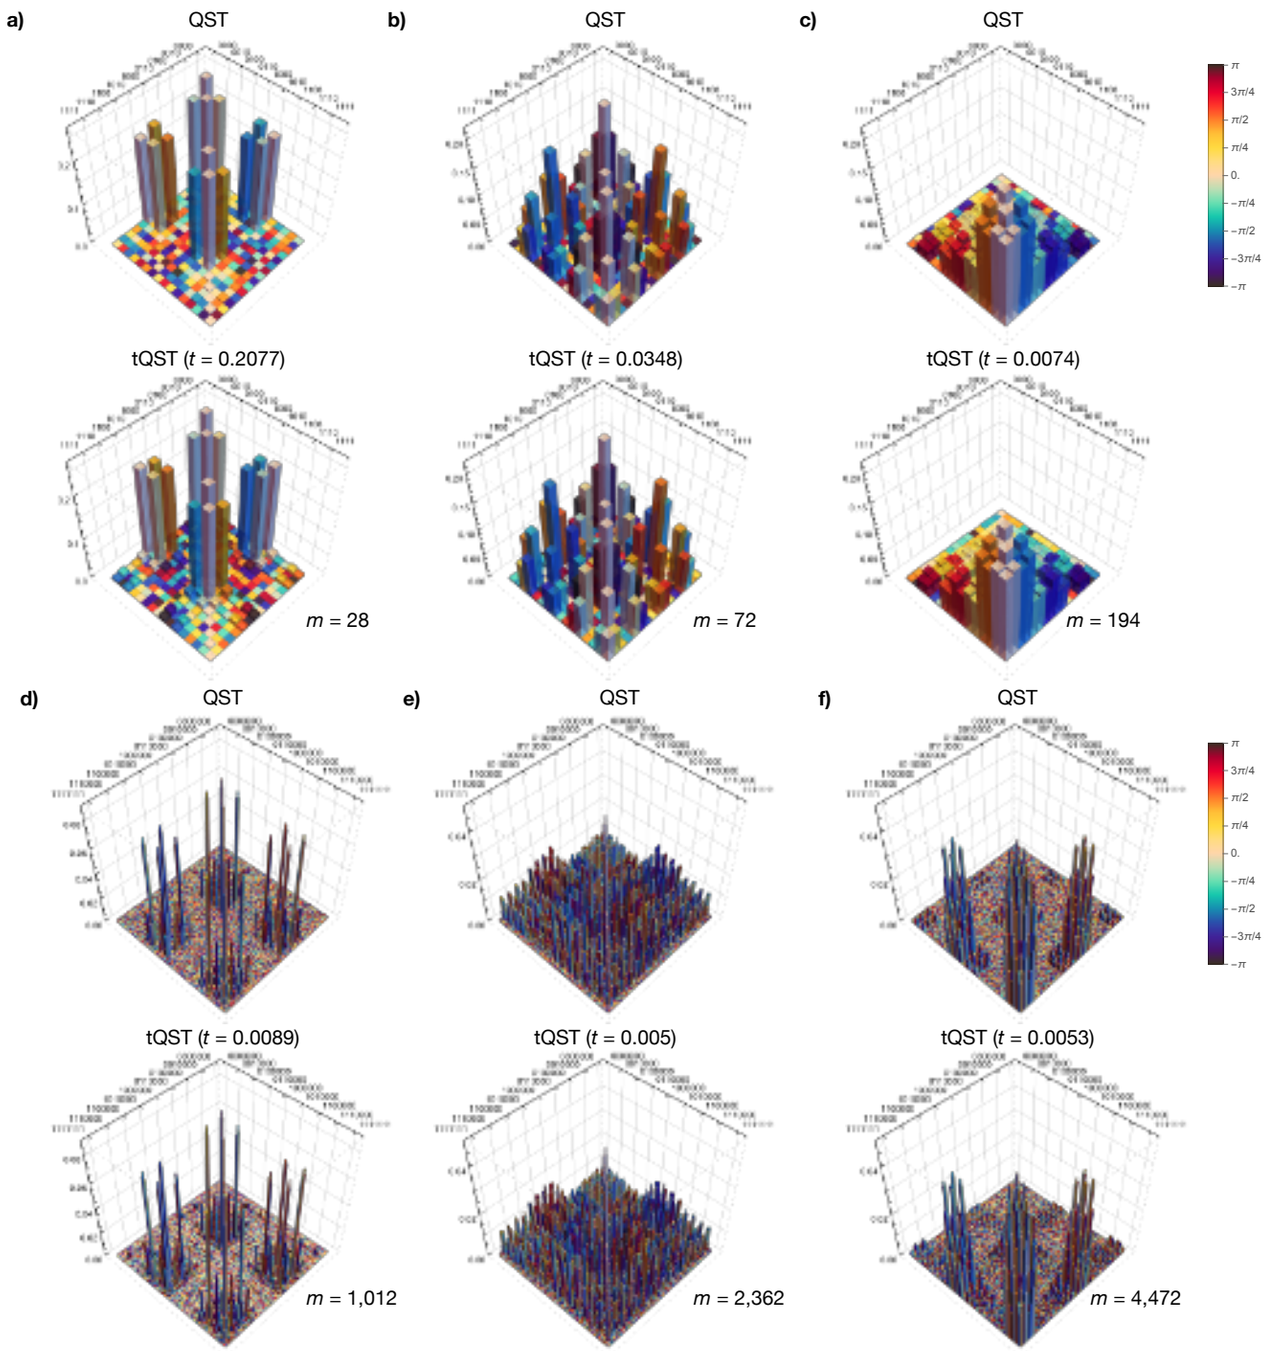

FIG. 3. tQST on random circuits implemented in a IBMQ processor. a), b), c): 4-qubit states generated from random circuits with a diagonal filling of 25% , 50%, and 75%, respectively. d), e), f): same as before for 7-qubit states. For more details see Sec. I in the Supplementary Material.

time we program the system to generate a target quantum state by constructing the corresponding quantum circuit. We first reconstruct the state density matrix by performing conventional QST as implemented by IBMQ, which uses linear inversion on the outcomes of an overcomplete set of 6^n observables. This process generally yields a non-positive reconstructed state, which is subsequently rescaled to be positive semi-definite.[35] We then reconstruct the same state by using tQST, where the threshold t is chosen by considering the typical signal-to-noise ratio (SNR) of the system to avoid unnecessary measurements (Methods, Sec. IV B). In this case, we use a maximum likelihood estimation to obtain the (pos-

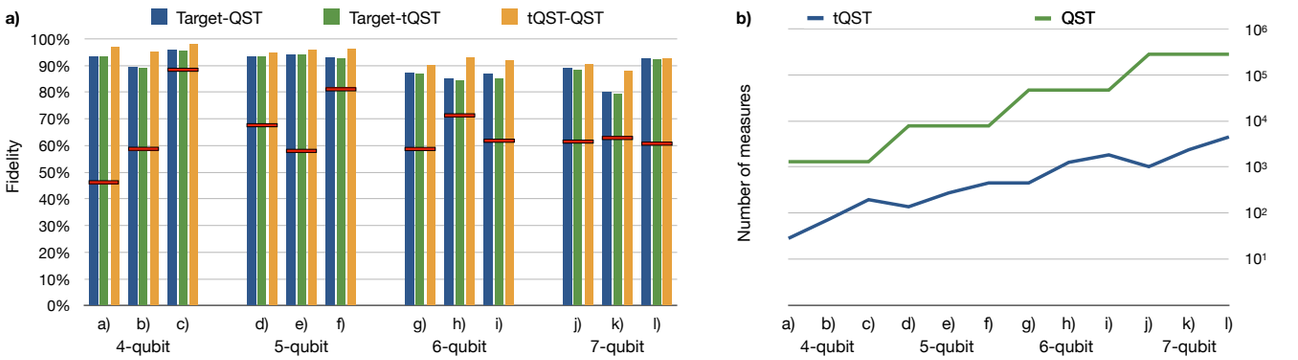

FIG. 4. Quantitative comparison between the reconstructed QST and tQST density matrices of the random circuits illustrated in Fig.1 of the Supplementary Material. a) Comparison of QST and tQST density matrix fidelities. The red line represents the tQST fidelity bound corresponding to the specified circuit threshold. b) Scaling of the required measurements with the number of qubits for both QST and tQST reconstructions.

itive semi-definite) density matrix from the measured expectation values (Methods, Sec. IV D). To compare the results we compute the fidelities between the target state and the reconstructed ones, and also that between the two reconstructed states. This choice is motivated by the fact that while not exhaustive, fidelity remains a widely recognized metric.

In our analysis, we generated 300 random states for a number of qubits from 4 to 7 and sorted them according to the diagonal filling, *i.e.*, the percentage of expected non-vanishing diagonal elements. In Fig. 3 we show representative results for filling percentages of 25% (a,d), 50% (b,e), and 75% (c,f) for the case of 4 and 7 qubits (other representative results for 5 and 6 qubits are shown in Fig.1 of the Supplementary Material). Each panel of the figure shows the circuit to generate the state, the density matrix reconstructed with conventional QST, and the one obtained with tQST. In this last case, we also report the value of the threshold t and the corresponding number of measurements used in the reconstruction. Since we choose t according to the typical SNR of the system, such a number is also related to the sparsity of the state representation. For the cases of 4, 5, and 6 qubits, we observe up to a 100-fold reduction in the number of required measurements compared to IBMQ-QST. For the case of 7 qubits, such reduction is ~ 300 for the case of Fig. 3d.

As evident from the state representation, despite the significant difference in the number of measurements, the state reconstructed with IBMQ-QST and tQST are very similar. A more quantitative analysis is obtained by calculating the fidelity between the two reconstructed states, which in all cases is about 90% or above, as shown in Fig. 4. The same figure also reports the fidelity of the tQST-reconstructed state with the target one. Remarkably, this is always comparable (within errors) to the fidelity between the IBMQ-QST reconstructed state and the target one (see Supplementary Material, Tab. S1). Thus no advantage is obtained with IBMQ-QST by performing more measurements than those set by the threshold in tQST. This suggests that when the threshold is determined by the SNR of the system, the tQST protocol can extract all the amount of information accessible with conventional QST, yet with a smaller number of measurements.

Qubits	Threshold (t)	Measures	Fidelity
8	0.053	312	91.5%
9	0.047	584	91.9%
10	0.042	1,114	91.2%
11	0.038	2,158	91.4%
12	0.035	4,228	91.4%
13	0.032	8,348	91.3%
14	0.030	16,556	91.3%

TABLE I. tQST reconstruction of a W state from synthetic data where an *ad-hoc* noise similar to the one observed in Ref. [2] has been introduced. The number of the needed off-diagonal measures for the state reconstruction scales as $n^2 - n$ with the number of qubits, and therefore it is subdominant with respect to the number of diagonal measurements (2^n) necessary for the determination of the threshold.

The IBMQ processor `lagos` limits our analysis to 7 qubits. Yet, it is interesting to further investigate the performance of tQST by extrapolating the analysis to a larger number of qubits. In this analysis we considered simulated data with a SNR analogous to that of the IBMQ system and followed the same strategy for the threshold choice. We considered the case of W states and increased the number of qubits up to 14, limited now only by our hardware. In Table I we report the number of qubits, the threshold t, the number of measurements, and the fidelity with respect to the target state. In all the cases we obtained fidelities exceeding 90%. We stress that in the cases of 14 qubits, the 16,556 measurements required by tQST make the reconstruction in principle experimentally accessible today. On the contrary, conventional QST would necessitate at least $4^{14} = 268,435,456$ measurements that, at the moment, appear as a prohibitive number.

III. CONCLUSIONS

We have proposed and implemented a protocol for QST that allows one to reconstruct the density matrix of any quantum state with a number of measurements that can be considerably smaller than that required by conventional QST. Such an approach does not make any *a priori* assumptions on the state and employs a threshold t to control the amount of information used in the state reconstruction. The presence of such a threshold allows one to trade the minimum fidelity with which the state can be reconstructed with the amount of resources to perform the characterization, thus reducing the number of measurements significantly. In addition, the threshold can be set to take into account the experimental limitations that may lead to unnecessary measurements. Our protocol was implemented on the IBMQ system `lagos` to characterize random states up to 7 qubits. In all the considered cases the fidelity achieved with tQST was compatible within the experimental uncertainty with that obtained by using conventional QST but with a smaller amount of measurements (in some cases ~ 300 times smaller). This suggests that our protocol is able to efficiently access all the information that can be extracted from the system. Finally, by using synthetic data, we pushed the approach to our computational limit and performed the characterization of W states up to 14 qubits, reaching a fidelity larger than 90% with only $\sim 16,000$ expectation values, four orders of magnitudes less than what would be required by conventional QST.

Our protocol is a flexible and practical approach for the full characterization of large quantum systems, including those based on atoms or photons. For example, the reconstruction of the 8-qubit state that was done in Ref. [2] took hours and required days to handle the outcomes of some 656,000 measurements, but our approach would require only 312 measurements, thus reducing the experimental efforts significantly. Similarly, in the case of photon-based systems[4] for which the detection rate is often limited by loss, our tomographic approach is expected to allow for the full reconstruction of states exceeding 10 photons. For these reasons, we believe that tQST will be particularly useful to the whole quantum community.

IV. MATERIALS AND METHODS

A. n-qubit projectors

The real (respectively, imaginary) part of an off-diagonal element of a density matrix ρ, can be obtained from:

$$\rho_{ij}^{\mathrm{re,im}} = \mathrm{tr}\left(\mathscr{O}_{ij}^{\mathrm{re,im}}\rho\right), \tag{2}$$

where $\mathscr{O}_{ij}^{\mathrm{re}}$ (respectively, $\mathscr{O}_{ij}^{\mathrm{im}}$) is an operator with: $1/2$ (respectively, $i/2$) at the entry ij; $1/2$ (respectively, $-i/2$) at the entry ji; and zero otherwise. How to experimentally conduct a series of measurements corresponding to these operators remains unknown. Instead, one can determine a set of 4^n projectors, denoted as $P = |\psi\rangle\langle\psi|$, where $|\psi\rangle$ is separable: $|\psi\rangle = \otimes_{k=1}^{n} |\phi_k\rangle$ with $|\phi_k\rangle$ an eigenstate of the Pauli matrices. The projectors associated with the real or imaginary part of the matrix element ij of the density matrix are then given by

$$P_{ij}^{\mathrm{re,im}} = \underset{P}{\mathrm{argmin}} \|\mathscr{O}_{ij}^{\mathrm{re,im}} - P\|_2, \tag{3}$$

where $\|\circ\|_2$ represents the Frobenius norm.

A given set of projectors is said to be *tomographically complete* if there exists a unique set of scalars $a_K^{\mathrm{re,im}}$ (with $K = \{ij\}$ a multi-index) such that the density matrix can be reconstructed, that is

$$\rho = \sum_K a_K^{\mathrm{re,im}} P_K^{\mathrm{re,im}}. \tag{4}$$

Equivalently, suppressing the "re" and "im" superindices, one gets

$$\mathrm{tr}\left(P_L^\dagger \rho\right) = \sum_K a_K \mathrm{tr}\left(P_L^\dagger P_K\right) = \sum_K M_{LK} a_K. \tag{5}$$

which means that the matrix $M_{LK} = \mathrm{tr}\left(P_L^\dagger P_K\right) = |\langle\psi_L|\psi_K\rangle|^2$ has to be invertible and the scalars a_K are uniquely determined as $a_K = \left(M^{-1}\right)_{KL} \mathrm{tr}\left(P_L^\dagger \rho\right)$.[10] Notice that, owing to the presence of equivalent minima in Eq. (3), it is not guaranteed that a complete set exists; however, if a complete set exists, then the matrix M is positive definite.

We found that in the case of an n-qubit system, a tomographically complete set of projectors associated to the elements of a n-qubit density matrix can be obtained from the eigenvectors of the Pauli matrices. In the conventional polarization notation, these eigenvectors are:

$$\begin{aligned} &|H\rangle; \quad |D\rangle = \frac{1}{\sqrt{2}}(|H\rangle + |V\rangle); \quad |R\rangle = \frac{1}{\sqrt{2}}(|H\rangle + i|V\rangle), \\ &|V\rangle; \quad |A\rangle = \frac{1}{\sqrt{2}}(|H\rangle - |V\rangle); \quad |L\rangle = \frac{1}{\sqrt{2}}(|H\rangle - i|V\rangle). \end{aligned} \tag{6}$$

In the 1-qubit case one possibility is: $P_{11}^{\mathrm{re}} = |H\rangle$; $P_{22}^{\mathrm{re}} = |V\rangle$; $P_{12}^{\mathrm{re}} = |D\rangle$; and $P_{12}^{\mathrm{im}} = |R\rangle$; or, in a more suggestive form

$$\pi_1 \equiv \begin{bmatrix} |H\rangle & |D\rangle + i|R\rangle \\ 0 & |V\rangle \end{bmatrix}. \tag{7}$$

This is a 2×2 table structured in such a way that the projector associated with measuring the real or imaginary component of the element ρ_{ij} of the 1-qubit density matrix is given by the real or imaginary part of the entry (ij) in (7). For example, information on the imaginary part of ρ_{12} is encoded in the matrix element $\langle R|\rho|R\rangle$. Entries with a value of "0" (which indeed acts as the zero element for the recursive operations below) do not need to be explicitly determined, as the density matrix is Hermitian. Consequently, we will assume, without loss of generality, that $j \geq i$ whenever we aim to determine P_{ij}. It should be noticed that the implementation of the projectors in Eq. (7) requires only local measurements on each qubit, as

in traditional QST approaches such as that implemented by IBMQ or in Ref. [10].

A set of tomographically complete separable projectors for $n > 1$ can be then constructed through the following recursion relation:

$$\pi_n = \begin{bmatrix} |H\rangle\,\pi_{n-1} & |D\rangle\,\pi_{n-1} + i\,|R\rangle\,\overline{\pi}_{n-1} \\ 0 & |V\rangle\,\pi_{n-1} \end{bmatrix}. \tag{8}$$

where we have defined

$$\overline{\pi}_1 \equiv \begin{bmatrix} |H\rangle & 0 \\ |D\rangle - i\,|R\rangle & |V\rangle \end{bmatrix}, \tag{9}$$

and analogously for all n. So in the 2-qubit case, Eq. (8) in conjunction with the results

$$|X\rangle\,\pi_1 = \begin{bmatrix} |XH\rangle & |XD\rangle + i\,|XR\rangle \\ 0 & |XV\rangle \end{bmatrix}, \tag{10a}$$

$$|X\rangle\,\overline{\pi}_1 = \begin{bmatrix} |XH\rangle & 0 \\ |XD\rangle - i\,|XR\rangle & |XV\rangle \end{bmatrix}, \tag{10b}$$

(with $X = H, V, D, R$), yields

$$\pi_2 = \begin{bmatrix} |HH\rangle & |HD\rangle + i\,|HR\rangle & |DH\rangle + i\,|RH\rangle & |DD\rangle + i\,|DR\rangle \\ 0 & |HV\rangle & |RR\rangle + i\,|RD\rangle & |DV\rangle + i\,|RV\rangle \\ 0 & 0 & |VH\rangle & |VD\rangle + i\,|VR\rangle \\ 0 & 0 & 0 & |VV\rangle \end{bmatrix}. \tag{11}$$

Information about of, *e.g.*, the real (imaginary) part of ρ_{23} is then encoded in the matrix element $\langle RR|\,\rho\,|RR\rangle$ ($\langle RD|\,\rho\,|RD\rangle$). The outlined procedure yields a total of 4^n projectors, and we have numerically verified that up to $n = 14$ the corresponding matrix M is invertible. In Eq. (8), the notation $|H\rangle\,\pi$ signifies a table where the entries result from the product of the ket $|H\rangle$ and the kets contained in π. Therefore, π_n is a table that has twice the number of rows and columns compared to π_{n-1}.

The recursive structure outlined in Eq. (8) can be leveraged to reduce the necessity of generating the entire set of projectors upfront. Instead, we can generate projectors on-demand, specifically for elements of the density matrix that need to be measured to achieve faithful reconstruction given a certain threshold.

To this end, we divide π_n into four quadrants, with "1" referring to the upper-left quadrant and "4" indicating the lower-right quadrant. Each of the quadrants 2 and 3 is further divided into an upper ("u") and a lower ("l") triangular part. The real part of the elements along the diagonal is assigned to the u portion, while the imaginary part is assigned to the l portion. This subdivision is pictorially represented in Fig. 5.

To determine the projector corresponding to the density-matrix element $\rho_{i,j}$, we initially locate it within π_n which immediately determines the projector associated with the first qubit, that is $|H\rangle$, $|D\rangle$, $|R\rangle$, or $|V\rangle$ according to its position: 1, $2u$, $2l$, or 4, respectively. We then continue splitting the quadrant where the element is found until we reach a resulting quadrant size of 2×2. At each splitting step, if the element falls into quadrants 1 or 4, the projector associated to the next qubit is $|H\rangle$ (for 1) or $|V\rangle$ (for 4). Conversely, if it falls into quadrants 2 or 3, the projector choice depends on its position in the previous splitting step. For an element in an upper quadrant, we select $|D\rangle$; and, for an element in a lower quadrant, we select $|R\rangle$ unless the previous quadrant was either $2l$ or $3l$, in which case the choice is reversed. Table II summarizes these steps.

Previous quadrant	New quadrant	1-qubit projector
any	1	$\|H\rangle$
any	4	$\|V\rangle$
any except $2l$, $3l$	$2u$, $3u$	$\|D\rangle$
any except $2l$, $3l$	$2l$, $3l$	$\|R\rangle$
$2l$, $3l$	$2u$, $3u$	$\|R\rangle$
$2l$, $3l$	$2l$, $3l$	$\|D\rangle$

TABLE II. Table determining the choice of the 1-qubit projector at each recursive step in Eq. (8). In the first step, "previous quadrant" is always "any". See text and Fig. 5 for details on the procedure.

As an illustration, Fig. 5 shows the determination of two projectors for a 3- and a 4-qubit system. More specifically, in the 3-qubit case, we consider the matrix element $\rho^{\rm im}_{3,5}$; its successive locations in π_3 are described by the sequence "$2l3l4$" which, according to Table II, corresponds to the local projector $|RDV\rangle$. The projector associated with $\rho^{\rm re}_{4,9}$ in a 4-qubit system, is instead determined by the locations in π_4 described by the sequence "$2l3u12u$"; using Table II, one then finds $|RRHD\rangle$.

B. Choice of threshold

The selection of an appropriate threshold value is dependent on the specific physical system used to implement the qubits (*e.g.*, noise level), the amount of available resources (*e.g.*, time requirements), and the desired quantum state to be generated.

In the case of IBMQ quantum processors considered here, it is possible to evaluate a convenient circuit-specific threshold

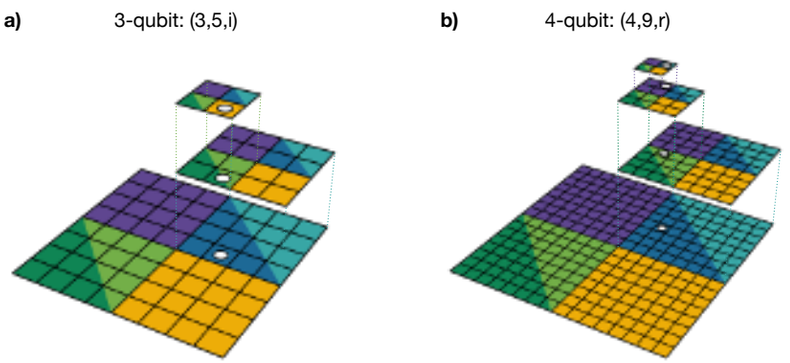

FIG. 5. Example determination of the projectors associated to two density matrix elements of 3- and 4-qubit systems using the recursive procedure outlined in the main text. At each state, the density matrix is divided into 4 different quadrants: quadrant "1" is colored in violet, quadrant "2" in blue, quadrant "3" in green, and quadrant "4" in golden yellow. Quadrants "2" and "3" are further subdivided into an upper (lighter) and a lower (darker) part. Each plane represents a successive iteration of the algorithm outlined in the main text, which proceeds from bottom to top. a) The 3-qubit matrix element $\rho^{\rm im}_{3,5}$ is located in quadrants "2*l*3*l*4" and corresponds to the $|RDV\rangle$ projector according to Tab. II b) 4-qubit matrix element $\rho^{\rm re}_{4,9}$ is located in quadrants "2*l*3*u*12*u*" therefore corresponding to the projector $|RRHD\rangle$ according to Table II.

using the IBMQ simulator available in the `qiskit` package. First, one simulates the unitary evolution of a ground-state initialized quantum register according to the circuit itself (we have used $n_s = 10,000$ shots). Measuring all qubits yields the expected diagonal counts in the absence of errors, which can be separated into zero and non-zero counts. Second, one uses the IBMQ simulator (which includes the effect of noise) to run the circuit a number of times (100 in our case) and record: the maximum value of the counts among the expected zero elements of the diagonal, $c_0^{\max}$; and the minimum value of the counts for the smallest expected non-zero diagonal element, $c_{>0}^{\min}$. Third, one defines, in a conservative way: the *noise threshold* as $t_0 = c_0^{\max} + n\sqrt{c_0^{\max}}$; and the *signal threshold* as $t_{>0} = c_{>0}^{\min} - n\sqrt{c_{>0}^{\min}}$. The square root terms consider the variability of the counts $c_0^{\max}$, $c_{>0}^{\min}$ each time the circuit is simulated; the n factor takes finally into account that, for the quantum processors considered, the noise increases with the number of qubits n. Then, we use as the circuit-specific (normalized) threshold the quantity

$$t = \max(t_0, t_{>0})/n_s, \tag{12}$$

which discards those diagonal entries most affected by noise. For the depth 3 quantum circuits analyzed herein, we generally have $t_{>0} \gtrsim t_0$; tQST works best whenever $t_{>0} \gg t_0$, whereas $t_{>0} \sim t_0$ indicates an unfavorable signal-to-noise ratio.

C. Fidelity lower bound

Fixing a threshold t given a set of diagonal elements $\{\rho_{ii}\}$, establishes a lower bound on the fidelity achievable through a tQST reconstruction of ρ. Let $\mathscr{C}_t = \{\{ij\} \mid \sqrt{\rho_{ii}\rho_{jj}} < t\}$, and consider the estimator matrix $\rho_t = \{0 \quad \forall\ \{ij\} \in \mathscr{C}_t; \sqrt{\rho_{ii}\rho_{jj}}$ otherwise$\}$ (note that $\rho = \rho_t$ if $t = 0$). Then, one has the inequalities[36,37]

$$1 - \sqrt{F(\rho,\rho_t)} \le \frac{1}{2}\|\delta_t\|_1, \tag{13a}$$

$$\|\delta_t\|_1 \le 2\sqrt{\min(r(\rho), r(\rho_t))}\,\|\delta_t\|_2, \tag{13b}$$

with the fidelity defined according to $F(\rho,\rho_t) = \mathrm{Tr}\left[\sqrt{\sqrt{\rho}\rho_t\sqrt{\rho}}\right]$, $\|\circ\|_1$ the trace distance, $\delta_t = \rho - \rho_t$, and $r(\rho)$ the rank of the density matrix. As t increases, the state purity of ρ_t decreases, so that $r(\rho) \le r(\rho_t)$. Indeed, the estimator ρ_t is a density matrix with the off-diagonal elements below threshold set to zero. Physically, this corresponds to making the density matrix more mixed and increasing its rank with respect to ρ. Additionally, $\|\delta_t\|_2 \le \sqrt{\sum_{\mathscr{C}_t}\rho_{ii}\rho_{jj}}$. Thus, we obtain the lower bound

$$F_{\rm bound}(\{\rho_{ii}\}, t) = \left(1 - \sqrt{r(\rho)\sum_{\mathscr{C}_t}\rho_{ii}\rho_{jj}}\right)^2. \tag{14}$$

In our investigations, we noticed that the actual tQST fidelity can be significantly greater than the one established by (14).

D. Maximum likelihood reconstruction

After measuring the counts N_K associated with the density matrix elements $|\rho_K| \ge t$ using the corresponding projectors, the density matrix is reconstructed by minimizing the function

$$L = \sum_K \left(\frac{n_K - N_K}{2\sqrt{n_K}}\right)^2; \qquad n_K = \langle P_K|\,\rho\,|P_K\rangle, \tag{15}$$

which is the negative log-likelihood function assuming that the noise on the counts has a Gaussian probability distribution.[10] In order to minimize L, it is necessary to

parametrize the density matrix ρ, in such a way that it fulfills the physical conditions of being Hermitian and positive semi-definite. A general approach is to write $\rho = T^\dagger T$ where T is a triangular matrix. In this case, the number of parameters to be determined grows as 4^n. If one has reasons to believe that ρ describes a high-purity state, then one can express it as $\rho = V^\dagger V$ with $\dim(V) = r(\rho_t) \times 2^n$. In this case, the parameters needed for the reconstruction will scale more favourably with the system dimension.

SUPPLEMENTARY MATERIAL

The supplementary material contains a more detailed analysis of the states reconstructed on the IBM quantum processor, a comparison between tQST and the compressed sensing approach, and an outline on the extension of tQST to qudit systems.

ACKNOWLEDGMENTS

We thank Prof. Daniel F.V. James of the University of Toronto for useful discussions.

Funding: This work is supported by PNRR MUR project PE0000023-NQSTI.

AUTHOR DECLARATIONS

Conflict of interests

The authors have no conflicts to disclose.

Author contributions

D. Binosi: Data curation (equal); Formal analysis (equal); Investigation (equal); Methodology (equal); Software (equal); Writing - original draft (equal); Writing - review and editing (equal). **G. Garberoglio**: Data curation (equal); Formal analysis (equal); Investigation (equal); Methodology (equal); Software (equal); Writing - original draft (equal); Writing - review and editing (equal). **D. Maragnano**: Conceptualization (equal); Data curation (equal); Formal analysis (equal); Investigation (equal); Methodology (equal); Software (supporting); Writing - original draft (equal); Writing - review and editing (equal). **M. Dapor**: Formal analysis (equal); Funding acquisition (equal); Investigation (equal); Methodology (equal); Project administration (equal); Writing - original draft (equal); Writing - reviewing and editing (equal). **M. Liscidini**: Conceptualization (equal); Formal analysis (equal); Funding acquisition (equal); Investigation (equal); Methodology (equal); Project administration (equal); Writing - original draft (equal); Writing - reviewing and editing (equal).

DATA AVAILABILITY

The data that support the findings of this study are available from the corresponding authors upon reasonable request.

REFERENCES

[1]S. Cao, B. Wu, F. Chen, M. Gong, Y. Wu, Y. Ye, C. Zha, H. Qian, C. Ying, S. Guo, *et al.*, "Generation of genuine entanglement up to 51 superconducting qubits," Nature **619**, 738–742 (2023).
[2]H. Häffner, W. Hänsel, C. Roos, J. Benhelm, D. Chek-al Kar, M. Chwalla, T. Körber, U. Rapol, M. Riebe, P. Schmidt, *et al.*, "Scalable multiparticle entanglement of trapped ions," Nature **438**, 643–646 (2005).
[3]T. Monz, P. Schindler, J. T. Barreiro, M. Chwalla, D. Nigg, W. A. Coish, M. Harlander, W. Hänsel, M. Hennrich, and R. Blatt, "14-qubit entanglement: Creation and coherence," Phys. Rev. Lett. **106**, 130506 (2011).
[4]X.-L. Wang, L.-K. Chen, W. Li, H.-L. Huang, C. Liu, C. Chen, Y.-H. Luo, Z.-E. Su, D. Wu, Z.-D. Li, H. Lu, Y. Hu, X. Jiang, C.-Z. Peng, L. Li, N.-L. Liu, Y.-A. Chen, C.-Y. Lu, and J.-W. Pan, "Experimental ten-photon entanglement," Phys. Rev. Lett. **117**, 210502 (2016).
[5]M. Zhang, L.-T. Feng, Z.-Y. Zhou, Y. Chen, H. Wu, M. Li, S.-M. Gao, G.-P. Guo, G.-C. Guo, D.-X. Dai, *et al.*, "Generation of multiphoton quantum states on silicon," Light: Science & Applications **8**, 41 (2019).
[6]W.-B. Gao, C.-Y. Lu, X.-C. Yao, P. Xu, O. Gühne, A. Goebel, Y.-A. Chen, C.-Z. Peng, Z.-B. Chen, and J.-W. Pan, "Experimental demonstration of a hyper-entangled ten-qubit schrödinger cat state," Nature **6**, 331–335 (2010).
[7]M. Agnew, J. Leach, M. McLaren, F. S. Roux, and R. W. Boyd, "Tomography of the quantum state of photons entangled in high dimensions," Phys. Rev. A **84**, 062101 (2011).
[8]C. Reimer, S. Sciara, P. Roztocki, M. Islam, L. Romero Cortés, Y. Zhang, B. Fischer, S. Loranger, R. Kashyap, A. Cino, *et al.*, "High-dimensional one-way quantum processing implemented on d-level cluster states," Nature Physics **15**, 148–153 (2019).
[9]L. S. Madsen, F. Laudenbach, M. F. Askarani, F. Rortais, T. Vincent, J. F. Bulmer, F. M. Miatto, L. Neuhaus, L. G. Helt, M. J. Collins, *et al.*, "Quantum computational advantage with a programmable photonic processor," Nature **606**, 75–81 (2022).
[10]D. F. V. James, P. G. Kwiat, W. J. Munro, and A. G. White, "Measurement of qubits," Phys. Rev. A **64**, 052312 (2001).
[11]D. Gross, Y.-K. Liu, S. T. Flammia, S. Becker, and J. Eisert, "Quantum state tomography via compressed sensing," Phys. Rev. Lett. **105**, 150401 (2010).
[12]D. Ahn, Y. S. Teo, H. Jeong, F. Bouchard, F. Hufnagel, E. Karimi, D. Koutný, J. Řeháček, Z. Hradil, G. Leuchs, and L. L. Sánchez-Soto, "Adaptive compressive tomography with no a priori information," Phys. Rev. Lett. **122**, 100404 (2019).
[13]D. Ahn, Y. S. Teo, H. Jeong, D. Koutný, J. Řeháček, Z. Hradil, G. Leuchs, and L. L. Sánchez-Soto, "Adaptive compressive tomography: A numerical study," Phys. Rev. A **100**, 012346 (2019).
[14]T. Xin, D. Lu, J. Klassen, N. Yu, Z. Ji, J. Chen, X. Ma, G. Long, B. Zeng, and R. Laflamme, "Quantum state tomography via reduced density matrices," Phys. Rev. Lett. **118**, 020401 (2017).
[15]R. Derka, V. Buzek, G. Adam, and P. Knight, "From quantum bayesian inference to quantum tomography," arXiv preprint quant-ph/9701029 (1997).
[16]H.-H. Lu, K. V. Myilswamy, R. S. Bennink, S. Seshadri, M. S. Alshaykh, J. Liu, T. J. Kippenberg, D. E. Leaird, A. M. Weiner, and J. M. Lukens, "Bayesian tomography of high-dimensional on-chip biphoton frequency combs with randomized measurements," Nature Communications **13**, 4338 (2022).
[17]C. Ferrie, "Self-guided quantum tomography," Phys. Rev. Lett. **113**, 190404 (2014).
[18]R. J. Chapman, C. Ferrie, and A. Peruzzo, "Experimental demonstration of self-guided quantum tomography," Phys. Rev. Lett. **117**, 040402 (2016).
[19]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).

[20]S. Aaronson, "Shadow tomography of quantum states," in *Proceedings of the 50th annual ACM SIGACT symposium on theory of computing* (2018) pp. 325–338.
[21]H.-Y. Huang, R. Kueng, and J. Preskill, "Predicting many properties of a quantum system from very few measurements," Nature Physics **16**, 1050–1057 (2020).
[22]G. Struchalin, Y. A. Zagorovskii, E. Kovlakov, S. Straupe, and S. Kulik, "Experimental estimation of quantum state properties from classical shadows," PRX Quantum **2**, 010307 (2021).
[23]H. Zhu, Y. S. Teo, and B.-G. Englert, "Minimal tomography with entanglement witnesses," Phys. Rev. A **81**, 052339 (2010).
[24]D. Maragnano, D. Binosi, G. Garberoglio, C. Cusano, M. Dapor, and M. Liscidini, "Threshold quantum state tomography," in *2023 23rd International Conference on Transparent Optical Networks (ICTON)* (2023) pp. 1–4.
[25]Z. Hradil, "Quantum-state estimation," Phys. Rev. A **55**, R1561–R1564 (1997).
[26]M. S. Kaznady and D. F. V. James, "Numerical strategies for quantum tomography: Alternatives to full optimization," Phys. Rev. A **79**, 022109 (2009).
[27]J. M. Lukens, K. J. H. Law, A. Jasra, and P. Lougovski, "A practical and efficient approach for bayesian quantum state estimation," New Journal of Physics **22**, 063038 (2020).
[28]R. Blume-Kohout, "Optimal, reliable estimation of quantum states," New Journal of Physics **12**, 043034 (2010).
[29]B. Qi, Z. Hou, L. Li, D. Dong, G. Xiang, and G. Guo, "Quantum state tomography via linear regression estimation," Scientific reports **3**, 3496 (2013).
[30]W. Dür, G. Vidal, and J. I. Cirac, "Three qubits can be entangled in two inequivalent ways," Phys. Rev. A **62**, 062314 (2000).
[31]C. A. Riofrio, D. Gross, S. T. Flammia, T. Monz, D. Nigg, R. Blatt, and J. Eisert, "Experimental quantum compressed sensing for a seven-qubit system," Nature communications **8**, 15305 (2017).
[32]D. Binosi, G. Garberoglio, D. Maragnano, M. Dapor, and M. Liscidini, "tresholdqst: A Python library to implement threshold quantum state tomography for qubits." (2024).
[33]Qiskit contributors, "Qiskit: An open-source framework for quantum computing," (2023).
[34]N. Kanazawa, D. J. Egger, Y. Ben-Haim, H. Zhang, W. E. Shanks, G. Aleksandrowicz, and C. J. Wood, "Qiskit experiments: A python package to characterize and calibrate quantum computers," Journal of Open Source Software **8**, 5329 (2023).
[35]J. A. Smolin, J. M. Gambetta, and G. Smith, "Maximum Likelihood, Minimum Effort," Phys. Rev. Lett. **108**, 070502 (2012), arXiv:1106.5458 [quant-ph].
[36]C. A. Fuchs and J. Van De Graaf, "Cryptographic distinguishability measures for quantum-mechanical states," IEEE Transactions on Information Theory **45**, 1216–1227 (1999).
[37]J. Haah, A. W. Harrow, Z. Ji, X. Wu, and N. Yu, "Sample-optimal tomography of quantum states," in *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing* (2016) pp. 913–925.