

Vector Field Visualization of Single-Qubit State Tomography

Adrien Suau
LIRMM, University of Montpellier
Montpellier, France
CERFACS, Toulouse, France
0000-0002-2412-7298

Marc Vuffray
Theoretical Division
Los Alamos National Laboratory
Los Alamos, NM 87545, USA
0000-0001-7999-9897

Andrey Y. Lokhov
Theoretical Division
Los Alamos National Laboratory
Los Alamos, NM 87545, USA
0000-0003-3269-7263

Lukasz Cincio
Theoretical Division
Los Alamos National Laboratory
Los Alamos, NM 87545, USA
0000-0002-6758-4376

Carleton Coffrin
Advanced Network Science Initiative
Los Alamos National Laboratory
Los Alamos, NM 87545, USA
0000-0003-3238-1699

Abstract—As the variety of commercially available quantum computers continues to increase so does the need for tools that can characterize, verify and validate these computers. This work explores using quantum state tomography for characterizing the performance of individual qubits and develops a vector field visualization for presentation of the results. The proposed protocol is demonstrated in simulation and on quantum computing hardware developed by IBM. The results identify qubit performance features that are not reflected in the standard models of this hardware, indicating opportunities to improve the accuracy of these models. The proposed qubit evaluation protocol is provided as free open-source software to streamline the task of replicating the process on other quantum computing devices.

Index Terms—Quantum computing, Quantum state tomography, Qubit benchmark

I. INTRODUCTION

The emergence of commercial Quantum Computing (QC) hardware has made tools for quantum characterization, verification, and validation (QCVV) [1] more important than ever. Especially in the era of Noisy Intermediate-Scale Quantum (NISQ) [2] devices, QCVV methods provide the means for QC users to measure and quantify the performance of quantum hardware platforms and enables consistent comparisons across different hardware architectures. The scope of QCVV is broad and ranges from testing individual quantum operations (e.g., error rates of one- and two-qubit gates [3]), verifying small circuits (e.g., Quantum State Tomography [4], Randomized Benchmarking [5], [6], Gate Set Tomography [7]), to full system-level protocols (e.g., quantum volume estimation [8], random quantum circuits [9]). Over the years these QCVV tools have become an invaluable foundation for benchmarking and measuring progress of quantum processors [10], culminating in multiple quantum supremacy demonstrations [11], [12].

In many ways quantum state tomography provides a gold standard for QCVV in that, it can theoretically provide an exact reconstruction of the full quantum state of QC hardware, with sufficient data. At its most basic level, quantum

state tomography provides a protocol for combining multiple observations to uniquely identify the state of a quantum system (i.e., a density matrix). On small quantum systems, where data collection and result computations are feasible, quantum state tomography provides a precise measure of how accurately a quantum hardware device can execute a desired quantum computation. The strength of quantum state tomography for QCVV is that it can provide a comprehensive picture of hardware performance. The weakness of quantum state tomography is that it often requires a prohibitively large amount of data and the interpretation of the results can be difficult. Indeed, the objective of many other QCVV protocols (e.g. Randomized Benchmarking [5], [6], Gate Set Tomography [7]) is to provide trade offs in data collection and result detail resulting in more scalable QCVV alternatives to quantum state tomography.

This work explores the use of quantum state tomography for conducting QCVV on commercially available QC platforms. We focus on quantum state tomography of single-qubits to minimize the data collection requirements and to develop a protocol that can be executed in parallel for all of the qubits in a given hardware platform. The core contributions of this work are a *Vector Field Visualization* of single-qubit quantum state tomography and an open-source software tool for data collection, state reconstruction and result presentation. Through experiments on QC hardware available in IBM’s Q-Hub, we show that the proposed method can identify qubit performance features that are not easily identified with captured with a single value and provide clear signatures that distinguish qubit performance from both perfect and noisy simulators of this hardware.

This work begins by introducing the foundations of quantum state tomography for a single-qubit in Section II and reviews the maximum likelihood estimation method [4] for reconstructing a quantum state from measurements. Section III then proposes the vector field visualization for presenting the single-qubit quantum state tomography results and illustrates

how this visualization can be leveraged to provide unique insights into qubit performance. Section IV investigates how the results from different quantum state tomography algorithms can be combined with the vector field visualization to identify signatures of data corruption. The details of the open-source software are provided in Section V and the paper concludes with a discussion of the usefulness of the proposed protocol and future work in Section VI.

II. SINGLE-QUBIT STATE TOMOGRAPHY

The state of a quantum system composed of n qubits is fully described by a $2^n \times 2^n$ Hermitian matrix ρ , the so-called density matrix. Density matrices are positive semi-definite, normalized matrices, i.e., $\langle \phi | \rho | \phi \rangle \geq 0$ for all $|\phi\rangle \in \mathbb{C}^{2^n}$, and $\text{Tr}[\rho] = 1$. A measurement outcome of a quantum system is represented by a real random variable, A , associated with a Hermitian matrix, customarily also denoted by the same letter. Measuring A is linked in expectation to the density matrix through the well-known formula $\mathbb{E}[A | \rho] = \text{Tr}[\rho A]$. This expression implies that only projections of the density matrix can be observed, as the trace is an inner product over Hermitian matrices.

The task of reconstructing density matrices from repetitive observations of these projections is called Quantum State Tomography (QST). Since the space of density matrices has $4^n - 1$ real parameters, exact QST can only be done if one records at least $4^n - 1$ different projections. Therefore, general QST remains prohibitive for quantum systems of moderate to large size due to the exponential growth in the number of required observations. However, for single-qubits, QST is tractable and provides a complete description of the quantum system, making it a powerful tool for conducting QCVV of individual qubits in QC hardware.

A. Maximum-Likelihood Quantum State Tomography

There exists multiple methods (i.e., statistical estimators) that can be used for QST. Each of these estimators comes with its own advantages and disadvantages such as improved reconstruction quality with respect to specific metrics, ease of implementation, or computational benefits. Some popular choices for QST include linear regression based methods [13], nuclear norm constrained reconstructions [14] and the Maximum-Likelihood Estimator (MLE) [15]. In this work, we adopted the MLE method as it is widely used in practice, leverages fundamental statistical theory principles and can be easily implemented for small quantum systems. However, a sensitivity study on simulated data suggested that all of these methods produce similar results under the specific data collection settings used in this work.

The Maximum-Likelihood approach for QST consists in finding the density matrix ρ that will maximize the probability of realized measurements. Our observables are measures k described by an ensemble of projectors $\{P_k\}$ that is the union of all measurement bases or more precisely the Projection-Valued Measures (PVMs) that we consider. What is recorded, and what serves as an input to the MLE algorithm, are the

number of times, n_k , that a particular measure k has been observed. The log-likelihood functions is then expressed in terms of our statistics as follows,

$$\ln \mathbb{P}[\{n_k\} | \rho] = \sum_k n_k \ln \text{Tr}[\rho P_k]. \quad (1)$$

The reconstructed density matrix ρ_{out} is the output of the following concave maximization problem on the positive semi-definite cone,

$$\rho_{\text{out}} = \arg \max_{\substack{\rho \succcurlyeq 0 \\ \text{Tr}[\rho]=1}} \sum_k n_k \ln \text{Tr}[\rho P_k] \quad (2)$$

Due to the limitations of off-the-shelf optimization software, in practice, enforcing the positive semi-definite (PSD) constraint, $\rho \succcurlyeq 0$, may require a specialized optimization algorithm. One possible approach to solving (2) consists in using local gradient ascent algorithm interleaved with eigendecomposition based projections onto the SDP cone. However, we will see that for single-qubit QST, the SDP constraint has a convenient simplification.

B. Specializing to Single-Qubit State Tomography

The Bloch vector representation of quantum states enables us to significantly simplify QST for an individual qubit. In this representation, density matrices are encoded as vectors $\vec{a} \in \mathbb{R}^3$ by $\rho = \frac{1}{2}(I + \vec{a} \cdot \vec{\sigma})$ where $\vec{\sigma}$ is the vector of Pauli matrices. The PSD constraint that ρ must satisfy is enforced through the requirement that its corresponding vector lies within the unit sphere, i.e. $\|\vec{a}\| \leq 1$. Projectors are described in a similar fashion by $P_{\vec{u}} = \frac{1}{2}(I + \vec{u} \cdot \vec{\sigma})$, where $\|\vec{u}\| = 1$. Any PVM is composed of only two projectors $P_{\vec{u}}$ and $P_{-\vec{u}} = I - P_{\vec{u}}$, and therefore can be identified by a single unit vector \vec{u} . Since we choose to perform the same number of measurements N in each PVMs, we only need to record the empirical probability $p_{\vec{u}} = n_{\vec{u}}/N$ of measuring the observable associated with \vec{u} . The MLE estimator from Eq. (2) is then simplified into the following program,

$$\vec{a}_{\text{out}} = \arg \max_{\|\vec{a}\| \leq 1} \sum_{\vec{u}} p_{\vec{u}} \ln(1 + \vec{a} \cdot \vec{u}) + (1 - p_{\vec{u}}) \ln(1 - \vec{a} \cdot \vec{u}). \quad (3)$$

The maximization problem in Eq. (3) is amenable to standard optimization software for it is a simple concave non-linear problem in \mathbb{R}^3 with the cumbersome SDP constraint from Eq. (2) replaced with a unit sphere constraint.

C. Single-Qubit State Tomography Experiment Design

Our main goal is to reconstruct the state of a single qubit programmed to be in the pure state $\rho_{\text{in}} := R_{\theta, \phi} |0\rangle\langle 0| R_{\theta, \phi}^\dagger$, where the rotation matrix is implemented through elementary rotations on the y and z axis $R_{\theta, \phi} = R_z(\phi)R_y(\theta)$ as depicted in Figure 1. The programmed density matrix ρ_{in} is represented on the Bloch sphere with the unit vector $\vec{a}_{\text{in}}(\theta, \phi) = (\sin(\theta) \cos(\phi), \sin(\theta) \sin(\phi), \cos(\theta))$. Using the MLE estimator from Eq. (3), we will assess the quality of the




Fig. 1: The state preparation procedure used to initialise a single-qubit quantum state where the elementary rotations are defined by $R_y(\theta) = \exp(-iY\theta/2)$ and $R_z(\phi) = (-iZ\phi/2)$.

state preparation for a specific qubit by comparing $\vec{a}_{\text{in}}(\theta, \phi)$ with the corresponding \vec{a}_{out} for multiple values of ϕ and θ .

In our experiments, we are limited to measurements in the computational basis, that is to say, to the PVM associated with the vector $(0, 0, 1)$. To remedy to this issue, we rotate the qubit before measuring it with the matrix $R_{\alpha, \beta}^\dagger$ to emulate a PVM on $\vec{u} = (\sin(\alpha) \cos(\beta), \sin(\alpha) \sin(\beta), \cos(\alpha))$. The rotation matrix is again implemented with elementary rotations $R_{\alpha, \beta}^\dagger = R_y(-\alpha)R_z(-\beta)$. The complete circuit is shown in Figure 2.




Fig. 2: The state preparation and tomography procedure used to initialise a single-qubit quantum state.

Conducting QST on a single qubit requires measurements from at least 3 different linearly independent PVMs. One natural choice are the Pauli PVMs along canonical axes (i.e., $(\alpha, \beta) \in \{(0, 0), (\pi/2, 0), (\pi/2, \pi)\}$). However, having more than the minimum number of PVMs can be beneficial in reducing the statistical error in state reconstruction. In this work we leverage a tetrahedral set of PVMs using $(\alpha, \beta) \in \{(0, 0), (\cos^{-1}(\frac{-1}{3}), 0), (\cos^{-1}(\frac{-1}{3}), \frac{2\pi}{3}), (\cos^{-1}(\frac{-1}{3}), \frac{-2\pi}{3})\}$ to provide a balance of mitigating statistical error and data collection requirements. The Bloch sphere vectors representing both the Pauli and tetrahedral PVMs are depicted in Figure 3.

The experimental procedure for conducting single-qubit QST proposed in this work uses the tetrahedral PVMs with the MLE model from (3) as follows. For a given state $\vec{a}_{\text{in}}(\theta, \phi)$ on the Bloch sphere, four variants of the tomography circuit (Figure 2) are executed, one for each (α, β) combination in the tetrahedral PVMs. State measurement statistics are collected for each of these circuits and converted into empirical probabilities $p_{\vec{u}}$ providing the information required for posing the MLE problem (3). The optimal solution to (3), \vec{a}_{out} , encodes the most likely density matrix that was implemented by the qubit at the state $\vec{a}_{\text{in}}(\theta, \phi)$.

An important subtlety in using a QST workflow in practice is to quantify how the accuracy of the empirical probability impacts the solution quality of model (3). On a quantum hardware platform one only has access to a finite number of measurements, N (a.k.a., *shots*), to estimate the empirical probabilities $p_{\vec{u}}$. If the number of measurements is not




Fig. 3: The Bloch sphere representation of PVMs that can be considered for single-qubit state tomography. The Pauli operators (left) provide a minimal set of PVMs while the tetrahedral PVMs used in this work (right) provides data redundancy and improved reconstruction accuracy. In this figure, $P_0 = |0\rangle$, $P_1 = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$, $P_2 = \frac{1}{\sqrt{2}}(|0\rangle + i|1\rangle)$, $T_0 = |0\rangle$ and $T_k = \frac{1}{\sqrt{3}}|0\rangle + \sqrt{\frac{2}{3}}e^{i\frac{2}{3}k\pi}|1\rangle$ for $k \in \{1, 2, 3\}$.

sufficient, finite sample errors can yield significant artifacts in \vec{a}_{out} . We performed a sensitivity study to quantify this statistical error on simulated data. For a given prepared state $\vec{a}_{\text{in}}(\theta, \phi)$, we sample $N = 20,000$ observations for each of the tetrahedral PVMs, and then reconstruct \vec{a}_{out} with our QST MLE algorithm. We repeat this procedure 10^4 times to accumulate an empirical histogram of the statistical error for each state, (θ, ϕ) , and estimate the 99-th percentile of the Euclidean distance error ϵ (i.e., $\mathbb{P}[\|\vec{a}_{\text{out}} - \vec{a}_{\text{in}}\| \leq \epsilon] = 0.99$). We find empirically that $\epsilon \leq 0.02$ for all angles (θ, ϕ) . Therefore, throughout this work we standardized on $N = 20,000$ observations per circuit to ensure that statistical fluctuations will only contribute to an error in the second digit, with high confidence.

III. VECTOR FIELD VISUALISATION OF SINGLE-QUBIT STATE TOMOGRAPHY

To illustrate the potential usefulness of single-qubit QST, this work investigates two questions arising in QCVV: (1) how accurate is single-qubit state preparation; (2) how consistent is the state preparation quality throughout the Bloch sphere. The first question amounts to quantifying the difference between the *ideal* quantum state, $\vec{a}_{\text{in}}(\theta, \phi)$, and the *reconstructed* quantum state resulting from the QST protocol, \vec{a}_{out} . The second question consists of repeating the state tomography protocol for a wide variety of possible state preparations and investigating how the reconstructed states vary. Throughout this section we develop the idea of a vector field visualisation for presenting the results of single-qubit QST to visually investigate these two questions. A variety of examples are then used to illustrate the usefulness of the proposed approach.

For a given single-qubit state, $\vec{a}_{\text{in}}(\theta, \phi)$, the tomography procedure described in Section II yields a density matrix ρ_{out} in the form of a vector in $\vec{a}_{\text{out}} \in \mathbb{R}^3$. Many metrics can be computed from \vec{a}_{out} such as quantum state purity ($\text{Tr}[\rho_{\text{out}}^2]$) or fidelity ($\text{Tr}[\sqrt{\sqrt{\rho_{\text{in}}}\rho_{\text{out}}}\sqrt{\rho_{\text{in}}}]$). A natural way to visualize

\vec{a}_{out} is a point within the Bloch sphere using the relation $\rho = \frac{1}{2}(I + \vec{a} \cdot \vec{\sigma})$ from Section II. However, this representation is difficult to interpret without an interactive visualization as many points within the sphere are co-located in standard orthographic projections, such as Figure 3. In particular, it is hard to distinguish if a given point is “on the front” or “on the back” of the sphere, it is hard to distinguish a pure state, that is represented on the surface of the Bloch sphere, from a mixed state, that is represented strictly within the Bloch sphere.

This work carefully combines three visualization tools to address the challenge of presenting this data. The first idea is to embed the single qubit states on the two-dimensional plane using established spherical projection methods. In this work we leverage the *Robinson projection* to place the QST results on a two-dimensional plot.¹ The second idea is to leverage a heatmap on this 2-dimensional representation to visualize a key metric of interest, such as the reconstructed state’s purity or fidelity. Throughout this work the color of the plot is used to present the state’s purity. The third idea is to use arrows to illustrate the locations difference between the *ideal* and *reconstructed* states, with each arrow starting at the *ideal* state and ending at the corresponding *reconstructed* state, projected onto the surface of the Bloch sphere. This communicates the rotational error that occurs in the *reconstructed* state. Overall, we call this representation of single-qubit QST the Vector Field Visualisation (VFV).


A. Vector Field Visualisation Examples

To verify the correctness of the proposed VFV approach, and to illustrate its usefulness, we begin with a series of QST studies of 200 quantum states spaced approximately equidistantly around the Bloch sphere using the `ibm_lagos` computer. The results are presented in Figure 4. As a validation exercise, we begin by performing the full QST protocol on an ideal quantum simulator in Figure 4a. This procedure yields reconstructed states with very high quality (purity 0.999, on average) with no rotational errors, confirming the correctness of the workflow’s implementation. The slight variability in these results illustrates the small errors that occur due to finite samples.


The second experiment, Figure 4b, consists of performing QST on data from a simulator of a noisy quantum computer. IBM’s Q-Hub tracks various performance properties of their qubits (e.g., gate and measurement errors) and provides an interface to perform noisy simulations using these properties. As expected, this procedure yields less consistent reconstructed states (purity 0.983, on average) with no rotational errors. Most notably, the simulated noisy qubit performance shows no rotational errors and is very homogeneous regardless of the quantum state that is being inspected.

The third and most interesting experiment, Figure 4c, consists of performing the state tomography protocol with


¹Note that any 2-dimensional projection of a 3-dimensional sphere is bound to imperfectly represent some of the features. The Robinson projection is a compromise and is neither angle-preserving nor equal-area, but a balance designed to reduce the overall distortion.



(a) Data obtained on an ideal simulator.



(b) Data obtained on a noisy simulator configured with the calibration data of `ibm_lagos` at the time of hardware data collection.



(c) Data obtained on `ibm_lagos`.

Fig. 4: Vector Field Visualisations of the reconstructed states using the single-qubit state tomography protocol from Section II using 20,000 shots on qubit 1 of `ibm_lagos`. The heatmap indicates the purity of the reconstructed quantum state and the horizontal red line in the color scale indicates the average purity of the states.

real data from the `ibm_lagos` computer. Three interesting observations can be made when comparing this result to the two simulations: (1) This is the first dataset where the vector field arrows are visible, indicating a type of rotational error that is not captured the simulations; (2) The purity of the reconstructed states is more heterogeneous than the simulators; (3) the purity of the reconstructed states is similar in average to the noisy simulator (0.983 compared to 0.972) but has a notably higher standard deviation (0.004 compared to 0.013). All of these observations indicate the potential for a state-dependent error model to better capture this qubit’s perfor-




Fig. 5: Quantum state-preparation circuit used to visualise the effect of state degradation over time. The original protocol is obtained by setting the delay time $t = 0$. Increasing t shows the errors that arise in idle open-quantum systems. On IBM Quantum chips, t is given as a multiple of a specific time dt . On `ibm_lagos`, $dt = \frac{2}{9}$ ns.

mance. Overall, the value of VFV approach is highlighted by the distinct motifs that are clear in each of these figures and may provide inspiration for new qubit performance measures and noise mitigation schemes.


B. Visualisation of State Degradation

One of the primary uses of data presentation tools like the VFV is to examine features of qubit performance that are not captured by current QC simulators. This can provide inspiration for which features could make simulators more faithful proxies of real-world QC hardware. To illustrate the point this section explores the impacts of adding a `delay(t)` instruction into the state preparation protocol, as shown in Figure 5. In this circuit the state preparation procedure used in the original protocol is now followed by a delay during which the qubit experiences the effects of an open-quantum system before the QST measurements are performed. Note that this delayed QST experiment would produce identical results in the case of *closed-system* simulations that are considered in Figure 4. This type of experiment is only interesting on QC hardware or simulations of open-quantum systems [16], [17].


Results of the time delay experiments are shown in Figure 6. Two notable observations can be made from these results. First we can see that the average of the state’s purity degrades steadily, starting with an average value of 0.972 and degrading to 0.959 and then 0.945. The second observation is that the rotational error of the state is also increasing steadily and non-uniformly with time; notice how the rotational bias at the top of the VFV is very different from the bottom. It is also suprising to see a systematic rotational shift from left to right that appears to be state-dependent. Although this study is only a proof-of-principle, it highlights the potential usefulness of VFV for designing models of open-quantum systems and provides some intuition for what decoherence *looks like* on this QC hardware.

IV. SIGNATURES OF SINGLE-QUBIT DATA CORRUPTION


In this section we explore how different QST reconstruction methods can be combined to identify additional issues with the performance of individual qubits. In particular, we compare reconstructed states using the MLE method with those produced by the Linear Regression (LR) method [13]. In short, LR QST identifies a density matrix that minimizes the difference between the observed and predicted measurement



(a) Data obtained on `ibm_lagos` with a delay of $t = 0$ dt.



(b) Data obtained on `ibm_lagos` with a delay of $t = 800$ dt.



(c) Data obtained on `ibm_lagos` with a delay of $t = 1600$ dt.

Fig. 6: Vector Field Visualisations of the reconstructed states of the state preparation with delay circuit using 20,000 shots on qubit 1 of `ibm_lagos` using three different delay values. The heatmap indicates the purity of the reconstructed quantum state and the horizontal red line in the color scale indicates the average purity of the states.

probabilities. For a single qubit system, the LR method solves the following optimization task on the Bloch sphere,

$$\vec{a}_{\text{out}}^{\text{LR}} = \arg \min_{\|\vec{a}\| \leq 1} \sum_{\vec{u}} (1 + \vec{u} \cdot \vec{a} - 2p_{\vec{u}})^2. \quad (4)$$

Using the statistical error analysis from Subsection II-C, we find that the 99-th percentile of the Euclidean distance error of LR is less than 0.02 for 20,000 observations, which is comparable to the MLE method.

The key insight of this section is that the difference in purity between the reconstructed states of MLE and LR should be very small, but in practice is it not always the case.

Specifically, when using identical input statistics, the 99-th percentile of the Euclidean distance between the state estimates of MLE and LR is also less than 0.02 with 20,000 observations (i.e., $\mathbb{P} [\|\hat{a}_{\text{out}}^{\text{LR}}\| - \|\hat{a}_{\text{out}}^{\text{MLE}}\| \leq 0.02] \geq 0.99$). However, in practice, we observe that some qubits produce data where the differences in state reconstructions by MLE and LR cannot be reasonably explained by statistical fluctuations, as shown in Figure 7. This suggests that the the measured statistics of these qubits are corrupted after the state preparation occurs. It is likely that the elementary rotations $R_y(-\alpha)$ and $R_z(-\beta)$ used for changing measurement basis introduce *state-dependant* errors that are incompatible with the QST models considered in this work. Nevertheless, the analysis in Figure 7 indicates that combining multiple tomography methods can be a useful and effective tool for identifying signatures of data corruption.

V. OPEN-SOURCE SOFTWARE IMPLEMENTATION


In the last couple of years, IBM’s Qiskit python package has emerged as a de facto standard for gate-based QC and can be used to access a wide variety of hardware platforms, even beyond those provided by IBM. As such, the libraries developed in this work leverage Qiskit for data collection on QC hardware. Even though Qiskit is the only supported framework for the moment, particular attention has been given to enable extending to other frameworks in the future.

The single-qubit QST protocol presented in this work is organized into two packages: `sqst` and `sqmap`. The first package, `sqst` (short for Single-Qubit Tomography), implements all the functions and interfaces to perform a single-qubit QST efficiently. In particular, it provides various single-qubit tomography basis and quantum state reconstruction methods. `sqst` also provides ways to parallelize a single-qubit experiment over all the available qubits in QC hardware for a quick assessment of the qubit performance across a full-chip. The second package, `sqmap` (short for Single-Qubit Map), provides various ways of visualising single-qubit tomography results obtained with `sqst`, such as the VFV figures presented in this work.


All of the QST procedures and plotting facilities used in this paper are available as open-source at <https://github.com/nelimee/sqst> and <https://github.com/nelimee/sqmap> and can be used freely by anyone to benchmark their QC platform. Installation of these packages can be done with the official Python package manager `pip`.

VI. CONCLUSION


As the variety of quantum computing platforms continues to increase so does the need of tools to inspect their performance characteristics. In this work we have demonstrated that quantum state tomography of individual qubits is a viable approach for inspecting qubit performance, although similar procedures are unlikely to scale to much larger systems due to the notable data collection requirements. Through the careful design of data collection, state reconstruction and result visualization, this work illustrates that the proposed QCVV procedures can highlight elusive patterns in qubit performance that are difficult



(a) Data obtained on `ibmq_belem` and reconstructed with MLE.



(b) Data obtained on `ibmq_belem` and reconstructed with LR.



(c) Reconstructed quantum states where the purity difference between MLE and LR reconstruction methods is above the 0.02 threshold.

Fig. 7: A comparison of the quantum state reconstruction methods MLE and LR on qubit 3 of `ibmq_belem`, where the reconstruction methods do not entirely agree. Each post-processing method is given the same raw data from `ibmq_belem`.

to capture with simpler metrics. The results indicate that there is room for improvement on the simulation models that are currently used and highlight the importance of modeling open-quantum system effects at medium time scales. A careful comparison different tomography methods also indicates that more general models for quantum state tomography should be considered to better capture the exotic effects that can be observed in current QC platforms. We hope that the proposed QCVV procedure and vector field visualization will be a valuable tool for the quantum computing community in evaluating qubit performance and have provided the implementation as open-source to support that aim.

ACKNOWLEDGMENT

This work was partly supported by the U.S. DOE through a quantum computing program sponsored by the Los Alamos National Laboratory (LANL) Information Science & Technology Institute and the Laboratory Directed Research and Development (LDRD) program of LANL under project numbers 20210116DR, 20210114ER and the Center for Non Linear Studies. This research was also partly supported by the U.S. Department of Energy (DOE), Office of Science, Office of Advanced Scientific Computing Research, under the Accelerated Research in Quantum Computing (ARQC) program. This research used quantum computing resources provided by the LANL Institutional Computing Program, which is supported by the U.S. Department of Energy National Nuclear Security Administration under Contract No. 89233218CNA000001. Adrien Suau also thanks TotalEnergies for their general support of this work. Adrien Suau also thank Gabriel Staffelbach and Aida Todri-Sanial for interesting and lively discussions around the paper.

REFERENCES

- [1] J. Eisert, D. Hangleiter, N. Walk, I. Roth, D. Markham, R. Parekh, U. Chabaud, and E. Kashefi, “Quantum certification and benchmarking,” *Nature Reviews Physics*, vol. 2, no. 7, pp. 382–390, Jul 2020. [Online]. Available: <https://doi.org/10.1038/s42254-020-0186-4>
- [2] J. Preskill, “Quantum Computing in the NISQ era and beyond,” *Quantum*, vol. 2, p. 79, Aug. 2018. [Online]. Available: <https://doi.org/10.22331/q-2018-08-06-79>
- [3] K. Wright, K. M. Beck, S. Debnath, J. M. Amini, Y. Nam, N. Grzesiak, J.-S. Chen, N. C. Pisenti, M. Chmielewski, C. Collins, K. M. Hudek, J. Mizrahi, J. D. Wong-Campos, S. Allen, J. Apisdorf, P. Solomon, M. Williams, A. M. DuCore, A. Blinov, S. M. Kreikemeier, V. Chaplin, M. Keesan, C. Monroe, and J. Kim, “Benchmarking an 11-qubit quantum computer,” *Nature Communications*, vol. 10, no. 1, p. 5464, Nov 2019. [Online]. Available: <https://doi.org/10.1038/s41467-019-13534-2>
- [4] Z. Hradil, “Quantum-state estimation,” *Phys. Rev. A*, vol. 55, pp. R1561–R1564, Mar 1997. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevA.55.R1561>
- [5] E. Magesan, J. M. Gambetta, and J. Emerson, “Scalable and robust randomized benchmarking of quantum processes,” *Phys. Rev. Lett.*, vol. 106, p. 180504, May 2011. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.106.180504>
- [6] E. Knill, D. Leibfried, R. Reichle, J. Britton, R. B. Blakestad, J. D. Jost, C. Langer, R. Ozeri, S. Seidelin, and D. J. Wineland, “Randomized benchmarking of quantum gates,” *Phys. Rev. A*, vol. 77, p. 012307, Jan 2008. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevA.77.012307>
- [7] E. Nielsen, J. K. Gamble, K. Rudinger, T. Scholten, K. Young, and R. Blume-Kohout, “Gate set tomography,” 2020.
- [8] A. W. Cross, L. S. Bishop, S. Sheldon, P. D. Nation, and J. M. Gambetta, “Validating quantum computers using randomized model circuits,” *Phys. Rev. A*, vol. 100, p. 032328, Sep 2019. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevA.100.032328>
- [9] S. Boixo, S. V. Isakov, V. N. Smelyanskiy, R. Babbush, N. Ding, Z. Jiang, M. J. Bremner, J. M. Martinis, and H. Neven, “Characterizing quantum supremacy in near-term devices,” *Nature Physics*, vol. 14, no. 6, pp. 595–600, Jun 2018. [Online]. Available: <https://doi.org/10.1038/s41567-018-0124-x>
- [10] J. Gambetta and S. Sheldon, “Cramming more power into a quantum device,” Published online at <https://www.ibm.com/blogs/research/2019/03/power-quantum-device/>, 2019, accessed: 03/28/2021.
- [11] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. S. L. Brandao, D. A. Buell, B. Burkett, Y. Chen, Z. Chen, B. Chiaro, R. Collins, W. Courtney, A. Dunsworth, E. Farhi, B. Foxen, A. Fowler, C. Gidney, M. Giustina, R. Graff, K. Guerin, S. Habegger, M. P. Harrigan, M. J. Hartmann, A. Ho, M. Hoffmann, T. Huang, T. S. Humble, S. V. Isakov, E. Jeffrey, Z. Jiang, D. Kafri, K. Kechedzhi, J. Kelly, P. V. Klimov, S. Knysh, A. Korotkov, F. Kostritsa, D. Landhuis, M. Lindmark, E. Lucero, D. Lyakh, S. Mandrà, J. R. McClean, M. McEwen, A. Megrant, X. Mi, K. Michielsen, M. Mohseni, J. Mutus, O. Naaman, M. Neeley, C. Neill, M. Y. Niu, E. Ostby, A. Petukhov, J. C. Platt, C. Quintana, E. G. Rieffel, P. Roushan, N. C. Rubin, D. Sank, K. J. Satzinger, V. Smelyanskiy, K. J. Sung, M. D. Trevithick, A. Vainsencher, B. Villalonga, T. White, Z. J. Yao, P. Yeh, A. Zalcman, H. Neven, and J. M. Martinis, “Quantum supremacy using a programmable superconducting processor,” *Nature*, vol. 574, no. 7779, pp. 505–510, Oct 2019. [Online]. Available: <https://doi.org/10.1038/s41586-019-1666-5>
- [12] Y. Wu, W.-S. Bao, S. Cao, F. Chen, M.-C. Chen, X. Chen, T.-H. Chung, H. Deng, Y. Du, D. Fan, M. Gong, C. Guo, C. Guo, S. Guo, L. Han, L. Hong, H.-L. Huang, Y.-H. Huo, L. Li, N. Li, S. Li, Y. Li, F. Liang, C. Lin, J. Lin, H. Qian, D. Qiao, H. Rong, H. Su, L. Sun, L. Wang, S. Wang, D. Wu, Y. Xu, K. Yan, W. Yang, Y. Yang, Y. Ye, J. Yin, C. Ying, J. Yu, C. Zha, C. Zhang, H. Zhang, K. Zhang, Y. Zhang, H. Zhao, Y. Zhao, L. Zhou, Q. Zhu, C.-Y. Lu, C.-Z. Peng, X. Zhu, and J.-W. Pan, “Strong quantum computational advantage using a superconducting quantum processor,” *Phys. Rev. Lett.*, vol. 127, p. 180501, Oct 2021. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.127.180501>
- [13] B. Qi, Z. Hou, L. Li, D. Dong, G. Xiang, and G. Guo, “Quantum state tomography via linear regression estimation,” *Scientific reports*, vol. 3, no. 1, pp. 1–6, 2013.
- [14] D. Gross, Y.-K. Liu, S. T. Flammia, S. Becker, and J. Eisert, “Quantum state tomography via compressed sensing,” *Phys. Rev. Lett.*, vol. 105, p. 150401, Oct 2010. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.105.150401>
- [15] Z. Hradil, “Quantum-state estimation,” *Phys. Rev. A*, vol. 55, pp. R1561–R1564, Mar 1997. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevA.55.R1561>
- [16] H. Chen and D. A. Lidar, “Hoqst: Hamiltonian open quantum system toolkit,” 2020.
- [17] J. Johansson, P. Nation, and F. Nori, “Qutip: An open-source python framework for the dynamics of open quantum systems,” *Computer Physics Communications*, vol. 183, no. 8, pp. 1760–1772, 2012. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0010465512000835>

LA-UR-22-23640