scientific reports

OPEN

Error rate reduction of single-qubit gates via noise-aware decomposition into native gates

Thomas J. Maldonado[1,2✉], Johannes Flick[3], Stefan Krastanov[4,5] & Alexey Galda[6,7,8]

In the current era of Noisy Intermediate-Scale Quantum (NISQ) technology, the practical use of quantum computers remains inhibited by our inability to aptly decouple qubits from their environment to mitigate computational errors. In this paper, we introduce an approach by which knowledge of a qubit's initial quantum state and the standard parameters describing its decoherence can be leveraged to mitigate the noise present during the execution of a single-qubit gate. We benchmark our protocol using cloud-based access to IBM quantum processors. On `ibmq_rome`, we demonstrate a reduction of the single-qubit error rate by 38%, from 1.6×10^{-3} to 1.0×10^{-3}, provided the initial state of the input qubit is known. On `ibmq_bogota`, we prove that our protocol will never decrease gate fidelity, provided the system's T_1 and T_2 times have not drifted above 100 times their assumed values. The protocol can be used to reduce quantum state preparation errors, as well as to improve the fidelity of quantum circuits for which some knowledge of the qubits' intermediate states can be inferred. This paper presents a pathway to using information about noise levels and quantum state distributions to significantly reduce error rates associated with quantum gates via optimized decomposition into native hardware gates.

Four decades after the conception of a quantum computer (QC)[1], its far-reaching computational potential remains abundantly clear[2]. Among the various physical systems whose quantum properties can be harnessed for computation[3–6], superconducting transmon qubits have demonstrated promise in their ability to realize scalable QCs[7–9] and have accordingly been made available to the public via cloud-based services offered by private companies such as IBM, Rigetti Computing, and Amazon. Despite the recent increase in availability, in the current era of Noisy Intermediate-Scale Quantum (NISQ) technology[10], our ability to utilize these machines to their full potential remains significantly inhibited by the computational errors that arise from interactions between the physical qubits and their environment.

While these detrimental interactions can be suppressed through the development of noise-resilient quantum hardware, the effective noise present during circuit execution can also be mitigated by optimizations in the compilation process[11]. Examples of software-based optimization protocols have been demonstrated across the full quantum computing stack, from high-level circuit depth compression via quantum-assisted quantum compiling[12] down to the optimization of individual native gates-the default quantum operations calibrated by the hardware provider-through the use of pulse-level control[13–15]. Some notable examples include noise tailoring via randomized compiling[16], dynamical decoupling of idle qubits[17], optimized state preparation via active reset[18,19], and measurement via excited state promoted readout[20,21].

In this paper, we pioneer a software-based optimization protocol for fidelity improvements of general single-qubit gates by leveraging knowledge of the qubit decoherence parameters to generate an optimized noise-aware decomposition into native hardware gates. The optimization of native gates themselves is a complementary task and a powerful noise-mitigation tool in its own right, but it requires pulse-level control, a level of hardware access both unfamiliar and inaccessible to many users of NISQ devices. The goal of this paper is to demonstrate a reduction of single-qubit error rates without the need for this lower level of access. By optimizing the decomposition of

[1]Department of Electrical and Computer Engineering, Princeton University, Princeton, NJ 08544, USA. [2]Department of Physics, Harvard University, Cambridge, MA 02138, USA. [3]Center for Computational Quantum Physics, Flatiron Institute, 162 Fifth Avenue, New York, NY 10010, USA. [4]Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA 02139, USA. [5]John A. Paulson School of Engineering and Applied Sciences, Harvard University, Cambridge, MA 02138, USA. [6]Menten AI, Inc., San Francisco, CA 94111, USA. [7]James Franck Institute, University of Chicago, Chicago, IL 60637, USA. [8]Computational Science Division, Argonne National Laboratory, Lemont, IL 60439, USA. ✉email: maldonado@princeton.edu

single-qubit gates without improving the native gates themselves, we demonstrate the efficacy of a protocol that is straightforward to implement at the gate level and requires minimal knowledge of the underlying Hamiltonian governing the qubit dynamics during gate execution. Accordingly, it can be easily adapted for use in QCs based on arbitrary physical systems. Our results demonstrate that it is possible to significantly improve the fidelity of single-qubits gates by leveraging knowledge of the qubit's initial state, along with its characteristic coherence times T_1 and T_2. We perform two randomized benchmarking (RB)[22] experiments: (i) on the `ibmq_rome`[23] quantum processor, we empirically determine the reduction of the single-qubit error rate offered by our optimization technique when the initial state of the input qubit is known, and (ii) on the `ibmq_bogota`[24] quantum processor, we analyze the sensitivity of our approach to the accuracy and drifts of the device's calibrated T_1 and T_2 coherence times. Our results demonstrate that it is possible to reduce the single-qubit error rate by up to 38% and that the approach is extremely robust against drifts and miscalibrations of $T_{1,2}$ coherence times, providing measurable fidelity improvements even when the $T_{1,2}$ values are up to 2 orders of magnitude different from the true values, i.e. off by a factor of 0.1 to 100. While we demonstrate our approach on two five-qubit IBM transmon devices, our optimization protocol is hardware-agnostic and assumes the two most prominent channels of Markovian noise in NISQ devices: relaxation and dephasing.

Native gates. We start by introducing IBM's native gate set and describing how the compiler of quantum circuits decomposes single-qubit gates into native gates. Throughout this paper, we use the terms single-qubit gate, rotation, and operation interchangeably. An arbitrary rotation U on the Bloch sphere can be parameterized by its Euler angles[25]. Concretely, $\forall U \in SU(2), \exists \alpha, \beta, \delta \in [0, 2\pi), \gamma \in [0, \pi]$ such that

$$U(\beta, \gamma, \delta) = e^{i\alpha} R_z(\beta) R_y(\gamma) R_z(\delta),\tag{1}$$

where R_z and R_y are rotations about the z- and y-axes, respectively. Since a global phase enacted on a quantum state has no physical effect, the value of α is irrelevant. Effectively, Eq. (1) defines the decomposition of an arbitrary single-qubit gate into three rotations, two about the z-axis and one about the y-axis. In this paper, we work exclusively with the native gates used by IBM in their standard decomposition framework[26]: rotations about the x-axis by integer multiples of $\pi/2$ and rotations about the z-axis by an arbitrary angle. We note that, in principle, IBM's native gate set can be expanded using pulse-level control[13,14], but in hopes of making our protocol easier to implement for those without this lower level of access, we maintain the standard native gate set described above. The decomposition outlined in Eq. (1) can be rewritten in terms of these native gates R_z and $R_x(\pm\pi/2)$:

$$U(\beta, \gamma, \delta) = e^{i\alpha} R_z(\beta) R_x(-\pi/2) R_z(\gamma) R_x(\pi/2) R_z(\delta).\tag{2}$$

Thus, any single-qubit gate can be applied via the sequential application of IBM's native gates defined by Eq. (2).

Before proceeding to the noise model, we make a brief note on the physical implementation of the native gate set on IBM quantum processors[26]. IBM quantum systems are built using fixed-frequency superconducting transmon qubits, wherein the qubits are manipulated using microwave pulses. The $R_x(\pm\pi/2)$ gates in the decomposition defined by Eq. (2) are implemented using calibrated microwave pulses, while the R_z rotations are realized as zero-duration "virtual" gates by adding a phase offset in software[27]. For the purposes of the noise model outline below, it is important that microwave pulses of fixed shape and duration are applied to qubits only to implement the $R_x(\pm\pi/2)$ native gates.

Noise model. While there are a number of noteworthy sources of single-qubit decoherence, including leakage[15] and non-Markovian noise[28], we limit our attention to amplitude damping (relaxation) and dephasing, respectively characterized by IBM's publicly reported T_1 and T_2 coherence times. All single-qubit noise specifications for `ibmq_rome`[23] and `ibmq_bogota`[24] were provided by IBM through Qiskit[29] and are tabulated in the Supplementary Information.

Based on the physical implementation of IBM's native gates discussed above, we model the noisy application of an $R_x(\pm\pi/2)$ gate as an instantaneous rotation, followed by decay and dephasing over time t_* equal to the gate duration. We emphasize that this model is an approximation, and though it is not necessarily exact, it captures enough of the noise dynamics for the purpose of this study. This is mathematically realized via the initial application of an $R_x(\pm\pi/2)$ unitary, followed by the appropriate Kraus operators. Accordingly, we model the noisy application of a single-qubit gate parameterized by Euler angles (β, γ, δ) by applying these Kraus operators after each instance of $R_x(\pm\pi/2)$ in Eq. (2); a pictorial representation can be found in Fig. 1. In this model, the noisy application of a single-qubit gate with Euler angles (β, γ, δ) will transform an initially pure state with Bloch sphere coordinates (θ, ϕ) into a mixed state with the following density matrix:

$$\rho_{(\beta,\gamma,\delta,\theta,\phi)} = \begin{bmatrix} a & b \\ b^* & 1-a \end{bmatrix},\tag{3}$$

$$a = \frac{1}{2}\left[\left(-\sin\gamma\cos(\phi+\delta)\sin\theta + \cos\gamma\cos\theta\right)(1-\lambda_A)^{3/2}\sqrt{1-\lambda_P} + 1 + \lambda_A\right],\tag{4}$$

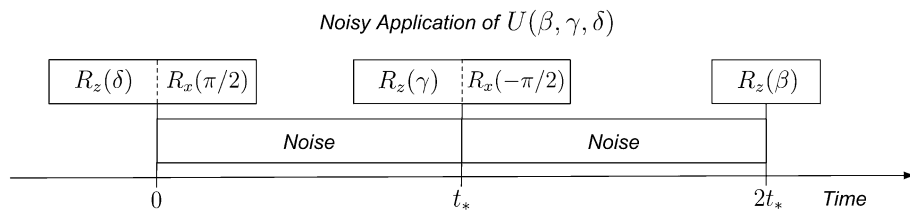*Noisy Application of $U(\beta, \gamma, \delta)$*

Figure 1. Noisy gate application: a single-qubit gate $U(\beta, \gamma, \delta)$ is decomposed into native gates according to Eq. (2). Each R_z gate is applied via a noiseless frame change. Each $R_x(\pm\pi/2)$ gate is applied via a microwave pulse and modeled as an instantaneous rotation, followed by decoherence over time t_* equal to the duration of the microwave pulse.

$$b = \frac{e^{-i\beta}}{2}\left[\left(\cos(\phi+\delta)\cos\gamma\sin\theta + \sin\gamma\cos\theta\right)(1-\lambda_A)(1-\lambda_P)\right.$$
$$\left. - i\left(\sin(\phi+\delta)\sin\theta(1-\lambda_A) + \lambda_A\right)\sqrt{1-\lambda_A}\sqrt{1-\lambda_P}\right]. \tag{5}$$

The derivation of this expression is provided in the Supplementary Information.

The variable λ_A is equal to the probability of a spontaneous emission during the application of an $R_x(\pm\pi/2)$ gate, and λ_P is equal to the probability of a spontaneous phase flip during the application of an $R_x(\pm\pi/2)$ gate. Both parameters are defined as functions of the system's T_1 and T_2 times, respectively, along with the $R_x(\pm\pi/2)$ gate duration t_*:

$$\lambda_A = 1 - e^{-t_*/T_1}, \tag{6}$$

$$\lambda_P = 1 - e^{-t_*/T_2}. \tag{7}$$

We note that all time-dependent terms in Eqs. (3–5) are functions of $t_*/T_{1,2}$.

Optimization.   Using the noise model given by Eqs. (3–5), we now outline the protocol by which the fidelity of an arbitrary single-qubit gate can be improved. Suppose we wish to implement the target operation $U(\beta_t, \gamma_t, \delta_t)$ acting on the initially pure state with Bloch sphere coordinates (θ, ϕ), represented below by $|\psi(\theta,\phi)\rangle$. Our protocol amounts to maximizing the following fidelity over the Euler angles (β, γ, δ):

$$F(\beta_t, \gamma_t, \delta_t, \beta, \gamma, \delta, \theta, \phi) = \langle\psi(\theta,\phi)|U(\beta_t, \gamma_t, \delta_t)^\dagger \rho_{(\beta,\gamma,\delta,\theta,\phi)}U(\beta_t, \gamma_t, \delta_t)|\psi(\theta,\phi)\rangle. \tag{8}$$

We find the optimal Euler angles $(\beta', \gamma', \delta')$ via gradient descent over the parameters (β, γ, δ). In the presence of noise, the native gate decomposition of $U(\beta', \gamma', \delta')$ will map the initial state $|\psi(\theta,\phi)\rangle$ to the target state $U(\beta_t, \gamma_t, \delta_t)|\psi(\theta,\phi)\rangle$ with higher fidelity than the default decomposition of $U(\beta_t, \gamma_t, \delta_t)$. We note that Eq. (8) has an explicit closed form and that in general, $U(\beta', \gamma', \delta') \neq U(\beta_t, \gamma_t, \delta_t)$, i.e., the optimized decomposition is not constrained to perform the target operation perfectly in the absence of noise. Throughout this study, the gradient descent was performed in Python using the function `scipy.optimize.minimize` from the SciPy[30] library with the method `L-BFGS-B`. For the parameters used in both experiments performed on IBM's hardware, the gradient descent to optimize a single gate could be performed in approximately 0.2 s on a standard computer.

To provide some intuition for how the optimized operation improves the fidelity, we begin by noting that the effect of amplitude damping is most pronounced on the south pole of the Bloch sphere (excited state), and the effect of phase damping is most pronounced on the equator of the Bloch sphere (equal superposition states). Thus, to best map the initial state to the target state, the optimizer finds a trajectory through the Bloch sphere that most aptly avoids these noisy regions. For the sake of visualization, we have included an example of an optimized (blue) and an unoptimized (red) trajectory through the Bloch sphere in Fig. 2, wherein the optimized trajectory tends towards the coherent north pole more than its unoptimized counterpart. Figure 2 was generated using QuTiP[31,32].

Results
Error rate reduction on `ibmq_rome`.   We now experimentally validate the protocol outlined above on qubit 3 of the `ibmq_rome` device by empirically determining the reduction in the single-qubit error rate achieved by the optimizer. At the time of the experiment, the qubit had the following characteristics: $T_1 = 46.4\,\mu s$, $T_2 = 105\,\mu s$, pulse duration $t_* = 35.6$ ns, and therefore damping probabilities $\lambda_A = 7.7 \times 10^{-4}$ and $\lambda_P = 3.4 \times 10^{-4}$ (see Eqs. (6–7)). Because the fidelity improvement offered by the optimizer is relatively small, we use an RB experiment to detect the improvement in fidelity by accumulating it over a long sequence of gates.

We begin by generating a circuit composed of a sequence of randomized single-qubit gates $C = (G_1, ..., G_N)$ acting on the initial state $|\psi_0\rangle = |0\rangle$. To randomize each gate $G_i \in C$, we sample its axis of rotation uniformly from the surface of the Bloch sphere and its angle uniformly from the interval $[0, 2\pi)$. We note that our method
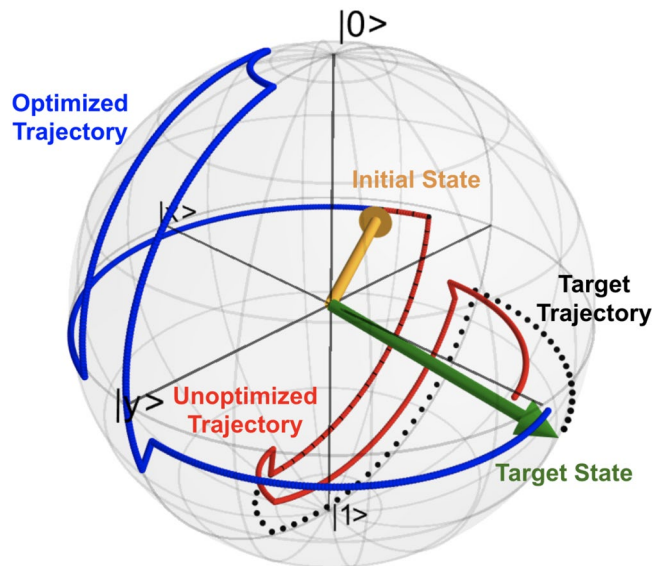
Figure 2. Optimized vs. unoptimized trajectories: an optimized trajectory (blue) and unoptimized trajectory (red) are mapped through the Bloch sphere in the presence of simulated noise corresponding to coherence times 100x shorter than the coherence times of `ibmq_rome` qubit 3 on its date of use (07/14/20). Also depicted is the initial state (orange), target trajectory (dotted black), and target state (green). To maximize fidelity, the optimized trajectory evolves the qubit through an intermediate state (after the $R_x(\pi/2)$ pulse) that avoids regions of the Bloch sphere (south pole and equator) that are more susceptible to noise.

for randomly generating single-qubit gates is one of many and that another popular approach is to sample from the Clifford Group[33]. We then optimize each gate G_i acting on the initial state $|\psi_{i-1}\rangle = G_{i-1}...G_1|0\rangle$. The result is a new circuit C_{opt} composed of the optimized versions of the gates in the circuit C. For each circuit C and C_{opt}, we denote the subsequence composed of the first d gates by C^d and C_{opt}^d, respectively. We measure the unoptimized fidelity after d gates by applying C^d, then applying the 2-pulse native gate decomposition of $(C^d)^{-1}$, and then measuring the probability of collapse to $|0\rangle$. Similarly, we measure the optimized fidelity by applying C_{opt}^d, followed by $(C^d)^{-1}$, and then measuring the probability of collapse to $|0\rangle$. In the presence of noise, the subsequence that maps the $|0\rangle$ state closer to the target state $|\psi_d\rangle = C^d|0\rangle$ will return the higher probability of measuring $|0\rangle$ after $(C^d)^{-1}$ is applied. We generate 10 circuits, each consisting of $N = 246$ randomized rotations. For each of the 10 circuits, we measure the unoptimized and optimized fidelities at circuit depths increasing by 7, $d \in \{1, 8, 15, ..., 246\}$. The empirically obtained fidelities at each depth d are then averaged over the 10 circuits. Readout errors were mitigated for all measurements made on `ibmq_rome` and `ibmq_bogota` by inverting a calibration matrix[34] composed of IBM's publicly reported readout error probabilities.

The results from simulating the execution of the circuits using our noise model and from executing the circuits on `ibmq_rome` qubit 3 are displayed in Fig. 3a,b, respectively. In both plots, the blue points represent the average fidelities of the optimized circuits, and the red points represent the average fidelities of the unoptimized circuits. The accumulation of noise in both circuits is reflected in the decrease in fidelity with circuit depth. We fit the data with the ansatz $f(x) = \frac{1}{2}(1 + \exp -ax)$ because it satisfies the limiting cases $f(0) = 1$ and $\lim_{x \to \infty} f(x) = \frac{1}{2}$. Thus, the error rate of one randomized single-qubit gate is given by the following expression:

$$\text{error rate} = 1 - \text{fidelity} = 1 - f(1) \approx \frac{a}{2}. \tag{9}$$

In the simulation and on the hardware, the optimized circuits outperformed the unoptimized circuits, thereby experimentally validating the optimization protocol put forth in this paper. Both the optimized and unoptimized circuits on the hardware have lower fidelities than in the simulation. We attribute this to two factors. First, our noise model only accounts for relaxation and dephasing, when in reality there are other noise channels present, such as sources of non-Markovian noise and leakage. Second, we made the approximation that a noisy single-qubit gate can be modeled as coherent evolution followed by decoherent evolution, when in reality the two occur simultaneously. We believe these two factors are the largest contributors to the lower fidelities observed on the hardware. Finally, from direct calculation of the optimized and unoptimized error rates defined by Eq. (9), we conclude that on `ibmq_rome` qubit 3, our optimization protocol reduces the error rate of a single-qubit gate acting on a known initial state by 38%, from 1.6×10^{-3} to 1.0×10^{-3}. The unoptimized error rate that we report here does not agree with IBM's reported error rate of 3.4×10^{-4}. We attribute this discrepancy to differences in methodology when calculating error rates, as well as to other sources of error not included in the model, such as coherent or calibration errors. The discrepancy is also likely attributable to a bias introduced by the 10 randomized gate sequences used in the experiment[35]. Nonetheless, we maintain our reported error rate reduction as
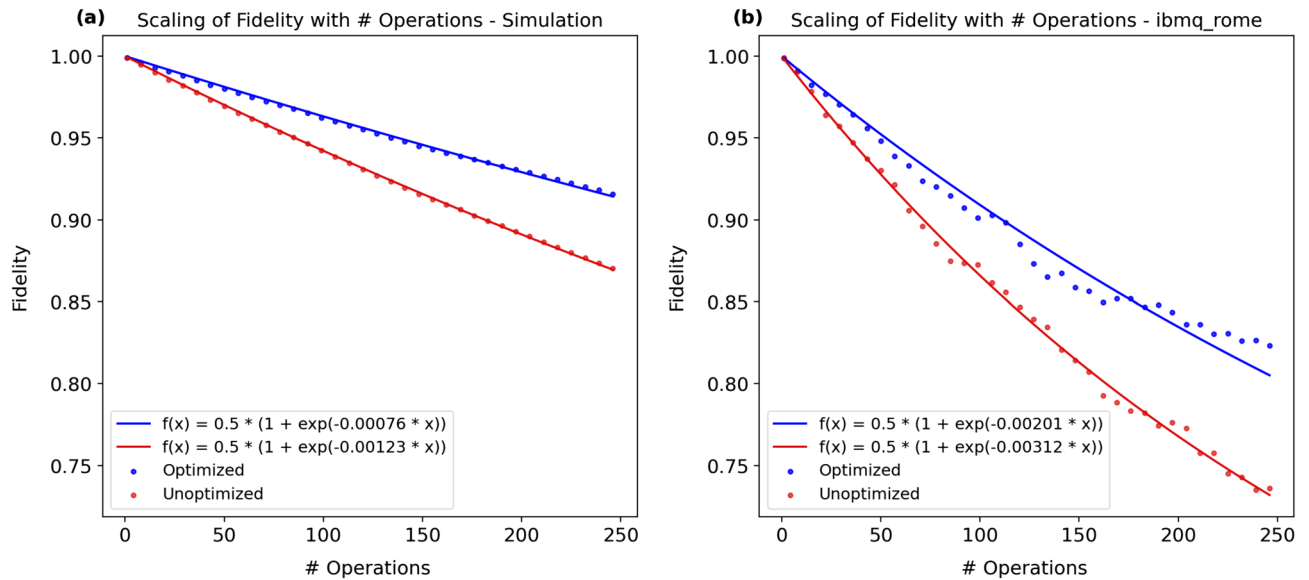
 4

Figure 3. Scaling of fidelity with number of operations: the fidelity (vertical axis) representing the overlap between the state output by the noisy application of each circuit (unoptimized and optimized) and the target state output by the noiseless application of the unoptimized circuit is plotted against the circuit depth (horizontal axis) at which the fidelity was measured. Each data point is the average fidelity of 10 randomized gate sequences with 8,192 shots per measurement. Measurements were taken at circuit depths increasing by 7, $d \in \{1, 8, 15, ..., 246\}$.

an estimate for the degree to which the noise present during the execution of a single-qubit gate can be mitigated by leveraging knowledge of the initial state of the input qubit and a description of the noise present during gate execution. An experiment of identical structure was carried out on Rigetti's Aspen-8[36] device and is detailed in the Supplementary Information.

Single-qubit state preparation. As an illustrative example application of the optimization protocol outlined above, we now analyze its ability to improve the preparation fidelity of a single-qubit state. Suppose that we wish to implement the target operation $U(\phi_t, \theta_t, 0)$ mapping the initial state $|0\rangle$ to the target state with Bloch sphere coordinates (θ_t, ϕ_t):

$$|\psi(\theta_t, \phi_t)\rangle = U(\phi_t, \theta_t, 0)|0\rangle. \tag{10}$$

Without loss of generality, we set $\delta = \theta = \phi = 0$ and reduce Eq. (8) to the following:

$$F(\theta_t, \phi_t, \beta, \gamma) = \langle \psi(\theta_t, \phi_t)|\rho_{(\beta, \gamma, 0, 0, 0)}|\psi(\theta_t, \phi_t)\rangle. \tag{11}$$

We find the optimal Euler angles $(\beta', \gamma', 0)$ via gradient descent over the parameters β and γ. In the presence of noise, the native gate decomposition of $U(\beta', \gamma', 0')$ will map the initial state $|0\rangle$ to the target state $|\psi(\theta_t, \phi_t)\rangle$ with higher fidelity than the default decomposition of $U(\phi_t, \theta_t, 0)$.

We now proceed by analyzing the relationship between the improvement in preparation fidelity offered by the optimizer and the amount of noise in the system. Since λ_A and λ_P are typically of comparable magnitude, we consider the improvement offered by optimization in the presence of noise parameterized by $\lambda = \lambda_A = \lambda_P$. For a fixed noise level λ, we randomly sample the target state $|\psi(\theta_t, \phi_t)\rangle$ uniformly from the surface of the Bloch sphere. We then find the optimal angles β' and γ' and simulate the application of $U(\phi_t, \theta_t, 0)$ and $U(\beta', \gamma', 0)$ on the input $|0\rangle$. Finally, we calculate the increase in fidelity to the target state $|\psi(\theta_t, \phi_t)\rangle$ and repeat this 100 times to find the average increase in preparation fidelity from optimization. The results from simulating the state preparation over a range of possible noise levels λ are displayed in Fig. 4.

Knowledge of the initial state. The above optimization protocol requires knowledge of the initial state to achieve an improvement in fidelity, as reflected in the fidelity function's explicit dependence on the initial state's Bloch sphere coordinates (see Eq. (8)). However, perfect knowledge of the initial state is not necessarily required for optimization. Provided a probability density function $p(\theta, \phi)$ for the distribution of the initial state over the Bloch sphere, we can optimize the target operation $U(\beta_t, \gamma_t, \delta_t)$ by maximizing the expected fidelity:

$$\langle F(\beta_t, \gamma_t, \delta_t, \beta, \gamma, \delta) \rangle = \int_0^{2\pi} \int_0^{\pi} F(\beta_t, \gamma_t, \delta_t, \beta, \gamma, \delta, \theta, \phi)p(\theta, \phi)d\theta d\phi. \tag{12}$$

The fidelity function $F(\beta_t, \gamma_t, \delta_t, \beta, \gamma, \delta, \theta, \phi)$ is defined in Eq. (8). We find the optimal Euler angles $(\beta', \gamma', \delta')$ via gradient descent over the parameters (β, γ, δ). In the presence of noise, the native gate decomposition of
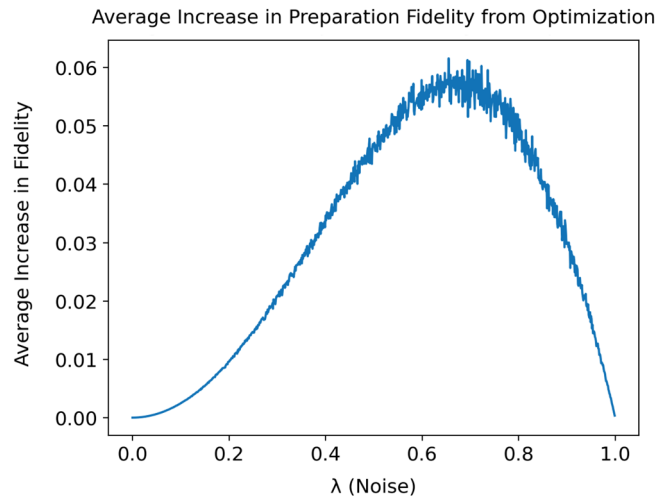
Figure 4. Average increase in preparation fidelity from optimization: the average increase in preparation fidelity (vertical axis) is plotted against the simulated noise level λ (horizontal axis). For each of the 1000 evenly spaced noise levels, the increase in fidelity is averaged from 100 randomly sampled target states.

$U(\beta', \gamma', \delta')$ will on average map an initial state $|\psi(\theta, \phi)\rangle$ sampled from the distribution defined by $p(\theta, \phi)$ to the target state $U(\beta_t, \gamma_t, \delta_t)|\psi(\theta, \phi)\rangle$ with higher fidelity than the default decomposition of $U(\beta_t, \gamma_t, \delta_t)$.

We proceed by analyzing the expected fidelity improvement offered by our protocol as a function of the initial state uncertainty. We have already shown a nontrivial improvement in expected fidelity provided perfect knowledge (i.e., minimal uncertainty) of the initial state $|0\rangle$ in Fig. 4. The other extreme corresponds to no knowledge (i.e., maximal uncertainty) of the initial state and is represented by the distribution in which all states are equally likely:

$$p(\theta, \phi) = \frac{1}{4\pi} \sin\theta . \tag{13}$$

The expected fidelity in this case is given by

$$\langle F(\beta_t, \gamma_t, \delta_t, \beta, \gamma, \delta) \rangle = \frac{1}{4\pi} \int_0^{2\pi} \int_0^{\pi} F(\beta_t, \gamma_t, \delta_t, \beta, \gamma, \delta, \theta, \phi) \sin\theta \, d\theta \, d\phi . \tag{14}$$

Further analysis of Eq. (14) shows

$$\frac{\partial \langle F \rangle}{\partial \beta}\bigg|_{(\beta, \gamma, \delta) = (\beta_t, \gamma_t, \delta_t)} = \frac{\partial \langle F \rangle}{\partial \gamma}\bigg|_{(\beta, \gamma, \delta) = (\beta_t, \gamma_t, \delta_t)} = \frac{\partial \langle F \rangle}{\partial \delta}\bigg|_{(\beta, \gamma, \delta) = (\beta_t, \gamma_t, \delta_t)} = 0 . \tag{15}$$

Regardless of the target operation and the amount of noise, the expected fidelity achieves a local maximum at $(\beta, \gamma, \delta) = (\beta_t, \gamma_t, \delta_t)$. We deduce that our protocol requires *some* knowledge of the initial state to improve expected fidelity. To further analyze this dependence, we examine the effect of maximizing the expected fidelity for a new probability density function:

$$p(\theta, \phi) = \begin{cases} \frac{\sin\theta}{2\pi(1-\cos\theta_{max})} & \text{if } \theta < \theta_{max} , \\ 0 & \text{if } \theta \geq \theta_{max} . \end{cases} \tag{16}$$

The initial state is now uniformly distributed over the portion of the Bloch sphere with polar angle less than θ_{max}. Accordingly, as θ_{max} approaches 0, we recover the case of state preparation, and for $\theta_{max} = \pi$, we recover the case of maximal uncertainty. By varying θ_{max} and examining its effect on the improvement in expected fidelity achieved by maximizing

$$\langle F(\beta_t, \gamma_t, \delta_t, \beta, \gamma, \delta) \rangle = \int_0^{2\pi} \int_0^{\theta_{max}} \frac{F(\beta_t, \gamma_t, \delta_t, \beta, \gamma, \delta, \theta, \phi) \sin\theta}{2\pi(1-\cos\theta_{max})} d\theta \, d\phi , \tag{17}$$

We quantify the degree to which an arbitrary single-qubit gate can be optimized as a function of the initial state uncertainty. To visualize this dependence, we begin by fixing the noise parameter $\lambda = \lambda_A = \lambda_P$ and the maximum polar angle θ_{max}. We proceed to randomly generate a target rotation $U(\beta_t, \gamma_t, \delta_t)$ by sampling the axis of rotation uniformly from the surface of the Bloch sphere and sampling the angle of rotation uniformly from the interval $[0, 2\pi)$. We then find the optimal Euler angles $(\beta', \gamma', \delta')$ via gradient descent of Eq. (17) over the parameters (β, γ, δ). Finally, we sample the initial state $|\psi(\theta, \phi)\rangle$ uniformly from the portion of the Bloch sphere with polar angle less than θ_{max}, simulate the application of $U(\beta', \gamma', \delta')$ and $U(\beta_t, \gamma_t, \delta_t)$ in the presence of noise

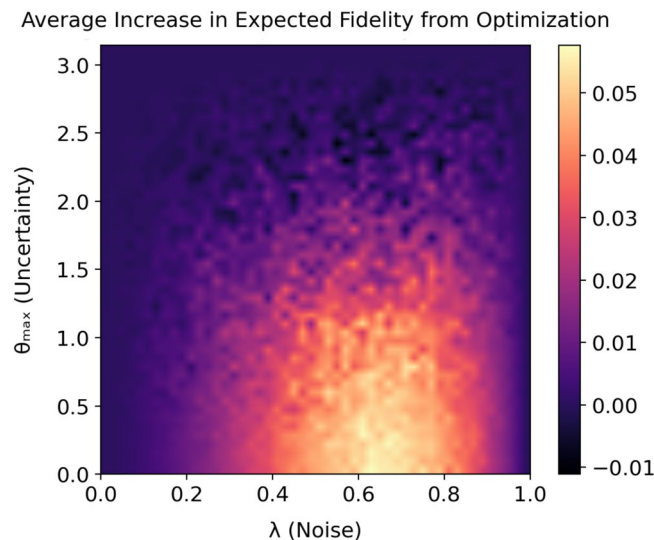Average Increase in Expected Fidelity from Optimization

Figure 5. Average increase in expected fidelity from optimization: the average increase in expected fidelity (color bar) is plotted against the simulated noise level λ (horizontal axis) and the uncertainty in the initial state θ_{max} (vertical axis). $\theta_{max} = 0$ corresponds to perfect knowledge of the initial state $|0\rangle$ and therefore recovery of Fig. 4; $\theta_{max} = \pi$ corresponds to no knowledge of the initial state and therefore no fidelity improvement. Each of the 50x50 data points was generated from 100 randomly sampled target operations and initial states.

Coherence Time Drifts - ibmq_bogota
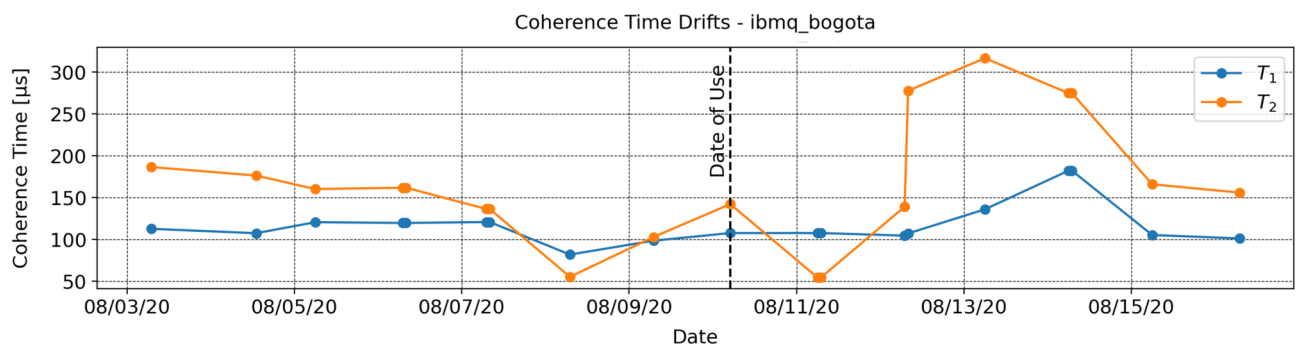
Figure 6. Coherence time drifts on `ibmq_bogota` qubit 2: IBM's publicly available coherence times on `ibmq_bogota` qubit 2 (vertical axis) are plotted against the date of calibration (horiztonal axis) up to one week before and after it was used in this study (08/10/20). Some of the reported coherence times violate $T_2 \leq 2 \cdot T_1$, thereby suggesting calibration errors.

parameterized by λ, and calculate the increase in fidelity to the target state $U(\beta_t, \gamma_t, \delta_t)|\psi(\theta, \phi)\rangle$. We note that this represents the increase in expected fidelity due to the random sampling of the initial state. We repeat this for 100 randomized rotations and calculate the average increase in expected fidelity. The results over a range of possible noise levels λ and polar angles θ_{max} are displayed in Fig. 5. For $\theta_{max} = 0$, we have perfect knowledge of the initial state $|0\rangle$ and accordingly recover Fig. 4. For $\theta_{max} = \pi$, we have no knowledge of the initial state and accordingly see no improvement in fidelity. Intuitively, the more we know about the initial state of the input qubit, the more we can fine-tune our optimization of the target operation to achieve a higher expected fidelity.

Sensitivity to coherence time drifts on `ibmq_bogota`. The optimization protocol outlined in this paper reduces the error rate of a single-qubit gate by employing the T_1 and T_2 times characterizing the qubit's decoherence. These noise parameters are obtained empirically and are therefore subject to error as a result of coherence time drifts between measurements (see Fig. 6). In the RB experiment outlined below, we demonstrate that our technique can improve single-qubit gate fidelity even when the assumed noise parameters are reasonably inaccurate.

In reality, the coherence times assumed by the optimizer are fixed, and the system's coherence times are free to drift. However, since the system's coherence times are not controllable parameters, we simulate the effect of a drift by intentionally providing the optimizer with inaccurate $T_{1,2}$ times. By varying the assumed $T_{1,2}$ times and measuring the fidelity of the optimized decomposition, we quantify the sensitivity of the optimizer to inaccurate
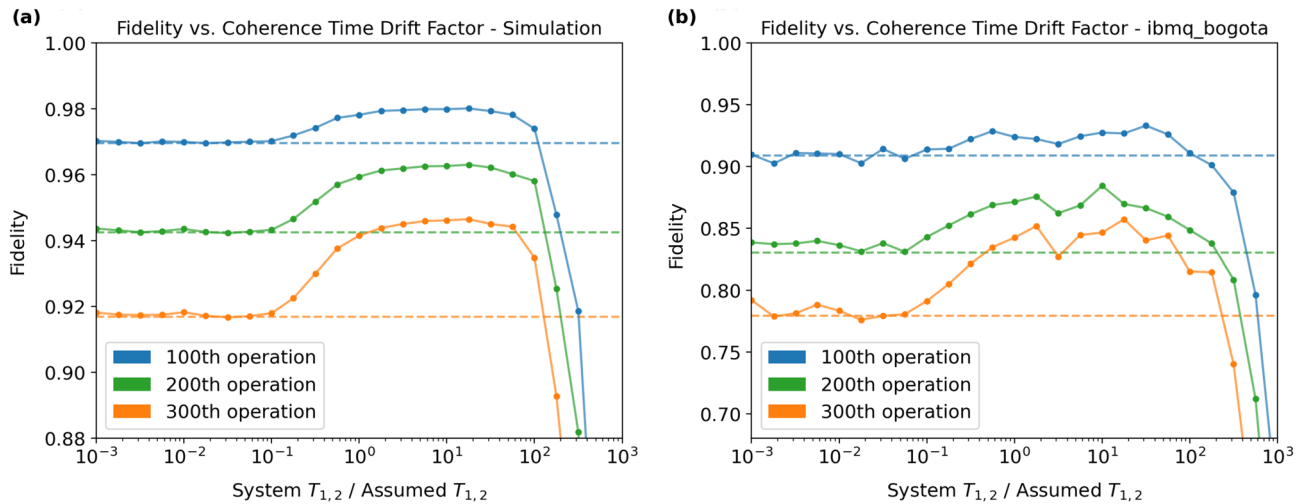
Figure 7. Fidelity vs. Coherence Time Drift Factor (System $T_{1,2}$/Assumed $T_{1,2}$): the fidelity (vertical axis) representing the overlap between the state output by the noisy application of the optimized circuit and the target state output by the noiseless application of the unoptimized circuit is plotted against the factor (horizontal axis) by which the system $T_{1,2}$ times differ from their assumed values during optimization. Horizontal dashed lines represent represent unoptimized fidelities. Each data point is the average fidelity of 10 randomized gate sequences with 16,384 shots per measurement. Measurements were taken at circuit depths increasing by 100, $d \in \{100, 200, 300\}$.

coherence times. Though T_1 and T_2 can generally drift by different factors, we simplify our experiment by simulating drifts of T_1 and T_2 from their assumed values by the same coherence time drift factor k:

$$(\text{system } T_{1,2}) = k \times (\text{assumed } T_{1,2}). \tag{18}$$

We justify this simplification in two ways. First, with limited reservable time on `ibmq_bogota` and limited computational resources for optimization, the experiment described in this section could only be made computationally feasible by reducing the scan of possible drifts from 2D to 1D. Second, though T_1 and T_2 do not generally drift by a common factor, they are indeed correlated, since $T_2 \leq 2 \cdot T_1$. We note that in many systems, $T_2 \ll T_1$, and T_2 is thus only weakly correlated with T_1. Nonetheless, up to one week before and after `ibmq_bogota` qubit 2 was used in this study, T_1 and T_2 were moderately correlated with a correlation coefficient of 0.68 (see Fig. 6). Albeit an approximation, the dimensionality reduction that made this experiment computationally feasible is thus statistically motivated. For each scaled pair of coherence times assumed by the optimizer, we perform the same RB experiment used to generate Fig. 3, now optimizing each gate in a 300-gate circuit and measuring the fidelity at depths of 100, 200, and 300 gates. The fidelities are once again averaged over 10 randomized gate sequences.

The results from simulating the execution of the circuits using our noise model and from executing the circuits on `ibmq_bogota` qubit 2 are displayed in Fig. 7a,b, respectively. Each color corresponds to a particular circuit depth, as indicated by the legend. For each circuit depth, the corresponding dashed line represents the fidelity obtained without any optimization. The accumulation of noise in all circuits is reflected in the decrease in fidelity with circuit depth. As with the first experiment, the fidelities observed on the hardware are lower than those predicted by the simulation. We once again attribute this to noise channels unaccounted for by our noise model, as well as to the approximation that noisy single-qubit gates can be modeled as coherent evolution followed by decoherent evolution.

As the factor k approaches 0, the optimizer assumes that there is no noise in the system. For any target rotation, in the absence of noise, the optimal rotation is the target rotation itself. Thus, as k approaches 0, we expect the fidelity offered by the optimizer to approach the unoptimized fidelity. This holds true in both the simulation data and the hardware data in Fig. 7. On the other hand, as k approaches ∞, the optimizer assumes an unrealistically large amount of noise in the system. The gradient descent is therefore performed over an assumed landscape unrepresentative of the system. The optimal rotation found by the optimizer will thus become uncorrelated with the target rotation, and the resulting fidelity will drop to 0.5 on average. This limit is also correctly captured in both the simulation data and the experiments performed on the quantum processor; however, since this regime is impractical, we have omitted much of this fidelity drop-off for the sake of visualization. In between the extremes, we expect to see a fidelity greater than the unoptimized fidelity. Once again, this is correctly reflected in both the simulation and experimental data. Most importantly, when the initial state of the input qubit is known, the optimizer will only decrease gate fidelity if the coherence times have drifted above about 100 times their assumed values. Since such drastic drifts are unrealistic (see Fig. 6), we conclude that the optimizer is extremely unlikely to decrease gate fidelity and is therefore robust against coherence time drifts.

Discussion

The optimization protocol outlined in this work reduces the error rate of single-qubit gates by leveraging knowledge of the initial state of the input qubit, along with the level of decoherence in the system, defined via T_1 and T_2 coherence times. On the `ibmq_rome` quantum processor, we proved that the protocol can reduce single-qubit error rates by 38%, from 1.6×10^{-3} to 1.0×10^{-3}, provided the initial state of the input qubit is known. On `ibmq_bogota`, we showed that the protocol always increases gate fidelity, provided the T_1 and T_2 times have not drifted below 0.1 or above 100 times their assumed values. The protocol can improve the expected fidelity of a single-qubit gate provided *some* knowledge of the initial state of the input qubit-the more localized the initial state distribution, the more we can improve the expected fidelity. The technique can be applied as a means to improve the fidelity of state preparation, as well as to improve the fidelity of quantum circuits for which some knowledge of intermediate states of qubits can be inferred; the former is well-suited for use in variational quantum eigensolvers (VQE), while the latter is likely only applicable in the near-term. Our optimization technique is not limited to the native gate set of IBM quantum hardware and can be adapted for optimization over an arbitrary native gate set.

In the noise model employed in this work, we used the approximation that native gates are applied instantaneously and are followed by a period of decoherence over the time interval equal to the gate duration. This approximation can be avoided by deriving a noise model from master equations involving the system's Hamiltonian during gate execution (e.g., the Linbladian for Markovian environments), which falls outside of the scope of this work. If the exact initial state of the input qubit is unknown, the optimization protocol can be modified to instead maximize the expected fidelity, averaged over the distribution of possible initial states. Our technique can also be adapted to optimize n-qubit gates by parameterizing all possible decompositions into an expanded native gate set (including entangling gates), and then proceeding as usual with the noise model and expected fidelity maximization for a specified distribution of n-qubit initial states. Since this work focused specifically on the optimization of single-qubit gates, gradient descent was sufficient for fidelity maximization; however, since the degrees of freedom required to parameterize an n-qubit gate are exponential in n, efficient alternatives to gradient descent are likely necessary to maximize the fidelity of larger n-qubit operations with more practical run-times.

In the near-term, while the execution of quantum circuits can still be simulated by classical computers, one can track the state of each qubit and use our protocol to optimize an arbitrary gate embedded within a quantum circuit. While the target qubit to said gate will generally be entangled with other qubits in the circuit, its reduced density matrix can be recovered by tracing over the degrees of freedom introduced by the qubits with which it is entangled. Though this entanglement will cause the qubit's initial state to be mixed, the optimization protocol can be easily adapted to accommodate mixed initial states, as is alluded to in the Supplementary Information. We emphasize that in the near-term, the protocol can be used to optimize all single-qubit gates on NISQ devices, beyond those involved in state preparation. This work presents a pathway to using information about noise levels and quantum state distributions to significantly reduce error rates associated with quantum gates via optimized decomposition into native gates.

Data availability

The datasets generated during and/or analyzed during the current study are available from the corresponding author on reasonable request.

References

1. Feynman, R. P. Simulating physics with computers. *Int. J. Theor. Phys* **21**, 467–488 (1982).
2. Shor, P. W. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM J. Comput.* **26**, 1484–1509. https://doi.org/10.1137/s0097539795293172 (1997).
3. Loss, D. & DiVincenzo, D. P. Quantum computation with quantum dots. *Phys. Rev. A* **57**, 120–126. https://doi.org/10.1103/PhysRevA.57.120 (1998).
4. Cirac, J. I. & Zoller, P. Quantum computations with cold trapped ions. *Phys. Rev. Lett.* **74**, 4091–4094. https://doi.org/10.1103/PhysRevLett.74.4091 (1995).
5. Nakamura, Y., Pashkin, Y. A. & Tsai, J. S. Coherent control of macroscopic quantum states in a single-cooper-pair box. *Nature* **398**, 786–788. https://doi.org/10.1038/19718 (1999).
6. Kane, B. E. A silicon-based nuclear spin quantum computer. *Nature* **393**, 133–137. https://doi.org/10.1038/30156 (1998).
7. Devoret, M. H. & Schoelkopf, R. J. Superconducting circuits for quantum information: An outlook. *Science* **339**, 1169–1174. https://doi.org/10.1126/science.1231930 (2013).
8. Wendin, G. Quantum information processing with superconducting circuits: A review. *Rep. Prog. Phys.* **80**, 10601. https://doi.org/10.1088/1361-6633/aa7e1a (2017).
9. Krantz, P. *et al.* A quantum engineer's guide to superconducting qubits. *Appl. Phys. Rev.* **6**, 021318. https://doi.org/10.1063/1.5089550 (2019).
10. Preskill, J. Quantum computing in the NISQ era and beyond. *Quantum* **2**, 79. https://doi.org/10.22331/q-2018-08-06-79 (2018).
11. Jurcevic, P. *et al.* Demonstration of quantum volume 64 on a superconducting quantum computing system. *Quantum Sci. Technol.* **6**, 025020. https://doi.org/10.1088/2058-9565/abe519 (2021).
12. Khatri, S. *et al.* Quantum-assisted quantum compiling. *Quantum* **3**, 140. https://doi.org/10.22331/q-2019-05-13-140 (2019).
13. Gokhale, P., Javadi-Abhari, A., Earnest, N., Shi, Y. & Chong, F. T. Optimized quantum compilation for near-term algorithms with openpulse. In *2020 53rd Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*, 186–200, https://doi.org/10.1109/MICRO50266.2020.00027 (2020).
14. Carvalho, A. R. R., Ball, H., Biercuk, M. J., Hush, M. R. & Thomsen, F. Error-robust quantum logic optimization using a cloud quantum computer interface (2020). arXiv:2010.08057.
15. Werninghaus, M. *et al.* Leakage reduction in fast superconducting qubit gates via optimal control. *npj Quantum Inf.* **7**, 1–6 (2021).
16. Wallman, J. J. & Emerson, J. Noise tailoring for scalable quantum computation via randomized compiling. *Phys. Rev. A* **94**, 052325. https://doi.org/10.1103/physreva.94.052325 (2016).

17. Pokharel, B., Anand, N., Fortman, B. & Lidar, D. A. Demonstration of fidelity improvement using dynamical decoupling with superconducting qubits. *Phys. Rev. Lett.* **121**, 220502. https://doi.org/10.1103/physrevlett.121.220502 (2018).
18. Magnard, P. et al. Fast and unconditional all-microwave reset of a superconducting qubit. *Phys. Rev. Lett.* **121**, 060502. https://doi.org/10.1103/PhysRevLett.121.060502 (2018).
19. Egger, D. et al. Pulsed reset protocol for fixed-frequency superconducting qubits. *Phys. Rev. Appl.* **10**, 044030. https://doi.org/10.1103/PhysRevApplied.10.044030 (2018).
20. Mallet, F. et al. Single-shot qubit readout in circuit quantum electrodynamics. *Nat. Phys.* **5**, 791–795. https://doi.org/10.1038/nphys1400 (2009).
21. Elder, S. S. et al. High-fidelity measurement of qubits encoded in multilevel superconducting circuits. *Phys. Rev. X* **10**, 011001. https://doi.org/10.1103/PhysRevX.10.011001 (2020).
22. Knill, E. et al. Randomized benchmarking of quantum gates. *Phys. Rev. A* **77**, 012307. https://doi.org/10.1103/PhysRevA.77.012307 (2008).
23. ibmq_rome v1.1.0, IBM Quantum team. Retrieved from https://quantum-computing.ibm.com (2020).
24. ibmq_bogota v1.0.0, IBM Quantum team. Retrieved from https://quantum-computing.ibm.com (2020).
25. Nielsen, M. A. & Chuang, I. *Quantum Computation and Quantum Information* (Cambridge University Press, 2002).
26. Cross, A. W., Bishop, L. S., Smolin, J. A. & Gambetta, J. M. Open quantum assembly language. *arXiv preprint* arXiv:1707.03429 (2017).
27. McKay, D. C., Wood, C. J., Sheldon, S., Chow, J. M. & Gambetta, J. M. Efficient z gates for quantum computing. *Phys. Rev. A* **96**, 022330. https://doi.org/10.1103/PhysRevA.96.022330 (2017).
28. Breuer, H.-P., Laine, E.-M., Piilo, J. & Vacchini, B. Colloquium: Non-Markovian dynamics in open quantum systems. *Rev. Mod. Phys.* **88**, 021002. https://doi.org/10.1103/RevModPhys.88.021002 (2016).
29. Aleksandrowicz, G. et al. Qiskit: An open-source framework for quantum computing. https://doi.org/10.5281/zenodo.2562110 (2019).
30. Virtanen, P. et al. SciPy 1.0: Fundamental algorithms for scientific computing in python. *Nat. Methods* **17**, 261–272. https://doi.org/10.1038/s41592-019-0686-2 (2020).
31. Johansson, J. R., Nation, P. D. & Nori, F. Qutip: An open-source python framework for the dynamics of open quantum systems. *Comput. Phys. Commun.* **183**, 1760–1772 (2012).
32. Johansson, J. R., Nation, P. D. & Nori, F. Qutip: An open-source python framework for the dynamics of open quantum systems. *Comput. Phys. Commun.* **184**, 1234–1240 (2013).
33. Helsen, J., Xue, X., Vandersypen, L. M. K. & Wehner, S. A new class of efficient randomized benchmarking protocols (2019). arXiv:1806.02048.
34. Nachman, B., Urbanek, M., de Jong, W. A. & Bauer, C. W. Unfolding quantum computer readout noise. *npj Quantum Inf.* **6**, 1–7. https://doi.org/10.1038/s41534-020-00309-7 (2020).
35. Mavadia, S. et al. Experimental quantum verification in the presence of temporally correlated noise. *npj Quantum Inf.* **4**, 1–9. https://doi.org/10.1038/s41534-017-0052-0 (2018).
36. Rigetti Computing Quantum Cloud Services, http://qcs.rigetti.com, accessed August, (2020).

Acknowledgements

Author contributions

T.M., J.F., S.K., and A.G. conceived the work, T.M. carried out the calculations, and all authors contributed to the interpretation of the results and the writing of the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary Information The online version contains supplementary material available at https://doi.org/10.1038/s41598-022-10339-0.

Correspondence and requests for materials should be addressed to T.J.M.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.