

Optimized Bayesian system identification in quantum devices

Thomas M. Stace¹,^{*} Jiayin Chen,¹ Li Li¹, Viktor S. Perunicic,¹ Andre R. R. Carvalho,¹ Michael Hush,^{1,*} Christophe H. Valahu^{2,3}, Ting Rei Tan^{2,3} and Michael J. Biercuk^{1,2,3}

¹*Q-CTRL, Sydney, NSW 2000, Australia*

²*School of Physics, University of Sydney, Sydney, NSW 2006, Australia*

³*ARC Centre of Excellence for Engineered Quantum Systems, University of Sydney, Sydney, NSW 2006, Australia*

 (Received 20 January 2023; revised 19 May 2023; accepted 13 September 2023; published 9 January 2024)

Identifying and calibrating quantitative dynamical models for physical quantum systems is important for a variety of applications. Here we present a closed-loop Bayesian learning algorithm for estimating multiple unknown parameters in a dynamical model, using optimized experimental “probe” controls and measurement. The estimation algorithm is based on a Bayesian particle filter, and is designed to autonomously choose informationally optimized probe experiments with which to compare to model predictions. We demonstrate the performance of the algorithm in both simulated calibration tasks and in an experimental single-qubit ion-trap system. Experimentally, we find that, with 60 times fewer samples, we exceed the precision of conventional calibration methods, delivering an approximately 93 times improvement in efficiency (as quantified by the reduction of measurements and resets required to achieve a target residual uncertainty and multiplied by the increase in accuracy). In simulated and experimental demonstrations, we see that successively longer pulses are selected as the posterior uncertainty iteratively decreases, leading to an exponential improvement in the accuracy of model parameters with the number of experimental queries, and a commensurate increase in the per-query time.

DOI: [10.1103/PhysRevApplied.21.014012](https://doi.org/10.1103/PhysRevApplied.21.014012)

I. INTRODUCTION

Accurately and efficiently identifying parameters in quantitative dynamical models that describe quantum systems is an open challenge in the development of quantum technologies [1]. This is an important task broadly known as system identification, which enables various applications, including system tuneup, error budgeting, and control design. More specifically, in the context of quantum computing, possession of an accurate and predictive system model can be used to develop high-fidelity error-robust quantum logic gates [2]; the accuracy of such models will be a key determinant [3] in our ability to reduce gate errors substantially below fault-tolerant error-correction thresholds [4,5].

Increasing the accuracy of parameter estimates typically involves the performance of comprehensive—but resource intensive—experimental processes [6] that can conflict with the fundamental premise that in real experiments quantum measurements are generally expensive. This is a consequence of both the exponential growth of the parameter space with system size, and fundamental measurement-induced noise, requiring repeated state preparation and evolution cycles to build up statistics of measurement

outcomes. Furthermore, measurements in real quantum devices are typically up to an order of magnitude slower than other quantum gates [7–11]. Finding measurement-efficient routines enabling high-fidelity parameter estimates is therefore critical for the advancement of useful experimental methods in parameter estimation.

Here we introduce a closed-loop learning-control algorithm that iteratively finds information-maximizing control pulses [12,13] in order to improve the accuracy of uncertain model parameters, such as Hamiltonian coefficients, with high efficiency in the number of measurements employed. The approach we adopt adaptively adjusts an interrogating control pulse waveform (duration, amplitude, phase, etc) in a manner similar to techniques used for high-dimensional control-parameter optimization [14]; here, informationally optimized probe waveforms are chosen based on a user-defined system model and the prior state of knowledge of the system parameters in a Bayesian framework. Importantly, this approach can accommodate realistic system constraints in the optimization process, such as control band limits, known linear or time-invariant transfer functions, and nonlinearities. The algorithm tracks the distribution of model-parameter values with a sample population, which is filtered and repopulated using a particle filter as new experimental data become available. This results in a sequence of posterior populations

*michael.hush@q-ctrl.com

[15,16] that incorporate the measured data and enable rapid convergence. We demonstrate that this algorithm delivers an exponential reduction in parameter uncertainty over iterations in both simulated and experimental tests using trapped atomic ions. In these tests the algorithm learns to make a commensurate increase in the probe duration in order to optimize the trade-off of probe sensitivity and uncertainty due to aliasing, and fully utilizes its freedom to design control probes for high-dimensional systems.

In Sec. II we introduce the context for system identification, including an illustrative two-parameter system model that also serves as a testbed for subsequent benchmarking. We then use the Fisher information [1,17] to elucidate important factors in system identification of uncertain model parameters. Section III describes in detail our optimized Bayesian system identification (OBSID) protocol, illustrated in Fig. 1, which is the central innovation in this manuscript. In Sec. IV we demonstrate the algorithm on simulated one- and two-qubit systems, described by models with up to five parameters. We also validate the protocol in an experimental calibration of a ion-trap system, and achieve parameter estimates in agreement with standard metrological approaches but requiring about 20 times fewer experimental measurements. We conclude with a discussion of OBSID applications and usage.

II. SYSTEM IDENTIFICATION CONTEXT

Formally, by *system identification* we mean the measurement and calibration of initially uncertain parameters, \mathbf{g} , in a quantitative model for some experimental system. This may include terms in a Hamiltonian, decay rates in a dissipative system, or calibration constants relating control pulses produced by a signal generator to the field experienced by a physical system, including, e.g., distortion or crosstalk.

To make explicit the ingredients for system identification, we assume the following.

(A1) *Initialization*—the experimental system can be reliably initialized into a specific reference state $|0\rangle$, such as the ground state or a separable product state.

(A2) *Control fields*—the system evolves under application of externally specified time-dependent control fields, $\mathbf{c} = \{c_1(t), c_2(t), \dots\}$ of duration T , which can be modulated by a programmable signal generator. For example, several qubits may be simultaneously addressed by several control fields.

(A3) *Measurements*—projective measurements are performed in orthogonal subspaces $\{\Pi_0, \Pi_1, \dots\}$, returning empirical probabilities $\hat{\mathbf{P}}(\mathbf{c}) = \{\hat{P}_0, \hat{P}_1, \dots\}$ and their corresponding uncertainties $\sigma = \{\sigma_0, \sigma_1, \dots\}$. We assume that

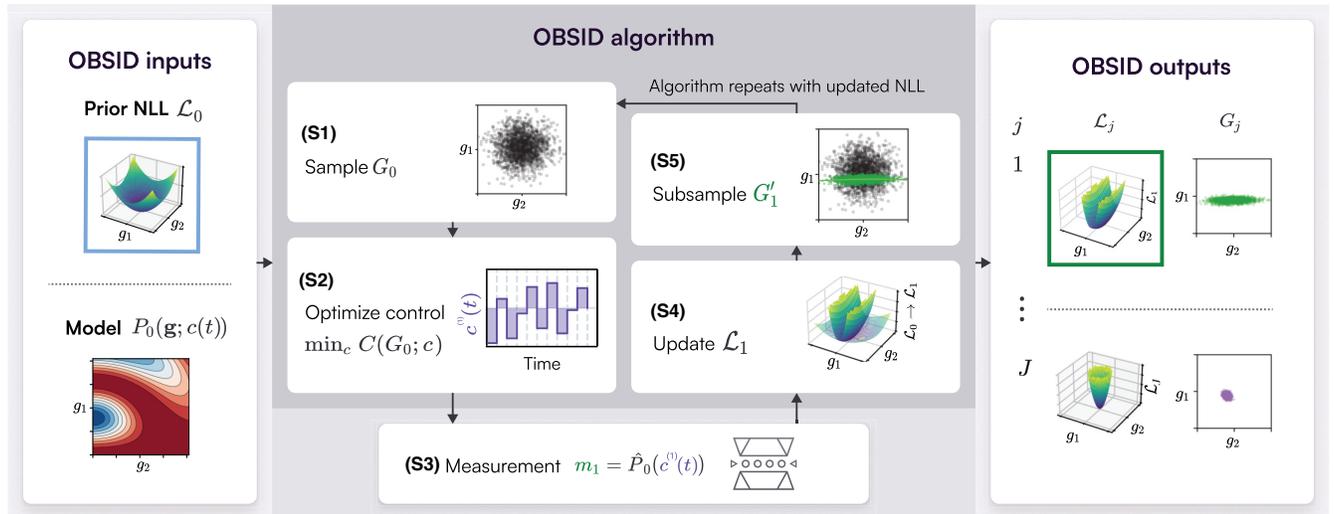


FIG. 1. A loop of the OBSID algorithm, shown in the dark gray box. The inputs to the OBSID algorithm are the prior negative loglikelihood (NLL) \mathcal{L}_0 (blue square), representing the prior uncertainty over the system parameters $\mathbf{g} = \{g_1, g_2, \dots\}$, and the system model. The prior NLL is used in (S1) to sample a prior population of model parameters, G_0 (gray points). (S2) An optimized control pulse $c^{(1)}(t)$ is found using the response $P_0(\mathbf{g}, c)$ predicted by the system model to maximize the anticipated parameter information gain, based on the cost C averaged over G_0 . (S3) The optimized control is used to drive an experiment. The experimental measurement outcome m_1 is used to (S4) update the posterior NLL, $\mathcal{L}_1 = \mathcal{L}_0 + \delta\mathcal{L}_1$, and (S5) subsample a posterior population, $G'_1 \subset G_0$ (green points), which is used to compute the posterior sample mean $\bar{\mathbf{g}}_1$ and covariance matrix Σ_1 . The updated NLL \mathcal{L}_1 (green square) becomes the prior for the next iteration of OBSID (light arrow). The output from the OBSID algorithm is the sequence of posterior NLLs \mathcal{L}_j and parameter populations G_j determined after each measurement. The algorithm completes after it reaches a target accuracy, iteration limit J_{\max} , or other termination criteria. Further details are given in Sec. III.

measurement outcomes are normally distributed around their mean $\hat{P}_i \sim \mathcal{N}_{P_i, \sigma_i^2}$.

(A4) *Model*—evolution of the system under investigation is well described by a quantitative model with p model parameters $\mathbf{g} = \{g_1, g_2, \dots, g_p\}$ that determine the evolution. For a given choice of \mathbf{g} and time-dependent control pulses \mathbf{c} , we use the model to compute probabilities $\mathbf{P}(\mathbf{g}, \mathbf{c}) = \{P_0, P_1, \dots\}$ for subspaces $\{\Pi_0, \Pi_1, \dots\}$. Comparing the experimental measurement $\hat{\mathbf{P}}(\mathbf{c})$ with model predictions $\mathbf{P}(\mathbf{g}, \mathbf{c})$ provides information about \mathbf{g} .

(A5) *Prior uncertainty*—the model parameters have a prior uncertainty that we assume is a multivariate normal distribution $\mathbf{g} \sim \mathcal{P}_0 = \mathcal{N}_{\bar{\mathbf{g}}_0, \Sigma_0}$, with prior mean $\bar{\mathbf{g}}_0$ and covariance matrix Σ_0 .

Here, we write an abstract distribution over the p -dimensional model parameter space as $\mathbf{g} \sim \mathcal{P}_j$, and the corresponding probability density as $\text{PDF}_j(\mathbf{g})$.

Assumptions (A1) to (A5) are practically reasonable, and are sufficient to define the goals of OBSID. A key insight from these assumptions is that, since we have a choice about what probe pulse to deliver to the experiment, there is an opportunity to optimize this choice of control in order to maximize the information gained about \mathbf{g} .

Within OBSID (Fig. 1), at each iteration of the algorithm, we numerically search for control pulses \mathbf{c} that maximize the anticipated information gain of proposed experiments, averaged over the prior parameter uncertainty. We then apply the optimized control to the physical (or simulated) system and measure the response $\hat{\mathbf{P}}$. This is compared with the model prediction \mathbf{P} , in order to update a posterior distribution over model parameters under a Bayesian framework.

We note that various of these assumptions can be relaxed if required. For example, mixed initial states in (A1) and open evolution in (A4) can be accommodated, albeit with additional computational cost. In (A3), generalized quantum measurements can be accommodated if necessary, as well as non-normal distributions in (A3) and (A5).

A. Illustrative model for probe-control optimization

We introduce an illustrative model that highlights the role of the control probe pulse in the efficiency of the system-identification procedure. The physical model we build upon embodies many salient features in OBSID, and is also used later in some of the one- and two-qubit benchmarking results we report in Sec. IV. It is a two-parameter model, $\mathbf{g}^{(1q)} = \{\Delta, \Omega\}$, for a single qubit driven by one control field, $c(t)$, with Hamiltonian

$$H^{(1q)}(t) = -\Delta Z/2 + \Omega(c(t)L + c(t)^*L^\dagger)/2, \quad (1)$$

where Δ is the detuning, $\Omega c(t)$ is the Rabi frequency, Ω is the Rabi calibration factor, Z is a Pauli operator, and

L is the corresponding lowering operator for a two-level system. We note that $c(t)$ has arbitrary units specified at the output of a signal generator (e.g., “volts”).

For this single-qubit model, we assume that the qubit is initialized in the ground state of the undriven Hamiltonian, $|0\rangle$, and measured in the undriven energy eigenbasis $\Pi_0 = \{|0\rangle\}$ and $\Pi_1 = \{|1\rangle\}$, so that $P_0 = |\langle 0|\psi(t)\rangle|^2$, and $P_1 = 1 - P_0$. Indeed, for the rest of this paper, we assume that system measurements are two-state projectors, in which we measure the *return* probability P_0 and the complement P_1 .

Using this model, Figs. 2(a)–2(c) illustrate the prototypical system response, $P_0(\mathbf{g}^{(1q)}; c(T))$ that reflects assumptions (A1) to (A5). Here T is the control pulse duration, assuming a Rabi-type control pulse of unit amplitude, $c(t) = 1$. We plot the response for three different pulse durations, $T = 1, 2$ and 8 “seconds” (using a normalized energy basis), in each case starting from the ground state $|0\rangle$. Shown in Figs. 2(a)–2(c) as a purple diamond is a putative (but arbitrary) “true” system parameter, $\mathbf{g}_t^{(1q)} = \{\Delta_t, \Omega_t\}$ Hz. For a given pulse, measurement of the system response will yield a measured probability \hat{P}_0 , indicated by the green contours that pass through $\mathbf{g}_t^{(1q)}$, up to measurement noise. Illustrating assumption (A5), the white 1σ -uncertainty disk in each panel represents prior uncertainty in the parameters.

This model gives us a simple example to illustrate the effect of different pulse choices (e.g., durations) on the posterior uncertainty. For the shortest pulse shown in Fig. 2(a), P_0 is a relatively slowly varying function over parameter space. For longer pulses [Figs. 2(b) and 2(c)], P_0 becomes increasingly oscillatory. The slowly varying response landscape of the shortest pulse has a unique contour cutting the uncertainty disk, highlighted in green. However, because the response is slowly varying (i.e., shallow gradient), the posterior distribution will be loosely clustered around the green cutting contour, with relatively large spread off the contour. That is, the response has low sensitivity, so the posterior distribution will have relatively large posterior uncertainty.

By contrast, for the long-duration control pulse shown in Fig. 2(c), the rapidly varying landscape has high sensitivity to the parameter values and, consequently, the posterior distribution would be relatively tightly clustered around the consistent cutting contours (green lines). However, since there are multiple consistent contours cutting through the uncertainty disk, the posterior distribution will still have wide support across the disk. For an experimentalist, this is a manifestation of the familiar $\text{mod } 2\pi$ uncertainty associated with a periodic function such as the measured population under Rabi oscillations.

The response at the intermediate duration shown in Fig. 2(b) with $T = 2$ is ideal (in this example). Its sensitivity is as large as possible, giving a posterior distribution that is tightly clustered around the green contour, while still

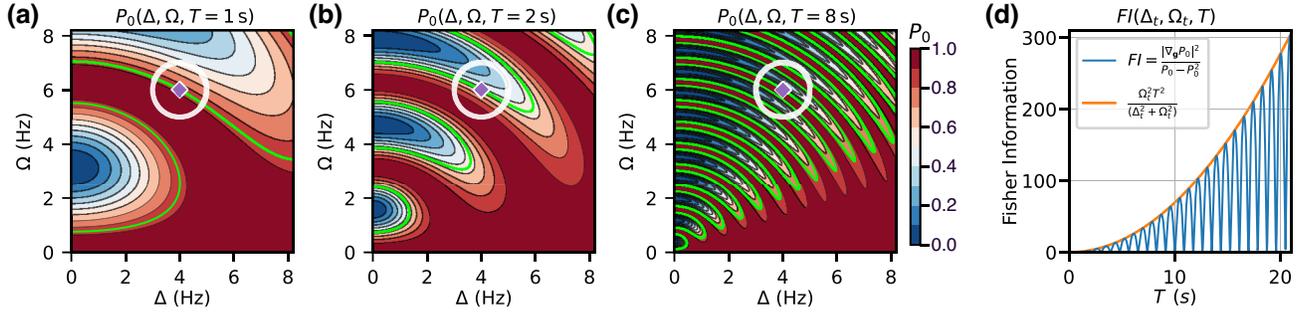


FIG. 2. Contours of the return probability landscape $P_0(\Delta, \Omega; T)$ for the one-qubit, two-parameter model from Eq. (1), driven by a uniform (Rabi-type) control pulse $c(t) = 1$ of duration (a) $T = 1$, (b) $T = 2$, and (c) $T = 8$ “seconds”; note that the units here are entirely nominal. In each plot, we indicate a specific point at $\mathbf{g}_t^{(1q)} = \{\Delta_t, \Omega_t\}$ Hz (purple diamond). For a system whose true parameters were $\mathbf{g}_t^{(1q)}$, the green contour cutting through $\mathbf{g}_t^{(1q)}$ represents the set of parameter values that would be consistent with an experimental measurement, \hat{P}_0 , using a control pulse of the corresponding pulse duration, T . The white 1σ -uncertainty disk pictured is helpful for illustrating the competing effects that lead to an optimal pulse duration, as described in the main text. Panel (d) shows the Fisher information for this model, which grows as T^2 with an oscillatory modulation.

having a unique contour cutting the disk. This illustrates a pulse for which the anticipated posterior uncertainty will be small. Qualitatively, the intermediate pulse duration represents a “Goldilocks” choice with highest information gain about the parameters, given the indicated uncertainty region.

More generally, the trade-off identified above between the response sensitivity and the density of consistent contours within a prior uncertainty region illustrates the possibility for varying the duration (and shape) of the control signal \mathbf{c} to maximize the information gain about the model parameters \mathbf{g} . Iterating, as the uncertainty region shrinks, the control pulse should be adapted accordingly.

B. Sensitivity, uncertainty, and Fisher information

To help formalize these observations, we introduce the Fisher information (FI), which is a widely used proxy for the sensitivity of a probability distribution with respect to its parameters [17]. Typically, larger FI corresponds to higher sensitivity with respect to the model parameters, and correspondingly higher potential information gain about unknown parameters. The FI also provides intuition about the “geometry” of information in the p -dimensional model parameter space.

For a multiparameter problem, with the two-state response probabilities $\{P_0, P_1 \equiv 1 - P_0\}$, the Fisher information matrix (FIM) [1] is given by

$$\text{FIM}(\mathbf{g}; \mathbf{c}) = \vec{f}(\mathbf{g}; \mathbf{c}) \boxtimes \vec{f}(\mathbf{g}; \mathbf{c}), \quad (2)$$

where $\vec{f}(\mathbf{g}; \mathbf{c}) = (P_0 - P_0^2)^{-1/2} \nabla_{\mathbf{g}} P_0(\mathbf{g}; \mathbf{c})$ is the FI eigenvector, and $a \boxtimes b$ denotes the outer product of a and b . For a two-state response, the FIM is a rank-1 matrix. The FI is

the unique nonzero eigenvalue of the FIM

$$\text{FI}(\mathbf{g}; \mathbf{c}) = |\vec{f}(\mathbf{g}; \mathbf{c})|^2. \quad (3)$$

For later discussion, we refer to the unit vector

$$\tilde{f}(\mathbf{g}; \mathbf{c}) = \vec{f}(\mathbf{g}; \mathbf{c}) / |\vec{f}(\mathbf{g}; \mathbf{c})| = \nabla_{\mathbf{g}} P_0 / |\nabla_{\mathbf{g}} P_0| \quad (4)$$

as the *information direction*. This defines the direction at each point in parameter space along which measurement of $\hat{P}_0(\mathbf{c})$ reveals information, and is locally orthogonal to contours of constant P_0 . It follows that, for a p -dimensional model space, we need at least p independent control fields, $\{\mathbf{c}^{(1)}, \mathbf{c}^{(2)}, \dots, \mathbf{c}^{(p)}\}$, to generate a spanning set of information directions, $\{\tilde{f}(\mathbf{g}; \mathbf{c}^{(1)}), \dots, \tilde{f}(\mathbf{g}; \mathbf{c}^{(p)})\}$, that constrain the posterior distribution over all p parameters. These statements can be generalized to multistate measurement outcomes.

The FI is a function of the control pulse envelope, and it is generally very strongly dependent on the overall control pulse duration T . This is illustrated in Fig. 2(d), which shows the FI as a function of T for the single-qubit model generated by Eq. (1) with $c(t) = 1$. The envelope grows as T^2 , with an oscillatory modulation.

The quadratic growth of FI with pulse duration is a fairly generic property [18]. For unitary evolution, when $g_i \in \mathbf{g}$ are energetic scales in a Hamiltonian, then

$$\text{FI}(\mathbf{g}, T) = T^2 \vartheta(\mathbf{g}, T), \quad (5)$$

where ϑ is a bounded and generally oscillatory function of pulse duration and the Hamiltonian parameters (see the supporting discussion in Appendix A). It follows that large FI will be obtained by increasing $T \rightarrow \infty$. However, given the foregoing “Goldilocks” argument that the optimal *finite* duration, $T^{(\text{opt})} < \infty$, should depend on the prior parameter

uncertainty, we see that maximizing FI alone is not sufficient to define informationally optimal control pulses given finite prior uncertainty.

This qualitative discussion gives rise to a rough estimate of the optimal control pulse duration, assuming that parameters in \mathbf{g} are energy scales in a Hamiltonian. Firstly, we assume that the prior parameter uncertainty is specified by a covariance matrix, Σ , that determines an uncertainty ellipse in parameter space [19] (e.g., illustrated in Fig. 2 as a white disk). Then the length of the major axis of the ellipse is given by the largest eigenvalue λ^{maj} of the deviation matrix $\Sigma^{1/2}$, which corresponds to the maximally uncertain combination of model parameters. Secondly, both P_0 and the FI oscillate in parameter space with a period $|\delta\mathbf{g}| \sim 1/T$ (and taking $\hbar = 1$ so that energy and time are reciprocal units).

Optimal pulses tend to be those for which the oscillation period in parameter space matches the major uncertainty, i.e., $|\delta\mathbf{g}| \approx \lambda^{\text{maj}}$, corresponding to the ‘‘Goldilocks’’ situation depicted in Fig. 2(b). It follows that the duration of the optimal pulse is $T^{\text{(opt)}} \sim 1/\lambda^{\text{maj}} \gtrsim \text{Tr}(\Sigma^{-1/2})$. This is an important estimate for two reasons: it gives intuition for how optimal pulse durations relate to the size of the parameter uncertainty ellipse, and it gives a practically useful order of magnitude for initializing numerical optimization over T in OBSID.

III. OPTIMIZED BAYESIAN SYSTEM IDENTIFICATION PROTOCOL

A. OBSID algorithm overview

The OBSID protocol starts from a prior uncertainty in the model parameters, \mathcal{P}_0 , and with the aid of the system model that computes \mathbf{P} , finds a control pulse that is likely to optimally improve our knowledge of the model parameters. We encode the anticipated knowledge gain from an arbitrary pulse in terms of a cost function, C . Using this optimized control, an experiment is performed, and the measurement results are incorporated into a posterior distribution, \mathcal{P}_1 , over model parameters. Iterating this yields a sequence of probability distributions \mathcal{P}_j , quantified by their density functions, $\text{PDF}_j(\mathbf{g})$, that iteratively and autonomously localizes the model parameters.

A loop of the OBSID algorithm is illustrated in Fig. 1. The inputs to the algorithm are the system model with which to compute $P_0(\mathbf{g}, \mathbf{c})$, and the initial prior distribution, \mathcal{P}_0 over parameters \mathbf{g} , represented as a negative-loglikelihood (NLL) function, $\mathcal{L}_0(\mathbf{g}) = -\ln(\text{PDF}_0(\mathbf{g}))$, shown as a two-dimensional surface inside the blue square in Fig. 1. Numerically, we track the sequence of posterior distributions, \mathcal{P}_j , using a Bayesian population filter [15], which form the output of the algorithm.

The algorithm itself, shown inside the dark gray box in Fig. 1, loops over the following sample, optimize, measure, and inference steps, initialized with $j = 1$, and terminating

after reaching a desired accuracy, iteration limit J , or other terminal conditions.

(S1) Sample a prior population, $G_{j-1} = \{\mathbf{g}_s\}_{s=1}^S$, from the prior distribution $\text{PDF}_{j-1}(\mathbf{g}) = \exp(-\mathcal{L}_{j-1}(\mathbf{g}))$, so that $\mathbf{g}_s \sim \mathcal{P}_{j-1}$.

(S2) Minimize a population cost function $C(G_{j-1}, \mathbf{c})$ over pulses \mathbf{c} , to find an optimal pulse $\mathbf{c}^{(j)} = \text{argmin}_{\mathbf{c}} C(G_{j-1}, \mathbf{c})$. The optimization uses the model predictions, $P_0(\mathbf{g}_s, \mathbf{c})$, at each $\mathbf{g}_s \in G_{j-1}$, and also model gradients $\nabla_{\mathbf{g}} P_0(\mathbf{g}_s, \mathbf{c})$.

(S3) Run an experiment with the optimal pulse $\mathbf{c}^{(j)}$ to measure $m_j = \hat{P}_0(\mathbf{c}^{(j)})$, with uncertainty σ_j determined by the measurement statistics.

(S4) Update the posterior NLL over $\mathbf{g}_s \in G_{j-1}$, so that $\mathcal{L}_j(\mathbf{g}_s | m_j) = \mathcal{L}_{j-1}(\mathbf{g}_s) + \delta\mathcal{L}_j(\mathbf{g}_s, m_j)$, where

$$\begin{aligned} \delta\mathcal{L}_j(\mathbf{g}, m_j) &= -\ln(\delta L_j(m_j | \mathbf{g})), \\ &\equiv (P_0(\mathbf{g}, \mathbf{c}^{(j)}) - m_j)^2 / (2\sigma_j)^2, \end{aligned} \quad (6)$$

assuming that $\delta L_j(m | \mathbf{g}) = \mathcal{N}_{P_0(\mathbf{g}, \mathbf{c}^{(j)}), \sigma_j}(m)$.

(S5) Return

(1) a posterior likelihood subsample $G'_j \subset G_{j-1}$, i.e., retain $\mathbf{g}_s \in G_{j-1}$ with probability $\delta L_j(\mathbf{g}_s, m_j)$,

(2) the posterior sample mean, $\bar{\mathbf{g}}_j(m_j) = \text{E}_S(\mathbf{g} | m_j)$,

(3) the posterior sample covariance $\Sigma_j(m_j) = \text{E}_S((\mathbf{g} - \bar{\mathbf{g}}_j) \boxtimes (\mathbf{g} - \bar{\mathbf{g}}_j) | m_j)$,

where the sample expectation of x conditioned on m is $\text{E}_S(x | m) = \sum_s x_s \delta L_j(\mathbf{g}_s, m)$, and we have defined the sample-normalized likelihood

$$\tilde{\delta L}_j(\mathbf{g}, m) = \delta L_j(m | \mathbf{g}) / N_j(m), \quad (7)$$

where $N_j(m) = \sum_s \delta L_j(m | \mathbf{g}_s)$ and δL_j is defined implicitly in Eq. (6).

(S6) Increment j or terminate.

The iterative outputs of the loop constitute a sequence of posterior quantities $\{\bar{\mathbf{g}}_j(m_j), \Sigma_j(m_j), G_j\}_{j=1}^J$, which increasingly localize the system parameters as measurement data are accumulated and processed.

B. Comments on implementation

Step (S1) is implemented using importance sampling to sample a population from the prior NLL, \mathcal{L}_{j-1} [20]. This is depicted as a population point cloud in Fig. 1. Importance sampling is implemented by generating a trial presample $T_{j-1} = \{\mathbf{g}_s^T\}$ from a proposal distribution, which we take to be the multinormal distribution so that $\mathbf{g}_s^T \sim \mathcal{T}_{j-1} = \mathcal{N}_{\bar{\mathbf{g}}_{j-1}, \Sigma_{j-1}}$, and then assigning an importance

score

$$i_s \propto \text{PDF}_{j-1}(\mathbf{g}_s^T) / \mathcal{N}_{\hat{\mathbf{g}}_{j-1}, \Sigma_{j-1}}(\mathbf{g}_s^T) \in (0, 1]. \quad (8)$$

We retain $\mathbf{g}_s^T \in G_{j-1} \subseteq T_{j-1}$ with probability i_s . If necessary, we repeat this procedure until the resampled posterior population reaches a predefined size, typically a few thousand sample points.

Step (S2) finds a control pulse that optimizes the anticipated information gain, as quantified in a cost function, C . In Sec. III C below, we describe cost functions that are suitable for OBSID. We have implemented the OBSID algorithm, including the cost function evaluation, in a graph-based machine-learning environment [21,22], building on *TensorFlow* [23], which facilitates automatic differentiation for the purpose of efficient computation and gradient-based optimization of C . This environment is designed for automated quantum control and learning applications, and it efficiently scales to handle controls with many optimizable parameters. It also straightforwardly accommodates soft constraints in the cost function that encode desirable properties of control pulses, such as limited bandwidth, slew rate, etc. This computational approach delivers greater flexibility to efficiently search the control space relative to previous discrete optimizations constrained to changing only the number of repetitions of a fixed probe control in related experiments [16].

Step (S3) requires an interface to a physical (or simulated) experiment. The interface should communicate pulse specifications provided by OBSID to the experimental apparatus' control system, which then implements the pulse with high fidelity on the physical target system. Care needs to be taken that the control system does not produce unmodeled distortion (e.g., nonlinearities, crosstalk, or dead time in buffers) in the pulse as received by the physical target. If such effects are present, they should be included in the system model's parameterization. The interface waits for the experiment and measurement process to complete, and then receives the measurement outcome m_j along with the standard measurement error σ_j .

Step (S4) is a direct application of Bayes' rule to compute the posterior NLL given the measurement outcome m_j . Internally, the NLL is stored as an interpolating function that can be evaluated anywhere inside the convex hull of the sample points in G_{j-1} ; this implements a consistent probability density estimator for amortizing the computation of previous predictive calculations and experimental measurements, as described in [24]. The right-hand side of Eq. (6) implements assumption (A3) that the measurement statistics are normally distributed. However, this can be generalized to other statistical distributions, e.g., the binomial distribution for a finite number of Bernoulli trials, if necessary.

Step (S5) uses rejection sampling [20] based on the acceptance likelihood δL_j to generate a posterior subpopulation G'_j , shown as green points in Fig. 1, and computes the posterior sample mean and covariance matrix, which are reported as part of the output of the protocol.

Lastly, the posterior NLL from (S4) becomes the prior NLL to start the next iteration of the OBSID loop, illustrated in the green square in Fig. 1.

C. Cost functions

The cost function that is minimized in (S2) is a critical part of the protocol. We have implemented two suitable cost functions. The first, C_{APC} , is based on the *anticipated posterior covariance* (APC), which was described in [16]. The second, C_{MFI} , is based on a modified Fisher information that penalizes highly oscillatory responses, and is described in Appendix A.

In practice, we find that both cost functions work comparably well. However, C_{APC} is more straightforward to describe and is more generalizable, so for benchmarking OBSID in later sections, we only present results based on C_{APC} , which we describe presently. The existence of at least two cost functions is conceptually useful as a foil to abstract the OBSID protocol from any specific choice of cost function. This abstraction opens avenues to finding other cost functions with desirable properties, such as reduced computational overhead [12,13], or enhanced robustness. We comment on alternative cost functions in Appendix B.

The motivation for the APC is that, for a given pulse $\mathbf{c}^{(j)}$, the experiment will return some measurement outcome, $m_j = \hat{P}_0(\mathbf{c}^{(j)})$, with statistical uncertainty σ_j . Of all the prior parameter samples, $\mathbf{g}_s \in G_{j-1}$, those that are consistent with m_j are those for which

$$P_0(\mathbf{g}_s, \mathbf{c}^{(j)}) \approx m_j \pm \sigma_j. \quad (9)$$

More precisely, under assumption (A3), the relative likelihood of a prior sample $\mathbf{g}_s \in G_{j-1}$ point to be represented in a posterior population G_j is given by

$$\delta L_j(\mathbf{g}_s, m_j) = e^{-(P_0(\mathbf{g}_s, \mathbf{c}^{(j)}) - m_j)^2 / (2\sigma_j)^2}, \quad (10)$$

consistent with Eq. (9). Rejection sampling with this relative likelihood produces a statistically consistent posterior subsample, G'_j . The posterior sample covariance Σ_j depends implicitly on $\mathbf{c}^{(j)}$, and so the APC cost function is designed to reward control pulses that minimize Σ_j , in some measure.

Since we must choose $\mathbf{c}^{(j)}$ prior to the measurement that determines Σ_j , we average the posterior sample covariance over all possible measurement outcomes [12] to produce

the *anticipated* posterior covariance

$$\Xi_j = \int dm Q_j(m) \Sigma_j(m), \quad (11)$$

where $Q_j(m) = N_j(m) / \int dm' N_j(m')$ is the likelihood of anticipated measurement outcomes, and $N_j(m)$ is defined in step (S5). The APC matrix Ξ that was introduced in Ref. [16] is positive definite and depends implicitly on pulse $\mathbf{c}^{(j)}$.

A well-chosen pulse is one for which the eigenvalues of Ξ are small, with a particular bias towards minimizing the largest eigenvalue, which is the major posterior uncertainty. A convenient cost function that captures this is

$$\begin{aligned} C_{\text{APC}}(\mathbf{c}^{(j)}) &= \text{Tr}(A \cdot \Xi_j), \\ &= \int dm Q_j(m) \text{Tr}(A \cdot \Sigma_j(m)), \end{aligned} \quad (12)$$

where A is a diagonal preconditioning matrix that rescales parameter uncertainties if required (e.g., if the model parameters carry different units, or are of greatly differing orders of magnitude). For benchmarking in the rest of this paper, we take $A = \mathbb{I}$. In practice, we approximate the nested integrals in Eqs. (11) and (12) over m and m' as a discretized sum over the range of possible measurement outcomes represented across the prior parameter sample population G_{j-1} .

IV. PERFORMANCE BENCHMARKING

In this section, we benchmark the performance of OBSID in estimating Hamiltonian parameters for both simulated and experimental systems. The scope for parameterizing control pulses is essentially unlimited, for example, superpositions of smooth basis functions, or piecewise continuous (PWC) functions. In these demonstrations we use PWC controls with a user-defined segment count, and with variable segment amplitudes and durations, so that the overall pulse shape and duration are included in the optimization. In this section, we first demonstrate OBSID on the simulated single-qubit calibration of a two-parameter Hamiltonian model described by Eq. (1). We then validate this with the experimental calibration of an ion-trap system. Finally, we demonstrate the calibration of a simulated two-qubit system characterized by a five-parameter model, with two independent control fields.

A. Simulated single-qubit parameter estimation

We begin by demonstrating parameter estimation for the pair $\mathbf{g}^{(1q)}$ used to specify the model introduced in Eq. (1) via numerical simulation of the physical model. The simulator calculates the time evolution of the system under an arbitrary control and generates measurement statistics,

m_j and σ_j , replicating experimentally realistic finite sampling errors. We assume that the true parameter values are $\mathbf{g}_t^{(1q)} = \{\Delta_t, \Omega_t\} = \{4, 6\}$ Hz (noting that in all simulation units are entirely nominal). We assume independent normally distributed priors $\Delta_0 \sim 4.1 \pm 0.5$ Hz and $\Omega_0 \sim 6.2 \pm 0.5$ Hz (here $x \sim \mu \pm \sigma$ signifies that $x \sim \mathcal{N}_{\mu, \sigma^2}$).

1. Automated, optimized Rabi-Ramsey measurements

Measuring $\mathbf{g}^{(1q)} = \{\Delta, \Omega\}$ is a standard single-qubit calibration scenario, which requires a minimum of two independent parameterizations. This is conventionally accomplished by a combination of experiments that use a uniform Rabi-type pulse,

$$c_{\text{Rabi}}(t) = \Theta(T_{\text{Rabi}} - t), \quad (13)$$

where $\Theta(x)$ is the unit-step function, and experiments that use a Ramsey-type pulse,

$$c_{\text{Ramsey}}(t) = \frac{\pi}{2\Omega} (\delta(t) - \delta(t - T_{\text{Ram}})), \quad (14)$$

which begins and ends with approximate $\pm\pi/2$ pulses. These two control parameterizations are shown schematically in Fig. 3(a). Rabi-type experiments are sensitive to the combination $(\Omega^2 + \Delta^2)^{1/2} \equiv \Omega_{\text{eff}}$, while Ramsey-type experiments are sensitive to the free evolution determined by Δ . In combination, the interleaved experiments are sufficient to localize both parameters.

We first illustrate the performance of OBSID in a simulated scenario that optimizes this conventional single-qubit tuneup procedure. At each iteration, OBSID autonomously optimizes over Rabi or Ramsey pulse types, choosing the pulse type and duration that minimizes C_{APC} . The simulated experiment is then addressed with the optimal pulse, and returns $m_j = \hat{P}_0(c^{(j)})$ and the simulated sample uncertainty σ_j .

Figures 3(b) and 3(c) show the outputs from the first two iterations of the OBSID loop, with the prior population at each iteration shown as gray points, and the posterior population after importance resampling as green points. Dark contours show the predicted response P_0 over the parameter space, and the contour consistent with the ‘‘measured’’ value m_j returned from the simulated experiment is shown in green. Because this is a simulation, we know the true parameter values $\mathbf{g}_t^{(1q)}$ (shown as a purple diamond), which are contained inside the convex hull of the posterior population (green points), illustrating the ability of OBSID to rapidly converge on an effective estimate of $\mathbf{g}^{(1q)}$.

Figure 3(d) shows the optimal duration and pulse type chosen at each OBSID iteration. The algorithm autonomously determines a procedure that alternates between Ramsey- and Rabi-type pulses. The effect of the alternating pulse types chosen by OBSID is seen in Figs. 3(b) and 3(c). The first iteration selected a Ramsey-type

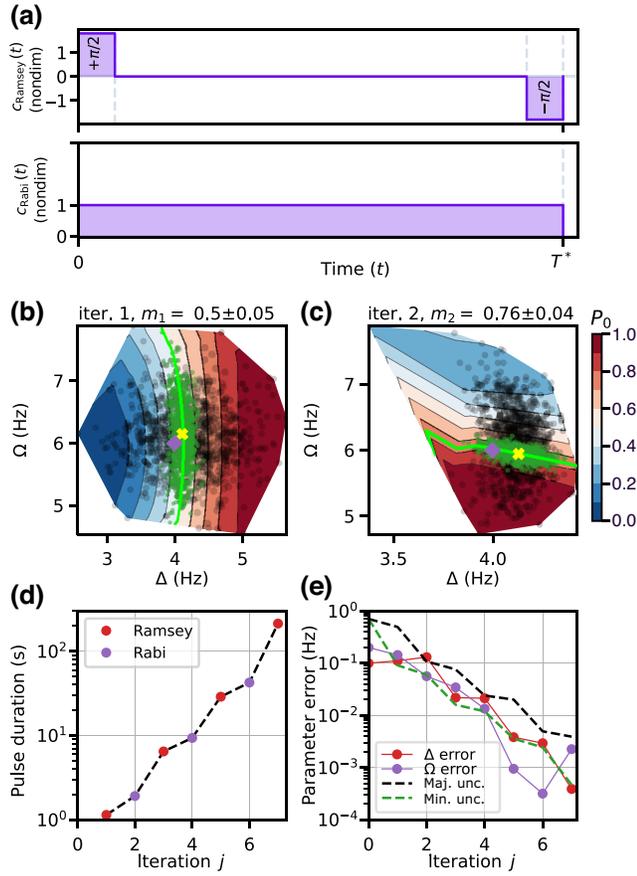


FIG. 3. Results of simulated single-qubit calibration using conventional “Ramsey”- and “Rabi”-type control pulses. Panel (a) shows the pulse envelopes for the two control types, in which the duration $T^{(j)}$ is optimizable. The first and last segments in the Ramsey control pulse are (approximate) $\pm\pi/2$ pulses. Panels (b) and (c) show the first two iterations of the prior population (black circles) and posterior population (green dots), superimposed on the predicted response P_0 (contours) for the optimally chosen pulse. The green contour is the measured value $m_j = \hat{P}_0$ (from simulation). The posterior mean (yellow cross) and true parameter value (purple diamond) are also shown. The posterior population (green) at iteration 1 becomes the prior population for iteration 2 (black). (d) Pulse durations and types selected to optimally estimate single-qubit parameters, and (e) the iterative improvement in the absolute parameter errors, $|\bar{\mathbf{g}}_j - \mathbf{g}_t|$, and the uncertainty characterized by the eigenvalues of $\Sigma_j^{1/2}$.

pulse, which principally localizes Δ and whose insensitivity to Ω is evident in the orientation of the contours of Fig. 3(b); the second iteration chose a Rabi-type pulse that localizes along a contour of constant Ω_{eff} , which is locally a linear combination of Δ and Ω in the neighborhood of the “true” value $\mathbf{g}_t^{(1q)}$.

The estimated posterior uncertainty after each iteration is measured by the largest (major) eigenvalue of the posterior sample deviation matrix, $\lambda_j^{\text{maj}} = \text{MaxEval}[\Sigma_j(m_j)^{1/2}]$.

This is shown in Fig. 3(e) (dashed black line), and is a statistical upper bound for the uncertainty of all marginals. We also plot the absolute error between the estimated population mean and the known values used in the simulation, shown as points. We expect that the absolute error in the parameters $|\bar{\mathbf{g}}_j - \mathbf{g}_t^{(1q)}| \lesssim \lambda_j^{\text{maj}}$. This is borne out in practice.

Returning to Fig. 3(d) we also observe that the optimized pulse duration grows by an approximately constant factor from iteration to iteration, consistent with the Fisher information tending to grow with T . This results in an overall exponential growth in the pulse duration with iteration count j . Thus, the optimal pulse duration varies approximately inversely with the major uncertainty, as discussed earlier. This is a general characteristic of all simulations we have performed.

2. Automated, optimized arbitrary pulse measurements

The example optimization over Rabi and Ramsey pulses uses predefined control envelopes that have well-understood sensitivity to the model parameters. However, the general approach in OBSID is capable of finding good control pulses directly, without relying on intuition or predefined control types. This becomes important for parameter estimation in multidimensional systems that are complex enough that well-understood calibration and control pulse families are not available.

We demonstrate this capability leveraging the same simulation, but now using a PWC control pulse whose segment amplitudes and total pulse duration are independently optimized. To demonstrate the flexibility of OBSID, we show the results using 10-segment pulses (using fewer segments also works in practice). During optimization, the segment amplitudes are constrained to the real interval $[-1, 1]$, and the maximum duration of each pulse is constrained by $T^{(j)} \leq 2T^{(j-1)}$, so that the overall duration does not grow too quickly. Figure 4(a) shows the pulses chosen by the OBSID protocol for the first five iterations. Clearly the pulse durations tend to grow with iteration.

Figures 4(b) and 4(c) show the prior and posterior sample populations, along with the model-predicted response in the first two OBSID iterations. Figure 4(d) shows the typical exponential growth in overall pulse duration, and (e) shows the major uncertainty, and the absolute parameter errors. These outcomes are qualitatively the same as in Fig. 3 for the Rabi-Ramsey calibration discussed above.

Again, we observe that the OBSID algorithm selects exponentially growing pulse durations, $T^{(j)}$, with iteration count [Figs. 4(d) and 4(e)]. Furthermore, the contours of the predicted response P_0 , which depend on $c^{(j)}$, characteristically rotate in parameter space from one iteration to the next: contours and the posterior population in the first iteration are broadly aligned to the Ω axis, while for

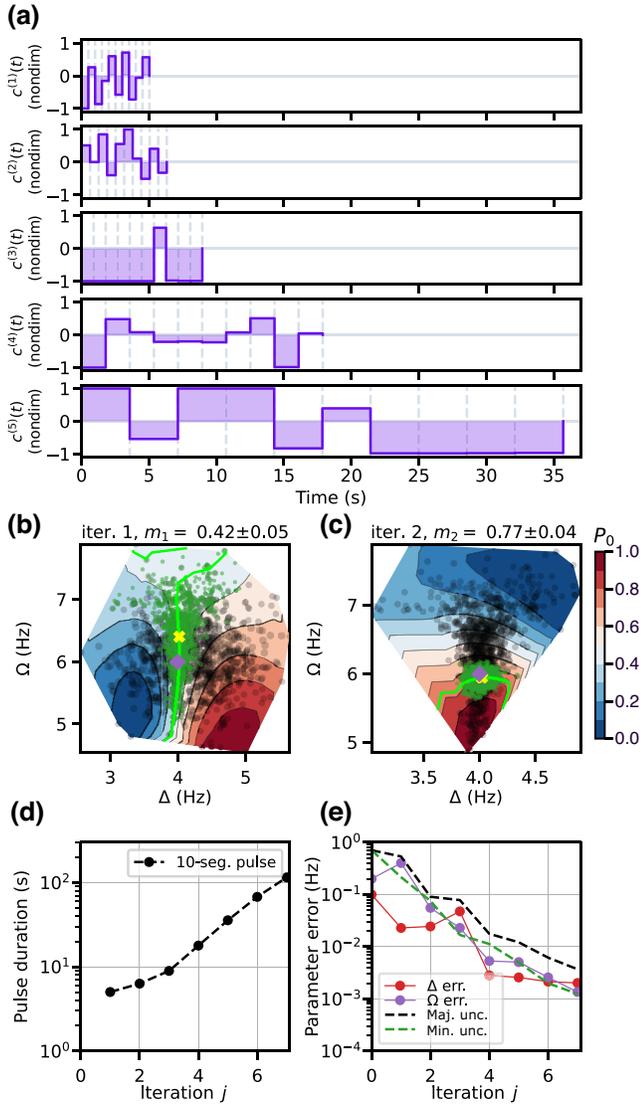


FIG. 4. Results of simulated single-qubit calibration using fully optimized 10-segment piecewise-continuous pulses. (a) The sequence of pulses generated in each iteration of OBSID, in which the segment amplitudes and the total pulse duration $T^{(j)}$ are all optimized. Panels (b) and (c) show populations for the first two iterations, showing prior (gray points) and posterior (green points) populations, superimposed on the predicted response P_0 (interpolated contours) for the optimally chosen pulses. (d) Pulse durations for each iteration, and (e) the iterative improvement in the parameter error and uncertainty.

the second iteration, they are more aligned with the Δ axis. This behavior, in which sequential iterations produce roughly orthogonal “cuts” through parameter space, is a consequence of the APC cost function that minimizes the major uncertainty.

The pulses selected by the OBSID algorithm achieve performance that is marginally better than results from using Rabi-Ramsey control pulses. Specifically, comparing panels (d) and (e) in Figs. 3 and 4 shows that after

several iterations, the arbitrary pulse optimization procedure returns pulses of somewhat shorter duration while maintaining similar accuracy in the resulting parameter estimates. Unlike the Rabi-Ramsey sequences, the effect of the 10-segment PWC pulses is not readily interpretable except by direct simulation.

B. Experimental single-qubit parameter estimation in a trapped-ion system

Having benchmarked the OBSID protocol for a single qubit using simulated experiments, we now validate it using a real trapped-ion experimental system. The qubit is encoded in the $^2S_{1/2}$ electronic ground state of a $^{171}\text{Yb}^+$ ion, where we assign the qubit states $|F=0, m_F=0\rangle \equiv |0\rangle$ and $|F=1, m_F=0\rangle \equiv |1\rangle$. Considering a 12.64-GHz microwave field near resonance with the $|0\rangle \leftrightarrow |1\rangle$ transition, the system’s dynamics are described by the same two-parameter Hamiltonian of Eq. (1). The detuning Δ corresponds to the frequency difference between the microwave control signal and the qubit’s resonance frequency, while Ω is the calibration factor relating the actual Rabi frequency experienced by the ion to the microwave driving voltage amplitude from a signal generator, so it has units of “Hz/V”.

The experimental apparatus consists of a linear Paul trap held at room temperature [25]. The degeneracy of the $F=1$ Zeeman levels is lifted by applying a static magnetic field of 0.44 mT. The native, uncompensated qubit coherence time is measured to be $T_2^* = 8.7(7)$ s with a Ramsey-type experiment. A set of laser beams addressing the $^2S_{1/2} \rightarrow ^2P_{1/2}$ transition near 369.5 nm is used for Doppler cooling, state preparation, and measurement [26]. State detection is performed by discriminating between the differing number of scattered photons collected on an avalanche photodiode for each of the two qubit states, when turning on the 369.5-nm laser.

The microwave control field is obtained from a synthesis chain in which a direct digital synthesizer (DDS; AD9910) is mixed with an arbitrary waveform generator (AWG; Keysight M8190A). The resulting signal is up-converted to 12.64 GHz. The DDS has a frequency resolution of 0.23 Hz and the clock speed of the AWG is 5.8 GHz; they both have an amplitude resolution of 14 bits. The 12.64-GHz signal is amplified and delivered to the ion by an antenna positioned 2.37 cm from the trap center. The amplifier chain in the synthesizer system has been independently calibrated to assure linearity between the amplitude specified in a pulse sequence and the amplitude experienced by the ion.

We use two-stage calibration measurements, implemented with DDS, to determine the best initial estimate of the “true” physical parameters that will later serve as the “ground truth” for the OBSID results. The true

detuning Δ_t is determined using 150 samples of a standard Ramsey-locking method [27] with up to a 50-ms Ramsey interrogation time. The Rabi frequency Ω_t is measured by fitting resonant Rabi oscillations over approximately 28 Rabi cycles. The oscillations were sampled with 300 uniformly spaced samples in time, so that there were about 11 samples per Rabi cycle. Each data point in these experiments used between 50 and 100 individual, binary-valued single-qubit fluorescence readouts to estimate \hat{P}_0 . In total, 450 experimental measurement samples were taken (i.e., 150 for Ramsey and 300 for Rabi), giving calibrated values $\Omega_t = 1249.1 \pm 0.1$ Hz/V and $\Delta_t = 500 \pm 0.4$ Hz, where the uncertainty is given as the standard error.

To initialize the OBSID protocol, we assume uncorrelated, normally distributed priors for each parameter, with $\bar{\Delta}_0 \sim 525 \pm 52.5$ Hz/V and $\bar{\Omega}_0 \sim 1311 \pm 131.1$ Hz. This was chosen to reflect a scenario in which the prior standard deviation is 10% of the prior mean parameter values. The prior uncertainty is large enough that it includes the pre-calibrated “ground-truth” parameter values within a region of high probability density, and so is statistically consistent with the conventionally calibrated system parameters. Such a circumstance is conventionally encountered experimentally in which a “coarse” spectroscopy scan is followed by experimental fine tuning to estimate experimental parameters. The OBSID waveform is generated with an AWG.

Figure 5(a) shows the first five PWC control pulses chosen by OBSID. The control pulse profiles specify the amplitudes input to the AWG to modulate the experimental drive, at a fixed carrier frequency.

Figures 5(b) and 5(c) show the prior and posterior populations of parameters for the first two OBSID iterations. As before, the posterior population mean (yellow cross) converges to the “ground truth” (purple diamond). Figure 5(d) shows the characteristic growth in control pulse duration, by roughly an order of magnitude over five iterations, consistent with the simulated results in Fig. 4(d).

Lastly, Fig. 5(e) illustrates the convergence between the OBSID posterior and the ground-truth values, along with the estimated major uncertainty from the population covariance matrix. Again, we see consistency between the internal estimate of the major parameter uncertainty and the error relative to the ground truth. Given the accuracy of the conventional calibration measurements, the floor in each of the absolute errors is the same as the calibrated uncertainty of 0.1 and 0.4 Hz for both Ω and Δ .

We see from Fig. 5(e) that the major uncertainty estimate from OBSID reaches ± 2.6 Hz after $j = 5$ samples. This is a statistical upper bound on the absolute error, which is 2.2 and 0.8 Hz for both Ω and Δ . The total iteration number j is limited by the AWG memory in our system due to a high sample rate, and does not pose a fundamental limit for the protocol. We compare OBSID to the conventional method of simultaneously estimating Ω and Δ , fitting a detuned

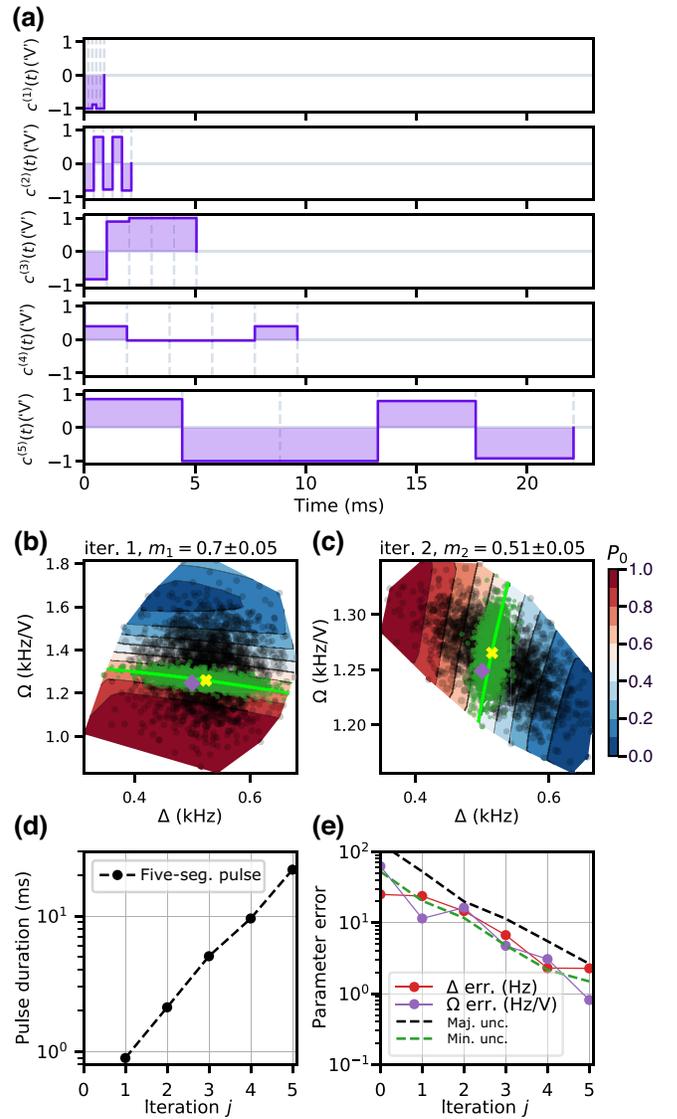


FIG. 5. Experimental OBSID results in a single-ion qubit calibration, using optimized five-segment pulses. (a) The sequence of pulses generated in the first five iterations of OBSID, in which the segment amplitudes and the total pulse duration $T^{(j)}$ are all optimized. Panels (b) and (c) show populations for the first two iterations, showing prior (gray points) and posterior (green points) populations, superimposed on the predicted response P_0 (interpolated contours) for the optimally chosen pulses. (d) Pulse durations for each iteration, and (e) the iterative improvement in the parameter accuracy. The floor in the absolute errors is set by the inferred uncertainty in the conventional calibration methods, which has a standard error of ± 0.3 Hz for both Δ and Ω .

Rabi flop, in order to make a quantitative comparison of protocol efficiency. We sample 300 equally spaced points up to OBSID’s longest duration, 22.1 ms, and use the same number of experimental shots at each sample as in OBSID. The resulting uncertainties are 3.4 and 7.9 Hz for Ω and Δ , which are greater than the absolute error obtained by OBSID for a 60 times larger number of samples (300 vs

5). Overall, this shows that, for a comparable number of samples, and using the smaller residual uncertainty on Ω as a bound, OBSID is approximately 93 times more efficient than standard parameter estimation techniques

The experimental results shown in Fig. 5 are qualitatively indistinguishable from the simulated results shown in Fig. 4. In both simulation and experiment, the OBSID protocol treats the simulation or experimental system as a black-box probability generator, and chooses informationally optimized pulses based on the model and the sample population. This experimental demonstration of OBSID illustrates one of its main strengths: OBSID is able to efficiently characterize a device without being provided knowledge of “good” calibration pulses used in standard techniques. This will become increasingly important in larger systems, which we treat next. Note that this strength of OBSID exists as long as a system has *structural identifiability* [28,29], which means that there is sufficient control and measurements available to observe information about all the parameters in the system. This is further discussed in Sec. V.

Beyond the demonstrations in Fig. 5 we have experimentally implemented other parameterizations of control pulses for use with OBSID in this experimental system. We present summary results for experimental implementations of these alternative pulse parameterizations in Appendix C.

(1) The Rabi-Ramsey pulse parameterization described in Sec. IV A 1, and illustrated schematically in Fig. 3(a). OBSID autonomously selects the optimal pulse type and pulse duration at each iteration.

(2) Optimized “bang-bang” control, which uses PWC pulses alternating between “on” (unit magnitude) and “off” (zero magnitude) segments. OBSID independently optimizes each of the segment durations, as well as the carrier phases $e^{i\phi(t)}$ during the “on” segments. This pulse parameterization has the advantage that, since the “on”-pulse carrier magnitude is constant, linearity of the transfer chain from the AWG to the target ion is not required.

In both cases, the results appear qualitatively the same as in Figs. 3–5.

C. Simulated two-qubit parameter estimation

We next demonstrate the operation of the OBSID protocol for a more complicated simulated five-parameter system, consisting of two coupled qubits, each of which is driven by an independent control field, specified by time-dependent control fields $c_q(t)$ for $q \in \{1, 2\}$. The Hamiltonian for this system is chosen to be

$$H^{(2q)} = \sum_q H_q^{(1q)}(t) + J(L_1^\dagger L_2 + L_2^\dagger L_1), \quad (15)$$

where $H^{(1q)}$ was introduced in Eq. (1). The Hamiltonian is characterized by the five parameters

$$\mathbf{g}^{(2q)} = \{\Delta_1, \Omega_1, \Delta_2, \Omega_2, J\}, \quad (16)$$

with detunings Δ_q , Rabi frequencies $\Omega_q c_q(t)$, and the interqubit coupling strength J . We note in passing that this Hamiltonian could be derived from a system of two qubits driven by control fields that have a common carrier frequency [30,31]; however, we use it here simply as a multiparameter, multiqubit demonstration of the SID technique.

As in the single-qubit case, we assume that the qubits are initialized in the ground state of the undriven Hamiltonian, $|\psi(0)\rangle = |0, 0\rangle$. We assume a two-state measurement process that discriminates between the initial state $\Pi_0 = \{|0, 0\rangle\}$ and the complementary subspace $\Pi_1 = \{|0, 1\rangle, |1, 0\rangle, |1, 1\rangle\}$ to yield the return probability $P_0 = |\langle 0, 0 | \psi(t) \rangle|^2$. This measurement model represents incomplete information for the two-qubit system, since only the return probability is distinguished. This can be generalized to include an informationally complete measurement, which will be the subject of future work.

For the purpose of simulation, we take the true parameter values to be $\mathbf{g}_t^{(2q)} = \{4.1, 5.5, 4, 6, 0.5\}$ Hz, and assume a prior deviation matrix given by $\Sigma_0^{1/2} = \text{diag}\{0.3, 0.3, 0.3, 0.3, 0.03\}$ Hz. We initialize the OBSID prior with a mean value sampled from $\bar{\mathbf{g}}_0 \sim \mathcal{N}_{\mathbf{g}_t^{(2q)}, \Sigma_0}$, and then set the prior to be $\mathcal{P}_0 = \mathcal{N}_{\bar{\mathbf{g}}_0, \Sigma_0}$, that is, the relative uncertainty in each parameter is around 7% of the true value.

Figure 6 shows the results obtained using OBSID for the two-qubit model, using PWC control pulses with 10 segments. Figure 6(a) shows the optimized pulse pairs, $c_1^{(j)}(t)$ and $c_2^{(j)}(t)$, chosen by OBSID for the first three iterations. In the pulse parameterization shown here, at each iteration, the duration of the two control pulses are constrained to be the same. This sequence of control-pulse pairs shows that OBSID is capable of optimizing both control pulses simultaneously. This would be particularly important in cases where calibrating crosstalk is necessary [25].

Figures 6(b) and 6(c) show the prior and posterior population samples obtained during the first two iterations of OBSID. To partially visualize the five-dimensional parameter space, we project onto two particular unit vectors. For the vertical axis of the projected plots, the direction is chosen to align with the maximum average information direction $\hat{\lambda}_{\text{grad}} \propto E_S(\tilde{f}(\mathbf{g}; \mathbf{c}^{(j)}))$, where the information direction \tilde{f} is defined in Eq. (4). The unit vector defining the horizontal axis is chosen to align with the maximum deviation between the population mean and the true value used in the simulation, $\hat{\lambda}_{\text{dev}} \propto \bar{\mathbf{g}}_j - \mathbf{g}_t^{(2q)}$. These unit

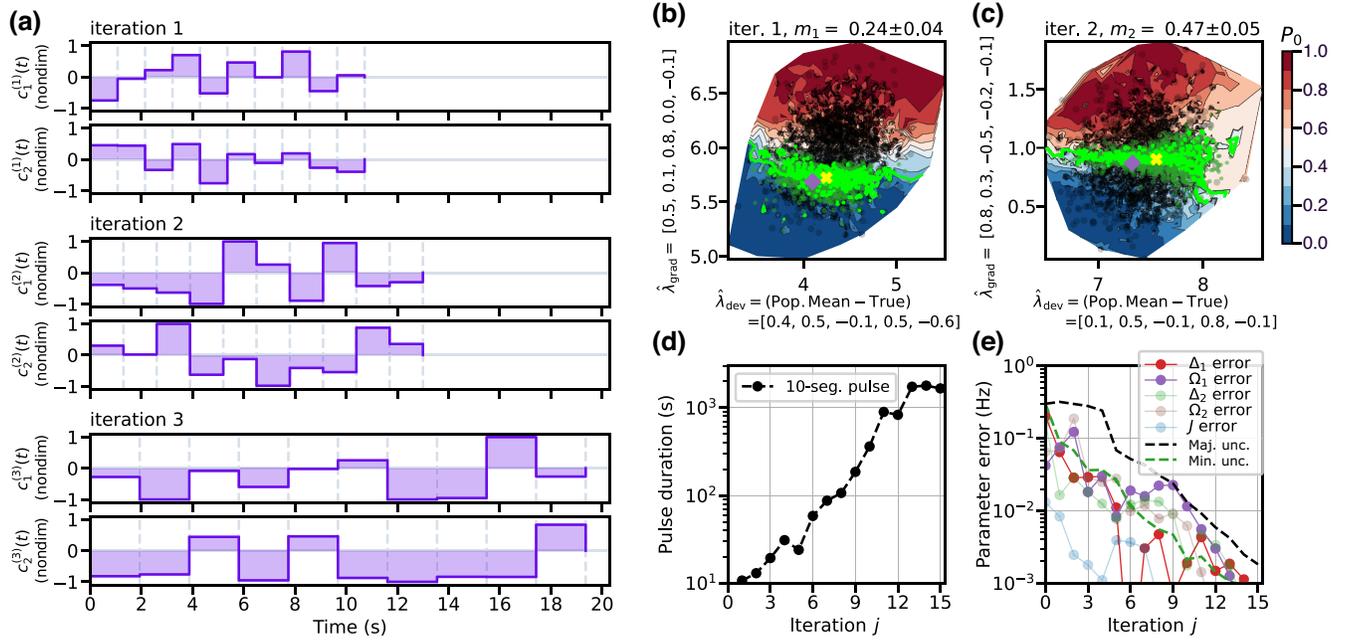


FIG. 6. Results of SID on a simulated two-qubit system with five parameters and two control fields. (a) The sequence of optimized control pulse amplitudes, $c_1^{(j)}(t)$ and $c_2^{(j)}(t)$, applied to each qubit for the first three iterations of OBSID. Panels (b) and (c) show scatter plots and contours of the response, P_0 , as a function of parameters. To represent the points in five-dimensional parameter space, we project onto two composite coordinates, defined by the unit vectors $\hat{\lambda}_{\text{dev}}$ along the deviation between the true value and the population mean (horizontal axis), and $\hat{\lambda}_{\text{grad}}$ along the direction of steepest gradient of P_0 (vertical axis). Both unit vectors vary with iteration count j , indicated by the numerical arrays shown in the axis labels. Panel (d) shows that pulse durations tend to grow with iteration, and (e) shows that the uncertainties, estimated from both the five-dimensional population covariances and absolute errors, improve consistently over iterations. We have highlighted the error in Δ_1 and Ω_1 relative to the simulated “true” value; fainter curves show that other parameter errors evolve similarly.

vectors evolve over OBSID iterations, and we list the unit-vector coefficients in the axis labels with respect to the parameter ordering given in Eq. (16).

Figures 6(b) and 6(c) illustrate that, as for the two-parameter case, OBSID tends to choose control pulses whose response, P_0 , varies by a large amount along some direction in parameter space, and without oscillations in that direction. Here, due to the dimensionality of the parameter space the axes in these graphs correspond to nontrivial cuts rather than the parameters themselves in previous figures (see the caption). Importantly, the true value is located within the convex hull of the prior and posterior populations. Furthermore, the posterior distributions projected in Figs. 6(b) and 6(c) are distinctly non-Gaussian.

Figure 6(d) again shows the tendency of the optimal control pulses to become longer with iteration count, which leads to a corresponding reduction in the estimated uncertainty in the parameters, shown in Fig. 6(e). The largest (smallest) estimated uncertainty is computed from the major (minor) eigenvalue of the population deviation matrix. The major uncertainty is seen to provide a good approximate bound on the absolute parameter error $|\hat{\lambda}_{\text{dev}}|$, shown as points in panel (e). For clarity, in Fig. 6(e),

we highlight the error in parameters Δ_1 and Ω_1 , with other parameter errors indicated by fainter lines. Experimentally, the true parameter values and absolute errors are unknown, but we reiterate that the posterior sample deviation provides a reliable estimate for the parameter uncertainty.

V. DISCUSSION

These results, including simulated and experimental tests, illustrate that OBSID is able to efficiently converge on estimates that closely approximate the true parameter values for a system. Comparing the single-qubit, two-parameter results in Fig. 4(e) with the two-qubit, five-parameter results in Fig. 6(e), we see empirically that the number of iterations to improve the parameter uncertainty by a fixed factor has scaled with the number of unknown parameters p .

A heuristic argument supporting this observation is as follows. Suppose that the prior uncertainty density was a spherically symmetric distribution in the p -dimensional model parameter space. Each iteration of OBSID compresses the prior distribution in some direction in parameter space by a factor $\alpha < 1$. [For example, by inspecting

panels (b) and (c) in all the OBSID outputs given earlier, the compression of the posterior population samples relative to the prior samples shows that $\alpha \approx 0.1$ to 0.2 .] Since OBSID preferentially compresses the most uncertain direction, after p iterations, it will have compressed the spherical prior approximately symmetrically in all p dimensions, and each will be compressed by approximately the same factor α . So after p iterations, the initially spherical uncertainty prior will have been compressed approximately uniformly. Iterating, it follows that the number of iterations to uniformly shrink the prior uncertainty scales with p .

A natural question to ask is whether the fixed factor (exponential) improvement of the parameter precision will continue with more iterations. We argue heuristically that the exponential reduction of the error is directly connected to the exponential increase in the length of the pulses. In most circumstances, there is a maximum pulse duration that can be achieved in the experiment, such as when the pulse duration approaches a T1 limit. The limit is rapidly approached because the pulse duration scales exponentially. If further experiments were completed at this maximum duration, improvements to estimates would largely be due to statistical averaging, which typically follows a power-law relationship $1/\sqrt{N}$ with the number of samples N . Hence, we recommend halting the OBSID algorithm once this maximum duration is reached, as further measurements will improve the precision according to a power law rather than an exponential scaling.

Most naturally occurring nonunitary (Markovian) evolution will limit the advantage of increasingly long pulse durations. This follows because dynamics governed by master equations typically reduce the distinguishability of states [32]. That is, the distinguishability of states governed by a master equation with a unique steady state will go to zero. Once a system has reached its steady state, measuring the dynamics will no longer provide any additional insight into the system parameters.

OBSID as we have designed and implemented it is extremely flexible. For standard calibration tasks, it is capable of optimizing the conventional Rabi-Ramsey-based calibration schedule for a single qubit that alternates between localizing the detuning and the Rabi frequencies. As shown in Figs. 3(b) and 3(c), these choices correspond to cuts through parameter space that are approximately oriented orthogonal to the corresponding parameter at each iteration. This has merit for “human interpretability,” in situations where it is desirable to calibrate model parameters individually. In this case, the pulse shapes in Fig. 3(a) that isolate a given parameter are known from conventional practice. For more general problems, the APC cost function could be straightforwardly adapted to select pulses that preferentially constrain parameters along the named parameter axes, that is, to include a “parameter-selective”

preference in the cost function. However, for highest efficiency, OBSID accesses the full p -dimensional parameter space, as illustrated in Figs. 4 and 6. In doing so, OBSID tends to introduce nonzero covariances, which correlates the uncertainty of different parameters.

The power of OBSID comes from increasing the control pulse duration T as the estimated uncertainty decreases. For suitable control-pulse parameterizations, this allows the difference in the dynamical states generated by different prior model parameter samples to grow in the system Hilbert space, amplifying the effect of small parameter uncertainty into large differences in Hilbert space evolution without introducing additional uncertainty due to oscillatory behavior.

We understand this through the FI, which ideally grows as T^2 . In such cases, for r repetitions of a measurement, the parameter uncertainty scales as $\sigma \propto 1/\sqrt{FI}r \sim 1/(T\sqrt{r})$. Since the total accumulated measurement time scales as $\tau_{\text{meas}} \sim Tr$, the overall measurement time resource is best utilized by choosing as long a pulse duration T as possible. OBSID does this dynamically: it increases T as quickly as possible, implicitly constrained by the prior uncertainty at each iteration.

The computational cost of optimizing population-averaged model predictions over control pulses is non-trivial. Our current implementation of OBSID builds on *TensorFlow* [23], a graph-based toolbox for machine learning optimization. This gives it great flexibility in searching for optimal pulses, and automatically distributes highly parallelizable tasks across computational resources [21,22]. For example, we have run the optimization tools on systems ranging from few-core laptops to high-performance workstations with many cores, with run times that scale with the resource availability. However, the optimization task is necessarily computationally demanding, typically taking tens of seconds to run on a 64-core workstation running at 4 GHz. Improving this performance by using alternative less-intensive cost functions is a subject of recent [12,24] and ongoing research.

Using OBSID effectively requires some care. It relies on statistical sampling, and so has some possibility to fail in any give run of system identification. This was noted in Ref. [16], and can be mitigated by including more sample points, at the cost of slower optimization or greater parallel resources. As the initial uncertainty in the priors increases, the statistical sampling issue becomes more challenging. This typically requires more samples in the particle filter, which increases the computational resources.

There is a more subtle issue, related to *statistical* and *structural identifiability* [28,29], which determine whether OBSID succeeds. Not all control parameterizations are suitable for effectively extracting model parameters with arbitrarily high accuracy. For example, if the control pulse parameterization has fewer free control parameters than

there are model parameters to fit, we expect OBSID will not be able to identify all system parameters, as described in Sec. II B. In this situation, the span of the information direction unit vectors does not cover the parameter space, and the control parameterization will not be able to structurally identify the system model parameters. In practice, when a system is not structurally identifiable for a specific control parameterization, the major uncertainty reported by OBSID will stall after several iterations, ceasing to improve with additional iterations. In Sec. D we describe an example of a control parameterization that is not structurally identifiable and that illustrates this stalling behavior. In practical usage of OBSID, simulations can be used to establish whether a system is structurally identifiable under a chosen control parameterization, before applying it to an experimental system.

Characterization of large quantum computers is an important application of OBSID. A direct simulation of the complete Hamiltonian of a large quantum computer on a classical computer is too complex. Therefore, OBSID must assume that characterizing the quantum computer can be broken down into subsystems. OBSID can estimate the parameters in the Hamiltonians of these subsystems *in parallel*. For example, the system identification of the Rabi rate and detuning of all the individual qubits on a quantum computer could be done in parallel, as each qubit has individual addressability. When run in this mode, the number of iterations required to parameterize all the qubits on a quantum computer should equal the number of iterations for an individual qubit.

The biggest practical challenge with running characterization of devices in parallel is individual addressability of the qubits. For example, in superconducting devices, engineering the electronics to ensure that there is no classical crosstalk between the microwave control lines is a challenging engineering feat. Fortunately, we have seen significant technical progress on this front and there is good experimental evidence that individual qubit parameters in superconducting quantum computers can be characterized in parallel simultaneously [33,34], and theoretical evidence in trapped ion devices [35].

Characterizing multiqubit coupling presents a more difficult challenge, but experimental evidence in superconducting devices suggests that this can be done in parallel with pairs of qubits [33]. When characterizing multiqubit coupling in this way, the scalability is limited by the maximum number of qubits that any given qubit can be coupled to, rather than the total number of qubits. Therefore, the total number of iterations required would typically be the amount of iterations required to identify one coupling, multiplied by the highest connectivity of the qubits. Investigating the viability of this subsystem approach to system identification is the most important next step for research on OBSID.

VI. CONCLUSIONS

We have described and demonstrated a general-purpose optimized Bayesian system identification protocol that uses model-based optimization over a statistically sampled prior population to choose control pulses that optimally extract information about unknown system parameters. Once the optimized pulse is applied experimentally, the system uses standard Bayesian update and resampling methods to produce a posterior distribution that is consistent with the new measurement information.

The implementation we have demonstrated in this paper is autonomous and hardware agnostic, and requires only that the system is well described by a quantitative numerical or analytical model. The typical performance of the scheme is demonstrated in several simulated systems, as well as in a single-ion experiment. We have shown experimentally that OBSID uses far fewer experiments to reach the same level of accuracy as conventional calibration methods; in the example presented, OBSID used just five different experiments to reach the accuracy of conventional single-ion calibration methods, which used 300 different experiments.

There are a variety of areas to explore, including quantifying the robustness and reliability of using OBSID in experimental systems with many parameters, and in open quantum systems with dissipative effects. There are also opportunities for optimizing the algorithmic and computational runtime performance using alternative cost functions.

ACKNOWLEDGMENTS

This work was supported by the U.S. Army Research Office/Laboratory for Physical Sciences Grant No. W911NF-21-1-0003. Experimental results reported here were undertaken at the University of Sydney and supported by the Intelligence Advanced Research Projects Activity Grant No. W911NF-16-1-0070, the Australian Research Council Centre of Excellence for Engineered Quantum Systems Grant No. CE170100009, and a private grant from H. and A. Harley. T.R. Tan was supported by the Sydney Quantum Academy Postdoctoral Research Fellowship.

APPENDIX A: QUADRATIC TIME DEPENDENCE IN THE FISHER INFORMATION

The quadratic growth in the FI with pulse duration can be derived straightforwardly in the case of a term, $g_i h_i$ with $g_i \in \mathbf{g}$, appearing in a time-independent Hamiltonian, $H = \dots + g_i h_i + \dots$. Then $P_0(T) = |\psi_0(T)|^2$, where $\psi_0(T) = \langle 0 | U(T) | 0 \rangle$ and $U(t) = e^{-itH}$. Using an integral expression for derivatives of operator exponentials (see Eq. 13

of Ref. [36]), we find that

$$\begin{aligned} \partial_{g_i} P_0(T) &= 2 \operatorname{Re}[\psi_0^*(T) \langle 0 | \partial_{g_i} U(T) | 0 \rangle], \\ &= 2 \operatorname{Re} \left[\psi_0^*(T) \langle 0 | U(T) \int_0^1 d\tau U^\dagger \right. \\ &\quad \left. (\tau T) \partial_{g_i} (-iTH) U(\tau T) | 0 \rangle \right] \\ &= 2 T \operatorname{Im}[\psi_0^*(T) \langle 0 | U(T) \bar{h}_i | 0 \rangle], \end{aligned} \quad (\text{A1})$$

where $\bar{h}_i = \int_0^1 d\tau h_i(\tau T)$ with $h_i(t) \equiv U^\dagger(t) h_i U(t)$ is the time-averaged expectation of $h_i(t)$ on the interval $t \in [0, T]$. In the last line, the term in square brackets is an oscillatory, bounded function of T , and the prefactor is linear in T . It follows that $(\partial_{g_i} P_0(T))^2$ scales as T^2 , and so $F = (\partial_{g_i} P_0(T))^2 / (P_0 - P_0^2) = T^2 \vartheta(\mathbf{g}, T)$, as stated in Eq. (5). This is consistent with specific simulations shown in Fig. 2(d) and in Ref. [18].

This property generalizes straightforwardly for time-dependent PWC controls whose segment durations scale with overall control duration. Briefly, we transform to a suitable rotating frame and make a rotating-wave approximation to remove explicit time dependence. In this frame, the evolution is given by a concatenation of unitary-evolution operators, each generated by a time-independent Hamiltonian (in the rotating frame) that depends linearly on g_i . The proportionality in Eq. (A1) will hold for each segment, and using the product rule for differentiation, we see that the growth of the FI will then also be quadratic in the total duration.

APPENDIX B: ALTERNATIVE COST FUNCTIONS

The Q -weighted averaging over m and m' implicit in Eq. (12) gives the statistically correct expectation value for the APC, but it also means that low-probability measurement outcomes are weakly represented. This could lead to some fragility in OBSID if these rare outcomes have a posterior distribution with high covariance, or other pathologies such as having very few samples represented in the prior population. We are free to modify the cost function to penalize the risk of such situations occurring.

We comment here briefly on alternative cost functions that are more robust, or enjoy other advantages. While we do not use them in the results reported in the main body of this paper, it is conceptually useful to abstract away from the specific choice of purely information-maximizing cost functions.

One way to make the cost function more robust is to use different aggregate measures over m . For example, a conservative ‘‘variance-averse’’ cost function would track the worst-case posterior covariance for any possible experimental measurement. In this case, the *maximum* posterior

covariance (MaxPC) cost would be

$$C_{\text{MaxPC}}(\mathbf{c}^{(j)}) = \max_m \operatorname{Tr}(A \cdot \Sigma_j(m)). \quad (\text{B1})$$

Clearly, a continuum of variations are possible using different norms to measure the distribution of $\Sigma_j(m)$ over m or the prior population G_{j-1} .

Other cost functions are also available based on approximations for the expected information gain [12,13]. In general, cost functions may be derived based on different uncertainty estimators [37]. We have developed a further alternative, based on a modified Fisher information, C_{MFI} , which relies on both the Cramer-Rao bound as an estimate of the sensitivity of the response P_0 and a bespoke penalty function designed to disfavor oscillatory responses. For completeness, we formulate this below.

The APC requires averaging over hypothetical measurement outcomes that are computationally intensive [12, 13,16]. Conversely, the MFI requires no such averaging, but does use gradients directly in the cost to evaluate the Fisher information, and in a penalty term that discounts oscillatory responses. These gradients are provided by automatic differentiation, so add a constant computational cost. However, optimizing MFI with gradient-based optimizers introduces an additional round of automatic differentiation, so that the cost gradients include second-order derivatives of the response, which add further computationally overhead as well. Both cost functions therefore have significant computational overheads. In both cases the computation is parallelizable over the prior population $\mathbf{g}_s \in G_{j-1}$, and so the computation benefits from high-performance multiprocessor computing resources.

1. Modified Fisher information cost function

Choosing control pulses that maximize the FI is a conceptually appealing approach to choosing maximally informative control pulses. Given a prior distribution $\text{PDF}_{j-1}(\mathbf{g})$ over the system parameters, we seek to maximize the expected FI

$$\bar{\text{FI}}(\mathbf{c}) = \int d^p \mathbf{g} \text{FI}(\mathbf{g}; \mathbf{c}) \text{PDF}_{j-1}(\mathbf{g}). \quad (\text{B2})$$

However, if the evolution is unitary, and g_i is a coefficient in the system Hamiltonian, then $\text{FI}(\mathbf{g}, T) \sim T^2 \vartheta(\mathbf{g}, T)$, where ϑ is some generally oscillatory function of pulse duration and the Hamiltonian parameters, as shown in Fig. 2. Optimizing \bar{F} will then yield arbitrarily long pulse sequences.

This reflects the fact that the FI is a point estimator of the sensitivity of a probability distribution to infinitesimal variations in a parameter, and this sensitivity can grow unboundedly (at the cost of longer pulse duration).

For parameter estimation however, we have a *finite* prior uncertainty (illustrated as the white disk in Fig. 2), and so

we have to discriminate amongst plausible model parameters distributed according to the prior \mathcal{P}_{j-1} . If the pulse duration T is too long then the predicted response $P_0(\mathbf{g}, \mathbf{c})$ will typically be highly oscillatory over the support of $\text{PDF}_{j-1}(\mathbf{g})$, so a measurement of $\hat{P}_0(\mathbf{c})$ will not effectively constrain the posterior range of plausible parameter values \mathbf{g} , as illustrated in Fig. 2(c). Instead, for a given prior covariance matrix Σ_{j-1} , the optimal pulse duration typically scales as $T^{(j)} \sim \text{Tr}(\Sigma_{j-1}^{-1/2})$, as discussed in Sec. II B.

We therefore define a modified Fisher information (MFI) designed to penalize oscillatory behavior in the response:

$$C_{\text{MFI}}(\mathbf{c}, \mathcal{P}_{j-1}) = -\bar{\text{FI}}(\mathbf{c})(1 - \alpha [M(\mathbf{c}, \mathcal{P}_{j-1}) - \beta]). \quad (\text{B3})$$

Here $[x] \equiv \max[x, 0]$ is the ramp function, M is a functional that penalizes oscillatory responses, $\beta > 0$ defines the tolerable penalty threshold, and $\alpha \sim 1$ scales the penalty.

The prior distribution \mathcal{P}_{j-1} is characterized by a prior mean $\bar{\mathbf{g}}_{j-1}$ and deviation matrix $\Sigma_{j-1}^{1/2}$ whose maximum eigenvalue $\lambda_{j-1}^{\text{maj}}$ is associated with the prior major uncertainty direction $\hat{\mathbf{y}}_{j-1}^{\text{maj}}$ in parameter space. A function that penalizes oscillations in P_0 along the direction of $\hat{\mathbf{y}}_{j-1}^{\text{maj}}$ is

$$M(\mathbf{c}, \mathcal{P}_{j-1}) = \int d^p \mathbf{g} \text{PDF}_{j-1}(\mathbf{g}) \Theta \times \left[-\frac{\hat{\mathbf{y}}_{j-1}^{\text{maj}} \cdot \nabla_{\mathbf{g}} P_0(\mathbf{g}; \mathbf{c})}{\hat{\mathbf{y}}_{j-1}^{\text{maj}} \cdot \nabla_{\mathbf{g}} P_0(\bar{\mathbf{g}}_{j-1}; \mathbf{c})} \right],$$

where $\Theta[x]$ is the unit-step function. Qualitatively, if the gradient of the response, $\nabla_{\mathbf{g}} P_0$, projected along $\hat{\mathbf{y}}_{j-1}^{\text{maj}}$ (i.e., the directional derivative of P_0) has a constant sign, either $+$ or $-$, over the support of $\text{PDF}_{j-1}(\mathbf{g})$, then the penalty will be zero. Changes in the sign of the directional derivative reveal oscillatory behavior, and are penalized accordingly.

For a population sampled from the prior distribution $\mathbf{g}_s \sim \mathcal{P}_{j-1}$, we estimate the integrals in $\bar{\text{FI}}$ and M as discrete averages over the sample populations:

$$\bar{\text{FI}}(\mathbf{c}) = \frac{1}{S} \sum_{s=1}^S \text{FI}(\mathbf{g}_s; \mathbf{c}), \quad (\text{B4})$$

$$M(\mathbf{c}, \mathcal{P}_{j-1}) = \frac{1}{S} \sum_{s=1}^S \Theta \left[-\frac{\hat{\mathbf{y}}_{j-1}^{\text{maj}} \cdot \nabla P_0(\mathbf{g}_s; \mathbf{c})}{\hat{\mathbf{y}}_{j-1}^{\text{maj}} \cdot \nabla P_0(\bar{\mathbf{g}}_{j-1}; \mathbf{c})} \right]. \quad (\text{B5})$$

We reiterate that these are implemented in a graph-based computational framework that allows for efficient automatic differentiation. This allows us to compute the gradients that appear explicitly in $\bar{\text{FI}}$ and M , and also to

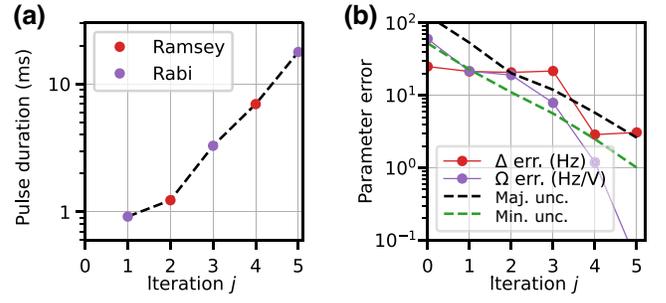


FIG. 7. Experimental OBSID results in a single-ion qubit calibration, using optimized Rabi or Ramsey pulse types, as illustrated in Fig. 3(a). (a) Pulse durations are shown at each iteration, with the autonomously selected pulse type indicated by color. (b) The iterative improvement in the parameter accuracy, indicated by the major uncertainty of the posterior population, and verified by the absolute error relative to conventionally calibrated parameter values. The floor in the absolute errors for Δ and Ω is set by the inferred uncertainty in the conventional calibration methods, which has a standard error of about 0.3 Hz for both Ω and Δ .

implement efficient gradient-based optimization of the cost function.

APPENDIX C: SUMMARY RESULTS OF ALTERNATIVE EXPERIMENTAL PULSE PARAMETERIZATIONS

As described in Sec. IV B, we have implemented other experimental control pulse sequences within the OBSID loop for calibration of the ion-trap experiment. Specifically, these are (1) optimizable Rabi-Ramsey pulses and (2) optimizable bang-bang control pulses. Here we present the results of experimental OBSID with these pulse parameterizations.

We note that the prior and posterior population plots look qualitatively similar to those shown in the main text, so we do not plot them here.

1. OBSID-selected Rabi-Ramsey pulses

Figure 7(a) shows the optimized pulse durations and types for the Rabi-Ramsey parameterizations illustrated in Fig. 3(a). The pulse type and the duration are autonomously chosen by OBSID for *in situ* experimental calibration of a single ion. The pulse durations grow by a factor $\lesssim 2$ per iteration.

Figure 7(b) shows the progression in the absolute error (points) and the major uncertainty (dashed black line) with iteration count. The major uncertainty provides a good estimate for the absolute error, although not a strict upper bound. At the last iteration, $j = 5$, the major uncertainty is ± 3 Hz, consistent with the absolute error in Δ .

The experimental results from the OBSID-selected Rabi-Ramsey sequence of pulses are very similar to those

in the equivalent simulated results shown in Fig. 3. In particular, at $j = 5$, comparing Figs. 3 and 7, we see that the pulse duration has grown by a factor of about 20 to 30 relative to the initial duration, and the major uncertainty has reduced by a similar factor.

2. OBSID-selected bang-bang pulses

Bang-bang control involves repeatedly switching a control field between an “on” state, with unit magnitude, and an “off” state, with zero magnitude. This is implemented here as a PWC with varying segment phases and durations, and allowing the carrier phase to vary from segment to segment.

This control type may be useful if there is a nonlinear transfer function between the control field generator (which in this experiment is an AWG) and the target system (which in this experiment is an ion). In this case, fixing the “on”-state magnitude removes the need to fully calibrate the nonlinear transfer function for a continuum of different pulse magnitudes.

Figure 8(a) shows the optimized complex-valued bang-bang pulses [represented as in-phase (I) and quadrature (Q) pairs] for the first three iterations of OBSID. In this example, the control pulses are five-segment PWCs, and the first segment is always “on,” with zero phase. The segment durations are optimized independently with a lower bound of zero.

Figure 8(b) shows that in this experimental run, the pulse durations grow by a factor of 2 to 3 per iteration, up to a maximum duration of about 60 ms at OBSID iteration $j = 5$.

The final pulse duration is somewhat longer in this example than the other experimental single-qubit pulse parameterizations reported in Figs. 5 and 7. There is a corresponding improvement in the final major uncertainty, which is shown in Fig. 8(c). We see that at iteration $j = 5$, the major uncertainty is ± 1 Hz, which is close to the ± 0.3 Hz floor set by standard experimental calibration methods used to independently determine the “true” parameter values for the system being measured. The absolute error in the parameters remains consistent with the major uncertainty as an estimated bound.

APPENDIX D: PRACTICAL SYSTEM IDENTIFIABILITY

There is a more subtle issue, related to statistical identifiability [28], which we first describe, and then illustrate with further simulations. The goal of OBSID is to accurately approximate the true parameter values \mathbf{g}_t that determine the system dynamics. At each iteration, OBSID improves its estimate by optimally discriminating between different plausible values \mathbf{g}_t in a prior sample population, and it achieves this by choosing pulses that minimize the posterior uncertainty. However, for some control-pulse

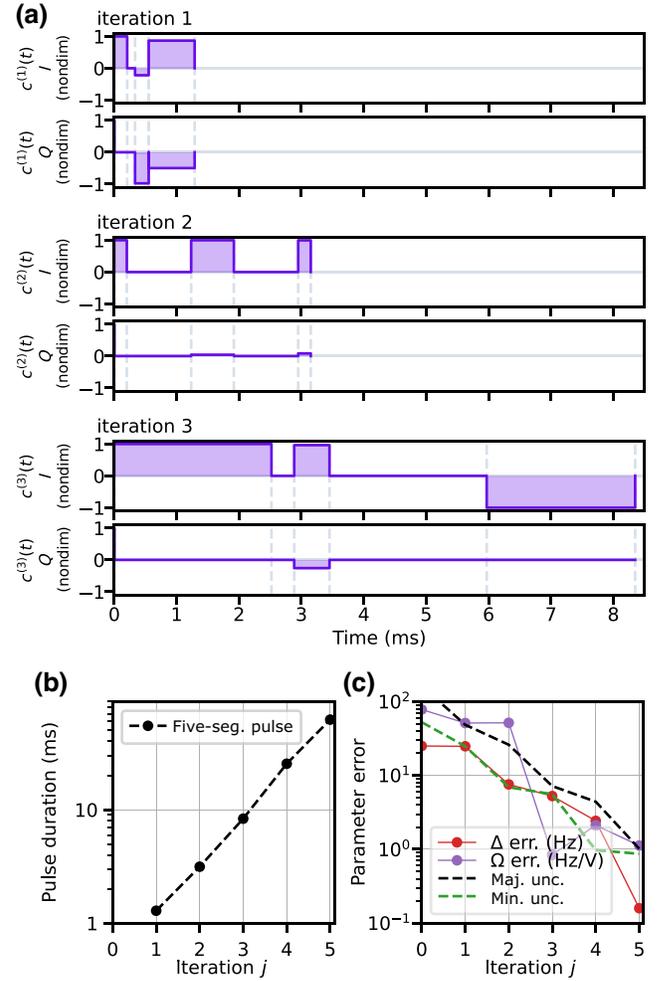


FIG. 8. Experimental OBSID results in a single-ion qubit calibration, using optimized five-segment bang-bang pulses, in which the control is either “on” (unit magnitude) with optimized phases or “off” (zero magnitude). (a) The sequence of complex-valued pulses [represented as in-phase (I) and quadrature (Q) components] generated in the first three iterations of OBSID, which optimizes over segment durations and phases. (b) Pulse durations for each iteration, and (c) the iterative improvement in the parameter accuracy. The floor in the absolute errors for Δ and Ω is set by the inferred uncertainty in the conventional calibration methods, which has a standard error of about 0.3 Hz for both Ω and Δ .

parameterizations, there may be a manifold of model parameter values that give statistically close measurement results for *all* available pulses within the parameterization. In this case, the OBSID protocol will eventually stall, with a distribution whose major uncertainty is independent of the pulse duration. At this point, OBSID will simply report a nearly stationary posterior parameter distribution.

In control theory and statistics, a parameterized probability distribution is called *identifiable* if empirical observations allow the parameters of the distribution to be identified uniquely, with no limitation on the number of

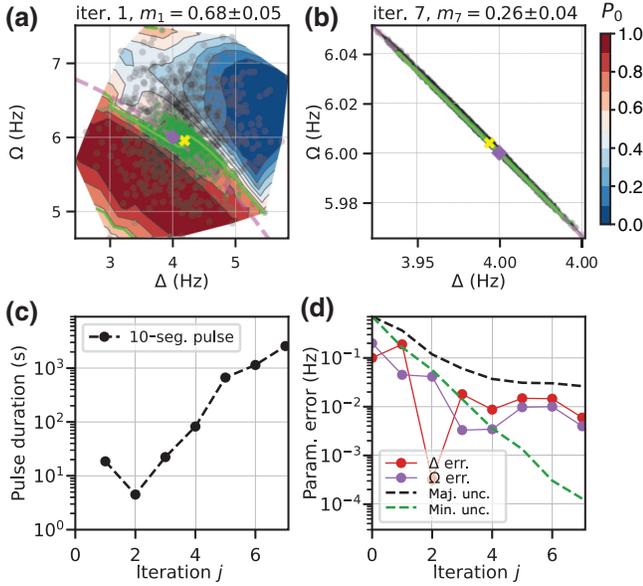


FIG. 9. A control example where OBSID fails to improve parameter estimates. The control pulse has uniform amplitude, but with $n = 10$ tunable phase segments $\phi_n(t)$, so that $c(t) = e^{i\phi_n(t)}$. Panels (a) and (b) show the prior populations and posterior populations after the first and seventh iterations, respectively. The magenta line indicates the contour $\Delta^2 + \Omega^2 = \Delta_t^2 + \Omega_t^2$, showing that the control successfully collapses the posterior distribution onto the contour, but fails to constrain the posterior distribution along the contour tangent direction. Panel (c) shows that OBSID chooses increasingly long pulses, as in previous examples; however, (d) shows that improvement in the major uncertainty “stalls” at a marginal uncertainty of around ± 0.02 Hz.

observations allowed. If more than one set of parameters yields the same distribution then it is not *structurally identifiable* [29]. In some cases, the response may be very shallow along some submanifold of parameter space, in which case the parameters may not be *practically identifiable* [29] given a finite number of measurements; this is the cause of OBSID stalling.

Practically, OBSID is limited to a finite number of empirical observations, and so practical statistical identifiability asks: for a given control pulse parameterization, and an upper limit on the number of measurements, r , can we improve the posterior parameter uncertainty by increasing the pulse duration (while tuning other control parameters, e.g., segment phases or amplitudes)? As we now demonstrate, there are poorly parameterized control pulses for which the model is not practically identifiable.

An example of this effect is a parameterization with a smaller number of free control parameters than the number of unknown system parameters, p . For example, with a finite number of observation repetitions, r , only using Rabi pulses with only a controllable duration parameter, $T_{\text{Rabi}}^{(\text{opt})}$, to identify the two parameters $\{\Delta, \Omega\}$

in Eq. (1) yields statistically significant information only about the combination $\Delta^2 + \Omega^2$. Consider two “nearby” models in parameter space, described by $\mathbf{g}_a = \{\Delta_a, \Omega_a\}$ and $\mathbf{g}_b = \{\Delta_b, \Omega_b\}$, with $\Delta_a \approx \Delta_b$ and $\Omega_a \approx \Omega_b$. These will give almost identical predictions for the Rabi response whenever $\Delta_a^2 + \Omega_a^2 = \Delta_b^2 + \Omega_b^2$. That is, within the two-dimensional parameter space, there is one-dimensional submanifold of models constrained by $\Delta^2 + \Omega^2 = \Delta_t^2 + \Omega_t^2$ that are statistically indistinguishable, given a finite r . It follows that, within this submanifold, Δ and Ω cannot be independently determined. The model is not practically identifiable with “Rabi-pulse-only” control.

We illustrate practical nonidentifiability with a generalization of the “Rabi-pulse-only” example. Suppose that we attempt to use OBSID to identify the single-qubit, two-parameter model in Eq. (1), using a pulse with uniform amplitude, but with controllable, time-dependent phases, so that $c(t) = e^{i\phi_n(t)}$, where the control phase $\phi_n(t)$ is a real-valued PWC with $n \geq 2$ segments. For this control parameterization, there are more control parameters available than unknown model parameters, which is a necessary condition for OBSID to succeed. Figures 9(a) and 9(b) show the prior and posterior population samples at the first and last of seven OBSID iterations; the final distribution is confined to a (nearly) one-dimensional manifold aligned with the contour $\Delta^2 + \Omega^2 = \Delta_t^2 + \Omega_t^2$, indicated by the dashed magenta line. This posterior distribution is strongly covariant, so that neither parameter can be accurately fixed. Figure 9(c) shows that the optimally chosen pulses still grow in duration; however, Fig. 9(d) illustrates that the major uncertainty saturates around iteration $j = 4$. Further iterations decrease the major uncertainty of this distribution slowly as $1/\sqrt{r}$, rather than quickly as $1/T^{(j)}$. This is a symptom that this “phase-only” control parameterization is unable to practically identify the system [29].

This example shows how OBSID can fail for poorly adapted control parameterizations that are unable to identify the system in practice. Importantly, it also illustrates a symptom of this failure, which is that major uncertainty stalls after several iterations of OBSID. The resolution is to extend the control parameterization to access more of the model parameter space. This self-diagnostic capability is a key advantage of the protocol. In the example described above, including control pulse amplitudes in the control parameterization would suffice.

-
- [1] S. G. Schirmer and F. C. Langbein, Ubiquitous problem of learning system parameters for dissipative two-level quantum systems: Fourier analysis versus Bayesian estimation, *Phys. Rev. A* **91**, 022125 (2015).
 - [2] A. R. R. Carvalho, H. Ball, M. J. Biercuk, M. R. Hush, and F. Thomsen, Error-robust quantum logic optimization using

- a cloud quantum computer interface, *Phys. Rev. Appl.* **15**, 064054 (2021).
- [3] F. Martínez-García, D. Vodola, and M. Müller, Adaptive Bayesian phase estimation for quantum error correcting codes, *New J. Phys.* **21**, 123027 (2019).
- [4] P. Aliferis, D. Gottesman, and J. Preskill, Accuracy threshold for postselected quantum computation, *Quantum Inf. Comput.* **8**, 181 (2008).
- [5] M. Kjaergaard, M. E. Schwartz, J. Braumüller, P. Krantz, J. I.-J. Wang, S. Gustavsson, and W. D. Oliver, Superconducting qubits: Current state of play, *Annu. Rev. Condens. Matter Phys.* **11**, 369 (2020).
- [6] E. Nielsen, J. K. Gamble, K. Rudinger, T. Scholten, K. Young, and R. Blume-Kohout, Gate set tomography, *Quantum* **5**, 557 (2021).
- [7] M. Takita, A. W. Cross, A. D. Córcoles, J. M. Chow, and J. M. Gambetta, Experimental demonstration of fault-tolerant state preparation with superconducting qubits, *Phys. Rev. Lett.* **119**, 180501 (2017).
- [8] Z. Chen, *et al.*, Exponential suppression of bit or phase errors with cyclic error correction, *Nature* **595**, 383 (2021).
- [9] J. M. Pino, J. M. Dreiling, C. Figgatt, J. P. Gaebler, S. A. Moses, M. S. Allman, C. H. Baldwin, M. Foss-Feig, D. Hayes, K. Mayer, C. Ryan-Anderson, and B. Neyenhuis, Demonstration of the trapped-ion quantum CCD computer architecture, *Nature* **592**, 209 (2021).
- [10] L. Postler, S. Heussen, I. Pogorelov, M. Rispler, T. Feldker, M. Meth, C. D. Marciniak, R. Stricker, M. Ringbauer, R. Blatt, P. Schindler, M. Müller, and T. Monz, Demonstration of fault-tolerant universal quantum gate operations, *Nature* **605**, 675 (2022).
- [11] S. Krinner, N. Lacroix, A. Remm, A. Di Paolo, E. Genois, C. Leroux, C. Hellings, S. Lazar, F. Swiadek, J. Herrmann, G. J. Norris, C. K. Andersen, M. Müller, A. Blais, C. Eichler, and A. Wallraff, Realizing repeated quantum error correction in a distance-three surface code, *Nature* **605**, 669 (2022).
- [12] A. Foster, M. Jankowiak, E. Bingham, P. Horsfall, Y. W. Teh, T. Rainforth, and N. Goodman, in *Advances in Neural Information Processing Systems*, Vol. 32, edited by H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (Curran Associates, Inc., La Jolla, United States, 2019). <https://proceedings.neurips.cc/paper/2019/file/d55cbf210f175f4a37916eafe6c04f0d-Paper.pdf>.
- [13] A. Foster, M. Jankowiak, M. O'Meara, Y. W. Teh, and T. Rainforth, in *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, Proceedings of Machine Learning Research, Vol. 108, edited by S. Chiappa and R. Calandra (PMLR, 2020) p. 2959. <https://proceedings.mlr.press/v108/foster20a.html>.
- [14] Y. Baum, M. Amico, S. Howell, M. Hush, M. Liuzzi, P. Mundada, T. Merkh, A. R. Carvalho, and M. J. Biercuk, Experimental deep reinforcement learning for error-robust gate-set design on a superconducting quantum computer, *PRX Quantum* **2**, 040324 (2021).
- [15] M. Arulampalam, S. Maskell, N. Gordon, and T. Clapp, A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking, *IEEE Trans. Signal Process.* **50**, 174 (2002).
- [16] L. Gerster, F. Martínez-García, P. Hrmo, M. W. van Mourik, B. Wilhelm, D. Vodola, M. Müller, R. Blatt, P. Schindler, and T. Monz, Experimental Bayesian calibration of trapped-ion entangling operations, *PRX Quantum* **3**, 020350 (2022).
- [17] P. Kok and B. W. Lovett, *Introduction to Optical Quantum Information Processing* (Cambridge University Press, Cambridge, England, 2010).
- [18] A. H. Kiilerich and K. Mølmer, Quantum Zeno effect in parameter estimation, *Phys. Rev. A* **92**, 032124 (2015).
- [19] N. R. Draper and I. Guttman, Confidence intervals versus regions, *J R Stat Soc Ser D (Stat)* **44**, 399 (1995).
- [20] S. T. Tokdar and R. E. Kass, Importance sampling: A review, *Wiley Interdiscip. Rev. Comput. Stat.* **2**, 54 (2010).
- [21] H. Ball, M. J. Biercuk, A. R. R. Carvalho, J. Chen, M. Hush, L. A. D. Castro, L. Li, P. J. Liebermann, H. J. Slatyer, C. Edmunds, V. Frey, C. Hempel, and A. Milne, Software tools for quantum control: Improving quantum computer performance through noise and error suppression, *Quantum Sci. Technol.* **6**, 044011 (2021).
- [22] Q-CTRL, Boulder Opal, (2022), <https://q-ctrl.com/boulder-opal> [Online].
- [23] M. Abadi, *et al.*, TensorFlow: Large-scale machine learning on heterogeneous systems, (2015), software available from tensorflow.org <https://www.tensorflow.org/>.
- [24] A. Stuhlmüller, J. Taylor, and N. Goodman, in *Advances in Neural Information Processing Systems*, Vol. 26, edited by C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Weinberger (Curran Associates, Inc., La Jolla, United State, 2013). <https://proceedings.neurips.cc/paper/2013/file/7f53f8c6c730af6aeb52e66eb74d8507-Paper.pdf>.
- [25] A. R. Milne, C. L. Edmunds, C. Hempel, F. Roy, S. Mavandia, and M. J. Biercuk, Phase-modulated entangling gates robust to static and time-varying errors, *Phys. Rev. Appl.* **13**, 024022 (2020).
- [26] S. Olmschenk, K. C. Younge, D. L. Moehring, D. N. Matsukevich, P. Maunz, and C. Monroe, Manipulation and detection of a trapped Yb⁺ hyperfine qubit, *Phys. Rev. A* **76**, 052314 (2007).
- [27] E. Peik, T. Schneider, and C. Tamm, Laser frequency stabilization to a single ion, *J. Phys. B: At., Mol. Opt. Phys.* **39**, 145 (2005).
- [28] E. D. Sontag, *Mathematical Control Theory: Deterministic Finite Dimensional Systems* (Springer Science & Business Media, Berlin, Germany, 2013), Vol. 6.
- [29] A. Raue, C. Kreutz, T. Maiwald, J. Bachmann, M. Schilling, U. Klingmüller, and J. Timmer, Structural and practical identifiability analysis of partially observed dynamical models by exploiting the profile likelihood, *Bioinformatics* **25**, 1923 (2009).
- [30] J. M. Chow, A. D. Córcoles, J. M. Gambetta, C. Rigetti, B. R. Johnson, J. A. Smolin, J. R. Rozen, G. A. Keefe, M. B. Rothwell, M. B. Ketchen, and M. Steffen, Simple all-microwave entangling gate for fixed-frequency superconducting qubits, *Phys. Rev. Lett.* **107**, 080502 (2011).
- [31] Y. Yang, B. Gong, and W. Cui, Real-time quantum state estimation in circuit QED via the Bayesian approach, *Phys. Rev. A* **97**, 012119 (2018).

- [32] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, England, 2000).
- [33] P. S. Mundada, A. Barbosa, S. Maity, Y. Wang, T. Merkh, T. Stace, F. Nielson, A. R. Carvalho, M. Hush, M. J. Biercuk, and Y. Baum, Experimental benchmarking of an automated deterministic error-suppression workflow for quantum algorithms, *Phys. Rev. Appl.* **20**, 024034 (2023).
- [34] P. V. Klimov, *et al.*, Optimizing quantum gates towards the scale of logical qubits, *ArXiv:2308.02321* (2023).
- [35] M. Kang, Q. Liang, M. Li, and Y. Nam, Efficient motional-mode characterization for high-fidelity trapped-ion quantum computing, *Quantum Sci. Technol.* **8**, 024002 (2023).
- [36] I. Najfeld and T. F. Havel, Derivatives of the matrix exponential and their computation, *Adv. Appl. Math.* **16**, 321 (1995).
- [37] Y. Li, L. Pezzè, M. Gessner, Z. Ren, W. Li, and A. Smerzi, Frequentist and Bayesian quantum phase estimation, *Entropy* **20**, 9 (2018).