

Resonant Coupling Parameter Estimation with Superconducting Qubits

J.H. Béjanin^{1,2}, C.T. Earnest^{1,2}, Y.R. Sanders^{3,4} and M. Mariantoni^{1,2,*}

¹*Institute for Quantum Computing, University of Waterloo, 200 University Avenue West, Waterloo, Ontario N2L 3G1, Canada*

²*Department of Physics and Astronomy, University of Waterloo, 200 University Avenue West, Waterloo, Ontario N2L 3G1, Canada*

³*Department of Physics and Astronomy, Macquarie University, Sydney, New South Wales 2109, Australia*

⁴*ARC Centre of Excellence in Engineered Quantum Systems, Macquarie University, Sydney, New South Wales 2109, Australia*

 (Received 28 August 2020; revised 12 July 2021; accepted 25 October 2021; published 30 November 2021)

Today's quantum computers are composed of tens of qubits interacting with each other and the environment in increasingly complex networks. To achieve the best possible performance when operating such systems, it is necessary to have accurate knowledge of all parameters in the quantum computer Hamiltonian. In this paper, we demonstrate theoretically and experimentally a method to efficiently learn the parameters of resonant interactions for quantum computers consisting of frequency-tunable superconducting qubits. Such interactions include, for example, those with other qubits, resonators, two-level systems, or other wanted or unwanted modes. Our method is based on a significantly improved swap spectroscopy calibration and consists of an *offline* data collection algorithm, followed by an *online* Bayesian learning algorithm. The purpose of the offline algorithm is to detect and coarsely estimate resonant interactions from a state of zero knowledge. It produces a *quadratic* speedup in the scaling of the number of measurements. The online algorithm subsequently refines the estimate of the parameters to accuracy comparable with that of traditional swap spectroscopy calibration but in constant time. We perform an experiment implementing our technique with a superconducting qubit. By combining both algorithms, we observe a reduction of the calibration time by 1 order of magnitude. Our method will improve present medium-scale superconducting quantum computers and will also scale up to larger systems. Finally, the two algorithms presented here can be readily adopted by communities working on different physical implementations of quantum computing architectures.

DOI: [10.1103/PRXQuantum.2.040343](https://doi.org/10.1103/PRXQuantum.2.040343)

I. INTRODUCTION

Quantum computing architectures based on different types of physical qubits have been investigated since the late 1990s [1]. Although quantum computers will be lacking fault tolerance for the near future [2,3], a variety of potential applications have been devised for such systems. Examples include solving optimization problems [4,5], machine learning [6–9], materials science and chemistry [10–12], and even commercializable technologies [13].

Small-scale and medium-scale quantum computers built of superconducting circuits [14–17] and comprising up

to a few tens of physical qubits have been in operation for the past several years [18–23]. The two main reasons for the success of superconducting technologies in the quantum computing arena are the pre-existing facilities for scaling up fabrication, due to methods being similar to silicon technology, and the favorable coherence-to-gate-time ratio. For state-of-the-art superconducting qubits with very short gate times of less than 100 ns (e.g., see Ref. [24]), such a ratio reaches values close to 1000. This means that thousands of one-qubit and two-qubit gates can be performed within the qubit lifetime, with fidelities in excess of 99.4% [22]. Very recently, a major milestone in quantum computing was reached—quantum supremacy—where a 53-qubit superconducting quantum computer was used to sample the output of a pseudo-random quantum circuit [25]. As a result of these advances, quantum computing has transitioned into the so-called noisy intermediate-scale quantum era [26].

Despite these successes, the high-fidelity operation of medium-scale and large-scale quantum computers is

*matteo.mariantoni@uwaterloo.ca

Published by the American Physical Society under the terms of the *Creative Commons Attribution 4.0 International* license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

accompanied by the daunting task of calibrating numerous physical qubits. In particular, calibrating tunable qubits requires the estimation of resonant interaction parameters (i.e., the frequency and the coupling strength) of both wanted and unwanted resonances. Wanted resonances include those between pairs of interacting qubits [27,28], qubits and resonators [18], and qubits and tunable couplers [29,30]. In this case, knowledge of the interaction parameters allows the implementation of two-qubit gates and readout. Unwanted resonance are mainly those between qubits and two-level systems (TLSs) [31–33] as well as substrate and box modes [34]. Microwave control crosstalk may also behave as an unwanted resonance, causing Rabi oscillations at the frequency of a nearby qubit. Such unwanted resonances must be avoided through proper frequency spacing. TLSs, especially, are a pervasive source of errors in superconducting architectures that must be remediated [35]. In summary, learning all the interaction parameters allows a comprehensive calibration and, thereby, minimization of coherent and incoherent errors.

In this paper, we study theoretically and demonstrate experimentally a data-efficient and automated method for identifying and estimating the parameters of resonant interactions based on *swap spectroscopy* [18,36]. We realize swap spectroscopy by performing energy relaxation time T_1 measurements of a frequency-tunable Xmon transmon qubit [24] at different qubit frequencies. The identification and estimation method is divided into two parts: an *offline* data collection algorithm [37] and an *online* Bayesian learning algorithm [38,39]. Both algorithms are based on the dynamics of interacting quantum systems. The former is used from a state of zero knowledge about a particular frequency range to coarsely identify resonance parameters within that range. The latter focuses on improving the estimate of those parameters. In this context, the term “online” means that measurements taken during the execution of the algorithm inform the subsequent ones. For the “offline” method, the execution of the entire algorithm is predetermined. The offline algorithm is applicable to any of the resonance types mentioned above, whereas the online algorithm can be used only for coherent resonances, either wanted or unwanted.

By means of our parameter-estimation method, we can shorten the calibration time of an Xmon transmon qubit significantly. The offline data collection algorithm makes it possible to reduce the scaling of the number of measurements by a square root when compared with a traditional swap spectroscopy calibration (i.e., a quadratic speedup). In our experiment, this algorithm takes approximately 30 min to detect resonances in a bandwidth of 10 GHz: 1 order of magnitude less time than with traditional methods. The online Bayesian learning algorithm runs in approximately 250 s per resonance, bringing the estimation accuracy to the same level as in high-resolution traditional swap spectroscopy.

To test our method and compare it against traditional swap spectroscopy, we characterize the resonances within a 1-GHz bandwidth of an Xmon transmon qubit. We additionally synthesize two resonance modes (RMs) that emulate the interaction with another qubit, a resonator, or a TLS. These modes are created by applying a coherent drive with a microwave source to the qubit under test. The synthesized resonance mode is a convenient and flexible tool to test our method since we can arbitrarily change its *resonance frequency* by tuning the source frequency as well as its *coupling strength* by changing the emitted source power.

Our method is not confined to the realm of superconducting quantum computing. It can easily be adopted by practitioners working on different physical implementations of quantum computing architectures, such as trapped ions and semiconductor qubits [1].

This paper is organized as follows: In Sec. II, we explain qubit calibration in frequency-tunable architectures. In Sec. III, we summarize the working principle of traditional swap spectroscopy, explaining why it is inefficient for the purpose of detecting resonances. In Sec. IV, we introduce the offline octave sampling algorithm (Sec. IV A) and demonstrate its experimental implementation and associated data analysis procedure to detect interactions between an Xmon transmon qubit and four resonance modes, including an incoherent one (Sec. IV B). In Sec. V, we explain the online Bayesian learning algorithm (Sec. V A) and demonstrate its performance in finding an accurate estimate for the parameters of a resonance (Sec. V B). In Sec. VI, we discuss additional concerns with the algorithms and the relevance of our methods for quantum computing. Finally, in Sec. VII, we provide an outlook and conclusions.

II. QUBIT CALIBRATION IN FREQUENCY-TUNABLE ARCHITECTURES

A fundamental requirement for the operation of a quantum computer is the proper calibration of the physical qubits in the system. This calibration includes many specific operations. One of the most basic tasks, for example, is to run a Rabi experiment on each qubit. This allows the determination of some experimental parameters needed to set up, for example, a π -pulse and perform a measurement. Once this first task is realized, further experiments can refine the knowledge of the pulse amplitude, rotation axis, measurement parameters, etc. Finally, a full calibration requires knowing the precise parameters of the system Hamiltonian and the interaction with the environment, allowing the systematic optimization of the fidelity of one-qubit and two-qubit gates as well as measurement.

In a frequency-tunable superconducting qubit architecture such as the Google architecture [25] or the one used

in this work, an additional degree of freedom must be considered during calibration: the qubit frequency f_q . Xmon transmon qubits are one example of tunable qubits [24]. In this design, an on-chip capacitive island made from aluminum is coupled in parallel to a superconducting quantum interference device (SQUID) comprising two Josephson tunnel junctions in parallel, forming a superconducting loop [17]. An Xmon transmon qubit is a quantum *anharmonic* oscillator, characterized by a non-equally-spaced ladder of quantum states. The frequency (i.e., energy) difference f_q between the ground state $|g\rangle$ and the first excited state $|e\rangle$ differs from that between $|e\rangle$ and the second excited state $|f\rangle$ by the so-called qubit anharmonicity α [40]. The qubit transition frequency f_q is controlled *in situ* by applying a local external magnetic flux that threads the SQUID, tuning the Josephson energy E_J and therefore the level separation.

Frequency tunability leads to a few distinct advantages for the operation of a quantum computer. For instance, it allows adjustable qubit-qubit interactions because the effective coupling strength between two qubits depends on the frequency difference between them. This enables the implementation of several types of two-qubit gates, such as the controlled-phase gate, which takes advantage of state $|f\rangle$ as an auxiliary state [18,27,28,41], as well as the \sqrt{i} SWAP and i SWAP gates [18]. In addition, setting the frequencies of spatially neighboring qubits away from each other helps avoid control crosstalk and frequency crowding issues, the latter being endemic in fixed-frequency systems [42].

Another advantage inherent to frequency-tunable architectures is related to energy relaxation. On-chip superconducting qubits interact with a distribution of TLSs, which are present in the various amorphous dielectric materials surrounding the qubit metallic structures (e.g., silicon and aluminum oxides). While the microscopic origin of TLSs is still debated [31], their effect on the qubit leads to either a T_1 reduction or T_1 and f_q time fluctuations [43]. In particular, the T_1 reduction is caused by the coherent or incoherent exchange of energy between a qubit and semiresonant TLSs. The ability to set the frequency of a qubit away from that of TLSs is therefore desirable and realizable only with tunability.

Calibrating qubits to implement two-qubit gates or to avoid TLSs is a *parameter estimation* problem. We need to determine the Hamiltonian parameters that define the resonant interactions between a qubit and another system. In all the aforementioned cases, two parameters must be found: the resonance frequency and the coupling strength of the interaction.

Historically, swap spectroscopy has been a prominent method to perform this kind of calibration. Unfortunately, traditional swap spectroscopy is inefficient in the amount of data it requires and therefore slow. This is inconvenient for multiple reasons. First, as the number of qubits

in a system grows, so does the number of calibrations that must be performed. This is particularly relevant to qubit-qubit coupling calibration, which cannot be performed in parallel on all qubits. Second, TLSs in the environment are known to fluctuate over time [33,35,43–45]. Similarly, f_q itself can shift in time. The identification of resonant interactions must therefore be repeated at regular intervals. We thus require a robust, accurate, and time-efficient method to identify the parameters associated with resonant interactions.

III. TRADITIONAL SWAP SPECTROSCOPY

Swap spectroscopy is an experimental method that allows exploration of the environment of a qubit at various frequencies by use of the qubit itself as a probe. Traditionally, swap spectroscopy has been used to select the operating frequency of qubits, making it possible to avoid TLSs or regions of low T_1 . Additionally, it has been used to explore resonant interactions, such as those with other qubits [28] or resonators [36]. Performing swap spectroscopy requires a minimally calibrated qubit and, thus, is suitable as a tune-up experiment.

In a swap spectroscopy experiment the qubit is initialized at the so-called *idle frequency*. A π -pulse is then applied to the qubit, energizing it from $|g\rangle$ to $|e\rangle$. At the end of the π -pulse, a flux pulse is applied to the SQUID to tune the qubit to a different frequency, the *probe frequency* $f_p(A)$, where A is the pulse amplitude. This procedure requires knowledge of the correspondence between f_q and A , which can be calibrated via regular pulse spectroscopy (see Appendix B). After time t , the flux pulse is turned off and the qubit is measured back at the idle frequency. This pulse sequence is illustrated in Fig. 1. Note that use of a flux pulse to set f_p presents advantages over quasistatically changing the idle qubit frequency by means of a direct current to the SQUID. Namely, it avoids recalibrating the π -pulse and measurement pulse at each qubit frequency.

In a traditional experiment, t and f_p are swept linearly over a desired range and the qubit is measured at each point, recording how many measurement shots correspond

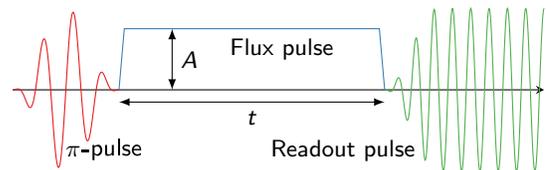


FIG. 1. Pulse sequence for a swap spectroscopy experiment. The initial π -pulse (red) excites the qubit, which is initialized at the so-called idle frequency. The flux pulse (blue) changes the qubit transition frequency f_q to the probe frequency $f_p(A)$ for duration t . The qubit is then measured (green) after being set back to its idle frequency.

to an excited or ground state, n_e or n_g , respectively. As a result of measuring the qubit in the energy basis, swap spectroscopy is insensitive to dephasing.

Figure 2(a) shows the result of a typical swap spectroscopy experiment, a *swap spectrum*, with data taken between 4.146 and 5.170 GHz and for times up to 500 ns. Resonant couplings appear as oscillations, or chevron patterns, of the measured average population $P_e = n_e/(n_e + n_g)$ in time. For example, on the far right of the spectrum it is possible to observe very fast oscillations, corresponding to a strong coupling g of approximately 40 MHz between the qubit and the measurement resonator (see Appendix A for details on the sample layout and experimental setup). To the left of the resonator we observe a slower oscillation, corresponding to a weaker interaction between the qubit and the synthesized resonance mode. Finally, at an even lower frequency, around 4.35 GHz, we observe a “streaky” structure. In this region, the qubit excitation is lost faster than elsewhere, and we cannot observe any oscillation. This structure is caused by an incoherent resonance, most likely a TLS.

The features observed in Fig. 2(a) demonstrate a selection of possible resonant interactions: strong interactions, where $g \gg 1/T_1$, resulting in multiple coherent oscillation cycles, and weak interactions, appearing as regions of lower T_1 . Neither of them is ideal for the operation of a qubit. In the case shown in Fig. 2, the best choice for the

qubit idle frequency is around 4.6 GHz, far away from any unwanted interactions.

The data in Fig. 2(a) gives us a rough idea about the parameters of any possible resonance modes coupled to the qubit within the measured spectrum. It is hard to tell, however, that there are *two* resonance modes at 4.8 GHz, or what the frequency of the oscillation for the resonator is. A more detailed scan, such as the one in Fig. 2(b), might be necessary to estimate the parameters with sufficient accuracy. Traditional swap spectroscopy, with data taken in a linear grid, is a possible method to detect and estimate resonance modes. We show in Sec. IV, however, that it is inefficient, and that there is a much better way to perform this task: octave sampling.

IV. OFFLINE OCTAVE SAMPLING

The offline octave sampling algorithm has a similar objective as swap spectroscopy; that is, to determine if there are any systems interacting resonantly with the qubit and provide an estimate of their coupling parameters. However, we want to achieve this purpose in a more efficient fashion by acquiring less data, therefore saving valuable experimental time. Note that the pulse sequence used to perform octave sampling is the same as for swap spectroscopy (see Fig. 1). The difference lies in how the spectrum is sampled. Whereas traditional swap spectroscopy

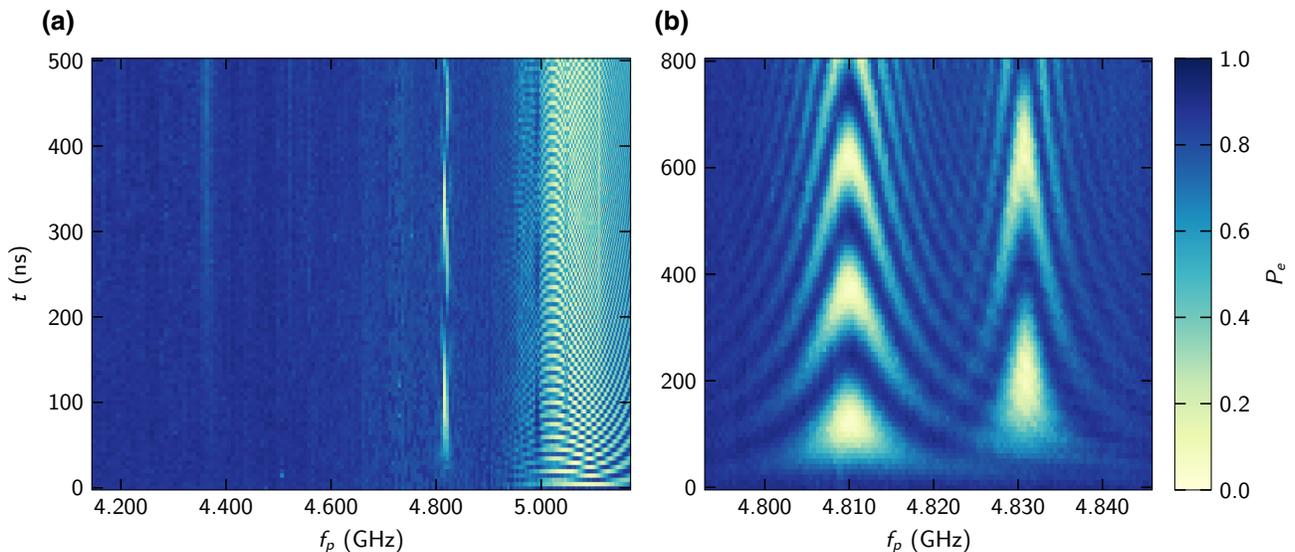


FIG. 2. Swap spectra for two frequency ranges. The x axis shows the probe frequency of the qubit, which is set by the amplitude of the flux pulse applied to the SQUID. The y axis indicates the length of the flux pulse before measurement and, therefore, corresponds to the interaction time with potential resonance modes. (a) Distinct features are visible in the full spectrum, including two chevron patterns around 4.8 GHz and one at 5.1 GHz. The resonance at 5.1 GHz looks aliased because of the low-resolution sampling of the time axis. A low- T_1 streak is also visible at 4.35 GHz, likely caused by an incoherent TLS. (b) Enlargement of the region with the slower coherent chevron patterns caused by synthesized resonance modes. For this experiment, the synthesizers are set at 4.8100 and 4.8314 GHz. A wide frequency scan is needed to see if and where there are resonance modes, as in (a), but a more detailed experiment, as in (b), is needed to properly estimate the resonance parameters.

samples the frequency-time space in a regular grid, octave sampling takes advantage of resonant dynamics to acquire as little data as possible.

A. Theoretical method

To explain the data collection strategy, we analyze the time dynamics of the systems at play. Since we are searching for resonant interactions with a qubit, we work in a single-excitation manifold ($|g\rangle \leftrightarrow |e\rangle$). Thus, even if a system is characterized by more than two energy levels (e.g., a resonator), we can still treat it as a two-level system because higher levels are never occupied. This is our working assumption throughout the rest of the paper.

Note that we can probe the environment of an anharmonic oscillator (e.g., a transmon) within a different single-excitation manifold. In that case, we can populate the second excited state and look for systems coupled to the $|e\rangle \leftrightarrow |f\rangle$ transition. This allows the calibration of certain two-qubit controlled-phase gates [18,22,27,28]. In either case, because we consider the exchange of a single excitation, the effective Hamiltonian remains unchanged.

After a rotating-wave approximation, the Hamiltonian of a qubit at the probe frequency f_p interacting with a resonance mode at a frequency f_{RM} reads

$$\hat{H} = \frac{hf_p}{2} \hat{\sigma}_{z,1} + \frac{hf_{\text{RM}}}{2} \hat{\sigma}_{z,2} + hg (\hat{\sigma}_1^+ \hat{\sigma}_2^- + \hat{\sigma}_1^- \hat{\sigma}_2^+), \quad (1)$$

where g is the coupling strength of the qubit–resonance-mode interaction, $\hat{\sigma}_{z,1(2)}$ are Pauli matrices for the qubit (1)

and resonance mode (2), $\hat{\sigma}_{1(2)}^+$ and $\hat{\sigma}_{1(2)}^-$ are raising and lowering operators for the qubit and resonance mode. We solve for the time evolution of the qubit when it is initialized in state $|e\rangle$ and with the resonance mode starting in $|g\rangle$. The theoretical probability of finding the qubit in the excited state after time t is then given by

$$\tilde{P}_e(t) = 1 - \left(\frac{2g}{\Omega} \right)^2 \sin^2(2\pi\Omega t/2), \quad (2)$$

where $\Omega^2 = \delta f^2 + 4g^2$, with $\delta f = f_p - f_{\text{RM}}$. The probability \tilde{P}_e of Eq. (2) is plotted in Fig. 3(a) as contours. Close to resonance, the excitation swaps between the qubit and the resonance mode with frequency Ω increasing at larger δf , resulting in the familiar chevron pattern. Both the width of the pattern, which we quantify by the full width at half maximum of the amplitude, $4g$, and Ω depend on g . Crucially, the width is proportional to g , while the period of the oscillation and the position in time of the first minimum are proportional to $1/g$. The goal is to detect a resonance mode by finding the *first minimum* of an oscillation, where $\tilde{P}_e \sim 0$ because the excitation has swapped into the resonance mode.

With these observations in mind, we choose to divide the frequency-time space into rows of *bins*, within which we take a constant number n_s of swap spectroscopy measurements. Instead of naively sampling the spectrum in a uniform grid, we adapt the measurement based on the value of g that we are trying to detect. The coupling strength determines the *time* t at which we measure and

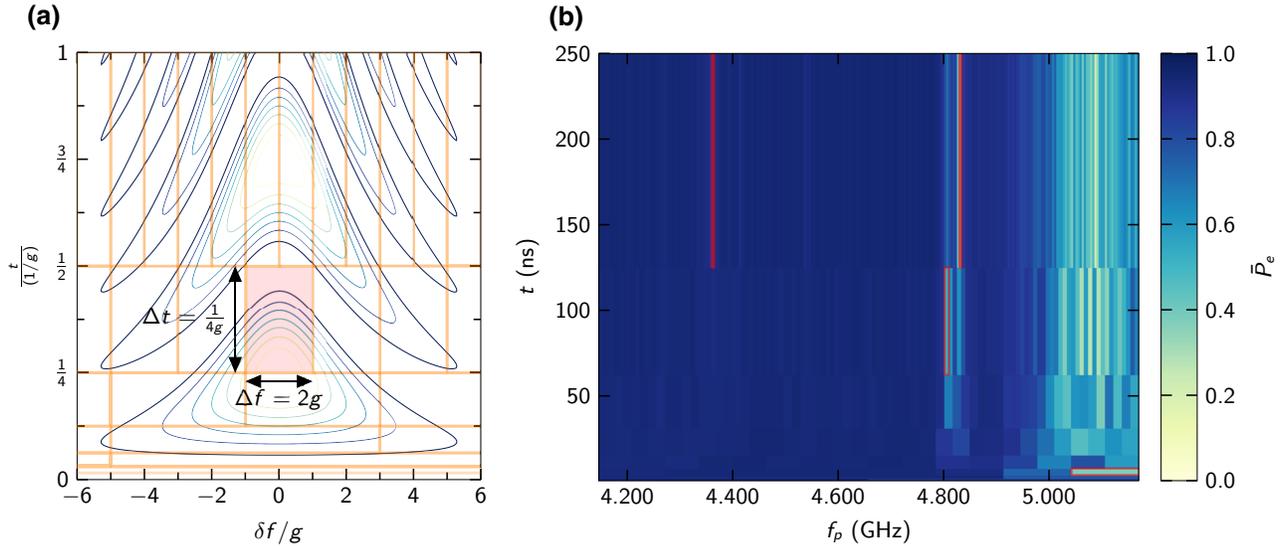


FIG. 3. Offline octave sampling. (a) Contour plot of the probability of finding the qubit in $|e\rangle$ [see Eq. (2)] as a function of time t and frequency detuning δf ; both axes are normalized by the coupling strength g . The highlighted box in the center indicates which portion of the chevron pattern is meant to be detected by the algorithm. (b) Swap spectroscopy experiment with octave sampling. Starting from a bin spanning the full measurement range, at each subsequent octave the bin width is halved and the bin length is doubled. The color of each bin represents the average value of the measured P_e over $n_s = 5$ samples. The red boxes indicate the resonances reported by the analysis explained in Sec. IV B.

the *bin size*. On the one hand, a resonance mode with large coupling strength g has a large width and a short period. For short time t , then, we choose the bins to be wide and short [see the bottom rows of bins in Fig. 3(a)]. On the other hand, a more weakly coupled resonance mode appears later in time, with a narrower frequency width and a longer period. In this case, the bins are longer and narrower [see the top rows of bins in Fig. 3(a)]. To cover the entire measurement space, all bins must be adjacent (without overlapping). This condition constrains the ratio of the number of bins in consecutive rows to be an integer. We choose this integer to be 2 because it is the only ratio for which a bin containing the first oscillation minimum (at $t = 1/4g$) does not contain any other later minima. For example, a factor of 3 would contain both the first oscillation minimum and the second oscillation minimum (at $t = 3/4g$), as can be deduced from Fig. 3(a). We refer to this method as *octave sampling* because consecutive bin rows are suited to detect resonances with a coupling strength ratio of 2.

To make this bin division systematic, we introduce the concept of a coupling octave with coupling strength g_m , where m is the octave number, ranging from 0 to m_f [37]. The final octave number m_f is determined by the desired frequency or time resolution, as discussed below. For each octave, the full frequency spectrum to be analyzed ranges between a minimum frequency f_{\min} and a maximum frequency f_{\max} . This range is divided into 2^m bins of equal size, with frequency width $\Delta f = 2g_m$ and time length $\Delta t = 1/4g_m$. The location in time of the bins' lower edge is $t = 1/4g_m$. One such bin, with $g_m = g$, is highlighted in red in Fig. 3(a). The highlighted bin is not centered on the oscillation minimum. This is because a low- P_e measurement in that area corresponds to a *range* of possible coupling strengths; namely, those for which $g_m/2 \leq g \leq g_m$. The resonance mode plotted in Fig. 3(a) is at the upper end of this range and, hence, at the lower edge of the bin.

The execution of the algorithm is determined by the total bandwidth $B = f_{\max} - f_{\min}$, which is the frequency width of the single bin spanning the whole spectrum at the zeroth octave. This bandwidth corresponds to a coupling octave $g_0 = B/2$ and, therefore, to a time length $\Delta t = 1/4g_0$. For the next octave, we divide the width of the bins by 2 such that the subsequent row has twice as many bins as in the previous step. The length of the bins in time is correspondingly doubled. An example of this division is shown by the orange grid in Fig. 3(a).

Thus, if we are given as inputs f_{\min} , f_{\max} , and m_f , the execution of the algorithm, starting from the zeroth octave $m = 0$, is as follows:

- (1) Divide the frequency range into 2^m bins, each with width $2g_m = B/2^m$.
- (2) Each bin spans the time values $1/4g_m \leq t \leq 1/2g_m$.

- (3) Take n_s swap spectroscopy samples within each of the 2^m bins, sampling uniformly in frequency and inverse uniformly in time. That is, for bin $k = 1, \dots, 2^m$, draw f_p and t as

$$f_p \sim \mathcal{U}(f_{\min} + 2(k-1)g_m, f_{\min} + 2kg_m),$$

$$t \sim \mathcal{U}(2g_m, 4g_m)^{-1},$$

where the notation $X \sim \mathcal{U}(a, b)$ signifies that X is drawn randomly from a continuous uniform distribution \mathcal{U} between a and b .

- (4) Increment m and start over for the next octave.

The total number of bins to be measured, N_{bins} , depends both on the size of the bandwidth B and the final octave number m_f . To set m_f , we can choose either a maximum time $t_{m_f} = 1/2g_{m_f}$ or a final frequency resolution $\Delta f_{m_f} = 2g_{m_f}$, or, and perhaps most usefully, a minimum coupling strength $g_{\min} = g_{m_f}/2$. Any of these quantities determine the number of bins for the final octave through the octave coupling g_{m_f} , and must be picked according to the goal of the experimenter. Then, following from the equation for the bin width above, we find $m_f = \lceil \log_2(B/\Delta f_{m_f}) \rceil = \lceil \log_2(B/2g_{m_f}) \rceil$ (the result of the logarithm is rounded up). Accordingly, N_{bins} can be calculated by summing the number of bins per octave:

$$N_{\text{bins}} = \sum_{m=0}^{m_f} 2^m$$

$$= 2^{m_f+1} - 1$$

$$= \frac{2B}{\Delta f_{m_f}} - 1. \quad (3)$$

The total number of points is thus given by $N_{\text{oct}} = n_s N_{\text{bins}}$.

In comparison, with the same frequency resolution, traditional swap spectroscopy divides the frequency axis into $B/\Delta f_{m_f}$ points and the time axis into $t_{m_f}/\Delta t_{\min}$ points, where Δt_{\min} is the time resolution [46]. While octave sampling reaches a time resolution of $1/B$, it would be unfair to the traditional method to use that number directly. Instead, we assume that $1/\Delta t_{\min}$ is on the order of hundreds of megahertz, allowing the detection of strong couplings such as those to other qubits or resonators. The total number of points is then given by

$$N_{\text{trad}} = \frac{t_{m_f}}{\Delta t_{\min}} \frac{B}{\Delta f_{m_f}}$$

$$= \frac{1}{\Delta t_{\min}} \frac{B}{\Delta f_{m_f}^2}. \quad (4)$$

The number of points scales as $\mathcal{O}(1/\Delta f_{m_f}^2)$ for the traditional method, whereas it scales as $\mathcal{O}(1/\Delta f_{m_f})$ for octave sampling: a *quadratic* improvement.

As a last remark, while the octave sampling strategy is designed according to the dynamics of *coherent* resonances, it is also well suited to handle *incoherent* resonances, such as the one at 4.35 GHz in Fig. 2(a). Although an incoherent resonance does not show clear oscillations, it still increases the qubit relaxation rate. This relaxation appears as a low-excitation region in the spectrum, where frequency width and time position obey similar scaling rules as explained above. Such regions of low excitation can be detected just as well as oscillation minima.

B. Experimental results

The result of an experimental implementation of the octave sampling algorithm is shown in Fig. 3(b), which is the efficient version of Fig. 2(a). Each bin is colored according to the average excitation probability \bar{P}_e measured over $n_s = 5$ samples. We are able to discern the same features as in Fig. 2(a), that is, distinct low-excitation regions, while acquiring much less data.

The frequency ranges from $f_{\min} = 4.146$ GHz to $f_{\max} = 5.170$ GHz such that $B = 1024$ MHz. We choose the final octave number m_f to be 8, allowing us to detect TLSs with $g \geq 1$ MHz. This corresponds to $g_{m_f} = 2$ MHz and a minimum bin width Δf_{m_f} of 4 MHz. With use of Eq. (3), $N_{\text{bins}} = 511$ and $N_{\text{oct}} = 2555$.

From Eq. (4) and given the resolution to detect oscillations up to $1/\Delta t_{\min} = 200$ MHz (corresponding to $g = 100$ MHz) in a time interval $t_{m_f} = 250$ ns, traditional swap spectroscopy requires $N_{\text{trad}} = 25600$ [47].

For the parameters used in this experiment, octave sampling requires 1 order of magnitude fewer points than traditional swap spectroscopy. It is worth noting that $N_{\text{oct}} \neq \sqrt{N_{\text{trad}}}$ due to the prefactors in Eqs. (3) and (4). Since octave sampling provides a quadratic scaling improvement, the reduction in the number of measurements grows for experiments with higher resolution. For example, if we increase m_f to 9, N_{oct} roughly doubles to 5115, whereas N_{trad} is quadrupled to 102 400.

Given the octave sampling results, we intend to determine if there are one or more resonance modes interacting with the qubit. If there are resonances, we also want an estimate of their coupling parameters f_{RM} and g . If there are no resonance modes at all, the qubit does not undergo any swap, and we should always measure it to be in $|e\rangle$ with $\bar{P}_e = 1$. Hence, a measurement of $\bar{P}_e < 1$ indicates energy loss due to a resonance mode interacting with the qubit.

In practice, however, other spurious experimental effects can lower the measured \bar{P}_e below the theoretical value of 1, even in the absence of a resonance mode. They include, for instance, the “bare” energy relaxation rate of the qubit,

TABLE I. Resonance modes detected after analysis of the octave sampling data; the corresponding bins are boxed in red in Fig. 3(b). The parameters listed result from a prominence value of 0.39, except for RM4. For RM4, a lower value of 0.09 (more sensitive) is necessary. For values above 0.72, no resonance is detected. The prominence threshold indicates the maximum tested value for which the resonance is detected; at that value, the reported parameters are slightly different.

Parameter	RM1	RM2	RM3	RM4
Octave number m	7	8	3	8
Bin center frequency (GHz)	4.806	4.832	5.106	4.364
Bin coupling range (MHz)	2–4	1–2	32–64	1–2
Prominence threshold	0.39	0.50	0.72	0.09

state preparation, measurement visibility, bin averaging, and statistical fluctuations. We therefore require an analysis method that will reliably detect and extract resonances from the octave data, while minimizing false positives.

The method used to analyze the octave data primarily relies on a peak-finding function meant to detect low-excitation bins in the spectrum. To avoid duplicate detection of the same resonance mode, we apply a procedure to combine peaks corresponding to the same resonance found in different octaves. We configure the sensitivity of the analysis by setting the minimum *prominence* value used for peak finding. Setting the prominence to a lower value (more sensitive) will detect more peaks, potentially leading to detection of fainter resonances. However, a low value could also generate false positives if the data is noisy. More details on this analysis method can be found in Appendix C.

The result of the octave analysis consists of the bin location—frequency and octave number—for each detected resonance. Since the octave number corresponds to a coupling strength range, the task of detecting resonances and finding their approximate coupling parameters is achieved.

We analyze the octave sampling data shown in Fig. 3(b) with different prominence values to provide an understanding of the sensitivity required to detect resonances. In that experiment, a qubit interacts with three known modes: two are synthesized with a microwave source and one is the on-chip readout resonator mode. The analysis detects those three modes at a prominence value of 0.39. When the prominence is decreased to 0.09, an additional resonance is detected at 4.35 GHz. Even at that sensitivity, no false positives are reported. The coupling parameters resulting from the analysis are reported in Table I. The bins corresponding to those detections are boxed in red in Fig. 3(b).

We purposely choose the two synthesized modes, RM1 and RM2, to be close in frequency to illustrate an important feature of our method: two distinct resonances are detected separately only if their frequency spacing is sufficiently large. In particular, the frequency separation must be at least twice as large as the largest of the two coupling

strengths. This ensures that the oscillation minima are separated by one bin width. If that condition is not met, the two resonances are located either in the same bin or in neighboring bins, resulting in the detection of a single peak. Here, RM1 and RM2 are separated by 20 MHz and the coupling strength of RM1 is at most 4 MHz. This means that there is at least one bin separating the two modes, allowing them to be detected independently.

The resonator mode, RM3, is detected at a low octave number. This is because it is characterized by a large coupling strength to the qubit (see Appendix A for the coupling capacitance). The last detected mode, RM4, is very faint. It corresponds to an incoherent weakly coupled TLS, as we also see in the traditional swap spectroscopy data in Fig. 2(a). This mode can be detected only by setting a sufficiently low prominence value. It is likely that RM4 could be detected with a less sensitive analysis if an additional octave were sampled.

For the three known modes, we coarsely estimate the frequency and coupling strength with a minimal amount of data. Obtaining more precise and accurate results necessitates the online estimation algorithm, which is explained in the next section.

V. ONLINE BAYESIAN LEARNING ALGORITHM

The offline octave sampling algorithm is data efficient and can be performed from a state of zero knowledge of the qubit's spectrum. However, it does not provide a very accurate estimate of the coupling parameters of a resonance mode. To increase accuracy, we can use the coarse estimate given by the offline method to execute an *online Bayesian learning algorithm* and refine the parameters in a very short time. This process relies on measuring a few dozen points of the qubit oscillation in the swap spectrum, using again the pulse sequence in Fig. 1. For the online algorithm to work, the qubit must undergo an oscillation; therefore, this method cannot be used to estimate the coupling parameters of an incoherent resonance mode. If it is unknown whether a mode is coherent or not, a traditional swap spectroscopy experiment has to be run first.

Given an initial probability distribution over the coupling parameters with a resonance mode, the online algorithm successively selects measurement settings to increase knowledge. After the result of a measurement is recorded, the distribution is updated according to Bayes' theorem and a new measurement setting is produced. This procedure is repeated iteratively until the distribution converges as desired [48].

A. Theoretical method

The online estimation algorithm is the experimental implementation of the theory proposed in Ref. [38]. It uses a particle filter method to efficiently represent the *prior*

and *posterior* distributions and compute Bayes' theorem at each iteration.

A particle distribution is a discretized representation of a probability distribution. The denser the distribution in a particular region of the parameter space, the higher the probability of those parameters. Here, each particle represents a two-tuple of the coupling parameters (f_{RM}, g) of a resonance mode. At the beginning of an iteration, we compute the means, $\mu(f_{\text{RM}}) = \langle f_{\text{RM}} \rangle$ and $\mu(g) = \langle g \rangle$, and standard deviations, $\sigma(f_{\text{RM}}) = \sqrt{\langle f_{\text{RM}}^2 \rangle - \langle f_{\text{RM}} \rangle^2}$ and $\sigma(g) = \sqrt{\langle g^2 \rangle - \langle g \rangle^2}$, of the prior distribution.

The next step is to perform a single measurement to determine the excited population P_e at a particular probe frequency f_p and time t . These *measurement settings* are heuristically selected to increase information gain [38]. In practice, t should scale inversely with $\sigma(g)$, while f_p should be within a factor of $\mu(g)$ on either side of $\mu(f_{\text{RM}})$.

We choose the following measurement settings:

$$f_p = \begin{cases} \mu(f_{\text{RM}}) + r_1 \mu(g) & \text{for } M \leq M_0, \\ \mu(f_{\text{RM}}) + cr_1 \sigma(f_{\text{RM}}) & \text{for } M > M_0, \end{cases} \quad (5)$$

$$t = \begin{cases} r_2 \tanh\left(\frac{a}{\sigma(g)t_{\text{max}}}\right) t_{\text{max}} & \text{for } M \leq M_0, \\ \frac{1+r_2}{2} \tanh\left(\frac{a}{\sigma(g)t_{\text{max}}}\right) t_{\text{max}} & \text{for } M > M_0, \end{cases} \quad (6)$$

where r_1 is picked from $\mathcal{U}(-1/2, 1/2)$, r_2 is picked from $\mathcal{U}(0, 1)$, $a = \pi/2$, $c = 5$, M is the iteration number, and we set $M_0 = 25$ (see below). These parameters are empirical constants determined in Ref. [38], although they are slightly adjusted for this experiment to have a larger distribution for f_p and t when $M > M_0$. Unlike in the method proposed in Ref. [38], we choose to limit t to a maximum value well under T_1 . This is done to prevent measurements from occurring after the oscillation has substantially decayed. For this purpose, we use the hyperbolic tangent function as it has linear behavior for small arguments, such that $\tanh(a/\sigma(g)t_{\text{max}})t_{\text{max}} \simeq a/\sigma(g)$ when $\sigma(g)$ is large.

After M_0 iterations, we modify the heuristic to accelerate convergence. Initially, we choose probe frequencies coarsely according to the value of g . Then, as our knowledge improves, $\sigma(f_{\text{RM}})$ decreases and can be used to select frequencies in a narrower range around $\mu(f_{\text{RM}})$. The factor c is used to avoid choosing measurement frequencies *too* narrowly. The time t is always weighted by approximately $1/\sigma(g)$, but we bias the selection to larger values after M_0 iterations.

The last step in the iteration is to apply Bayes' theorem to update our knowledge of the coupling parameters. We want to obtain the posterior distribution based on the measurement result $P_e(f_p, t)$. This is achieved in two sub-steps: (1) we compute the likelihood of obtaining the

measurement value given each particle's (f_{RM}, g) parameters; (2) we resample the distribution according to these likelihoods.

We compute the likelihood from the measurement result $P_e = n_e/n$, which is the proportion of n_e excited state outcomes for n individual measurement shots. Since the theoretical fraction we expect to measure is $\tilde{P}_e(f_p, t, f_{\text{RM}}, g)$ [given by Eq. (2) in the decoherence-free case [49]], we know that the result is a binomial random variable $n_e \sim \mathcal{B}(n, \tilde{P}_e)$. Accordingly, the likelihood of obtaining a particular measurement result given the measurement settings (f_p, t) and a particle with parameters (f_{RM}, g) is the probability mass function

$$\mathcal{L}(n_e | f_p, t, f_{\text{RM}}, g) = \binom{n}{n_e} \tilde{P}_e^{n_e} (1 - \tilde{P}_e)^{(n-n_e)}, \quad (7)$$

where $\binom{n}{n_e}$ is the binomial coefficient.

In effect, we are computing the probability that the measurement result corresponds to a resonance mode with coupling parameters (f_{RM}, g) . The next step is to resample the distribution to keep only those parameters that are most probable. Although this task can be achieved in a variety of ways, the general idea is to pick particles from the prior at random, weighted by the likelihood. To avoid duplicate particles in the posterior distribution, we add normally distributed random noise proportional to the covariance of the prior. The procedure chosen for this experiment is described in Appendix E.

The iteration process is visualized in Fig. 4, allowing us to understand more intuitively how the particle filter technique works. If the measurement is useful (i.e., the

resulting likelihood favors a subset of the prior), the posterior distribution is shrunk or filtered, improving knowledge of the parameters. Otherwise, if the likelihood does not discriminate the particles, the distribution is not modified significantly. After resampling, the next iteration starts with the last iteration's posterior as the prior [50].

The task of the online Bayesian learning algorithm is simpler than octave sampling because we already have a coarse estimate of the interaction parameters. The particle filter can therefore “fit” to the most likely parameters given the measurements. At the end of the final iteration, the parameters are given by the mean of the particle distribution. If the algorithm converges, the final particle “cloud” is small, resulting in an estimate that is accurate: both true and precise. If the algorithm does not converge, meaning that the final particle cloud is not tightly concentrated in a single region, it might be necessary to run the experiment again. Regenerating the initial particle distribution via the octave sampling method could also improve estimation performance.

Ideally, the uncertainty on the estimated parameters is given by the standard deviation of the final particle distribution. In practice, however, directly taking the standard deviation is generally not valid in an experimental context. This is because the algorithm does not take into account potential errors on the value of the measurement settings or any model inaccuracies. To obtain an uncertainty on the parameters, we instead rerun the full inference steps of the algorithm multiple times on the already-gathered data, starting from the same initial particle distribution. If the model and measurement settings were accurate, the results of rerunning the inference would be identical. However, since that is not the case, the estimates obtained

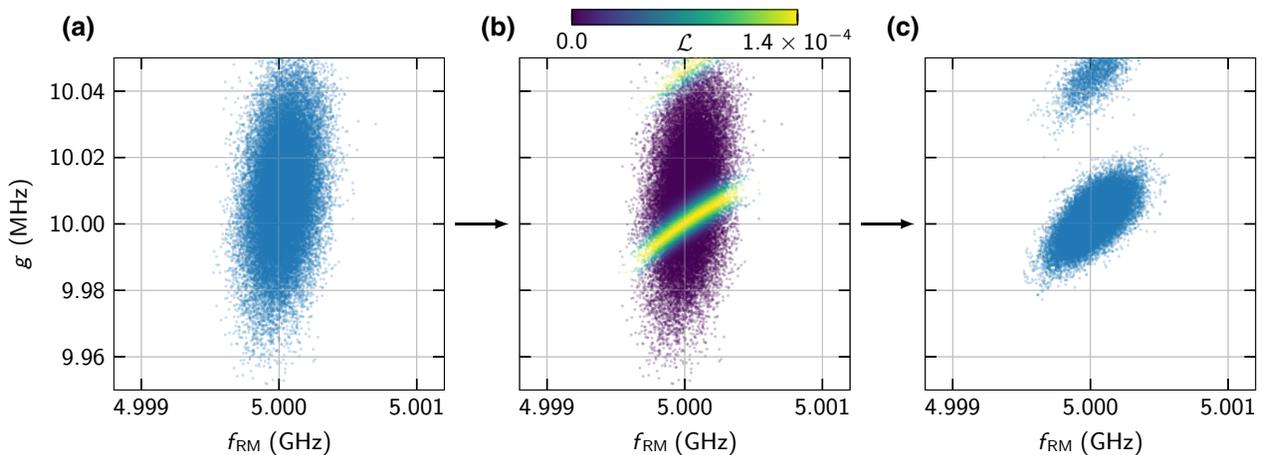


FIG. 4. Illustration of a simulated iteration of the online particle filter algorithm with 40 000 particles. (a) Given a prior distribution, we heuristically generate measurement settings f_p and t meant to increase information. (b) Following the measurement, the likelihood \mathcal{L} of the result is computed for each particle. The values shown on the scale are normalized. (c) We apply Bayes’ theorem to determine the posterior distribution. This task is achieved by resampling the particles according to their likelihoods. The distribution is split into two “clouds.” After resampling, the posterior distribution can be used as the next iteration’s prior.

TABLE II. Estimated parameters for the three resonance modes detected after running the online Bayesian learning algorithm. RM1 and RM2 are synthesized modes and RM3 corresponds to the qubit’s measurement resonator.

Parameter	RM1	RM2	RM3
f_{RM} (GHz)	4.8091(3)	4.8297(2)	5.033(12)
g (MHz)	2.78(3)	1.62(8)	37.9(3)

by rerunning the inference differ. The error we report is then the standard deviation of the different estimates. Since this procedure does not require acquisition of new data, it can be performed offline, after the experiment.

One possible cause of failure is the overestimation of g by an integer multiple. In that case, crests in the oscillation for the different frequencies partially overlap. To prevent such failures, the experiment can be run multiple times.

B. Experimental results

We run the online Bayesian learning algorithm on three distributions generated from the octave data, one for each detected resonance mode. The generation of those distributions is discussed in Appendix D. We do not execute the algorithm on RM4 since the resonance is incoherent and does not undergo the oscillations necessary to estimate the coupling parameters with this method.

For each mode, we perform 35 iterations of the algorithm, at which point the distribution has converged. The runtime of the algorithm for a single resonance mode is approximately 23 s. Afterwards, we rerun the full inference 200 times—enough for the resulting statistics to stop changing—and report the final means and standard deviations in Table II. As expected, the parameters of the synthesized modes and of the measurement resonator are correctly identified. The errors shown do not include systematic errors caused by an inaccurate flux-amplitude-to-frequency calibration (see Appendix B).

Here, the sources are set at 4.810 and 4.8305 GHz. These values are higher than the ones found in Table II by approximately 1 MHz, likely due to a systematic calibration error. In principle, it is possible to relate the power emitted by a synthesizer to the “coupling strength” of the mode. However, the attenuation and reflection of the signal between the source and the qubit make it impossible to accurately find such a relation. Instead, we fit the swap spectroscopy measurement in Fig. 2(b) with a nonlinear least squares regression, giving $g = 2.852(1)$ and $1.472(1)$ MHz for RM1 and RM2, approximately two standard deviations away from the results in Table II. The anticrossing frequency and the coupling strength of the resonator are estimated to be 5.032 GHz and 37.4 MHz by a full Hamiltonian fit (see Appendix A).

To test the performance of the estimation algorithm, we run it 1000 times on RM2 with slightly different initial distributions. The mean of each of the 1000 initial particle distributions is distributed uniformly at random within a 10-MHz range for f_{RM} and a 1.5-MHz range for g . Each individual particle distribution is uniform, with a width of 15 MHz in f_{RM} and 2.5 MHz in g .

We plot the convergence of the parameters in Fig. 5. As shown by the histograms, more than 99% of the runs converge successfully to properly estimate the frequency and coupling strength, with just a few failures. The average values after the 35th iteration are 4.8301(4) GHz for f_{RM} and 1.45(9) MHz for g .

The estimated true parameter values of the resonance, which we compute by fitting the dataset combining all 35 000 measurements, are 4.830 08 GHz for f_{RM} and 1.445 MHz for g . These values are shown with a red line in the histograms. We compare the results to a fit instead of the synthesized frequency to circumvent potential calibration inaccuracies. It is worth noting that the power for this experiment is slightly different from that of the experiment summarized in Table II.

Several experimental errors could cause the variation of a parameter in the likelihood model of Eq. (7). While the binomial likelihood accounts for variance in the qubit measurement, it assumes that all parameters are constant. Consequently, any parameter fluctuation or drift causes a discrepancy between the inference model and the physical system. For example, the qubit probe frequency f_p and energy relaxation time T_1 can be modified by strongly coupled TLSs [35,43–45]; f_p can additionally be perturbed by a nonconstant flux pulse amplitude (see Appendix B) or flux noise in the SQUID loop generated by, for example, voltage sources or on-chip magnetic defects. Depending on the nature of the resonance mode, f_{RM} and possibly g may vary due to, for example, changing temperature or strong coupling to TLSs. If the mode is another qubit, its frequency may be affected in the same manner as f_p . Finally, the measurement visibility itself could vary due to possible instrumentation issues.

The errors introduced above can be classified according to their timescale relative to the length of the online estimation experiment. A parameter that changes during the execution of the algorithm would lead to inconsistencies with subsequent measurements. This effect is expected to be taken into account by the variance of the repeated inferences, thereby increasing the calculated standard error. For example, if f_{RM} changes during the experiment, repeating the inference multiple times would likely lead to a bimodal distribution (i.e., a spread in the estimated value of f_{RM}).

If a parameter changes over a longer timescale, in our case greater than approximately 30 s, measurements taken during an experiment remain consistent and the standard error does not increase. If such an error is suspected,

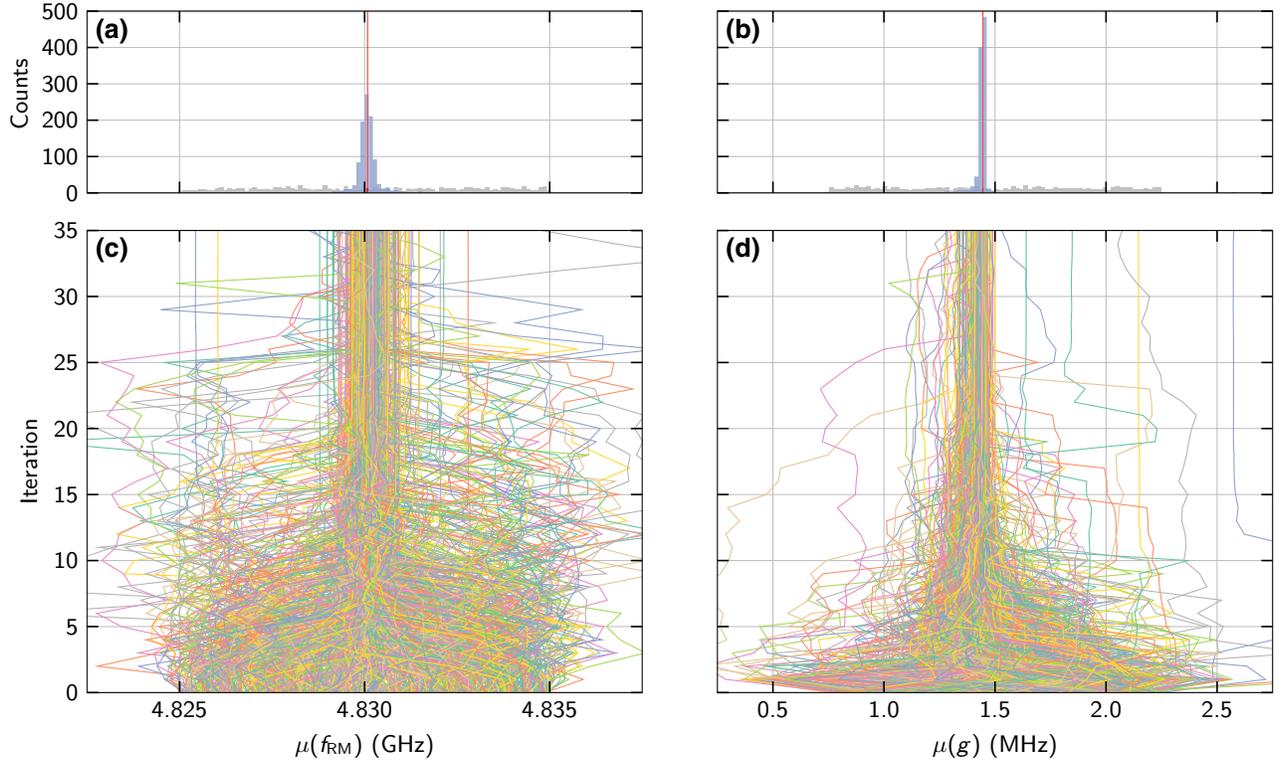


FIG. 5. Performance of the online estimation algorithm over 1000 runs with RM2. (a),(b) Histograms over the means of the initial and final particle distributions. (c),(d) Mean of the posterior particle distribution computed after each iteration’s measurement. As more iterations are made, the particle distribution converges toward the true value of the parameters. The zeroth iteration corresponds to the initial distribution. The true values of the parameters are identified by a red line in the histograms. The total runtime of the experiment is 6.4 h. Each individual run executes 35 iterations and thus 35 measurements, taking approximately 23 s. Most of the time (60%) is spent acquiring data. Each measurement comprises 786 shots at a repetition rate of 2 kHz. The remaining time is attributable to data transfer and processing. The total computation time for the estimation algorithm is approximately 1 s.

the calibration should be repeated at regular intervals, as mentioned in Sec. II.

VI. DISCUSSION

Both algorithms presented above depend on a few parameters that are crucial to their function. For the octave sampling algorithm, the choice of the frequency range to be measured is naturally determined by the properties of the device: superconducting qubits have a limited frequency range within which they operate optimally. For the device in this work, the upper end of the measurement range f_{\max} simply corresponds to the maximum attainable frequency. The lower limit f_{\min} is chosen to be as low as desired, keeping in mind that TLSs far below the operating frequency range of the qubit are not a cause for concern. In addition, given that the purpose of detecting interactions is to then *select* optimal operating frequencies, it might be sufficient to set f_{\min} to the lowest frequency where high-fidelity control and readout are achievable. Since we use a resonator for readout, the farther away the qubit is in

frequency, the lower the fidelity of the measurement. Other constraints (e.g., pulse control bandwidth) might dictate even tighter limits.

A second important parameter for the octave sampling method is n_s . In principle, a single high-quality (many shots) measurement of P_e at the center of the bin should be sufficient. This would be analogous to traditional swap spectroscopy. However, because of the efficiency of the octave method, we can afford to take a few more measurements per bin. This is what we chose to do by randomly distributing $n_s = 5$ measurements per bin. This redundancy increases the detection sensitivity and protects against possible statistical fluctuations in the measurement.

One more parameter worth discussing is the number of octaves to be measured m_f , which corresponds to t_{m_f} , or equivalently, Δf_{m_f} , as explained in Sec. IV A. Generally, this parameter should be determined by the requirements of the experiment for which the calibration is made. If a long gate sequence is needed (e.g., for randomized benchmarking), detecting weakly coupled resonance modes is important. This would not necessarily be the case for

shorter experiments, such as process tomography. A total time $t_{m_f} \sim t_{\text{exp}}$, where t_{exp} is the length of the experimental gate sequence, is therefore generally a good choice.

For the particle filter algorithm, the choice of a , c , and M_0 is discussed in Ref. [38]. Other parameters of interest include t_{max} and the number of particles to be used. The time t_{max} is used in Eq. (6) to restrict the maximum measurement time. This is necessary because the qubit eventually decays to the ground state. To obtain reliable results, t_{max} should be set well below T_1 . Another way to limit the maximum measurement time would be to replace $\tanh(a/\sigma(g)t_{\text{max}})t_{\text{max}}$ with $a/\sigma(g)$ in Eq. (6) (as in the original proposal) and simply stop the algorithm once a sufficiently small $\sigma(g)$ is reached.

The number of particles to be used is constrained mainly by the performance of the computer running the resampling procedure and, potentially, numerical accuracy issues [51]. As a rule of thumb, at least 10 000 particles should be used; in this work, we used 40 000.

In Sec. II, we explain that our method can be used not only for a simple qubit swap spectroscopy experiment but also to look for resonances with the $|e\rangle \leftrightarrow |f\rangle$ transition with a double-excitation protocol. The algorithms discussed in this work are very general and apply even to systems that do not involve a resonance mode. We can use the online and offline algorithms to efficiently detect the location and estimate the parameters of any qubit dynamics akin to a chevron pattern. This is the case, for example, with a whole class of parametric two-qubit gates, where instead of varying f_p , we vary the frequency of a flux drive applied to the SQUID of a qubit or tunable coupler [52–54].

Finally, we briefly discuss the problem of *choosing* qubit operating frequencies. Once the calibration showcased in this work is accomplished and all resonant couplings are identified, the next step is to use this information to optimize the performance of a quantum computer. This process depends on the quantum computing architecture. For an array of directly coupled superconducting qubits, we want to avoid crosstalk between neighboring qubits and minimize interactions with TLSs. We therefore need to choose the idle frequencies of all qubits at the same time, taking into account both wanted and unwanted couplings. Additional concerns apply for choosing the operating frequencies of two-qubit gates: we must consider the frequency *path* that the qubits will take during the gate. For example, it is undesirable for a qubit to cross through a resonance with a TLS. If such crossings are unavoidable (e.g., if the device is afflicted by many TLSs), knowing the coupling strength of each TLS helps select the optimal set of qubit interaction frequencies.

While this work does not explain the process needed to perform this optimization (see, e.g., Ref. [55] for an example), we emphasize that the runtime reduction of the offline and online algorithms when compared with traditional

swap spectroscopy presents several advantages: first, the calibration may be run more often; second, the calibration is affordable enough to be run on a larger spectrum, giving the frequency optimization process more information to work with.

VII. CONCLUSIONS

In conclusion, we explain two methods for the Hamiltonian parameter estimation of resonant couplings in the context of tunable superconducting qubits. The octave sampling technique can be run without prior knowledge of potential resonances in the environment of the qubit and allows efficient detection of coupled modes within a chosen parameter range. The online Bayesian algorithm can be performed or omitted depending on whether a more accurate estimate of the coupling parameters is desired. Using these algorithms reduces the number of measurements needed from $\mathcal{O}(1/\Delta f_{m_f}^2)$ to $\mathcal{O}(1/\Delta f_{m_f})$. This translates to a reduction in runtime by 1 order of magnitude in typical conditions.

We experimentally demonstrate both techniques on a superconducting Xmon transmon qubit and evaluate their performance. We are able to detect the resonance with the qubit’s measurement resonator, as well as with synthesized resonance modes and a naturally occurring weakly coupled TLS. We determine that the methods are efficient, reliable, and readily automated. We expect this type of calibration to be critical to the operation of large-scale quantum computers, superconducting or otherwise. Future work includes integrating the information we acquire by our methods into a comprehensive optimization process for selecting the operating frequency of each qubit in a quantum computer and implementing the calibration of a two-qubit gate with the online Bayesian algorithm.

ACKNOWLEDGMENTS

This research was undertaken thanks in part to funding from the Canada First Research Excellence Fund. We acknowledge the support of the Natural Sciences and Engineering Research Council of Canada, [Application Number: RGPIN-2019-04022]. We would like to acknowledge the Canadian Microelectronics Corporation Microsystems for the provision of products and services that facilitated this research, including CAD software. The authors thank the Quantum-Nano Fabrication and Characterization Facility at the University of Waterloo, where the sample was fabricated. The authors thank Christopher W. Warren for useful discussions.

APPENDIX A: SETUP AND SAMPLE

The sample is a chip comprising two independent Xmon transmon superconducting qubits that are fabricated with evaporated aluminum on a silicon wafer [32]. We use

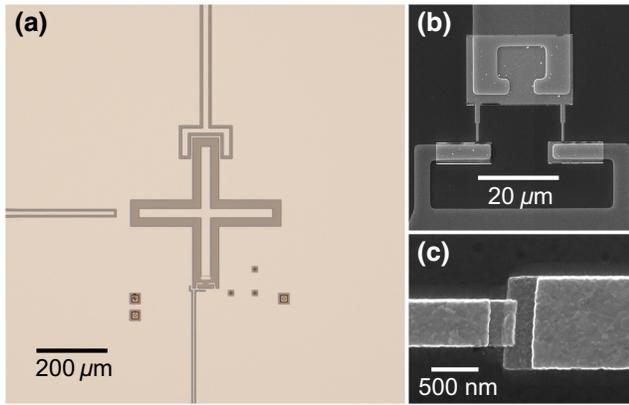


FIG. 6. Images of a sample identically fabricated to the one used in this work. (a) Optical image of the Xmon transmon qubit, with the drive line on the left, the measurement resonator coupler above, and the flux bias line below. The cross-shaped island constitutes a capacitor to ground, while the SQUID acts as a nonlinear inductor. (b) Scanning electron microscope image of the SQUID, which is located at the end of the bottom arm of the cross. (c) Scanning electron microscope image of a Josephson junction made with a Dolan bridge.

the qubit on the left side of the chip for the experiments shown in this work. Qubit rotations about an axis in the X - Y plane are controlled by a capacitive microwave drive, while flux pulses to set the frequency are done via a low-frequency line inductively coupled to the qubit's SQUID. Measurement is performed by means of a high-power readout scheme using a resonator capacitively coupled to the qubit. We read out the qubit state over 786 single-shot measurements to find P_e with visibility greater than approximately 90%. Images of an identically fabricated chip can be seen in Fig. 6. The package-to-chip connections are made with the quantum socket, a fully vertical packaging method [56].

The measurement schematics of the dilution refrigerator, including instrument and wiring details, can be found in Fig. 7. Compared with our previous setup in Ref. [56], carbon nanotube-based low-pass filters are placed on the microwave lines to prevent infrared noise [57]. In addition, coaxial lines are installed for flux biasing. Two lines are used to bias the qubit. One is meant for pure direct current coming from a battery, while the other is used for flux pulses made with an arbitrary waveform generator (AWG). These lines are joined with a custom-made Eccosorb-based bias tee, which acts as a low-pass filter for infrared shielding [58].

The parameters of the device used in the experiment are specified in Table III. Both T_1 and T_2 fluctuate over the range of accessible qubit frequencies due to the Purcell effect and flux noise sensitivity.

APPENDIX B: QUBIT FLUX PULSE

Performing swap spectroscopy requires the ability to set the qubit's frequency to a desired value f_p for a particular duration. In our experiment, this is done with a flux pulse applied to the SQUID with an AWG. As shown in Fig. 7, the flux pulse reaches the sample after going through multiple stages of filtering, attenuation, and connections. This means that the waveform is modified compared with what is generated by the AWG.

In addition, while we control the *amplitude* of the pulse, we are ultimately interested in the resulting *frequency* of the qubit. We therefore need a way to convert between the amplitude A of the flux pulse and the qubit probe frequency f_p . This can be done, for example, with pulse spectroscopy, where we send π -pulses to the qubit at different frequencies while it is detuned by a flux pulse of a particular amplitude. The qubit frequency for that amplitude can then be fit. This is repeated for many amplitudes to get a map between A and f_p .

The above considerations mean that the measurement settings (f_p, t) that we select may contain multiple kinds of potential errors. This must be taken into account when estimating the coupling parameters with either the offline algorithm or the online algorithm.

For example, if f_p is higher than the true probe frequency of the qubit due to some systematic error in the amplitude-frequency map, the result f_{RM} reported by the online algorithm is higher than the true value as well. Similarly, an error on the value of t leads to a wrong estimation of g . In practice, this kind of systematic error is not a major problem as long as the error is consistent between experiments. For our application, for example, it does not matter that the estimated frequency of the resonance is not truly 4.8305 GHz. The quantity of interest is the flux amplitude corresponding to the mode.

Another kind of error would occur if the flux pulse amplitude were not constant in time (e.g., due to filtering). This would result in a time-dependent probe frequency f_p over the length of the pulse. Assuming that the time dependence is not so strong that the qubit frequency changes by an amount of approximately g in less than a Rabi period, the consequences for the octave algorithm should be negligible. This is because the detection is made via the first minimum of the swap oscillation. The online estimation algorithm would be affected more strongly, since the model assumes that the qubit probe frequency is constant. The error on the parameters estimated by an experiment suffering from this issue, as computed by the repeated inference method discussed in Sec. V A, would be expected to capture this effect.

To increase accuracy, it is important to properly calibrate the flux pulses before the experiment, ensuring that the qubit frequency is stable. This calibration can be done, again, with pulse spectroscopy, by mapping out the qubit

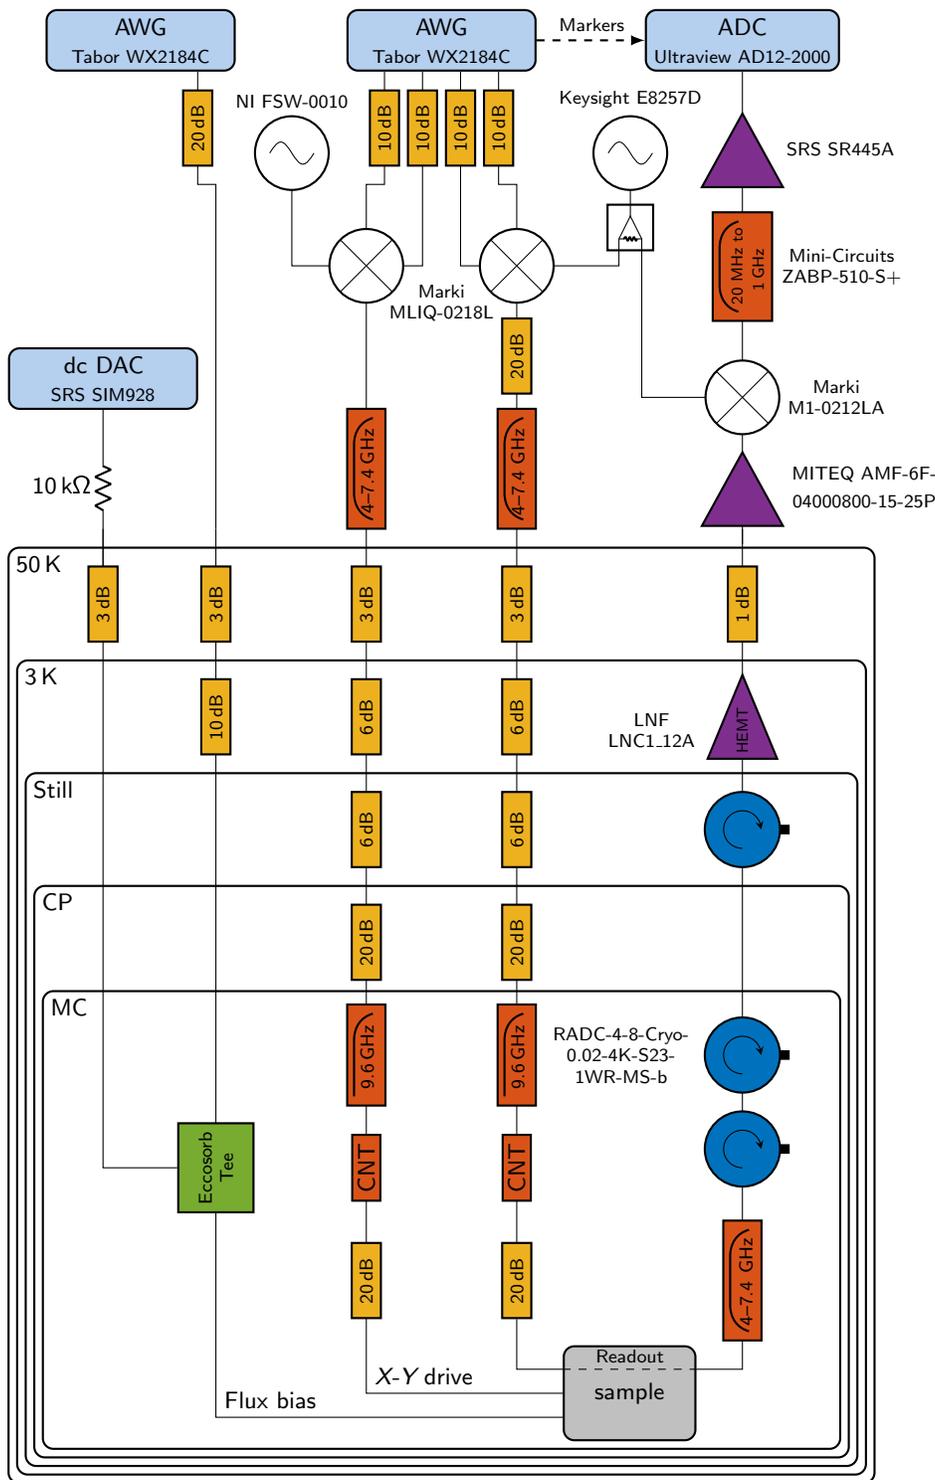


FIG. 7. Dilution refrigerator wiring schematics. The sample is mounted in an aluminum package at the mixing chamber stage of a dilution refrigerator, at a temperature of approximately 10 mK. We use a dedicated microwave line to drive the qubit. Two coaxial lines are joined by a custom Eccosorb tee at the mixing chamber for the flux biasing of the SQUID. We use a homodyne readout setup, with the in-phase quadrature mixers configured for image rejection. ADC, analog-to-digital converter; CNT, carbon nanotube; DAC, digital-to-analog converter; CP, cold plate; MC, mixing chamber.

frequency over the length of the pulse. Deviations can be subsequently corrected by modifying the shape of the pulse emitted by the AWG.

APPENDIX C: DETAILS ON OCTAVE ANALYSIS

The task of detecting resonances in data gathered with the octave sampling algorithm (the data plotted in Fig. 3)

is done in two substeps: (1) find peaks in each octave; (2) merge the octaves [59].

First, the average bin excitation is reversed to transform the low-excitation regions into peaks. Then, we find peaks in each octave individually. This is done by looking for points in the data that are *prominent* when compared with neighboring points. Each peak must be taller than the lowest point to the left and right by a certain amount,

TABLE III. Device parameters for the qubit used in the experiment. C_c is the coupling capacitance between the qubit and its readout resonator. The Hamiltonian parameters are determined via a fit of the resonator and qubit frequencies and the qubit anharmonicity. The relative error on each parameter is less than 1%. The resonator frequency shown is that of the “bare” resonator. If we include the effect of the coupling capacitance, the resonator frequency is 5.032 GHz.

Parameter	Value
E_J (GHz)	19.614(5)
E_C (MHz)	188.92(5)
f_{res} (GHz)	5.048 44(3)
C_c (fF)	3.371(5)
T_1 (μm)	10–25
T_2 (ns)	400–1000

the *prominence value*. Decreasing this value leads to more peaks being detected and therefore increases the detection sensitivity. The result of the peak-finding step for the highest octave of the data is shown in Fig. 8.

The second step is to check if the peaks detected in different octaves correspond to the same resonance. This is done by comparing the frequency position of the peaks found. If the frequencies are close, we assume that the peaks correspond to the same resonance and the detections are therefore merged. When detections from different octaves are being merged, the one from the lower octave is preferred. This is because the coupling strength is estimated from the first oscillation minimum, which corresponds to the lowest octave detected.

When the procedure is complete, we are left with a single peak location per resonance, along with the lowest octave where it is found. This gives a coarse estimate of the frequency and coupling strength of that resonance, as reported in Table I.

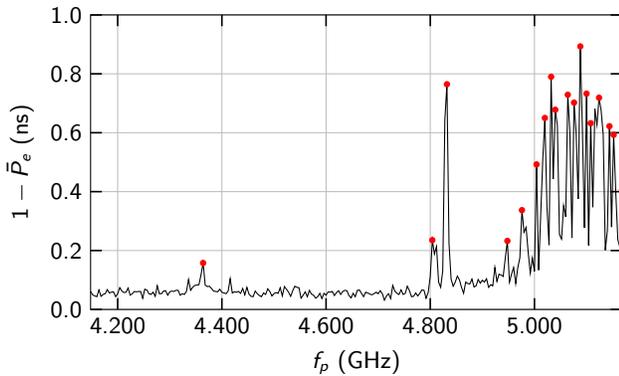


FIG. 8. Peaks found in the highest octave in Fig. 3(b) with the prominence value set to 0.09. Many peaks are found for RM3. They are merged when the detection analysis proceeds to lower octaves.

There are a few possible causes of errors with this procedure. The most likely error occurs when two detections corresponding to a single resonance are reported or, in contrast, when two resonances close in frequency are reported as one. The first issue is not very problematic, since further analysis done with the Bayesian online algorithm likely correctly reports that the two detections come from the same resonance. The second is solved by increasing the resolution of the octave data collection.

APPENDIX D: PARTICLE DISTRIBUTION GENERATION

While a prior particle distribution for the online Bayesian algorithm can be manually specified, it is much more convenient to generate one from available octave data. The first step in the process is to divide the octave data into multiple spectra such that we are left with one resonance mode per spectrum. To split the spectrum, we simply cut the data between each detected resonance (see Appendix C). If this cut happens to fall within a bin, we slice the bin into two sub-bins, keeping track of each part’s proportion.

Following this, we can restrict ourselves to a frequency-time spectrum containing a single resonance mode. To specify a discrete distribution representing our knowledge of the coupling parameters, we follow a procedure where we pick bins from the spectrum and generate a particle according to the bin parameters. We pick from bins according to a weight, which we choose to be proportional to $1 - \bar{P}_e$ but only for those above a particular threshold. We ignore bins below the threshold. Thus, bins for which the

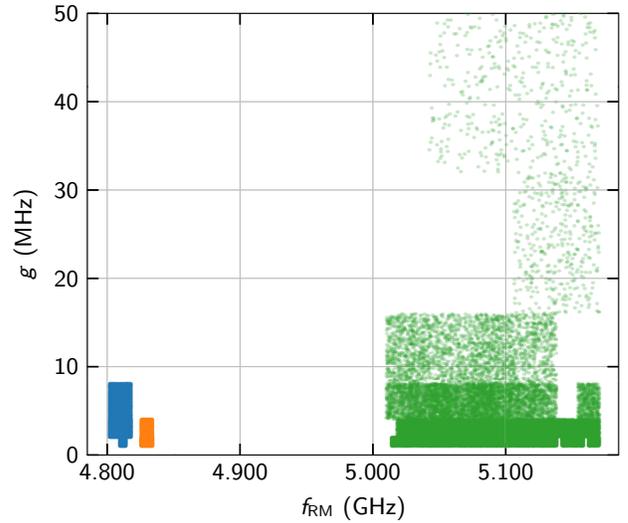


FIG. 9. Initial particle distribution generated from the octave data in Fig. 3(b) after splitting of the spectrum according to the detections. The distribution corresponding to RM1 is plotted in blue, the distribution corresponding to RM2 is plotted in orange, and the distribution corresponding to RM3 is plotted in green.

TABLE IV. Frequency and coupling parameter statistics of the particle distributions generated for RM1, RM2, and RM3. A plot of those distributions can be seen in Fig. 9.

Parameter	RM1	RM2	RM3
f_{RM} (GHz)	4.811(4)	4.830(2)	5.088(42)
g (MHz)	3(2)	2(1)	4(7)

average excitation \bar{P}_e is low have a high weight and have a higher chance of being chosen for generating particles.

Each particle represents a two-tuple of frequency–coupling–strength values. To generate a particle, we must choose those two values. The frequency of the generated particle is chosen uniformly at random within the frequency range of the bin drawn. The coupling strength is picked from the uniform distribution $\mathcal{U}(\Delta_m, 2\Delta_m)$, where Δ_m is the bin width. We repeat this procedure (draw a bin, generate a particle) to create as many particles as is desired for the distribution.

Exemplary particle distributions generated for RM1, RM2, and RM3 can be seen in Fig. 9, with their statistics tabulated in Table IV. These distributions can be used as the starting point for the online algorithm discussed in Sec. V. In our experiment we use distributions comprising 40 000 particles. This number can be adjusted depending on the capabilities of the computer executing the algorithm.

APPENDIX E: BAYESIAN RESAMPLING PROCEDURE

The exact resampling procedure is described in detail in Chap. 2 in Ref. [37]. We reproduce a shortened version here for completeness.

As inputs, we require the particle locations, as an array of frequency–coupling–strength tuples, in addition to their likelihoods, as computed in the main text. The likelihoods must be normalized to sum to unity, after which they become *weights*. This normalization ensures that the weight array is a valid discrete probability distribution that can be drawn from.

The algorithm then draws particles from the input distribution according to the weights and, from the position of the particles drawn, generates new ones by adding “noise.” This prevents there being duplicates in the output distribution, even if a particle from the input is drawn multiple times. The resampling algorithm is shown in pseudocode in Algorithm 1.

The particle at index i in the $\{\vec{x}_k\}$ array has its corresponding weight at the same index in the $\{w_k\}$ array. In addition, the amount of noise added to the position of the particle drawn is controlled by the spread of the input distribution, quantified by taking the covariance. For more details, see Ref. [37].

Input: Array of particle positions $\{\vec{x}_k\}$
Input: Array of particle weights $\{w_k\}$
Output: New particle positions $\{\vec{y}_k\}$
Function *resample*($\{\vec{x}_k\}, \{w_k\}$)

```

    a = 0.98
     $\vec{\mu} = \text{mean}(\{\vec{x}_k\})$ 
     $\Sigma = (1 - a^2)\text{cov}(\{\vec{x}_k\})$ 
    for i ∈ 1 : n do
        l = rand(Discrete( $\{w_k\}$ ))
         $\vec{\mu}_l = a\vec{x}_l + (1 - a)\vec{\mu}$ 
         $\vec{y}_i = \text{rand}(\text{Normal}(\vec{\mu}_l, \Sigma))$ 
    end
    return  $\{\vec{y}_k\}$ 

```

Algorithm 1. Particle resampling.

-
- [1] T. D. Ladd, F. Jelezko, R. Laflamme, Y. Nakamura, C. Monroe, and J. L. O’Brien, Quantum computers, *Nature* **464**, 45 (2010).
 - [2] D. Gottesman, in *Quantum Information Science and Its Contributions to Mathematics* (American Mathematical Society, 2010), p. 13.
 - [3] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, Surface codes: Towards practical large-scale quantum computation, *Phys. Rev. A* **86**, 032324 (2012).
 - [4] K. Kechedzhi, V. Smelyanskiy, J. R. McClean, V. S. Denchev, M. Mohseni, S. Isakov, S. Boixo, B. Altshuler, and H. Neven, in *Leibniz International Proceedings in Informatics (LIPIcs)*, Vol. 111, edited by S. Jeffery (Schloss Dagstuhl, Dagstuhl, Germany, 2018), p. 9:1.
 - [5] R. D. Somma, S. Boixo, H. Barnum, and E. Knill, Quantum Simulations of Classical Annealing Processes, *Phys. Rev. Lett.* **101**, 130504 (2008).
 - [6] E. Farhi and H. Neven, Classification with Quantum Neural Networks on Near Term Processors, *ArXiv:1802.06002* (2018).
 - [7] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, Barren plateaus in quantum neural network training landscapes, *Nat. Commun.* **9**, 4812 (2018).
 - [8] I. Cong, S. Choi, and M. D. Lukin, Quantum convolutional neural networks, *Nat. Phys.* **15**, 1273 (2019).
 - [9] S. Bravyi, D. Gosset, and R. König, Quantum advantage with shallow circuits, *Science* **362**, 308 (2018).
 - [10] A. Aspuru-Guzik, Simulated quantum computation of molecular energies, *Science* **309**, 1704 (2005).
 - [11] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O’Brien, A variational eigenvalue solver on a photonic quantum processor, *Nat. Commun.* **5**, 4213 (2014).
 - [12] C. Hempel, C. Maier, J. Romero, J. McClean, T. Monz, H. Shen, P. Jurcevic, B. P. Lanyon, P. Love, R. Babbush, A. Aspuru-Guzik, R. Blatt, and C. F. Roos, Quantum Chemistry Calculations on a Trapped-Ion Quantum Simulator, *Phys. Rev. X* **8**, 031022 (2018).
 - [13] M. Mohseni, P. Read, H. Neven, S. Boixo, V. S. Denchev, R. Babbush, A. Fowler, V. Smelyanskiy, and J. M. Martinis,

- Commercialize quantum technologies in five years, *Nature* **543**, 171 (2017).
- [14] J. Clarke and F. K. Wilhelm, Superconducting quantum bits, *Nature* **453**, 1031 (2008).
- [15] G. Wendin, Quantum information processing with superconducting circuits: A review, *Reports on Progress in Physics* **80**, 106001 (2017).
- [16] A. F. Kockum and F. Nori, in *Fundamentals and Frontiers of the Josephson Effect* (Springer, 2019), p. 703.
- [17] M. Mariani, The Energy of an Arbitrary Electrical Circuit, Classical and Quantum, *ArXiv:2007.08519* (2020).
- [18] M. Mariani, H. Wang, T. Yamamoto, M. Neeley, R. C. Bialczak, Y. Chen, M. Lenander, E. Lucero, A. D. O’Connell, D. Sank, M. Weides, J. Wenner, Y. Yin, J. Zhao, A. N. Korotkov, A. N. Cleland, and J. M. Martinis, Implementing the quantum von Neumann architecture with superconducting circuits, *Science* **334**, 61 (2011).
- [19] E. Lucero, R. Barends, Y. Chen, J. S. Kelly, M. Mariani, A. Megrant, P. O’Malley, D. Sank, A. Vainsencher, J. Wenner, T. White, Y. Yin, A. N. Cleland, and J. M. Martinis, Computing prime factors with a Josephson phase qubit quantum processor, *Nat. Phys.* **8**, 719 (2012).
- [20] A. D. Córcoles, E. Magesan, S. J. Srinivasan, A. W. Cross, M. Steffen, J. M. Gambetta, and J. M. Chow, Demonstration of a quantum error detection code using a square lattice of four superconducting qubits, *Nat. Commun.* **6**, 6979 (2015).
- [21] D. Risté, S. Poletto, M.-Z. Huang, A. Bruno, V. Vesterinen, O.-P. Saira, and L. DiCarlo, Detecting bit-flip errors in a logical qubit using stabilizer measurements, *Nat. Commun.* **6**, 6983 (2015).
- [22] R. Barends, *et al.*, Superconducting quantum circuits at the surface code threshold for fault tolerance, *Nature* **508**, 500 (2014).
- [23] J. S. Kelly, *et al.*, State preservation by repetitive error detection in a superconducting quantum circuit, *Nature* **519**, 66 (2015).
- [24] R. Barends, J. S. Kelly, A. Megrant, D. Sank, E. Jeffrey, Y. Chen, Y. Yin, B. Chiaro, J. Mutus, C. Neill, P. O’Malley, P. Roushan, J. Wenner, T. C. White, A. N. Cleland, and J. M. Martinis, Coherent Josephson Qubit Suitable for Scalable Quantum Integrated Circuits, *Phys. Rev. Lett.* **111**, 080502 (2013).
- [25] F. Arute, *et al.*, Quantum supremacy using a programmable superconducting processor, *Nature* **574**, 505 (2019).
- [26] J. Preskill, Quantum computing in the NISQ era and beyond, *Quantum* **2**, 79 (2018).
- [27] L. DiCarlo, J. M. Chow, J. M. Gambetta, L. S. Bishop, B. R. Johnson, D. I. Schuster, J. Majer, A. Blais, L. Frunzio, S. M. Girvin, and R. J. Schoelkopf, Demonstration of two-qubit algorithms with a superconducting quantum processor, *Nature* **460**, 240 (2009).
- [28] T. Yamamoto, M. Neeley, E. Lucero, R. C. Bialczak, J. S. Kelly, M. Lenander, M. Mariani, A. D. O’Connell, D. Sank, H. Wang, M. Weides, J. Wenner, Y. Yin, A. N. Cleland, and J. M. Martinis, Quantum process tomography of two-qubit controlled-Z and controlled-NOT gates using superconducting phase qubits, *Phys. Rev. B* **82**, 184515 (2010).
- [29] A. O. Niskanen, K. Harrabi, F. Yoshihara, Y. Nakamura, S. Lloyd, and J. S. Tsai, Quantum coherent tunable coupling of superconducting qubits, *Science* **316**, 723 (2007).
- [30] M. S. Allman, F. Altomare, J. D. Whittaker, K. Cicak, D. Li, A. Sirois, J. Strong, J. D. Teufel, and R. W. Simmonds, RF-SQUID-Mediated Coherent Tunable Coupling between a Superconducting Phase Qubit and a Lumped-Element Resonator, *Phys. Rev. Lett.* **104**, 177004 (2010).
- [31] C. Müller, J. H. Cole, and J. Lisenfeld, Towards understanding two-level-systems in amorphous solids: Insights from quantum circuits, *Reports on Progress in Physics* **82**, 124501 (2019).
- [32] C. T. Earnest, J. H. Béjanin, T. G. McConkey, E. A. Peters, A. Korinek, H. Yuan, and M. Mariani, Substrate surface engineering for high-quality silicon/aluminum superconducting resonators, *Supercond. Sci. Technol.* **31**, 125013 (2018).
- [33] M. Moeed, C. T. Earnest, J. Béjanin, A. Sharafeldin, and M. Mariani, Improving the time stability of superconducting planar resonators, *MRS Adv.* **4**, 2201 (2019).
- [34] T. G. McConkey, J. H. Béjanin, C. T. Earnest, C. R. H. McRae, Z. Pagel, J. R. Rinehart, and M. Mariani, Mitigating leakage errors due to cavity modes in a superconducting quantum computer, *Quantum Sci. Technol.* **3**, 034004 (2018).
- [35] P. V. Klimov, *et al.*, Fluctuations of Energy-Relaxation Times in Superconducting Qubits, *Phys. Rev. Lett.* **121**, 90502 (2018).
- [36] M. Mariani, H. Wang, R. C. Bialczak, M. Lenander, E. Lucero, M. Neeley, A. D. O’Connell, D. Sank, M. Weides, J. Wenner, T. Yamamoto, Y. Yin, J. Zhao, J. M. Martinis, and A. N. Cleland, Photon shell game in three-resonator circuit quantum electrodynamics, *Nat. Phys.* **7**, 287 (2011).
- [37] Y. R. Sanders, *Characterizing Errors in Quantum Information Processors*, Ph.D. thesis, school University of Waterloo (2016).
- [38] M. P. Stenberg, Y. R. Sanders, and F. K. Wilhelm, Efficient Estimation of Resonant Coupling between Quantum Systems, *Phys. Rev. Lett.* **113**, 210404 (2014).
- [39] C. E. Granade, C. Ferrie, N. Wiebe, and D. G. Cory, Robust online Hamiltonian learning, *New J. Phys.* **14**, 103013 (2012).
- [40] J. Koch, T. Yu, J. M. Gambetta, A. A. Houck, D. I. Schuster, J. Majer, A. Blais, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, Charge-insensitive qubit design derived from the Cooper pair box, *Phys. Rev. A* **76**, 042319 (2007).
- [41] G. Haack, F. Helmer, M. Mariani, F. Marquardt, and E. Solano, Resonant quantum gates in circuit quantum electrodynamics, *Phys. Rev. B* **82**, 024514 (2010).
- [42] P. Krantz, M. Kjaergaard, F. Yan, T. P. Orlando, S. Gustavsson, and W. D. Oliver, A quantum engineer’s guide to superconducting qubits, *Appl. Phys. Rev.* **6**, 021318 (2019).
- [43] J. H. Béjanin, C. T. Earnest, A. S. Sharafeldin, and M. Mariani, Interacting defects generate stochastic fluctuations in superconducting qubits, *Phys. Rev. B* **104**, 094106 (2021).
- [44] S. Schlör, J. Lisenfeld, C. Müller, A. Bilmes, A. Schneider, D. P. Pappas, A. V. Ustinov, and M. Weides, Correlating Decoherence in Transmon Qubits: Low Frequency

- Noise by Single Fluctuators, *Phys. Rev. Lett.* **123**, 190502 (2019).
- [45] J. J. Burnett, A. Bengtsson, M. Scigliuzzo, D. Niepce, M. Kudra, P. Delsing, and J. Bylander, Decoherence benchmarking of superconducting qubits, *npj Quantum Inf.* **5**, 54 (2019).
- [46] Although the concept of an octave number m does not apply to the traditional method, we use the values of Δf_{m_f} and t_{m_f} for comparison.
- [47] A time resolution $\Delta t_{\min} = 2.5$ ns corresponds to the 400-MHz sampling rate used for sampling a signal up to 200 MHz.
- [48] The source code developed for this work can be found online at <https://gitlab.com/DQMLab/TLSInfer.jl>.
- [49] In the experiment, we use a model for \tilde{P}_e that accounts for relaxation and measurement visibility. For relaxation, we use Eq. (17) in Ref. [38]. To account for measurement visibility, we clamp the theoretical probability between 0.05 and 0.95.
- [50] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PRXQuantum.2.040343> for an animation of this iterative process.
- [51] L. M. Murray, A. Lee, and P. E. Jacob, Parallel resampling in the particle filter, *J. Comput. Graph. Stat.* **25**, 789 (2016).
- [52] P. Bertet, C. J. P. M. Harmans, and J. E. Mooij, Parametric coupling for superconducting qubits, *Phys. Rev. B* **73**, 064512 (2006).
- [53] D. C. McKay, S. Filipp, A. Mezzacapo, E. Magesan, J. M. Chow, and J. M. Gambetta, Universal Gate for Fixed-Frequency Qubits via a Tunable Bus, *Phys. Rev. Appl.* **6**, 064007 (2016).
- [54] S. A. Caldwell, *et al.*, Parametrically Activated Entangling Gates Using Transmon Qubits, *Phys. Rev. Appl.* **10**, 034050 (2018).
- [55] P. V. Klimov, J. S. Kelly, J. M. Martinis, and H. Neven, The Snake Optimizer for Learning Quantum Processor Control Parameters, *ArXiv:2006.04594* (2020).
- [56] J. H. Béjanin, T. G. McConkey, J. R. Rinehart, C. T. Earnest, C. R. H. McRae, D. Shiri, J. D. Bateman, Y. Rohanizadegan, B. Penava, P. Breul, S. Royak, M. Zapatka, A. G. Fowler, and M. Mariantoni, Three-Dimensional Wiring for Extensible Quantum Computing: The Quantum Socket, *Phys. Rev. Appl.* **6**, 044010 (2016).
- [57] M. V. Moghaddam, C. W. S. Chang, I. Nsanzineza, A. M. Vadiraj, and C. M. Wilson, Carbon nanotube-based lossy transmission line filter for superconducting qubit measurements, *Appl. Phys. Lett.* **115**, 213504 (2019).
- [58] D. F. Santavicca and D. E. Prober, Impedance-matched low-pass stripline filters, *Meas. Sci. Technol.* **19**, 087001 (2008).
- [59] The source code implementing this analysis as well as the particle generation procedure can be found online at <https://gitlab.com/DQMLab/TLSInfer.jl> in the “Octave-Analysis” module.