

Real-Time Adaptive Tracking of Fluctuating Relaxation Rates in Superconducting Qubits

Fabrizio Berritta^{1,2,*}, Jacob Benestad³, Jan A. Krzywda⁴, Oswin Krause⁵, Malthe A. Marciniak^{1,2}, Svend Krøjer^{1,2}, Christopher W. Warren^{6,1,2}, Emil Hogedal⁶, Andreas Nylander⁶, Irshad Ahmad⁶, Amr Osman⁶, Janka Biznárová⁶, Marcus Rommel⁶, Anita Fadavi Roudsari⁶, Jonas Bylander⁶, Giovanna Tancredi⁶, Jeroen Danon^{6,3}, Jacob Hastrup^{1,2}, Ferdinand Kuemmeth^{1,7,8} and Morten Kjaergaard^{1,2,†}

¹*Center for Quantum Devices, Niels Bohr Institute, University of Copenhagen, 2100 Copenhagen, Denmark*

²*NNF Quantum Computing Programme, Niels Bohr Institute, University of Copenhagen, 2100 Copenhagen, Denmark*

³*Center for Quantum Spintronics, Department of Physics, Norwegian University of Science and Technology, NO-7491 Trondheim, Norway*

⁴*Lorentz Institute for Theoretical Physics and Leiden Institute of Advanced Computer Science, Universiteit Leiden, 2311 EZ Leiden, The Netherlands*

⁵*Department of Computer Science, University of Copenhagen, 2100 Copenhagen, Denmark*

⁶*Department of Microtechnology and Nanoscience, Chalmers University of Technology, SE-412 96 Gothenburg, Sweden*

⁷*Institute of Experimental and Applied Physics, University of Regensburg, 93040 Regensburg, Germany*

⁸*QDevil, Quantum Machines, 2750 Ballerup, Denmark*



(Received 19 June 2025; accepted 5 January 2026; published 13 February 2026)

The fidelity of operations on a solid-state quantum processor is fundamentally bounded by environmental decoherence. Characterizing environmental fluctuations is challenging because the acquisition time of nonadaptive experimental protocols limits temporal precision and can average out rapid features of the underlying dynamics. Here, we overcome this temporal-resolution limit by 2 orders of magnitude using a field-programmable gate-array powered classical controller that adaptively and continuously tracks the relaxation-time fluctuations of two fixed-frequency superconducting transmon qubits, which exhibit average relaxation times of approximately 0.17 ms and occasionally exceed 0.5 ms. We report events in which the relaxation time switches by nearly an order of magnitude over timescales of just tens of milliseconds, rather than minutes or hours as previously reported. Our real-time Bayesian estimation protocol estimates relaxation times within a few milliseconds, close to the decoherence timescale itself. Our statistical analysis further suggests that some of these fast fluctuations arise from two-level systems switching at rates up to 10 Hz, 4 orders of magnitude faster than earlier reports. These results redefine the timescales relevant for calibration in superconducting quantum processing units, establish a reference for rapid relaxation-rate characterization in device screening, and improve our understanding of fast relaxation dynamics.

DOI: [10.1103/gk1b-stl3](https://doi.org/10.1103/gk1b-stl3)

Subject Areas: Condensed Matter Physics,
Mesoscopics,
Quantum Physics

I. INTRODUCTION

Superconducting qubits [1–5] are among the main candidates for fault-tolerant quantum computation schemes, with quantum operations on these devices approaching error rates capable of demonstrating quantum error correction [6–9]. However, as the number of physical qubits increases in quantum processing units (QPUs), the QPU's performance is bounded by the lowest-performing outlier qubits [10–12]. Identifying such outliers can be nontrivial, as time-dependent fluctuations in physical qubit parameters may alter which qubits qualify as outliers at any given moment. This is complicated by the physical mechanisms for these variations

*Contact author: fabrizio.berritta@mit.edu

†Contact author: mkjaergaard@nbi.ku.dk

‡Present address: Research Laboratory of Electronics, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA.

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.

drifting over several different and competing timescales. In particular, the relaxation rate Γ_1 of a superconducting qubit directly limits the fidelity of quantum operations [13]. Γ_1 fluctuates unpredictably in the time domain [14–22] and also as a function of the qubit frequency [14–18,22,23]. One of the major contributions to energy relaxation in state-of-the-art superconducting qubits is frequently attributed to their (semi-)resonant interaction with environmental two-level system (TLS) defects [14,15,24], though their detailed microscopic origins remain unknown. It has been shown that TLS frequencies drift over repeated cooldowns, whereas the overall number of TLSs does not [20,25].

Previous works in transmon qubits have resolved Γ_1 with a time resolution of seconds [18,26,27] or minutes [28,29]. More efficient and scalable estimation methods are needed to (i) probe previously unexplored subsecond regimes of Γ_1 dynamics, which is necessary for understanding the underlying physics, and (ii) identify outlier qubits and time-dependent fluctuations [14–16,24] in large QPUs to ensure fast and reliable characterization and error mitigation. Modern field-programmable gate array (FPGA) advancements have facilitated *online* (during experimental data collection) Hamiltonian learning [30,31], which is a useful tool to probe drifts in qubit parameters through real-time estimation [31–40].

In this work, we investigate fast fluctuations of the relaxation rates of two long-lived transmon qubits (with $T_1 \equiv 1/\Gamma_1 \approx 0.17$ ms measured over several hours) on millisecond timescales, almost comparable to the relaxation times themselves. We use a commercial controller with an integrated FPGA that leverages single-shot readout and performs real-time Bayesian estimation of the relaxation time, overcoming the sampling limitations of the traditional method mentioned below by two orders of magnitude. Our adaptive estimation allows the controller to investigate the stochastic behavior of the relaxation rate on unprecedentedly short timescales, revealing events where the relaxation time switches by almost an order of magnitude on the timescale of tens of milliseconds, instead of minutes or hours [14–16]. We analyze the fluctuations of Γ_1 using the power spectral density and Allan deviation [41] and find that a simple Lorentzian model describes the observed fluctuations, allowing us to extract TLS switching rates as fast as 10 Hz, $10^4 \times$ faster than previously reported, and to monitor changes in the dominating TLS environment with subminute time resolution.

To appreciate the advantages of our approach, it is useful to recall the traditional protocol. The most common method for estimating Γ_1 consists of initializing the qubit to the excited state and measuring its state projectively after a fixed (*nonadaptive*) waiting time. For each waiting time τ_{wait} , the measurement outcome is averaged over many repetitions, and this is repeated for all waiting times. The fraction of times the system is measured in the excited state as a function of waiting time is then fitted to an exponential

decay $\propto e^{-\Gamma_1 \tau_{\text{wait}}}$. The drawbacks of such a nonadaptive method are that (i) it is not optimally efficient in terms of experiment time, (ii) prior estimates of the qubit’s Γ_1 are not used to track its subsequent temporal fluctuations, and (iii) implementing curve fitting directly on an FPGA is challenging due to the available numerical precision, which prevents interleaving such estimations with qubit operations in real time.

Bayesian parameter estimation [31,42] instead is a natural approach to implementing real-time optimization techniques [31,34,35,38–40] compatible with low-latency control hardware. By using an onboard probability distribution of the parameter estimate, the optimal experimental settings for each subsequent probing cycle may be chosen *adaptively* so as to maximize the knowledge obtained from every (single-shot) measurement. In this work the technique is validated with shot-by-shot interleaved estimations of the relaxation time using our Bayesian method and the traditional approach.

Our scheme sets a new reference for fast QPU characterization of relaxation rates, for probing relaxation dynamics at previously unexplored time resolution, and is well suited for stable interleaved QPU execution in the presence of Γ_1 fluctuations.

II. SETUP AND PROTOCOL

We use a 5-qubit superconducting array with tunable couplers operated at the mixing-chamber stage (below 10 mK) of a dilution refrigerator. The device design and fabrication is similar to Ref. [43]. We implement the decay-rate estimation protocol on two of the five available qubits. In the main text, we focus on a transmon qubit (Q_1) of frequency ≈ 4.13 GHz and anharmonicity ≈ -210 MHz, as shown in Fig. 1(a). The results from another qubit (Q_2) in the same device are presented in Supplemental Material [44]. In our qubits, the expected Purcell-limited relaxation time is ≈ 0.8 ms. A commercial controller (Quantum Machines OPX1000) applies microwave pulses for qubit control and single-shot readout, using dedicated XY control lines and readout resonators, respectively [see Fig. 1(a)] [1–5].

The fluctuating parameter Γ_1 is estimated on the FPGA from the probing sequence shown in Fig. 1(b). For each probe cycle (labeled $1 \leq i \leq N$), the qubit is initialized to the ground state $|0\rangle$ via active reset [47] and is then brought to the excited state $|1\rangle$ using an X_π pulse. After waiting for an adaptive waiting time τ_i , the controller assigns the qubit state as ground ($m_i = 0$) or excited ($m_i = 1$) by thresholding the demodulated dispersive readout signal. The Bayesian probability distribution $\mathcal{P}_i(\Gamma_1)$, detailed in Sec. III, is updated based on m_i . The controller then uses the updated estimate of $\hat{T}_1 \equiv 1/\langle \Gamma_1 \rangle$ based on $\mathcal{P}_i(\Gamma_1)$ and selects a new waiting time $\tau_{i+1} = c\hat{T}_1$, with c fixed, for the subsequent probing cycle.

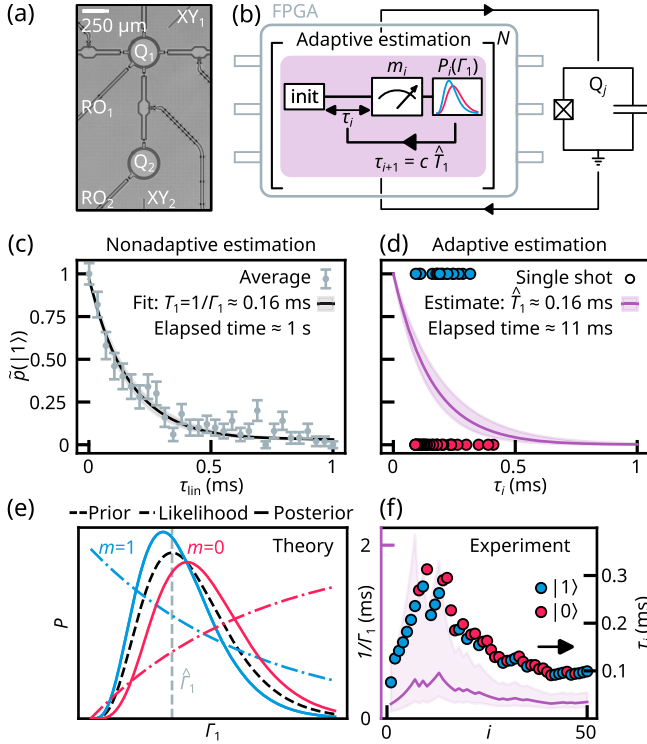


FIG. 1. Device and Bayesian adaptive decay-rate estimation. (a) Optical micrograph of transmon qubits (Q_1 and Q_2) nominally identical to the ones used in this work. Each transmon (Q_j) is individually controlled by microwave pulses (XY_j) and read out through independent resonators (RO_j). (b) Experimental scheme for adaptively estimating the qubit decay rate Γ_1 (purple box) on the FPGA in real time over N probe cycles. In each cycle, labeled i , the controller initializes Q_j to the excited state (init), waits a time τ_i adaptively chosen based on $\hat{T}_1 \equiv 1/\langle\Gamma_1\rangle$ from the previous Bayesian distribution $\mathcal{P}_{i-1}(\Gamma_1)$, then updates the probability distribution $\mathcal{P}_i(\Gamma_1)$ based on the measurement outcome m_i . (c) Example of a nonadaptive estimation of the relaxation time. From the normalized fraction $\tilde{p}(|1\rangle)$ of excited states as a function of linearly stepped probing waiting times τ_{lin} , Γ_1 is estimated by an exponential fit. The total elapsed time is ≈ 1 s. (d) Example of an adaptive estimation of \hat{T}_1 by Bayesian statistics implemented on the controller. Each circle is a single-shot measurement outcome $|0\rangle$ (red) or $|1\rangle$ (blue) which updates the current estimate of \hat{T}_1 and the subsequent adaptive waiting time τ_i . The total elapsed time for the entire estimation is ≈ 11 ms. (e) Evolution of the probability distribution \mathcal{P} during each probe cycle i of the estimation algorithm. The current estimate is \hat{T}_1 (gray dashed line). On each probing cycle, the estimate of \hat{T}_1 is updated according to two possible likelihood functions (dot-dashed lines), multiplied by the prior distribution (dashed). This yields the posterior distribution (solid), whose estimate is shifted left or right depending on the measurement outcome, while the uncertainty is reduced on average. (f) Convergence of \mathcal{P} as a function of the i th probe cycle. The line shows the resulting estimate \hat{T}_1 and the shaded area marks the 90% credible interval. As in panel (d), each circle is a single-shot outcome m_i with corresponding waiting time τ_i .

To contrast our estimation method with the standard (nonadaptive) approach, we first present in Fig. 1(c) a “standard T_1 experiment.” For clarity, the measurement outcomes of Figs. 1(c)–1(d) ignore readout errors. Specifically, in Fig. 1(c) we plot the normalized measured fractions of excited states, $\tilde{p}(|1\rangle)$ as a function of the linear waiting time τ_{lin} , along with the corresponding exponential fit, yielding a decay constant of $T_1 = 1/\Gamma_1 = (165 \pm 15) \mu\text{s}$. This experiment is based on 1890 single-shot measurements and spans approximately one second, comparable to the method used in Ref. [26].

In contrast, in Fig. 1(d), we present a representative run of our adaptive approach. The controller uses 50 single-shot measurement outcomes (experiment performed the day after the nonadaptive one), each of which is used to update the probability distribution on the controller iteratively. The updated distribution is then used to calculate an adaptive waiting time for the subsequent probing shot. After 50 single-shot measurements, the estimated value is $\hat{T}_1 \approx 159 \mu\text{s}$, with a 68% credible interval (CI) of $[131, 202] \mu\text{s}$, and a total elapsed time of only 11 ms ($\approx 69\hat{T}_1$). The purple curve illustrates the corresponding exponential decay, the shaded area indicating the CI.

We see that the resulting uncertainty from the adaptive method is slightly larger than the nonadaptive case, while the estimator \hat{T}_1 was obtained with a total estimation time that is two orders of magnitude shorter. Moreover, in Fig. 1(c), the chosen values of waiting times within the interval $[1, 1000] \mu\text{s}$ are appropriate for the estimated relaxation time in this specific example. However, if T_1 were to fluctuate drastically, a nonadaptive estimation on the same grid would result in a much greater uncertainty. The adaptive scheme is much more robust to these fluctuations as the waiting times are selected dynamically (for further details on how the adaptive method outperforms the nonadaptive one, see Supplemental Material [44]).

III. BAYESIAN ESTIMATION

We now describe our efficient and adaptive Bayesian estimation method for the decay rate implemented on the controller. The crux of the protocol is that all information about the current estimate of the probability distribution, $\mathcal{P}(\Gamma_1)$, is stored at any time, dynamically and with only a few parameters, on the controller. This on-controller parametrization allows each Bayesian update of $\mathcal{P}(\Gamma_1)$ to take only $\approx 2.2 \mu\text{s}$ (cf. the update time of $\approx 50 \mu\text{s}$ in Ref. [34], which used particle filtering for estimating decoherence rates in a nitrogen-vacancy center). Section IV presents our main results on Γ_1 estimation, which remain accessible without reference to the following implementation details.

An overview of our Bayesian estimation method was described in the previous section, however, we reiterate to be explicit as to what is executed at every step of the

protocol and its approximations. The estimation method involves a probing cycle where the qubit is first initialized in the excited state $|1\rangle$. This initialization is followed by a waiting time τ after which the state of the qubit is measured using dispersive readout. Since the fluctuations of $\Gamma_1(t)$ tend to be dominated by low-frequency noise [14,15], we approximate $\Gamma_1(t)$ to be quasistatic on the scale of tens of probing cycles (i.e., a few ms). In the following, we thus drop the time dependence of $\Gamma_1(t)$ for ease of notation.

We assume that after the initialization the probability of measuring an outcome $m \in \{0, 1\}$ corresponding to the states $|0\rangle$ and $|1\rangle$ is given by the likelihood function [dotted lines in Fig. 1(e)]

$$P(m|\Gamma_1, \tau) = 1 - m - (-1)^m [\beta + (1 - \alpha - \beta)e^{-\Gamma_1\tau}], \quad (1)$$

where Γ_1 is the parameter we want the controller to estimate, and α and β are misclassification probabilities for measuring $|0\rangle$ when at the beginning of the measurement the true state is $|1\rangle$ and measuring $|1\rangle$ when the true state is $|0\rangle$, respectively [48].

In the quasistatic approximation, by Bayes' rule, one obtains

$$\mathcal{P}_{i+1}(\Gamma_1) \propto \mathcal{P}_i(\Gamma_1)P(m_{i+1}|\Gamma_1, \tau_{i+1}), \quad (2)$$

where the prior $\mathcal{P}_i(\Gamma_1)$ describes the probability distribution for Γ_1 after the i th probing cycle, which depends on all the previously used waiting times and measurement outcomes, and the posterior $\mathcal{P}_{i+1}(\Gamma_1)$ describes the distribution after the subsequent cycle.

A simple parametrization of probability distributions is generally favorable in Bayesian approaches for on-FPGA optimization and computational speed, as demonstrated in Ref. [40]. In typical experiments, the distribution of estimated Γ_1 's fits well to a Gaussian [29], which suggests approximating the prior and posterior in Eq. (2) by a Gaussian for all i [40,50]. However, combining a Gaussian prior with the exponential likelihood function in Eq. (1) results in a posterior that is not well approximated by a Gaussian distribution, especially when the spread of the Gaussian becomes comparable to its mean, yielding also the unphysical ingredient of significant probability for negative values of Γ_1 .

Instead, we find that the gamma distribution is a convenient choice,

$$\mathcal{P}_i(\Gamma_1|k_i, \theta_i) = \frac{\theta_i^{k_i}}{\Gamma(k_i)} \Gamma_1^{k_i-1} e^{-\theta_i\Gamma_1}, \quad (3)$$

where k_i is the shape parameter, θ_i is the scale parameter, and $\Gamma(k)$ is the gamma function [for all positive integers $\Gamma(k) = (k-1)!$]. An example of a gamma-distributed prior with $k_i = 5$ is shown in Fig. 1(e) (black dashed line). A gamma-distributed prior yields a posterior that remains

exactly gamma-distributed in the absence of state-preparation and measurement errors ($\alpha = \beta = 0$) when $m = 1$. Therefore, we expect that the posterior [see solid lines in Fig. 1(e)] obtained via Eq. (2) using the prior in Eq. (3) is in all relevant cases still approximately gamma-distributed. The gamma distribution is also a convenient choice for the case when the standard deviation becomes comparable to the mean, since by definition it has zero weight at negative Γ_1 . For standard deviations small compared to the mean ($k_i \gg 1$) the gamma distribution approaches a Gaussian distribution.

In the following we will thus always approximate the probability distribution for Γ_1 after the i th probing cycle with the gamma distribution $\mathcal{P}_i(\Gamma_1|k_i, \theta_i)$. We emphasize that while the gamma distribution is often associated with stochastic waiting times, in this work it is not linked to the physical origin of Γ_1 [51]. Rather, it is implemented in the controller for its mathematical convenience and since it provides a good approximation of a Gaussian distribution after sufficiently many measurements. In summary, the use of gamma distributions is advantageous as it requires two parameters (k_i, θ_i) only, and once multiplied by the exponential likelihood function (2), the posterior remains close to a gamma distribution. The controller thus only needs to keep track of two parameters after each measurement, which reduces the time the Bayesian update takes to $\approx 2.2 \mu\text{s}$, as mentioned above.

Two steps are required to implement the estimation in the controller. The first step is to determine an adaptive waiting time τ based on the prior distribution. Then, one must approximate the resulting posterior $\mathcal{P}_{i+1}(\Gamma_1)$ to a gamma distribution. This can be done using equations implemented directly on the controller in real time.

A. Adaptive waiting time

Working within a Bayesian framework, a common approach is to choose parameters that minimize the expected value of a quantity that measures the inaccuracy of the Γ_1 estimate, such as the variance or Shannon information of its probability distribution, after a future measurement is obtained. This works well as long as one can find an analytical expression for the optimal experiment based on the chosen metric [52]. Here it turns out that even a simple metric like the expected posterior variance does not give a simple analytical update rule for choosing τ . Therefore, in our Bayesian framework, we make a similar heuristic choice for the adaptive parameter τ as in Ref. [34]: In each cycle, the controller uses the current estimate $\hat{T}_1 \equiv 1/\langle\Gamma_1\rangle = \theta/k$ based on the prior distribution and uses the adaptive waiting time:

$$\tau_{i+1} = c\hat{T}_{1,i}, \quad (4)$$

where c is fixed in each experiment and depends on the qubit cycle idle time t (e.g., initialization, readout), and

measurement error rates α and β . The optimal choice for the coefficient c is theoretically bound to the interval $c \in (0, 1.59)$ and is chosen based on binomial statistics, to reduce the estimation time locally and uncertainty given the current knowledge of Γ_1 (see Supplemental Material [44]).

B. Posterior approximation

The prior $\mathcal{P}_i(\Gamma_1)$ in Eq. (2) is assumed to be a gamma distribution as given by Eq. (3), illustrated by the black dashed line in Fig. 1(e). After measuring $m_{i+1} = \{0, 1\}$, the posterior distribution $\mathcal{P}_{i+1}(\Gamma_1)$ is obtained by inserting Eqs. (1), (3) in Eq. (2). Since the posterior distribution is usually not an exact gamma distribution, we use the method of moments and approximate $\mathcal{P}_{i+1}(\Gamma_1)$ with the gamma distribution which has the same mean $\mu_{i+1} = E[\Gamma_1 | m_{i+1}, \nu_{i+1}]$ and variance $\sigma_{i+1}^2 = E[\Gamma_1^2 | m_{i+1}, \nu_{i+1}] - E[\Gamma_1 | m_{i+1}, \nu_{i+1}]^2$ computed over $\mathcal{P}_{i+1}(\Gamma_1)$, where we use the notation $\nu_i \equiv (k_i, \theta_i, \tau_{i+1})$ [53]. The required gamma distribution then has the parameters

$$\theta_{i+1}^{-1} = f_m(k_i + 1, \theta_i, \tau_{i+1}) - f_m(k_i, \theta_i, \tau_{i+1}), \quad (5a)$$

$$k_{i+1}^{-1} = \frac{f_m(k_i + 1, \theta_i, \tau_{i+1})}{f_m(k_i, \theta_i, \tau_{i+1})} - 1, \quad (5b)$$

where

$$f_0(k, \theta, \tau) = \frac{k}{\theta} \frac{1 - \beta - (1 - \alpha - \beta) \left(\frac{\theta}{\theta + \tau}\right)^{k+1}}{1 - \beta - (1 - \alpha - \beta) \left(\frac{\theta}{\theta + \tau}\right)^k}, \quad (6a)$$

$$f_1(k, \theta, \tau) = \frac{k}{\theta} \frac{\beta + (1 - \alpha - \beta) \left(\frac{\theta}{\theta + \tau}\right)^{k+1}}{\beta + (1 - \alpha - \beta) \left(\frac{\theta}{\theta + \tau}\right)^k}. \quad (6b)$$

The approximated $\mathcal{P}_{i+1}(\Gamma_1)$ becomes then the new prior and the controller repeats this scheme in total N times per estimation repetition to obtain a sufficiently narrow distribution [54], during which Γ_1 is assumed to be quasistatic. If one were to relax the quasistatic approximation, the optimum choice of N would depend on details of the Γ_1 drift and the estimation efficiency [44,55].

In Fig. 1(f) we illustrate (i) the resulting evolution of $\mathcal{P}_i(\Gamma_1)$ as a function of the measurement number i for one representative estimation sequence, and (ii) the corresponding waiting times τ_i and measurement outcomes m_i . The misclassification probabilities $\alpha = 0.11$ and $\beta = 0.14$ are obtained by fitting an exponential decay from the non-adaptive method [e.g., see inset of Figs. 2(d) and S3(c) of Supplemental Material [44]]. The initial prior is defined by $(k_0, \theta_0) = (3, 450 \mu\text{s})$, with $N = 50$ and $\tau \approx 0.51 / \langle \Gamma_1 \rangle$ [56], based on previously measured relaxation rate fluctuations. The purple line shows the estimates \hat{T}_1 and the shaded area indicates their 90% credible interval, narrowing \hat{T}_1 down from 168 μs (90% CI: [80, 630] μs) to 191 μs

(90% CI: [133, 301] μs) after 50 single-shot measurements. The user can predefine the number of single-shot measurements based on the required target uncertainty, which is traded off against estimation speed. Upon measuring $m_i = 1$ (blue circle), the estimate $\hat{T}_{1,i}$ increases compared to its previous value, and the subsequent waiting time τ_{i+1} also increases. In contrast, after measuring $m_i = 0$ (red circle), $\hat{T}_{1,i}$ decreases, while τ_{i+1} decreases on average.

IV. RESULTS

A. Tracking high-frequency fluctuations by adaptive estimation

We now apply our fast adaptive method to characterize the qubit at a previously unexplored temporal resolution. We first task the controller to acquire a time trace of \hat{T}_1 by the adaptive estimation, with $N = 100$ probe cycles [57] for each estimation repetition. We plot the estimated \hat{T}_1 on the controller in Fig. 2(a). The 68% confidence interval of the moving mean is computed using the standard error of the standard deviations of the 5 Bayesian posterior distributions: The standard error is estimated by first calculating the moving mean of the posterior standard deviations and then dividing by the square root of the window size. The confidence bounds are then determined as the moving mean of the estimated \hat{T}_1 values ± 1 standard error under normality assumption.

In Fig. 2(b), we plot the estimated \hat{T}_1 around 1878 s of panel (a) (other windows are presented in Supplemental Material [44]). As shown, \hat{T}_1 switches between $\approx 500 \mu\text{s}$ and $\approx 100 \mu\text{s}$ on a timescale of tens to hundreds of milliseconds. We emphasize that the purple points are *independent* of each other since at the beginning of each estimation repetition the prior distribution is reset to $\mathcal{P}_0(\Gamma_1)$.

The fluctuations in panels Figs. 2(a) and 2(b) exhibit telegraphic noise with multiple stable points, which could be explained by the qubit interacting with an ensemble of TLSs [14,15,24] changing due to spectral diffusion [22,26] or background ionizing radiation [19]. We highlight that such fast fluctuations would not be measurable using the standard nonadaptive method, which has previously reported sampling times of a few seconds [18,26,27]. The dwell times, on the order of tens of milliseconds, are consistent with the state-switching dynamics observed in a TLS strongly coupled to a superconducting qubit used as a detector [58]. Our sampling interval, which is 2 orders of magnitude faster than previous works [18,26,27], reveals that significant T_1 fluctuations can occur at least 1 order of magnitude faster than previously reported in transmon qubits [18,26–29].

The estimation time is only tens of times longer than the average T_1 and can be integrated into adaptive quantum control strategies to error-mitigate the performance of QPUs. One may imagine interleaving the estimation protocol with quantum circuit operations, executing them only

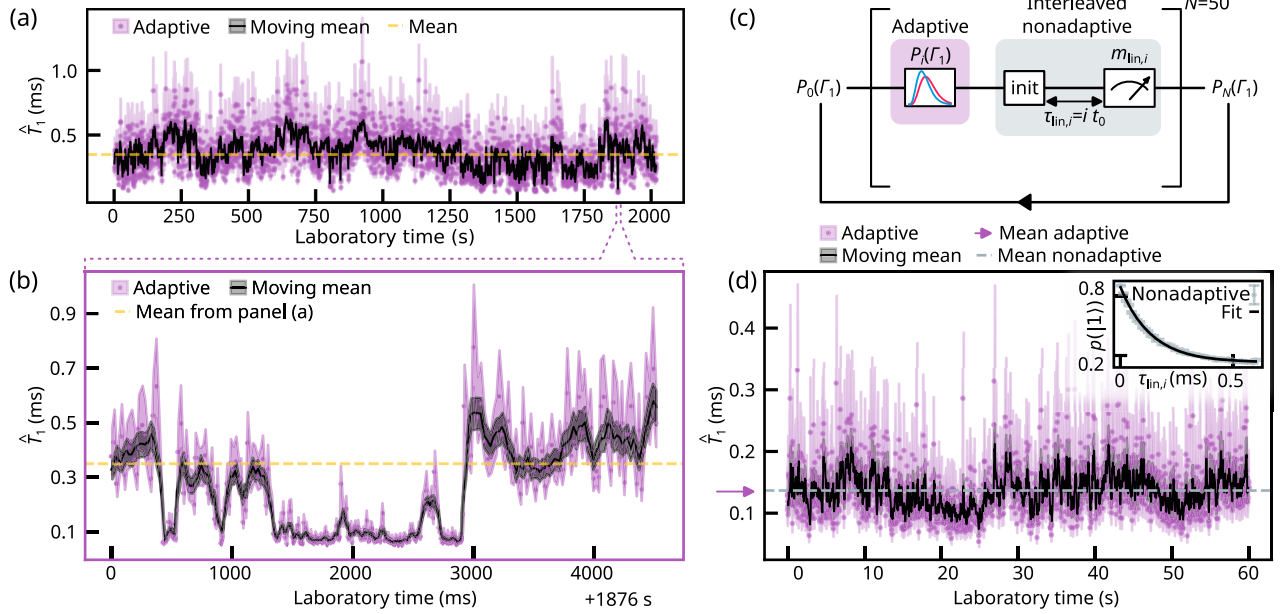


FIG. 2. Protocol for tracking and validation of the decay-rate fluctuations by adaptive estimation on the controller. (a) Experimental results for the adaptive tracking protocol with $N = 100$ to estimate \hat{T}_1 (purple dots, down sampled by $D = 30$) and its 68% credible interval (shaded area). Each purple point in this plot required an average estimation time of ≈ 20 ms. The dashed yellow line indicates the mean value of all estimates $\bar{T}_1 \approx 350$ μ s. The black line is a moving mean over 100 samples and the black shaded area is its 68% confidence interval (see main text). (b) Estimated \hat{T}_1 at ≈ 1876 s of panel (a) (see purple dashed lines), where \hat{T}_1 shows telegraphic switching with timescales on the order of tens or hundreds of milliseconds. (c) The interleaved estimation sequence for Γ_1 used to validate the adaptive protocol. Each of the N probe cycles, labeled i , consists of parts contributing to the adaptive (purple) and nonadaptive (gray) estimates. Each adaptive probe cycle is followed by the nonadaptive part of the cycle, where the qubit is again initialized in the excited state, the wait time is fixed to $\tau_{\text{in},i} = i\tau_0$, and the measurement outcome is stored for offline postprocessing. After the N probe cycles, the final adaptively obtained distribution $\mathcal{P}_N(\Gamma_1)$ is saved. (d) Experimental results for the adaptive tracking protocol, interleaved with nonadaptive measurements. Main panel: The estimate \hat{T}_1 (purple dots) and 68% credible interval (shaded area) corresponding to the final probability distribution $\mathcal{P}_{50}(\Gamma_1)$ of the 2000 adaptive estimates performed during the ≈ 60 s of the experiment. The black line shows a moving mean over 5 samples and the black shaded area is its 68% confidence interval. The purple arrow indicates the mean of all the adaptive estimates \bar{T}_1 . The dashed line is the value extracted from the fit shown in the inset. Inset: Experimental results for the nonadaptive estimate using linearly sampled waiting time $\tau_{\text{in},i}$. Error bars represent the standard error.

when the estimated qubit relaxation time T_1 exceeds a user-defined threshold to maintain a target gate fidelity. Since T_1 fluctuates over time, instead of running quantum circuits continuously, the system can pause its execution when the relaxation rate is too high, thus improving the overall QPU fidelity. Since both the interleaved estimation and qubit operations are performed on the same qubit, there exists a nontrivial relationship between the correlation time of the fluctuations being tracked, the efficiency of the estimation process, and the duration required for coherent operations between estimations. Overall, the results presented in this section demonstrate that our Bayesian estimation protocol performs real-time tracking of the decay rate of a superconducting qubit.

B. Protocol validation by interleaved adaptive and nonadaptive estimations

Next, we validate the protocol by programming the controller to perform interleaved measurements with the

nonadaptive method to verify that it correctly identifies the decay rate Γ_1 . The fluctuating parameter Γ_1 is estimated from the probing sequence shown in Fig. 2(c), where each probe cycle i begins with an adaptive cycle [see Fig. 1(b)] and is interleaved cycle-by-cycle with the nonadaptive one (gray) [59], the outcomes $m_{\text{in},i}$ of which are stored.

In Fig. 2(d) we characterize the experimentally found final posterior probability distribution $\mathcal{P}_{50}(\Gamma_1)$ from 2000 subsequent adaptive estimations. The 68% confidence interval of the moving mean is computed as in Fig. 2(a). In post-processing we calculate the 50 averages $\langle m_{\text{in},i} \rangle$ and fit them to an exponentially decaying curve, see the inset panel of Fig. 2(d), which yields $T_1 = (136.7 \pm 2.2)$ μ s (dashed gray line in the main panel). The 2000 adaptive estimations give on average $\bar{T}_1 = (135.0 \pm 0.9)$ μ s (purple arrow in the main panel), computed from the mean and standard error of the adaptive time trace shown in the main panel of Fig. 2(d). The main result is that the two values agree, and the reduced uncertainty of the adaptive method

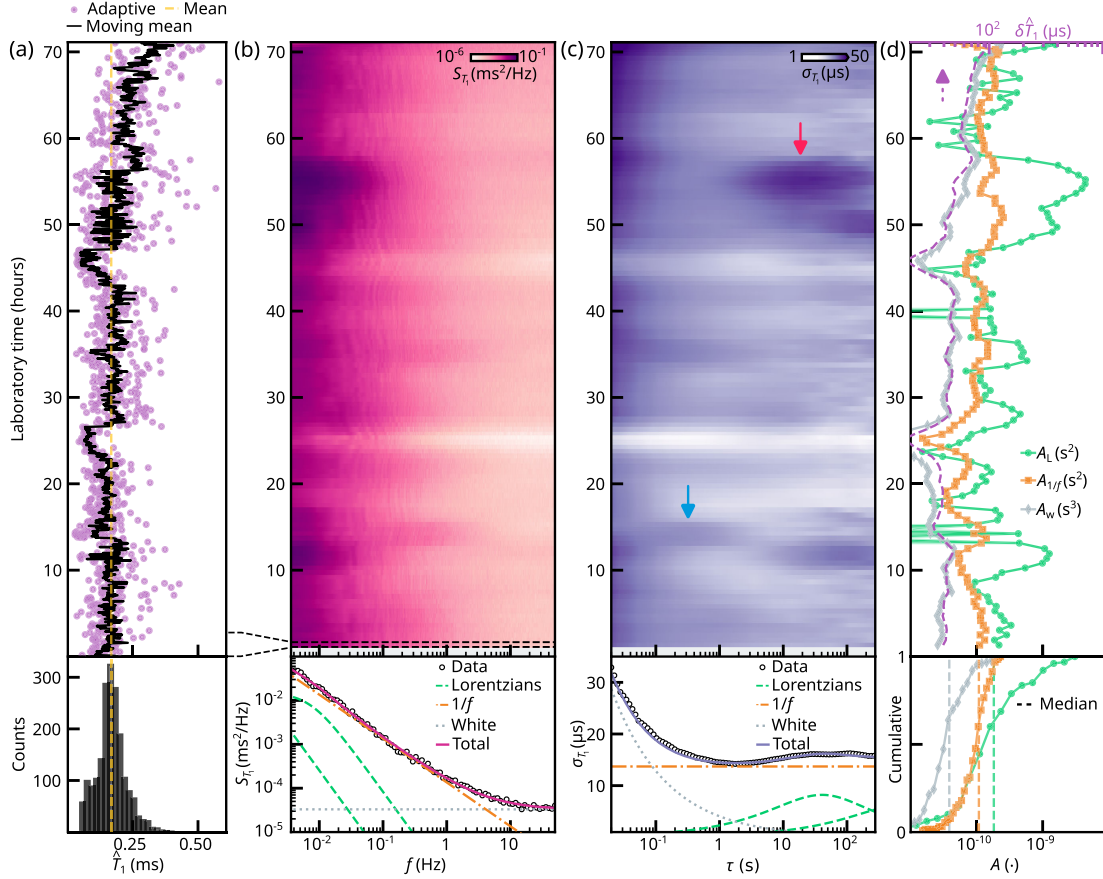


FIG. 3. Frequency and time domain analysis of \hat{T}_1 fluctuations on a 72-h timescale. (a) Estimated \hat{T}_1 (purple dots, down sampled by $D = 30000$) as a function of laboratory time by real-time adaptive tracking with $N = 49$ and sampling speed of ≈ 7 ms over 72 h. The dashed yellow line shows the mean value of all estimates $\bar{T}_1 \approx 168$ μs . The black line is a moving mean with a window of size 20 000 over the original estimates. Lower panel: Histogram of the moving mean. (b)–(c) Frequency and time domain analysis of the \hat{T}_1 fluctuations shown in panel (a), obtained from 2.8-h running windows with 80% overlap (black dashed lines): (b) Power spectral density. Lower panel: PSD of the full time trace. (c) Allan deviation on a logarithmic scale; arrows denote the enlarged regions of Fig. 4. Lower panel: Allan deviation of the full time trace. (d) Amplitudes of the Lorentzian (A_L), $1/f$ ($A_{1/f}$), and white noise (A_w) contributions, extracted from the full time trace, by a simultaneous PSD and Allan deviation fit to the analytical formulas from Table I. The purple dashed line is the standard deviation $\delta\hat{T}_1$ (see main text) of the Bayesian posterior distribution. Lower panel: Cumulative histogram of the fitted A_i and their medians (dashed lines).

results from the narrowing of the prior distribution and the adaptive waiting time chosen for the experiment.

The controller performs the same interleaved estimation procedure on another qubit for over 3 min (Q_2 , located on the same chip, see Supplemental Material [44]). In that case, the nonadaptive method yields $T_1 = (178 \pm 1.7)$ μs , which is again in good agreement with the mean adaptive estimate $\bar{T}_1 = (182.63 \pm 0.55)$ μs . Compared to Q_1 , we attribute the slightly larger discrepancy between these values to residual T_1 fluctuations occurring between the interleaved probing cycles.

C. Power spectral density and Allan deviation

To gain further insight into the physics of the fast fluctuations we observe, we calculate the power spectral density (PSD) and the Allan deviation [29,41,60] of a time

trace of \hat{T}_1 acquired over 72 h. The entire trace consists of $\approx 3.82 \times 10^7$ samples in total and is shown in Fig. 3(a). The controller uses $N = 49$ single-shot measurements per estimation repetition, with settings similar to those in Fig. 2(b) [61].

We compute the PSD [Fig. 3(b)] and Allan deviation [Fig. 3(c)] with a running 2.8-h-long window (with 80% overlap) in the top panels and for the full trace in the lower panels. Compared to the controller's fast sampling period of ≈ 7 ms, the relatively long window enables us to resolve Lorentzian processes in the Allan deviation within hundreds of seconds of observation time. Later, in Fig. 4, we use narrower windows to highlight faster dynamics.

In Fig. 3(b), top panel, we show the power spectral density $S_{\hat{T}_1}(f)$. The PSD of $\hat{T}_1(t)$ is defined as the Fourier transform of its autocorrelation function:

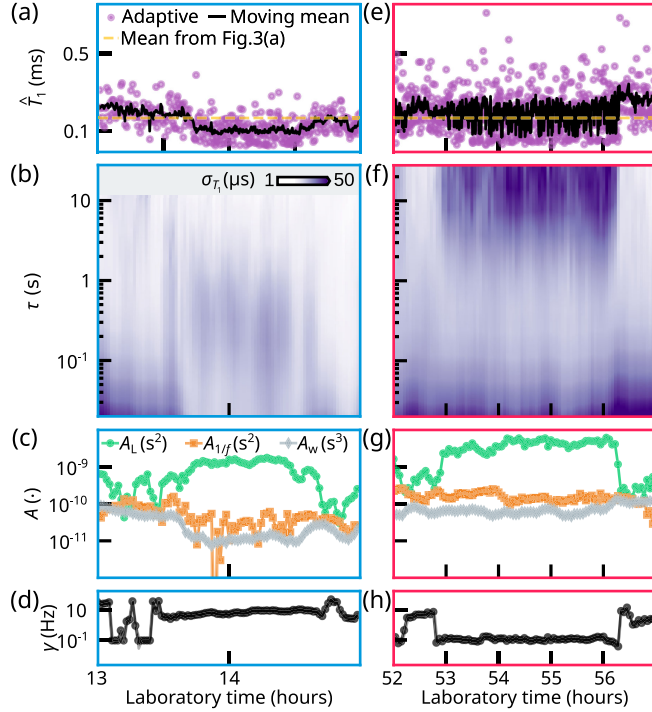


FIG. 4. Time domain analysis of \hat{T}_1 fluctuations on subminute observation times. Enlargements of the regions indicated by the arrows in Fig. 3(c). (a) Estimated \hat{T}_1 (purple dots, down sampled by $D = 3000$) as a function of laboratory time. The dashed yellow line represents the mean value of all estimates \bar{T}_1 of Fig. 3(a). The black line is a moving mean with a window size of 3000 over the original estimates. (b) Allan deviation in logarithmic scale of the \hat{T}_1 time trace with a running window of about 7 min and 80% overlap. (c) Amplitudes of the Lorentzian (A_L), $1/f$ ($A_{1/f}$), and white (A_w) noise contributions in each interval extracted from panel (a) by a simultaneous fit to the Allan deviation and PSD using analytical models from Table I. (d) Fitted switching rate $\gamma \approx 10$ Hz over more than 1 h of the Lorentzian component. (e)–(h) Same as (a)–(d) with a running window of 17 min and 80% overlap. The switching rate fit in panel (h) reveals a stable Lorentzian process with $\gamma \approx 100$ mHz for more than 3 h.

$$S_{\hat{T}_1}(f) = \int_{-\infty}^{+\infty} \langle \hat{T}_1(t) \hat{T}_1(t + \tau) \rangle e^{-2\pi i f \tau} d\tau. \quad (7)$$

The PSD exhibits increased amplitude at low frequencies over some laboratory times, but attributing this behavior to a TLS is not straightforward, as even a few TLSs can produce a smooth, featureless $1/f$ spectrum that conceals individual contributions. In contrast, the Allan deviation quantifies how much a signal switches on average over an observation time τ . If an individual TLS induces a random telegraph signal with a characteristic switching time, the Allan deviation of this signal exhibits a peak close to the average switching time of the TLS [29,60].

For our qubit parameter $\hat{T}_1(t)$, the Allan deviation is defined as

$$\sigma_{\hat{T}_1}(\tau) = \sqrt{\frac{1}{2} \langle (\bar{T}_1(t + \tau, \tau) - \bar{T}_1(t, \tau))^2 \rangle}, \quad (8)$$

where $\bar{T}_1(t, \tau) = (1/\tau) \int_t^{t+\tau} \hat{T}_1(t') dt'$ is the average over an interval of duration τ and $\langle \dots \rangle$ is the average over t . In Fig. 3(c), top panel, we plot $\sigma_{\hat{T}_1}(\tau)$ which is high at short observation times ($\tau \lesssim 100$ ms) due to high-frequency noise and sampling uncertainty. As τ increases, $\sigma_{\hat{T}_1}(\tau)$ decreases as the noise averages out. At longer τ , switching in \hat{T}_1 shows up as an increase in $\sigma_{\hat{T}_1}(\tau)$, which is clearly observed around τ in the range of tens of seconds at laboratory times of 11, 50, and 55 (red arrow) h. Another striking feature is the peak at 14 h (blue arrow) in the subsecond timescale. All the peaks shown in Fig. 3(c) are a clear signature of a Lorentzian noise process, as no power-law noise source can reproduce them, and they correspond to the rises in the PSD at smaller frequencies. The peaks of the Allan deviations are indications of individual TLSs that move in and out of resonance with the qubit [17,18]. To our knowledge, these features have not been previously observed at such short observation timescales and are accessible here because of the fast and adaptive estimation protocol. The PSD and Allan deviation are lower at 25 and 45 h of laboratory time, as a shorter T_1 generally corresponds to smaller fluctuations [26–29,62,63]. We tentatively attribute the shorter $\hat{T}_1 \approx 0.1$ ms to one or several TLSs being close to resonance with the qubit.

To quantify the impact of Lorentzian-type noise on the measured \hat{T}_1 fluctuations, it is common practice to fit the full time trace to a model including white noise, $1/f$ noise, and one or more Lorentzian noise sources [29]. We note that even if two Lorentzian processes fit the full trace, different TLS dynamics occur over laboratory time and it is actually sufficient for the fit to use only one Lorentzian in a shorter window slice. In the following, we present both analyses to support our claim.

In the bottom panels of Figs. 3(b) and 3(c) we plot the PSD and Allan deviation from the full trace shown in panel (a). Both curves are simultaneously fitted with a sum of white, $1/f$, and Lorentzian noise components, using the analytical expressions from Table I [29,41]. Each noise process has a corresponding amplitude coefficient, A_i , which is a free parameter for the Allan deviation $\sigma_{\hat{T}_1}(\tau)$ and the power spectral density $S_{\hat{T}_1}(f)$ simultaneously. The Lorentzian fit also includes the switching rate γ as a free parameter. To fit the full trace we use two Lorentzian processes and the fit parameters are given by $A_w = (3.30 \pm 0.06) \times 10^{-5} \text{ s}^3$, $A_{1/f} = (1.36 \pm 0.02) \times 10^{-4} \text{ s}^2$, $A_{L,1} = (1.8 \pm 0.1) \times 10^{-4} \text{ s}^2$, $\gamma_1 = (46 \pm 5) \text{ mHz}$, $A_{L,2} = (1.0 \pm 0.3) \times 10^{-4} \text{ s}^2$, and $\gamma_2 = (2 \pm 1) \text{ mHz}$.

Now, instead of the full trace, we fit each 2.8-hour-long time window of the top panels of Figs. 3(b) and 3(c) to show that only one Lorentzian process is sufficient to fit the data instead of two by selecting a shorter window. We plot

TABLE I. Power spectral density (PSD) and Allan deviation models for different noise processes. The amplitudes A_i and switching rate γ are free parameters for the simultaneous fit of the PSD and Allan deviation.

Noise	PSD $S_{\hat{T}_1}(f)$	Allan deviation $\sigma_{\hat{T}_1}(\tau)$
White	A_w	$\sqrt{A_w/\tau}$
$1/f$	$A_{1/f}f^{-1}$	$\sqrt{2A_{1/f}\ln 2}$
Lorentzian	$\frac{4A_L\gamma}{\gamma^2+(2\pi f)^2}$	$\sqrt{A_L\gamma\tau}\sqrt{2\gamma\tau+1-(e^{-\gamma\tau}-2)^2}$

the amplitudes of each noise process contribution in Fig. 3(d) and their cumulative histogram in the bottom panel. In the main panel of (d), the 68% confidence intervals of the fit parameters are smaller than the plotted symbols and are calculated as the square roots of the diagonal elements of the covariance matrix associated with the fit. The characteristic switching rate of the fitted Lorentzian component is shown in Supplemental Material [44], including the close agreement between the model and the data. The $1/f$ noise likely originates from the qubit interacting with an ensemble of TLSs [14,15,24] which dominate the low-frequency fluctuations of T_1 as observed from our fit. The presence of distinct Lorentzian components supports the existence of a few strongly coupled fluctuators, consistent with the telegraphic switching observed in Fig. 2(b).

In Fig. 3(d) we also plot the standard deviation $\delta\hat{T}_1 \approx \langle \Gamma_1 \rangle^{-2} \delta\Gamma_1 = \theta k^{-3/2}$ (purple dashed line) from the posterior distribution computed on the controller. We note that $\delta\hat{T}_1$ correlates very well with A_w (gray diamonds), which is the fitted amplitude of the white noise contribution according to our model. The strong correlation between the white noise power and the posterior standard deviation suggests that $\delta\hat{T}_1$ indeed reflects the true uncertainty in the estimate. In other words, if the white noise scales with $\delta\hat{T}_1$ above a known noise floor, it is strong indirect evidence that the Bayesian posterior's $\delta\hat{T}_1$ reliably captures the actual estimation error. To our knowledge, such correlation in real-time estimation methods has not been reported before.

To quantitatively extract the occurrence rate of large telegraphic switches in \hat{T}_1 on tens of milliseconds timescales [Fig. 2(b)], we split the 72-h-long time trace of Fig. 3(a) into independent 200-ms-long intervals. We find that $\approx 2.6\%$ of intervals exhibit on average changes in T_1 greater than 100 μs , roughly one event every 7.7 s (see Supplemental Material [44]).

As mentioned above, from Fig. 3(c) we observe dominant Lorentzian noise contributions for instance in the regions marked by the arrows. To highlight the fast noise dynamics pointed at by the blue arrow, we zoom in on the 2-h time period around hour 14 and plot the \hat{T}_1 estimates in Fig. 4(a). To increase the temporal resolution, we analyze

7-min intervals with 80% overlap and compute the corresponding Allan deviation for each, as shown in Fig. 4(b) with peaks around laboratory time of 14 h and within sub-second observation time τ . The fitted parameters are presented in Fig. 4(c) alongside the switching rate γ in Fig. 4(d).

The fitted amplitudes reveal distinct periods during which Lorentzian fluctuators dominate the T_1 dynamics. In particular, the analysis identifies a more than 1-h-long window dominated by a single fluctuator with a relatively stable switching rate of $\gamma \approx 10$ Hz, consistent with the timescales observed by Ref. [58]. Such fast dynamics would be extremely challenging to probe with traditional nonadaptive methods.

We also focus on the region marked by the red arrow of Fig. 3(c) over 5 h, see the corresponding \hat{T}_1 estimates in Fig. 4(e). Its Allan deviation is shown in Fig. 4(f), along with the fit parameters in panels (g) and (h). In this time span, the dominating TLS has a stable switching rate of about 100 mHz, slower than the 10 Hz presented in Fig. 4(d) but still 2 orders of magnitude faster than what is reported, for instance, in Ref. [29]. In panels (c), (d), (g), (h), the 68% confidence intervals of the fit parameters are smaller than the plotted symbols and are calculated as in Fig. 3(d).

In this section we have demonstrated how our adaptive Bayesian estimation method, about 2 orders of magnitude faster than conventional nonadaptive methods [18,26,27], confirms dominant Lorentzian noise processes in transmon qubits in a previously unexplored regime of quickly fluctuating relaxation times.

V. CONCLUSIONS AND OUTLOOK

This work presents the experimental demonstration of an adaptive Bayesian estimation protocol for the decay rate of two fixed-frequency transmon qubits. The scheme probes submillisecond relaxation times using only tens of single-shot measurements, with total acquisition times of a few milliseconds. Our approach achieves rapid adaptive estimation of the decay rate by integrating real-time Bayesian estimation with FPGA-based feedback control. The protocol has been validated by interleaved measurements with the standard nonadaptive method of extracting the decay rate from fitting an exponential decay curve.

The main results of our work are (i) our fast protocol improves the state-of-the-art estimation speed by 2 orders of magnitude without compromising accuracy [18,26,27], (ii) it reveals clear telegraphic changes in Γ_1 by almost one order of magnitude [Fig. 2(b)], with dwell times of tens of milliseconds, instead of minutes or hours, most likely caused by interaction with an environmental bath of TLSs, and (iii) temporal and spectral analyses of the estimated fluctuations are consistent with dominant Lorentzian processes, from which switching rates as fast as 100 mHz and 10 Hz are resolved in Fig. 4, 4 orders of magnitude faster

than previously reported. The estimation scheme is a powerful probe of dominant TLSs that switch in and out of resonance with the qubit on timescales of tens of seconds. Such timescales would be extremely difficult to observe with standard nonadaptive methods with the precision reported here, and they redefine the timescales relevant for calibration in superconducting QPUs, which traditionally operate on minute-to-hour cycles.

Our protocol uncovers new perspectives on materials characterization. High-throughput qubit screening targeting fast fluctuations previously required hours to accumulate sufficient statistics [16,29]. While long acquisition times remain necessary for characterizing slow drifts, our scheme collects fast fluctuations statistics within seconds. Average relaxation times of $T_1 \gtrsim 100 \mu\text{s}$ [64,65] have been shown in superconducting qubits, however achieving such results uniformly across a large wafer [10] and over time remains challenging. The overall performance is limited by the worst T_1 values, and T_1 tails can be improved through better fabrication that mitigates drops in T_1 . Our protocol provides a useful tool for qubit benchmarking and process control, enabling rapid optimization of superconducting qubits fabrication. Since the protocol employs only single-qubit gates, it can be readily extended to multiple qubits simultaneously, and its adaptiveness makes it appealing for the characterization of large qubit arrays with unknown decay rates.

In future work, the protocol could be interleaved with dephasing rate measurements [27] using existing real-time estimation techniques [34,40]. Potential correlations between decay and dephasing rates may provide insights into the microscopic origins of intrinsic decoherence channels at higher frequencies in Josephson junction qubits. Additionally, the estimation could be combined with spectral manipulation of TLSs via applied electric fields [21,23,65–67] or strain [66,68]. Our adaptive Bayesian technique could be employed for simultaneous relaxation time estimations across multiple qubits [18,63], to investigate temporal and spatial correlations at significantly higher frequency bandwidths than previously achieved. Spectral dynamics could also be explored in fixed-frequency transmons by implementing off-resonant microwave tones to drive ac Stark shifts [18,23]. Additionally, our method could be applied for the rapid detection of qubit decay induced by gamma and cosmic rays [14,15,69,70].

Potential modifications to the protocol include relaxing the assumption of single-shot readout [34], or mitigating state preparation and measurement errors by repeating probe cycles with the same waiting times [71,72], at the cost of a slower estimation rate. Furthermore, the estimation could be optimized by terminating it once a target total measurement time or desired uncertainty is reached, rather than using a fixed number of single-shot measurements.

Fluctuations in the decay rate degrade the stability of QPUs, introduce uncertainty in coherence benchmarking,

and hinder process optimization for superconducting qubits. Quantum error correction demonstrations are predominantly limited by the “worst” outlier qubits in a given processor [10,12]. Since state-of-the-art gate fidelities are limited by T_1 , the optimal gate duration depends on temporal variations of T_1 . The observed stochastic decay rate fluctuations suggest the following error mitigation approach: Continuous identification of the lowest-performing qubits, suggesting a shift from offline periodic recalibration every few hours to real-time adaptive recalibration at millisecond timescales for maintaining higher-fidelity gate operations in QPUs. More generally, the ability to monitor a qubit’s relaxation rate in real time opens new opportunities for the dynamic optimization and routing of quantum algorithms. A recently estimated T_1 value could be passed to a compiler’s routing algorithm to determine the optimal mapping between virtual and physical qubits, thereby boosting the benchmarking or diagnostics of useful QPUs, and could be further extended to other platforms.

Our adaptive Bayesian technique also offers an efficient and online Hamiltonian learning protocol for real-time estimation of decay rates beyond superconducting qubits. Our results support that TLSs are major contributors to rapid decay-rate fluctuations and a deeper understanding of TLSs is thus vital for further improving the performance of useful QPUs.

ACKNOWLEDGMENTS

F. B., M. A. M., S. K., C. W. W., J. H., and M. K. gratefully acknowledge support from the Novo Nordisk Foundation (Grant No. NNF22SA0081175, the NNF Quantum Computing Programme), the Villum Foundation through a Villum Young Investigator grant (Grant No. 37467), the European Union through an ERC Starting Grant (NovADePro, 101077479), the Innovation Fund Denmark (DanQ, Grant No. 2081-00013B), and the U.S. Army Research Office (Grant No. W911NF-22-1-0042). F. B., J. Benestad, J. A. K., J. D., and F. K. acknowledge funding from the European Union’s Horizon 2020 research and innovation programme under Grant Agreements No. 101017733 (QuantERA II), No. 101204890 (HORIZON-MSCA-2024-PF-01), and EUREKA Eurostars 3 (ECHIDNA), and from the Dutch National Growth Fund (NGF) as part of the Quantum Delta NL programme and the Research Council of Norway (RCN) under INTFELLES-Project No. 333990. O. K. received funding via the Innovation Fund Denmark for the project DIREC (9142-00001B). The device was fabricated at Myfab Chalmers. The Chalmers team was funded by the Knut and Alice Wallenberg Foundation through the Wallenberg Center for Quantum Technology (WACQT) and the EU Flagship on Quantum Technology HORIZON-CL4-2022-QUANTUM-01-SGA Project No. 101113946 OpenSuperQPlus100.

F. B. led the measurements and data analysis, and wrote the manuscript with input from all authors. F. B., J. Benestad, J. H., F. K., and M. K. performed the experiment with theoretical contributions from J. A. K., O. K., J. D. M. A. M. and S. K. developed experimental infrastructure. C. W. W. designed the device, which was fabricated by E. H., A. N., I. A., A. O., J. Biznárová, and M. R. under the supervision of A. F. R., J. Bylander, and G. T. The project was supervised by J. H., F. K., and M. K.

Any opinions, findings, conclusions, or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Army Research Office, the U.S. Government, the European Union, or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them.

DATA AVAILABILITY

The data that support the findings of this article are openly available [73].

-
- [1] 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).
- [2] H.-L. Huang, D. Wu, D. Fan, and X. Zhu, *Superconducting quantum computing: A review*, *Sci. China Inf. Sci.* **63**, 1 (2020).
- [3] A. Blais, A. L. Grimsmo, S. M. Girvin, and A. Wallraff, *Circuit quantum electrodynamics*, *Rev. Mod. Phys.* **93**, 025005 (2021).
- [4] Y. Y. Gao, M. A. Rol, S. Touzard, and C. Wang, *Practical guide for building superconducting quantum devices*, *PRX Quantum* **2**, 040202 (2021).
- [5] S. E. Rasmussen, K. S. Christensen, S. P. Pedersen, L. B. Kristensen, T. Bækkegaard, N. J. S. Loft, and N. T. Zinner, *Superconducting circuit companion—an introduction with worked examples*, *PRX Quantum* **2**, 040204 (2021).
- [6] E. Campbell, *A series of fast-paced advances in quantum error correction*, *Nat. Rev. Phys.* **6**, 160 (2024).
- [7] L. Caune, L. Skoric, N. S. Blunt, A. Ruban, J. McDaniel, J. A. Valery, A. D. Patterson, A. V. Gramolin, J. Majaniemi, K. M. Barnes, T. Bialas, O. Buğdaycı, O. Crawford, G. P. Gehér, H. Krovi *et al.*, *Demonstrating real-time and low-latency quantum error correction with superconducting qubits*, [arXiv:2410.05202](https://arxiv.org/abs/2410.05202).
- [8] Google Quantum AI and Collaborators, *Quantum error correction below the surface code threshold*, *Nature (London)* **638**, 920 (2025).
- [9] I. Besedin, M. Kerschbaum, J. Knoll, I. Hesner, L. Bödeker, L. Colmenarez, L. Hofele, N. Lacroix, C. Hellings, F. Swiadek, A. Flasby, M. B. Panah, D. C. Zanz, M. Müller, and A. Wallraff, *Realizing lattice surgery on two distance-three repetition codes with superconducting qubits*, [arXiv:2501.04612](https://arxiv.org/abs/2501.04612).
- [10] M. Mohseni, A. Scherer, K. G. Johnson, O. Wertheim, M. Otten, N. A. Aadit, K. M. Bresniker, K. Y. Camsari, B. Chapman, S. Chatterjee, G. A. Dagnew, A. Esposito, F. Fahim, M. Fiorentino, A. Khalid *et al.*, *How to build a quantum supercomputer: Scaling challenges and opportunities*, [arXiv:2411.10406](https://arxiv.org/abs/2411.10406).
- [11] D. Gao, D. Fan, C. Zha, J. Bei, G. Cai, J. Cai, S. Cao, F. Chen, J. Chen, K. Chen, X. Chen, X. Chen, Z. Chen, Z. Chen, Z. Chen *et al.*, *Establishing a new benchmark in quantum computational advantage with 105-qubit Zuchongzhi 3.0 processor*, *Phys. Rev. Lett.* **134**, 090601 (2025).
- [12] Z. Wei, T. He, Y. Ye, D. Wu, Y. Zhang, Y. Zhao, W. Lin, H.-L. Huang, X. Zhu, and J.-W. Pan, *Low-overhead defect-adaptive surface code with bandage-like super-stabilizers*, *npj Quantum Inf.* **11**, 75 (2025).
- [13] P. J. J. O'Malley, J. Kelly, R. Barends, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, A. G. Fowler, I.-C. Hoi, E. Jeffrey, A. Megrant, J. Mutus, C. Neill, C. Quintana *et al.*, *Qubit metrology of ultralow phase noise using randomized benchmarking*, *Phys. Rev. Appl.* **3**, 044009 (2015).
- [14] I. Siddiqi, *Engineering high-coherence superconducting qubits*, *Nat. Rev. Mater.* **6**, 875 (2021).
- [15] C. E. Murray, *Material matters in superconducting qubits*, *Mater. Sci. Eng.* **146**, 100646 (2021).
- [16] C. R. H. McRae, G. M. Stiehl, H. Wang, S.-X. Lin, S. A. Caldwell, D. P. Pappas, J. Mutus, and J. Combes, *Reproducible coherence characterization of superconducting quantum devices*, *Appl. Phys. Lett.* **119**, 100501 (2021).
- [17] 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).
- [18] M. Carroll, S. Rosenblatt, P. Jurcevic, I. Lauer, and A. Kandala, *Dynamics of superconducting qubit relaxation times*, *npj Quantum Inf.* **8**, 132 (2022).
- [19] T. Thorbeck, A. Eddins, I. Lauer, D. T. McClure, and M. Carroll, *Two-level-system dynamics in a superconducting qubit due to background ionizing radiation*, *PRX Quantum* **4**, 020356 (2023).
- [20] D. Colao Zanz, Q. Ficheux, L. Michaud, A. Orekhov, K. Hanke, A. Flasby, M. Bahrami Panah, G. J. Norris, M. Kerschbaum, A. Remm, F. Swiadek, C. Hellings, S. Lazăr, C. Scarato, N. Lacroix, S. Krinner, C. Eichler, A. Wallraff, and J. C. Besse, *Mitigating losses of superconducting qubits strongly coupled to defect modes*, *Phys. Rev. Appl.* **23**, 044054 (2025).
- [21] Y. Kim, L. C. Govia, A. Dane, E. v. d. Berg, D. M. Zajac, B. Mitchell, Y. Liu, K. Balakrishnan, G. Keefe, A. Stabile *et al.*, *Error mitigation with stabilized noise in superconducting quantum processors*, *Nat. Commun.* **16**, 8439 (2025).
- [22] S. Weeden, D. C. Harrison, S. Patel, M. Snyder, E. J. Blackwell, G. Spahn, S. Abdullah, Y. Takeda, B. L. T. Plourde, J. M. Martinis, and R. McDermott, *Statistics of strongly coupled defects in superconducting qubits*, [arXiv:2506.00193](https://arxiv.org/abs/2506.00193).
- [23] L. Chen, K.-H. Lee, C.-H. Liu, B. Marinelli, R. K. Naik, Z. Kang, N. Goss, H. Kim, D. I. Santiago, and I. Siddiqi,

- Scalable and site-specific frequency tuning of two-level system defects in superconducting qubit arrays*, arXiv:2503.04702.
- [24] C. Müller, J. H. Cole, and J. Lisenfeld, *Towards understanding two-level-systems in amorphous solids: Insights from quantum circuits*, *Rep. Prog. Phys.* **82**, 124501 (2019).
- [25] Y. Shalibo, Y. Rofe, D. Shwa, F. Zeides, M. Neeley, J. M. Martinis, and N. Katz, *Lifetime and coherence of two-level defects in a Josephson junction*, *Phys. Rev. Lett.* **105**, 177001 (2010).
- [26] P. V. Klimov, J. Kelly, Z. Chen, M. Neeley, A. Megrant, B. Burkett, R. Barends, K. Arya, B. Chiaro, Y. Chen, A. Dunsworth, A. Fowler, B. Foxen, C. Gidney, M. Giustina, R. Graff *et al.*, *Fluctuations of energy-relaxation times in superconducting qubits*, *Phys. Rev. Lett.* **121**, 090502 (2018).
- [27] 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).
- [28] C. Müller, J. Lisenfeld, A. Shnirman, and S. Poletto, *Interacting two-level defects as sources of fluctuating high-frequency noise in superconducting circuits*, *Phys. Rev. B* **92**, 035442 (2015).
- [29] 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**, 132 (2019).
- [30] C. E. Granade, C. Ferrie, N. Wiebe, and D. G. Cory, *Robust online Hamiltonian learning*, *New J. Phys.* **14**, 103013 (2012).
- [31] V. Gebhart, R. Santagati, A. A. Gentile, E. M. Gauger, D. Craig, N. Ares, L. Bianchi, F. Marquardt, L. Pezzè, and C. Bonato, *Learning quantum systems*, *Nat. Rev. Phys.* **5**, 141 (2023).
- [32] M. Spiecker, P. Paluch, N. Gosling, N. Drucker, S. Matityahu, D. Gusenkova, S. Günzler, D. Rieger, I. Tadmakov, F. Valenti *et al.*, *Two-level system hyperpolarization using a quantum Szilard engine*, *Nat. Phys.* **19**, 1320 (2023).
- [33] K. Reuer, J. Landgraf, T. Fösel, J. O’Sullivan, L. Beltrán, A. Akin, G. J. Norris, A. Remm, M. Kerschbaum, J.-C. Besse *et al.*, *Realizing a deep reinforcement learning agent for real-time quantum feedback*, *Nat. Commun.* **14**, 7138 (2023).
- [34] M. J. Arshad, C. Bekker, B. Haylock, K. Skrzypczak, D. White, B. Griffiths, J. Gore, G. W. Morley, P. Salter, J. Smith, I. Zohar, A. Finkler, Y. Altmann, E. M. Gauger, and C. Bonato, *Real-time adaptive estimation of decoherence timescales for a single qubit*, *Phys. Rev. Appl.* **21**, 024026 (2024).
- [35] F. Berritta, T. Rasmussen, J. A. Krzywda, J. van der Heijden, F. Fedele, S. Fallahi, G. C. Gardner, M. J. Manfra, E. van Nieuwenburg, J. Danon, A. Chatterjee, and F. Kuemmeth, *Real-time two-axis control of a spin qubit*, *Nat. Commun.* **15**, 1676 (2024).
- [36] N. Dumoulin Stuyck, A. E. Seedhouse, S. Serrano, T. Tantt, W. Gilbert, J. Y. Huang, F. Hudson, K. M. Itoh, A. Laucht, W. H. Lim, C. H. Yang, A. Saraiva, and A. S. Dzurak, *Silicon spin qubit noise characterization using real-time feedback protocols and wavelet analysis*, *Appl. Phys. Lett.* **124**, 114003 (2024).
- [37] N. R. Vora, Y. Xu, A. Hashim, N. Fruitwala, H. N. Nguyen, H. Liao, J. Balewski, A. Rajagopala, K. Nowrouzi, Q. Ji *et al.*, *ML-powered FPGA-based real-time quantum state discrimination enabling mid-circuit measurements*, arXiv:2406.18807.
- [38] F. Berritta, J. A. Krzywda, J. Benestad, J. van der Heijden, F. Fedele, S. Fallahi, G. C. Gardner, M. J. Manfra, E. van Nieuwenburg, J. Danon, A. Chatterjee, and F. Kuemmeth, *Physics-informed tracking of qubit fluctuations*, *Phys. Rev. Appl.* **22**, 014033 (2024).
- [39] J. Park, H. Jang, H. Sohn, J. Yun, Y. Song, B. Kang, L. E. A. Stehouwer, D. D. Esposti, G. Scappucci, and D. Kim, *Passive and active suppression of transduced noise in silicon spin qubits*, *Nat. Commun.* **16**, 78 (2025).
- [40] F. Berritta, J. Benestad, L. Pahl, M. Mathews, J. A. Krzywda, R. Assouly, Y. Sung, D. K. Kim, B. M. Niedzielski, K. Serniak, M. E. Schwartz, J. L. Yoder, A. Chatterjee, J. A. Grover, J. Danon, W. D. Oliver, and F. Kuemmeth, *Efficient qubit calibration by binary-search Hamiltonian tracking*, *PRX Quantum* **6**, 030335 (2025).
- [41] C. M. Van Vliet and P. H. Handel, *A new transform theorem for stochastic processes with special application to counting statistics*, *Physica (Amsterdam)* **113A**, 261 (1982).
- [42] R. C. Kurchin, *Using Bayesian parameter estimation to learn more from data without black boxes*, *Nat. Rev. Phys.* **6**, 152 (2024).
- [43] C. W. Warren, J. Fernández-Pendás, S. Ahmed, T. Abad, A. Bengtsson, J. Biznárová, K. Debnath, X. Gu, C. Križan, A. Osman, A. Fadavi Roudsari, P. Delsing, G. Johansson, A. Frisk Kockum, G. Tancredi, and J. Bylander, *Extensive characterization and implementation of a family of three-qubit gates at the coherence limit*, *npj Quantum Inf.* **9**, 44 (2023).
- [44] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/gk1b-stl3> for the experimental setup, additional examples of decay-rate estimation for both qubits, and theoretical details on the optimal waiting time based on binomial statistics and on the decay-rate estimation using gamma distributions. The Supplemental Material includes Refs. [45,46].
- [45] S. Gustavsson, F. Yan, G. Catelani, J. Bylander, A. Kamal, J. Birenbaum, D. Hover, D. Rosenberg, G. Samach, A. P. Sears *et al.*, *Suppressing relaxation in superconducting qubits by quasiparticle pumping*, *Science* **354**, 1573 (2016).
- [46] F. Leveraro, *Novel techniques for fast and precise estimation of qubit decay rates*, Master’s thesis, University of Copenhagen (2024).
- [47] D. Ristè, C. C. Bultink, K. W. Lehnert, and L. DiCarlo, *Feedback control of a solid-state qubit using high-fidelity projective measurement*, *Phys. Rev. Lett.* **109**, 240502 (2012).
- [48] Our model assumes that an exponential decay well approximates a relaxation time experiment. However, this assumption breaks down in the presence of strong coupling between the qubit and a TLS, which induces oscillations in the excited-state return probability. If the

- oscillation frequency is known, the protocol can be easily adapted by sampling at the maxima (or minima) of these oscillations. Otherwise, the model must be extended for multiparameter Bayesian estimation [34,49]. Such strong coupling regimes have not been observed in Q_1 and Q_2 during our experiments.
- [49] M. Valeri, V. Cimini, S. Piacentini, F. Ceccarelli, E. Polino, F. Hoch, G. Bizzarri, G. Corrielli, N. Spagnolo, R. Osellame *et al.*, *Experimental multiparameter quantum metrology in adaptive regime*, *Phys. Rev. Res.* **5**, 013138 (2023).
- [50] C. Ferrie, C. E. Granade, and D. G. Cory, *How to best sample a periodic probability distribution, or on the accuracy of Hamiltonian finding strategies*, *Quantum Inf. Process.* **12**, 611 (2013).
- [51] A gamma distribution models the statistics of the sum of independent decay processes, where the random variable represents the total decay time. In that sense, it is naturally related to decay dynamics. However, in our case, we measure the state m rather than the actual decay time τ . As a result, the relevant statistics are binomial rather than the exponential distribution typically associated with lifetime studies. Therefore, the appearance of the gamma distribution in this work is not due to physical origins, but rather because it closely resembles the true posterior.
- [52] Thus for adaptive methods, the Shannon information is rarely used in practice, as it tends to be computationally intractable.
- [53] For $(\alpha = \beta = 0)$, we have $E[\Gamma_1 | m_i, \nu_i] = k/\theta$ and $E[\Gamma_1^2 | m_i, \nu_i] = (k_i + k_i^2)/\theta_i^2$ (see Supplemental Material [44] for the case $\alpha, \beta \neq 0$).
- [54] We perform a numerical approximation of the Kullback-Leibler divergence (D_{KL}) to estimate how much the posterior distribution differs from our gamma-distribution approximation. The largest deviation occurs for the worst-case outcome $m = 1$ together with a small prior shape parameter, typically $k = 3$. In this regime the posterior is strongly skewed, and the approximation error reaches a maximum of roughly $D_{KL} \approx 0.09$ – 0.10 at intermediate ratios $\tau_{i+1}/\theta \gtrsim 1$. The accuracy improves systematically with increasing k . For example, at $k = 10$ the peak divergence drops below 0.006, and for $k = 20$ it falls below 0.001 across the entire range of τ_{i+1}/θ . See Supplemental Material [44] for further details on the validity of the approximation, also with respect to a normal distribution.
- [55] J. Benestad, J. Krzywda, E. van Nieuwenburg, and J. Danon, *Efficient adaptive Bayesian estimation of a slowly fluctuating overhauser field gradient*, *SciPost Phys.* **17**, 014 (2024).
- [56] Because of the numerical precision of the controller, the actual ratio between τ_{i+1} and $\hat{T}_{1,i}$ varies by approximately 10% in the vast majority of probe cycles.
- [57] Each probe cycle includes a 2.5 μs readout period, followed by an approximately 8 μs wait time to allow the resonator to cool down and fully deplete any residual photons from the readout pulse. Each Bayesian update takes ≈ 2.2 μs . The controller is programmed to start with an initial gamma distribution prior with $(k_0, \theta_0) = (3, 450 \mu\text{s})$. With this in mind, the adaptive waiting time in each cycle is set to $\tau_{i+1} \approx 0.51 \hat{T}_{1,i}$ (see Supplemental Material [44] for a motivation of the choice $c = 0.51$), where $\hat{T}_{1,i}$ is the estimator for T_1 after cycle i .
- [58] S. M. Meißner, A. Seiler, J. Lisenfeld, A. V. Ustinov, and G. Weiss, *Probing individual tunneling fluctuators with coherently controlled tunneling systems*, *Phys. Rev. B* **97**, 180505(R) (2018).
- [59] Here $(k_0, \theta_0) = (3, 450 \mu\text{s})$, $\alpha = 0.11$, $\beta = 0.14$. Considering also that the nonadaptive $\tau_{\text{in},i}$ is linearly stepped up to $\approx 600 \mu\text{s}$, the average idle time from the point of view of the adaptive scheme is $t \approx 345 \mu\text{s}$, mostly covered by the average $\langle \tau_{\text{in},i} \rangle \approx 300 \mu\text{s}$. The adaptive waiting time in each cycle is then set to $\tau_{i+1} \approx 0.98 \hat{T}_{1,i}$ (see Supplemental Material [44]).
- [60] F. Ye, A. Ellaboudy, D. Albrecht, R. Vudatha, N. T. Jacobson, and J. M. Nichol, *Characterization of individual charge fluctuators in Si/SiGe quantum dots*, *Phys. Rev. B* **110**, 235305 (2024).
- [61] $(k_0, \theta_0) = (3, 600 \mu\text{s})$, $\alpha = \beta = 0.12$, and $c \approx 0.53$.
- [62] M. Bal, A. A. Murthy, S. Zhu, F. Crisa, X. You, Z. Huang, T. Roy, J. Lee, D. v. Zanten, R. Pilipenko *et al.*, *Systematic improvements in transmon qubit coherence enabled by niobium surface encapsulation*, *npj Quantum Inf.* **10**, 43 (2024).
- [63] S. Kono, J. Pan, M. Chegnizadeh, X. Wang, A. Youssefi, M. Scigliuzzo, and T. J. Kippenberg, *Mechanically induced correlated errors on superconducting qubits with relaxation times exceeding 0.4 ms*, *Nat. Commun.* **15**, 3950 (2024).
- [64] M. Tuokkola, Y. Sunada, H. Kivijärvi, J. Albanese, L. Grönberg, J.-P. Kaikkonen, V. Vesterinen, J. Govenius, and M. Möttönen, *Methods to achieve near-millisecond energy relaxation and dephasing times for a superconducting transmon qubit*, *Nat. Commun.* **16**, 5421 (2025).
- [65] A. Dane, K. Balakrishnan, B. Wacaser, L.-W. Hung, H. J. Mamin, D. Rugar, R. M. Shelby, C. Murray, K. Rodbell, and J. Sleight, *Performance stabilization of high-coherence superconducting qubits*, [arXiv:2503.12514](https://arxiv.org/abs/2503.12514).
- [66] J. Lisenfeld, A. Bilmes, A. Megrant, R. Barends, J. Kelly, P. Klimov, G. Weiss, J. M. Martinis, and A. V. Ustinov, *Electric field spectroscopy of material defects in transmon qubits*, *npj Quantum Inf.* **5**, 105 (2019).
- [67] A. Bilmes, A. Megrant, P. Klimov, G. Weiss, J. M. Martinis, A. V. Ustinov, and J. Lisenfeld, *Resolving the positions of defects in superconducting quantum bits*, *Sci. Rep.* **10**, 3090 (2020).
- [68] G. J. Grabovskij, T. Peichl, J. Lisenfeld, G. Weiss, and A. V. Ustinov, *Strain tuning of individual atomic tunneling systems detected by a superconducting qubit*, *Science* **338**, 232 (2012).
- [69] J. M. Martinis, *Saving superconducting quantum processors from decay and correlated errors generated by gamma and cosmic rays*, *npj Quantum Inf.* **7**, 90 (2021).
- [70] P. M. Harrington, M. Li, M. Hays, W. Van De Pontseele, D. Mayer, H. D. Pinckney, F. Contipelli, M. Gingras, B. M. Niedzielski, H. Stickler *et al.*, *Synchronous detection of cosmic rays and correlated errors in superconducting qubit arrays*, *Nat. Commun.* **16**, 6428 (2025).
- [71] A. Sergeevich, A. Chandran, J. Combes, S. D. Bartlett, and H. M. Wiseman, *Characterization of a qubit Hamiltonian*

- using adaptive measurements in a fixed basis*, *Phys. Rev. A* **84**, 052315 (2011).
- [72] P. Cappellaro, *Spin-bath narrowing with adaptive parameter estimation*, *Phys. Rev. A* **85**, 030301(R) (2012).
- [73] F. Berritta, J. Benestad, J. Krzywda, O. Krause, M. A. M. Nielsen, S. Krøjer, C. Warren, E. Rehnman, A. Nylander, I. Ahmad, A. Osman, J. Biznárová, M. Rommel, A. Fadavi Roudsari, J. Bylander, G. Tancredi, J. Danon, J. Hastrup, F. Kuemmeth, and M. Kjaergaard, *Dataset underlying the manuscript: Real-time adaptive tracking of fluctuating relaxation rates in superconducting qubits [Data set]*, Zenodo (2026), [10.5281/zenodo.17345142](https://doi.org/10.5281/zenodo.17345142).