Editors' Suggestion

Physics-informed tracking of qubit fluctuations

Fabrizio Berritta,[1] Jan A. Krzywda,[2] Jacob Benestad,[3] Joost van der Heijden,[4] Federico Fedele,[1,5] Saeed Fallahi,[6,7] Geoffrey C. Gardner,[7] Michael J. Manfra,[6,7,8,9] Evert van Nieuwenburg,[2] Jeroen Danon,[3] Anasua Chatterjee,[1] and Ferdinand Kuemmeth[1,4,*]

[1] *Center for Quantum Devices, Niels Bohr Institute, University of Copenhagen, 2100 Copenhagen, Denmark*
[2] *Lorentz Institute and Leiden Institute of Advanced Computer Science, Leiden University, P.O. Box 9506, 2300 RA Leiden, The Netherlands*
[3] *Center for Quantum Spintronics, Department of Physics, Norwegian University of Science and Technology, NO-7491 Trondheim, Norway*
[4] *QDevil, Quantum Machines, 2750 Ballerup, Denmark*
[5] *Department of Engineering Science, University of Oxford, Oxford OX1 3PJ, United Kingdom*
[6] *Department of Physics and Astronomy, Purdue University, West Lafayette, Indiana 47907, USA*
[7] *Birck Nanotechnology Center, Purdue University, West Lafayette, Indiana 47907, USA*
[8] *Elmore Family School of Electrical and Computer Engineering, Purdue University, West Lafayette, Indiana 47907, USA*
[9] *School of Materials Engineering, Purdue University, West Lafayette, Indiana 47907, USA*

(Received 16 April 2024; revised 26 April 2024; accepted 7 June 2024; published 15 July 2024)

Environmental fluctuations degrade the performance of solid-state qubits but can in principle be mitigated by real-time Hamiltonian estimation down to timescales set by the estimation efficiency. We implement a physics-informed and an adaptive Bayesian estimation strategy and apply them in real time to a semiconductor spin qubit. The physics-informed strategy propagates a probability distribution inside the quantum controller according to the Fokker-Planck equation, appropriate for describing the effects of nuclear spin diffusion in gallium arsenide. Evaluating and narrowing the anticipated distribution by a predetermined qubit probe sequence enables improved dynamical tracking of the uncontrolled magnetic field gradient within the singlet-triplet qubit. The adaptive strategy replaces the probe sequence by a small number of qubit probe cycles, with each probe time conditioned on the previous measurement outcomes, thereby further increasing the estimation efficiency. The combined real-time estimation strategy efficiently tracks low-frequency nuclear spin fluctuations in solid-state qubits, and can be applied to other qubit platforms by tailoring the appropriate update equation to capture their distinct noise sources.

DOI: 10.1103/PhysRevApplied.22.014033

I. INTRODUCTION

Low-frequency environmental fluctuations cause decoherence in solid-state qubits [1–3]. Quantum error correction strategies [4] can detect and correct errors but demand an increased number of physical qubits. Conventional noise reduction techniques, such as dynamical decoupling [5,6] and active suppression of environmental fluctuations [7–9], are not universally effective and may not align with specific experimental goals.

Hamiltonian learning emerges as a promising solution for compensating for uncontrolled environmental effects and enhancing the qubit quality factor [10–16]. This approach leverages modern hardware capabilities to provide real-time feedback, but comes at the cost of dedicating time to estimate the fluctuating Hamiltonian parameters. Although several theoretical estimation schemes [17–24] have been proposed to boost the estimation efficiency, no experiment has yet demonstrated a *physics-informed* scheme within any qubit platform, where understanding of the physical processes driving the fluctuations is utilized to improve the estimations. Even the experimental adoption of real-time *adaptive* Bayesian strategies [25,26], where measurement parameters are chosen based on the previous measurements, is still missing in gate-defined spin qubits. This work reports a real-time physics-informed and adaptive Bayesian estimation of a qubit.

To demonstrate an adaptive and physics-informed estimation protocol, we employ a singlet-triplet (ST_0) qubit in GaAs. In nitrogen-vacancy centers in diamond [27] and semiconductor spin qubits [28], low-frequency noise

*Contact author: kuemmeth@nbi.dk

from spinful nuclear isotopes decreases qubit performance through hyperfine interactions. Isotopic purification techniques [29,30] mitigate this issue in group IV semiconductors such as silicon and germanium, though it comes with significant effort and does not remove low-frequency noise originating from other sources. For our demonstration we chose GaAs as its nuclear noise spectrum is well understood [31,32]. Our technique involves programming a commercial quantum controller, powered by an integrated field-programmable gate array (FPGA), to propagate the probability distribution of the effective nuclear fields on the dots in real time, using the Fokker-Planck (FP) equation [23,24]. This enables the dynamic tracking of the fluctuating nuclear field gradient across the qubit, which is the main source of decoherence in ST_0 qubits in GaAs [31,32].

The propagation of probability distributions on the quantum controller, here according to the FP equation, can be replaced by other update equations, e.g., a transition matrix for Markov processes [33], or machine-learning-based methods for signal prediction [34], the details of which depend on the specific nature of the qubit system.

Real-time capabilities of quantum controllers can also be used advantageously to choose optimal measurement parameters within an *adaptive* estimation sequence (in our case updating free-induction-decay times on the fly, based on previous measurement outcomes), which we will analyze separately below.

Our scheme can accomplish Hamiltonian learning for intermittent calibration of circuit parameters, making it ideal for the recurrent and reliable execution of quantum circuits against the impact of drift. Since the interleaved estimation and qubit operation take place in the same qubit, there is an intricate interplay between the correlation time of the fluctuations being estimated, the efficiency of the estimation procedure, and the time required for coherent operations between estimations. Optimizing and managing these timescales will be essential when going from single-qubit devices to multiqubit devices, making it even more valuable to estimate qubit and noise correlations times quickly and efficiently.

II. RESULTS

A. Device and Bayesian estimation

We employ the top-gated GaAs double quantum dot (DQD) array from [35] with one of its ST_0 qubits activated using gate electrodes as in Fig. 1(a). A dilution refrigerator provides a base temperature below 50 mK and a 200-mT in-plane magnetic field defines the z-direction. A commercial digital-to-analog converter [36] (FPGA-powered quantum controller [37]) applies low-frequency (high-frequency) baseband waveforms to the gate electrodes, and radio-frequency reflectometry off one ohmic contact of the sensor dot distinguishes the charge configurations of the DQD, allowing single-shot qubit readout [38]. Details about the experimental setup can be found in Ref. [15].

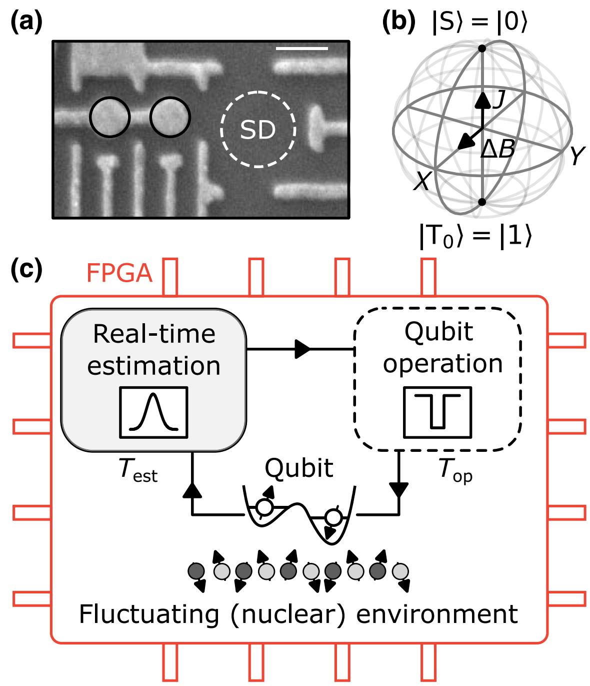

FIG. 1. Qubit implementation and estimation schedule. (a) Scanning electron micrograph of a GaAs double-dot device similar to the one used in this work [35] comprising a singlet-triplet qubit (black circles) next to a sensor dot (SD) used for qubit readout. Scale bar 100 nm. (b) Exchange coupling $J(\varepsilon)$ and Overhauser gradient $\Delta B \propto hf_B(t)$ drive rotations of the qubit around two orthogonal axes of the Bloch sphere, providing universal qubit control if the prevailing Overhauser frequency f_B can be estimated sufficiently efficiently. (c) Qubit schedule, alternating between periods T_{op} of quantum information processing (dashed box) and short periods T_{est} for efficiently learning the fluctuating environment (gray box).

The qubit operates in the $(1,1)$ and $(0,2)$ charge configuration, where the integers stand for the number of electrons in the left and right dot of the DQD. In the two-electron ST_0 basis, the Hamiltonian can be approximated in the regime of interest as

$$\mathcal{H}(t) = \frac{J(\varepsilon)}{2}\sigma_z + \frac{g^*\mu_B \Delta B(t)}{2}\sigma_x, \qquad (1)$$

where the σ_i represent the Pauli operators, g^* is the effective g-factor, and μ_B is the Bohr magneton. The energy $J(\varepsilon)$ characterizes the exchange interaction between the two electrons, which is tunable via the relative electrical detuning of the dots. By defining $\varepsilon = 0$ at the $(1,1)$–$(0,2)$ charge-state degeneracy, detuning is proportional to the difference in the effective onsite potentials on the two dots of the singlet-triplet qubit, where negative ε corresponds to the $(1,1)$ ground-state region. The field $\Delta B(t)$ denotes the z-component of the Overhauser gradient, which is the difference in effective magnetic fields on the two dots due to the hyperfine interaction of the electrons with approximately 10^5–10^6 spinful nuclei on each dot [31]. This gradient fluctuates slowly, and our goal is to efficiently estimate the corresponding Overhauser frequency $f_B(t) \equiv g^* \mu_B \Delta B(t)/h$ in real time on the quantum controller, using a physics-informed model with and without adaptive probe times.

A Bloch-sphere representation of the two contributions to $\mathcal{H}$ is sketched in Fig. 1(b). The qubit undergoes manipulation through voltage pulses applied to the plunger gates of the DQD, which effectively control the magnitude of $J(\varepsilon)$. Deep in the $(1,1)$ regime, where $J(\varepsilon) \ll |hf_B|$, the qubit is almost purely driven by the Overhauser gradient, whereas close to $\varepsilon = 0$ typically $J(\varepsilon) \gtrsim |hf_B|$.

After manipulation, the qubit is measured by projecting the unknown final spin state onto either the $(1,1)$ charge state ($|\mathrm{T}_0\rangle$) or the $(0,2)$ charge state ($|\mathrm{S}\rangle$), by tuning to positive ε. Each single-shot readout of the DQD charge configuration involves the generation, demodulation, and thresholding of a few-microsecond-long radio-frequency burst on the quantum controller [15].

The fluctuating frequency f_B is assessed on the quantum controller using a Bayesian estimation approach based on a series of N free-induction-decay experiments with evolution times t_i, where $i = 1, 2, \ldots, N$ [10–12,14–16]. Employing m_i to represent the outcome ($|\mathrm{S}\rangle$ or $|\mathrm{T}_0\rangle$) of the ith measurement, the likelihood function $P(m_i|f_B)$ is defined as the probability of obtaining m_i given a value of f_B,

$$P\left(m_i|f_B\right) = \frac{1}{2}\left[1 + m_i\left(\alpha + \beta\cos\left(2\pi f_B t_i\right)\right)\right], \quad (2)$$

where m_i takes a value of 1 (-1) if $m_i = |\mathrm{S}\rangle$ ($|\mathrm{T}_0\rangle$), and α and β are parameters accounting for the measurement error and axis of rotation on the Bloch sphere during a free-induction decay experiment [10]. In this work we use $\alpha = 0.28$ and $\beta = 0.45$ extracted from a series of separate free-induction decay (FID) experiments. Applying Bayes' rule to estimate f_B based on the series of measurements $m_N, \ldots, m_1$, which are assumed to be independent of each other, yields the final probability distribution $P_{\mathrm{final}}(f_B) \equiv P(f_B\,|m_N, \ldots, m_1)$ given by

$$P_{\mathrm{final}}\left(f_B\right) \propto P_0\left(f_B\right)\prod_{i=1}^{N}\left[1 + m_i\left(\alpha + \beta\cos\left(2\pi f_B t_i\right)\right)\right], \quad (3)$$

where $P_0(f_B)$ is the initial probability distribution assumed for f_B before the estimation starts. Equivalently, the measurement outcome m_i updates the Bayesian probability distribution according to $P_i(f_B) \propto P_{i-1}(f_B)P(m_i|f_B)$, up to a normalization factor, where the likelihood function $P(m_i|f_B)$ is given by Eq. (2). The final estimate of f_B is taken to be the expectation value $\langle f_B\rangle$, calculated over the final distribution $P_{\mathrm{final}}(f_B)$ after all N measurements have been performed. The estimation protocol can be repeated at user-defined times when the qubit is not in use for other operations.

Estimating low-frequency fluctuations is useful as outlined in the following example, depicted in Fig. 1(c). One starts by estimating the instantaneous magnitude of the slowly fluctuating field (the Overhauser frequency in our case), resulting in a strongly reduced uncertainty in this field. Subsequently, that knowledge is used to compensate for the random value of the field during coherent qubit operation, resulting in an increased qubit quality factor [15]. However, while operating the qubit for a period T_{op}, the field will again slowly drift, which can be captured by letting its distribution function evolve over time according to a known noise model [32]. For the Overhauser gradient, this amounts to a diffusion of its mean towards zero mean field and an increase of the uncertainty towards a maximum value that depends on the number and coupling strengths of the involved nuclear spins. Before such a stationary state is reached, the known dynamics of the probability distribution can be used to improve the feedback or make the next estimation more efficient. After a user-defined period T_{op}, qubit operations are momentarily halted and a new real-time estimation is initiated on the quantum controller. Its duration, approximately $T_{\mathrm{est}} \propto N$, depends on the desired estimation accuracy as discussed below. A series of estimation sequences, each resulting in an accurate distribution $P_{\mathrm{final}}(f_B)$, is what we refer to as qubit tracking.

B. Physics-informed tracking of the qubit frequency

This section describes how such "stroboscopic" physics-informed tracking of an Overhauser field is implemented on the quantum controller and to what extent it produces higher-quality estimates than obtainable via more commonly used estimation sequences [10,15]. The protocol is physics-informed in the sense that the assumed evolution of the distribution function in between two estimations is based on a physical model describing the nuclear spin dynamics in GaAs-based quantum dots.

The FPGA-based estimation of the Overhauser frequency f_B is illustrated in Fig. 2(a). One estimation sequence consists of N repetitions of an FID probe cycle. In each probe cycle, a singlet pair is initialized in $(0,2)$ and then detuned deeply into the $(1,1)$ region. At $\varepsilon \approx -40$ mV, the quantum controller lets the qubit evolve for a probe time $t_i = i\,t_0$, before thresholding the resulting qubit state and updating the probability distribution [$P_i(f_B) \propto P_{i-1}(f_B)P(m_i|f_B)$]. In this sequence, the probe times t_i are predetermined and linearly distributed by the probe time spacing $t_0 = 1$ ns. We assume that N is sufficiently small such that the Overhauser gradient remains constant during the sequence.

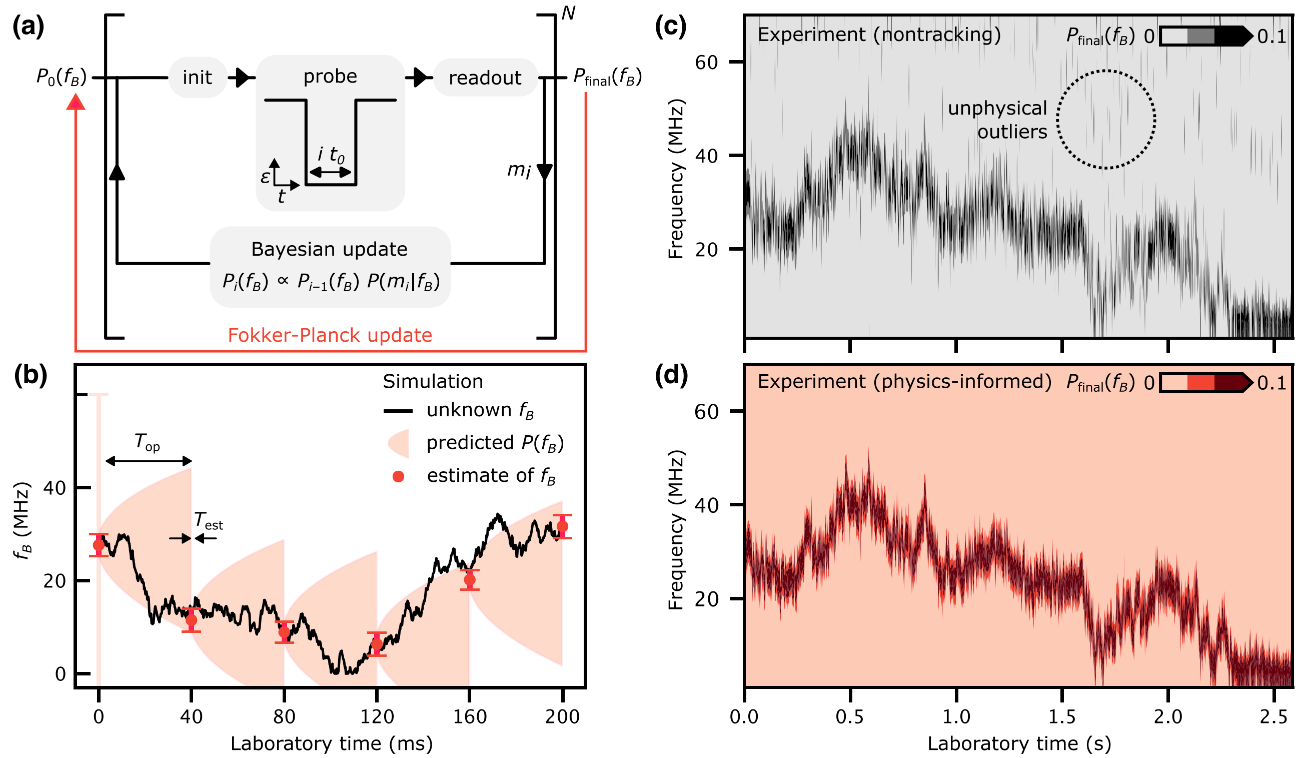

FIG. 2. Tracking the Overhauser frequency by anticipating nuclear spin diffusion on the quantum controller. (a) The physics-informed estimation sequence for f_B initializes the prior distribution $P_0(f_B)$ by evolving an older final distribution $P_{\rm final}(f_B)$ (Fokker-Planck update). For each of the N probe cycles, labeled i, the quantum controller initializes the qubit to the singlet state, performs an FID for time $t_i = it_0$, then updates the probability distribution $P_i(f_B)$ based on the measurement outcome m_i. After N probe cycles, the final distribution $P_{\rm final}(f_B)$ is saved. (b) Simulation of the unknown fluctuating Overhauser gradient (black) and five physics-informed estimation sequences, illustrating the tracking protocol. Every 40 ms, a sequence of FID probe cycles results in a final distribution with expected value $f_B^f = \langle f_B \rangle$ and error bar $2\sigma_f$ (red markers). The simulation assumes a uniform prior distribution $P_0(f_B)$ at $t = 0$, whereas subsequent priors $P_0(f_B)$ are based on the mean $\mu(t)$ and standard deviation $\sigma(t)$ propagated by the Fokker-Planck equation over period $T_{\rm op}$ (shaded in light red). (c) Experimental results for the nontracking reference protocol, using $P_0(f_B) \equiv P_{\rm uniform}(f_B)$ for each estimation sequence. (d) Experimental results for the physics-informed tracking protocol, obtained simultaneously with nontracking estimates in panel (c). The initial prior $P_0(f_B)$ for each column is $P_{\rm final}(f_B)$ from the previous column, propagated in time according to Eq. (5). Note the absence of multipeaked distributions $P_{\rm final}(f_B)$.

We model the dynamics of the Overhauser gradient as an Ornstein-Uhlenbeck (or drift–diffusion) process [32], driven by randomly occurring nuclear spin flips. The time dependence of the distribution function $P(f_B, t)$ resulting from such a process is governed by an FP equation [23,24,39], allowing the prediction of $P(f_B)$ in periods when the qubit is used for other operations ($T_{\rm op}$). Assuming that each final distribution $P_{\rm final}(f_B)$ is sufficiently characterized by its mean and variance, we instruct the quantum controller to approximate it by a Gaussian distribution [40]. Denoting the mean and variance immediately after estimation (time $t = 0$) as f_B^f and σ_f^2, respectively, the FP equation yields as solution for $t > 0$,

$$P(f_B, t) = \frac{1}{\sqrt{2\pi\sigma(t)^2}} \exp\left\{-\frac{[f_B - \mu(t)]^2}{2\sigma(t)^2}\right\}, \tag{4}$$

where

$$\mu(t) = f_B^f\, e^{-\Gamma t}, \tag{5a}$$

$$\sigma(t)^2 = \sigma_K^2 + \left[\sigma_f^2 - \sigma_K^2\right] e^{-2\Gamma t}. \tag{5b}$$

Here, σ_K is the steady-state root-mean-square value of the Overhauser field frequency (typically around 30–50 MHz [32]), while Γ reflects the slow relaxation rate of nuclear spin polarization (measured to be $\Gamma \approx 1.1$ Hz from autocorrelation). Notably, the inverse of Γ, denoted by $T_{\rm c} = \Gamma^{-1} \approx 0.91$ s, defines the timescale for the correlation of fluctuations in f_B; this establishes the time window within which an estimate of f_B is expected to remain useful.

In Fig. 2(b) we numerically simulate a fluctuating Overhauser gradient with $T_{\rm c} = 1$ s and $\sigma_K = 30$ MHz. The associated unknown frequency f_B (black trace) is assumed to be estimated every $T_{\rm op} = 40$ ms (red markers). The physics-informed evolution of probability distributions (shaded red areas, adapted from Ref. [24]) captures two properties expected for nuclear spin diffusion, namely the inclination of the average of the Overhauser gradient to drift back towards zero [Eq. (5a)], and a progressive expansion of the uncertainty in the gradient towards σ_K [Eq. (5b)]. Both processes take place on a timescale of T_c.

Initially, no knowledge of f_B is available, reflected by a uniform prior distribution $P_{\text{uniform}}(f_B)$ at $t = 0$ represented by the semitransparent error bar spanning the entire frequency range of the simulation (60 MHz) [41]. A number of FID cycles are performed until the updated probability distribution has a fitted $\sigma < 2$ MHz. This estimation sequence is assumed to take only a few hundred microseconds (i.e., much shorter than T_{op} and T_c), and we only plot the mean f_B^f and 95% confidence interval of the final distribution $P_{\text{final}}(f_B)$ (first red marker).

After the first estimation sequence, the Overhauser fields are left to evolve freely for T_{op}. During this time, the distribution function is assumed to be Gaussian; the time dependence of its mean and variance is given by Eq. (5). The evolution of the 95% confidence interval is indicated by the red shaded area in Fig. 2(b). At the end of T_{op} ($t = 40$ ms) the associated Gaussian distribution is characterized by $\mu(T_{\text{op}})$ and $\sigma(T_{\text{op}})$, and serves as the initial prior distribution for the next estimation sequence. Similarly, estimations at $t = 80$, 120, 160, and 200 ms use as prior the most recent Gaussian.

If T_{op} is smaller than T_c, the physics-informed $P_0(f_B)$ will remain somewhat constrained, providing a better prior compared to a uniform distribution and potentially requiring fewer FID experiments for a more accurate estimate of f_B. If T_{op} becomes comparable to or larger than T_c, prior knowledge about f_B becomes irrelevant and is not expected to improve the next estimation.

To experimentally test the benefits of physic-informed priors, we define a nontracking estimation scheme that always sets the initial distribution $P_0(f_B)$ to a uniform distribution $P_{\text{uniform}}(f_B)$ between 1 MHz and 70 MHz with 1 MHz resolution. Thus, as in previous works [10,15], each estimation sequence does not retain any memory of previous estimations. In parallel to this nontracking estimation, we instruct the quantum controller to also generate estimates based on the physics-informed initialization of $P_0(f_B)$, thereby improving the estimation accuracy as quantified below.

Figure 2(c) plots 1000 final probability distributions of the nontracking scheme, acquired over a span of 2.6 s using an $N = 31$ schedule with $T_{\text{est}} = 0.6$ ms and $T_{\text{op}} = 2$ ms. Specifically, each FID probe cycle lasts 20 μs, of which 5 μs is dedicated to qubit readout, 2.6 μs to initialize the qubit and discharge the bias tee with a zero-averaging pulse, and the remaining time is used to update the nontracking and physics-informed distributions $P_i(f_B)$ on the FPGA. Several estimation sequences result in a multipeaked probability distribution, with secondary peaks that randomly jump from one column to another. In simulations, such "outliers" also appear in the absence of measurement errors and appear to be a shortcoming of the algorithm, not an artifact of the device or the quantum controller. The known correlation time of the Overhauser field dynamics makes it improbable that the sudden jumps of the outliers represent the actual Overhauser field gradient, and similar jumps in previous work were associated with compromised qubit quality factors [cf. discussion of Fig. 2(b) of Ref. [15] in its supplemental material].

Figure 2(d) shows the physics-informed estimates $P_{\text{final}}(f_B)$, acquired concurrently with the nontracking estimates in Fig. 2(c). Strikingly, multipeaked probability distributions are absent, suggesting that the physics-informed model on the quantum controller suppresses unphysical jumps of the estimated Overhauser gradient (here with $T_c = 0.91$ s and $\sigma_K = 50$ MHz). By extracting the standard deviation from each column in Figure 2(d), we find that its average is reduced relative to the average standard deviation extracted from Fig. 2(c), suggesting an improved estimation accuracy.

Figure 3 compares the performance of the nontracking and physics-informed estimation sequences as a function of the number of FID probe cycles, for different choices of T_{op}. Each data point corresponds to an independent experiment comprising 10,000 repetitions of an estimation sequence. The plotted uncertainty is defined as the average standard deviation of the final probability distribution $P_{\text{final}}(f_B)$ of each of the 10,000 estimations. The shaded areas indicate the standard deviation of the associated 10,000 standard deviations. In our experiment, the true value of the real field, and thus the actual error in the estimation, is unknown, and therefore we rely on the uncertainty measure plotted as a reasonable metric. Indeed, low uncertainties at the end of T_{est} correlate with increased quality factors of controlled Overhauser rotations during T_{op} (see Supplemental Material [42]).

The uncertainty of the nontracking estimates in Fig. 3(a) does not depend on T_{op}. This is expected, as the prior distributions $P_0(f_B)$ in the nontracking scheme are always the uniform distribution $P_{\text{uniform}}(f_B)$, with no memory of the previous estimates. In contrast, the uncertainty of the physics-informed estimates decreases with decreasing T_{op}, for fixed number of measurements in the estimation sequence. This suggests that a narrower prior yields a more accurate estimate.

Remarkably, with as few as 10 probes the physics-informed estimates for $T_{\text{op}} = 1$ ms are more accurate than nontracking estimates based on 100 probes [in Fig. 3(a) the uncertainties are approximately 3 MHz and 5 MHz, respectively]. With increasing number of probe cycles, the uncertainty of nontracking estimates saturates near

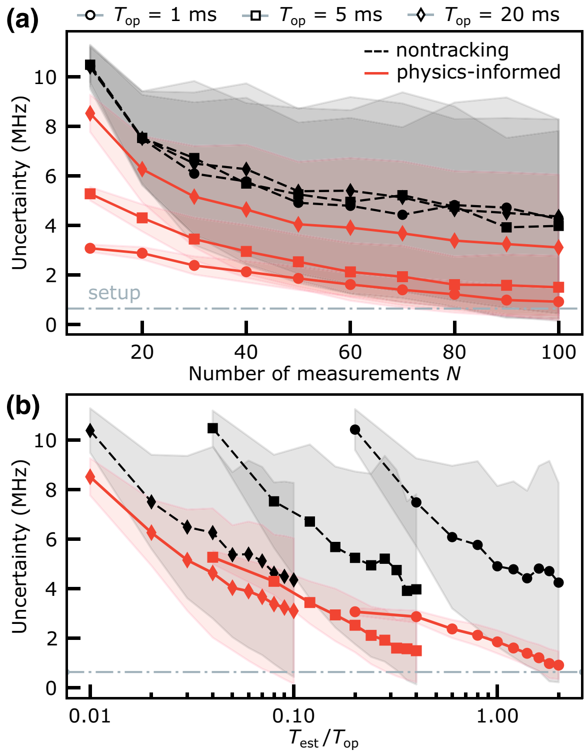

FIG. 3. Efficiency of the nontracking and physics-informed protocols. (a) Estimation uncertainty as a function of the number of FID probes in the estimation sequence, for the nontracking (black) and physics-informed (red) protocols. Symbols denote the average standard deviation of 10,000 $\langle f_B \rangle$ values, whereas shaded regions show their standard deviation, for different choices of operation time. (b) Uncertainty from (a) plotted as a function of the ratio T_{est}/T_{op}, where the estimation time is $T_{est} = N \cdot 20\,\mu$s. The dash-dotted gray line indicates the resolution limit imposed by our setup, see the main text.

5 MHz, whereas the physics-informed estimation uncertainty approaches the limitation imposed by our choice of frequency binning (0.8 MHz [43]).

The tradeoff between "qubit duty cycle" (T_{op}/T_{est}) and estimation accuracy is evident in Fig. 3(b). Here, we replot the uncertainties from (a) as a function of the estimation time $T_{est} = N \cdot 20\,\mu$s, where N is the number of qubit probes and 20 μs is the probe cycle duration. Depending on the desired Hamiltonian uncertainty, a maximum operation limit T_{op} and a significant qubit downtime (high T_{est}/T_{op} ratio) for estimation must be tolerated. The optimum choice of N depends on details of the noise spectrum and the estimation efficiency [24].

One may be tempted to pursue the lowest possible uncertainty while estimating the environmental fluctuations, but the operational benefits will depend on details such as the tolerable estimation uncertainty for a certain application and how long it is expected to survive given a specific environment. Because achieving lower uncertainties in general requires more qubit down time for estimation, quantum information processing applications may need to define a tolerated "error budget", which translates into a useful operation time T_{op} depending on the correlation time of the fluctuations T_c and a minimized estimation time T_{op} depending on the efficiency of the protocol.

So far, we have demonstrated an improved Hamiltonian learning protocol that tracks a slowly fluctuating environmental parameter, by instructing a quantum controller to generate physics-informed priors in real time. Next, we instruct the controller to adaptively choose the probe times, thereby reducing the length of the estimation sequences.

C. Adaptive Bayesian tracking of the qubit frequency

For the purpose of only monitoring fluctuating Hamiltonian parameters without interspersed qubit operation, nonadaptive Bayesian estimation is straightforward to execute because it does not require real-time feedback and could even be carried out *a posteriori*. However, numerical studies [17,18,20,21,23,24] suggest the beneficial use of adaptive estimation sequences in which the probe times t_i are chosen based on previous measurement outcomes, as experimentally realized in nitrogen-vacancy centers [25,26].

Previous experiments with gate-defined spin qubits employed nontracking and nonadaptive FID-based Bayesian estimation to probe the qubit frequency [10,15]. In this section, we supplement the generation of physics-informed time-evolved priors by the generation of adaptive probe times in real time, thereby reducing the number of required probes and showing a path towards much shorter estimation sequences.

Figure 4(a) illustrates the key difference of the adaptive estimation sequence, relative to that in Fig. 2(a): the free-evolution time t_i for the ith FID probe now depends on the previous Bayesian update as

$$t_i = \frac{1}{c\sigma_{i-1}}, \tag{6}$$

where σ_{i-1} is the standard deviation of the Gaussian-approximated probability distribution $P_{i-1}(f_B)$, except σ_0, which is the standard deviation of prior $P_0(f_B)$ based on the FP equation. The optimal numerical prefactor c is expected to depend on the experimental setup [21]. Intuitively, this choice for the free evolution times can be motivated by our desire that two oscillations with frequencies that differ by Δf develop a phase shift of π after time $t = 1/(2\Delta f)$. In other words, Eq. (6) maps a frequency range of width $c\sigma_{i-1}/2$ to a large phase contrast in the likelihood function.

Implementation of the estimation protocol of Fig. 4(a) on the quantum controller yields reliable estimates for f_B from only 10 probes per sequence, as shown in Fig. 4(b) for $T_{op} = 1$ ms and $c \approx 13$ [44]. This example demonstrates the estimation of a slowly fluctuating qubit frequency

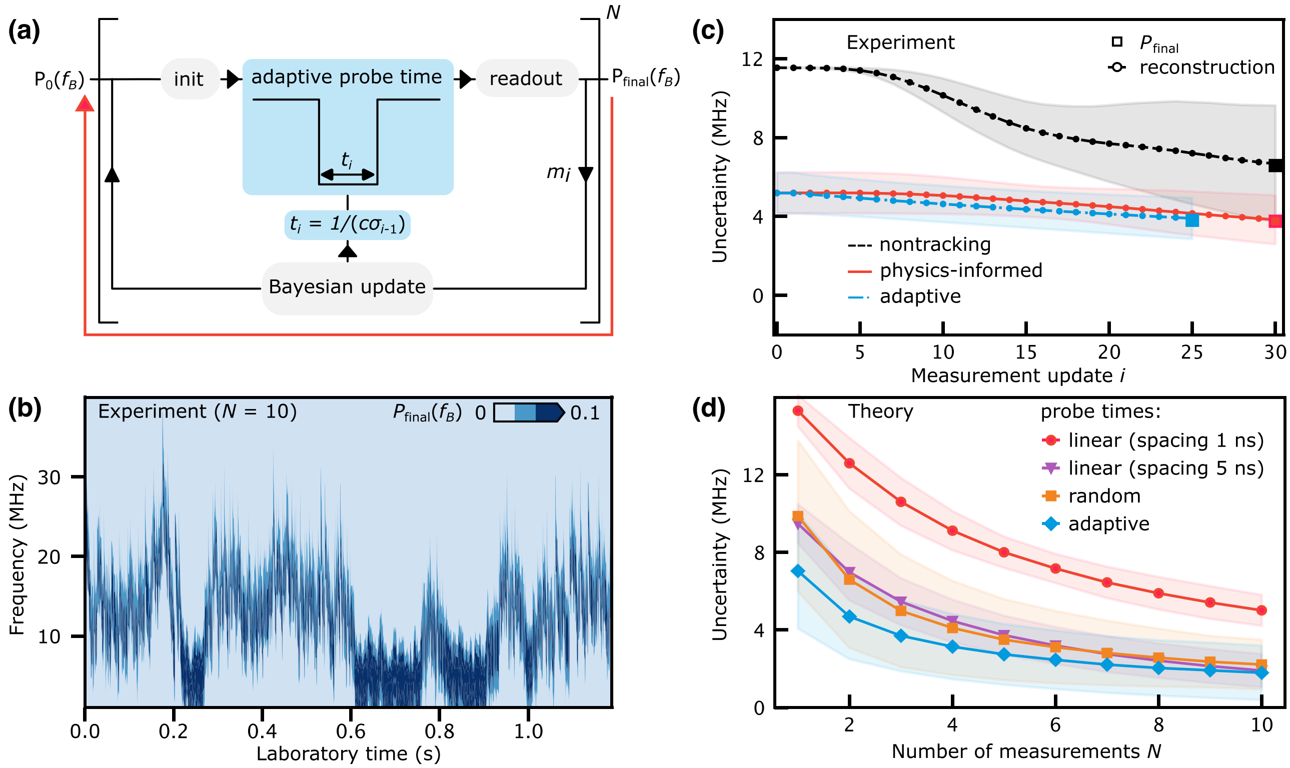

FIG. 4. Adaptive Bayesian tracking by real-time choice of qubit probe times. (a) In this adaptive Bayesian estimation sequence, probe times t_i are chosen based on the standard deviation σ_{i-1} of the previous Bayesian distribution. $P_0(f_B)$ is initialized based on the FP equation. (b) Adaptive tracking obtained from short estimation sequences ($N = 10$) for $T_{\text{op}} = 1$ ms. (c) Reconstructed uncertainty in the distribution function within an estimation sequence (defined in the text) as a function of the measurement update m_i. Squares at the end of the curves correspond to the experimental posterior distributions computed on the quantum controller. (d) Simulated uncertainty expected at the end of a short estimation sequence ($N \leq 10$) for different probe time protocols, including evenly distributed t_i (probe time spacing of 1 or 5 ns), adaptive probe times, and random probe times (see the main text). The initial prior distributions are assumed to be determined from the FP equation.

within 200 µs , which is one order of magnitude shorter and with better accuracy than previously reported [15]. Here, $c \approx 13$ was chosen empirically, and further improvements may be possible by better choices informed from numerical simulations; see the Supplemental Material [42].

Outliers appear to be absent both for the physics-informed [Fig. 2(c)] and adaptive tracking [Fig. 4(b)], likely for similar reasons, motivating a quantitative comparison based on experimental data and theoretical insights.

Figure 4(c) compares average uncertainties, inferred from experimental data in Fig. S2 of the Supplemental Material [42]. We choose $T_{\text{op}} = 5$ ms and perform 10,000 repetitions of three protocols, focusing on $N \leq 30$ to test whether short sequences benefit from adaptive probe cycles. The three squares at the end of the curves show the uncertainties σ (defined as in Fig. 3 and computed on the quantum controller from the posterior distributions P_{final}) for nontracking (black), physics-informed (red), and adaptive (blue) estimation sequences. For $N = 30$, the nontracking scheme yields an average $\sigma \approx 7.3$ MHz, while the physics-informed scheme yields $\sigma \approx 3.5$ MHz. The uncertainty of the adaptive scheme is similar, though obtained with fewer probes ($N = 25$).

To investigate how each probe cycle contributes information gain, we analyze how uncertainties evolve within a sequence (additional details can be found in Fig. S2 [42]). Specifically, we reconstruct the Bayesian probability updates $P_i(f_B)$ from our record of raw single-shot measurement outcomes m_i [45]. For each i, we plot the standard deviation of $P_i(f_B)$ (reconstruction), averaged over all 10,000 repetitions, as well as their standard deviation (shaded areas).

The nontracking method is clearly outperformed by the physics-informed and adaptive schemes. This is expected, as both the physics-informed and adaptive protocols use physics-informed prior distributions. Furthermore, the adaptive scheme has consistently lower uncertainty than the physics-informed scheme, though only marginally. Finally, we note that the uncertainties for the nontracking and physics-informed schemes barely decrease during the first few measurements ($i \lesssim 5$), as shown by the nearly flat curves in this range. In contrast, the adaptive scheme shows a negative slope already for the first measurement outcomes, indicating information gain and a narrowing of the probability distribution.

To explore the ultimate estimation efficiencies that can be expected for our spin-qubit system, unconstrained by coarse frequency binning and limited memory on the FPGA-powered controller, we now turn towards simulated Overhauser fluctuations, assumed to follow an Ornstein-Uhlenbeck process with $T_c = 1$ s and $\sigma_K = 40$ MHz, and simulate estimation sequences on a much finer and larger frequency grid (0–150 MHz with 0.25 MHz bin size) than currently possible in our experimental setup.

Figure 4(d) shows the resulting uncertainties and their standard deviations, assuming $T_{\mathrm{op}} = 5$ ms, for different distributions of probe times (see Figs. S3 and S4 [42] for further details). In the sequences with "linear" probe times, $t_i = i\, t_0$, we observe that the choice of the probe time spacing t_0 (1 and 5 ns are shown) has a drastic influence on the resulting accuracy. In the sequences with "random" probe times, t_i is randomly chosen from a uniform distribution between 1 ns and 50 ns. In the sequences with "adaptive" probe times, $t_i = 1/(c\sigma_{i-1})$, now with $c = 6$ and without rounding t_i to the temporal granularity of the quantum controller (see the Supplemental Material [42]).

The adaptive-probe-time sequence outperforms the linear sampling approach with $t_0 = 1$ ns, yielding uncertainties that are on average smaller by a factor of approximately 2.7, and is also superior to $t_0 = 5$ ns and random probe times, resulting in approximately 30% smaller uncertainties for short estimation sequences ($N \lesssim 5$). We therefore believe that adaptive estimation sequences will become crucial in applications that only permit a small number of probe cycles.

In summary, the results shown in Fig. 4 present an adaptive Bayesian estimation scheme implemented in a semiconductor-defined spin qubit.

The real-time capabilities of the quantum controller enable probe times t_i to be updated based on previous measurement outcomes $m_{i-1}, m_{i-2}, \ldots, m_1$, resulting in a small but measurable improvement compared to linearly spaced probe times. Our approach is substantiated by numerical simulations, indicating that high-quality estimates of the qubit frequency achieving only a few percent error (approximately 3 MHz uncertainty with a simulated dynamic range of about 150 MHz) should be possible with fewer than five qubit probe cycles.

III. OUTLOOK

We have implemented physics-informed and adaptive estimation sequences that allowed the efficient tracking of low-frequency fluctuations in a solid-state qubit. A quantum controller estimates in real time the uncontrolled magnetic field fluctuations in a gallium arsenide singlet-triplet spin qubit, yielding improved accuracy by temporally evolving a sufficiently recent probability distribution according to the FP equation. In addition, the adaptive choice of qubit probe times, based on the standard deviation of the updated probability distribution, allows for significantly shorter estimation sequences yielding similar or reduced uncertainties. Compared to previous experiments [15], this work extends the estimation bandwidth from a few hundred hertz to approximately 2.5 kHz, due a tenfold reduction of the estimation time and a reduced uncertainty.

While our work presents real-time adaptive tracking of a semiconductor spin qubit, determining optimal protocols compatible with constraints of the control hardware and application requirements remains an open question. We anticipate further progress by research that combines theoretical and hardware aspects.

Possibly useful modifications of the protocol could relax the assumption of single-shot readout [46] or mitigate state preparation and measurement errors by duplication of probe cycles [17,19]. Probe times can further be optimized by also taking into account the estimated qubit frequency, not just its uncertainty, and possibly it is advantageous to terminate an estimation sequence when reaching an accuracy target, rather than a predetermined length.

Fault-tolerant quantum computing based on quantum error correction will likely require qubits that are affected by limited amounts of Markovian noise. Therefore, real-time frequency tracking protocols may become important tools, as they suppress non-Markovian noise [16].

By properly modifying the tracking equation relevant to the specific noise source, this work offers an efficient, physics-informed, and adaptive Hamiltonian learning protocol for real-time estimation of low-frequency noise in solid-state qubits.

ACKNOWLEDGMENTS

We thank Torbjørn Rasmussen for help with the experimental setup. This work received funding from the European Union's Horizon 2020 research and innovation programme under grant agreements 101017733 (QuantERA II) and 951852 (QLSI), from the Novo Nordisk Foundation under Challenge Programme NNF20OC0060019 (SolidQ), from the Inge Lehmann Programme of the Independent Research Fund Denmark, from the Research Council of Norway (RCN) under INTFELLES-Project No 333990, as well as from the Dutch National Growth Fund (NGF) as part of the Quantum Delta NL program.

F.B. led the measurements and data analysis, and wrote the manuscript with input from all authors. F.B., J.v.d.H., A.C., and F.K. performed the experiment with theoretical contributions from J.A.K., J.B., J.D., and E.v.N. F.F. fabricated the device. S.F., G.C.G., and M.J.M. supplied the heterostructures. A.C. and F.K. supervised the project.

[1] L. Chirolli and G. Burkard, Decoherence in solid-state qubits, Adv. Phys. **57**, 225 (2008).
[2] J. Bergli, Y. M. Galperin, and B. L. Altshuler, Decoherence in qubits due to low-frequency noise, New J. Phys. **11**, 025002 (2009).
[3] G. Falci, P. J. Hakonen, and E. Paladino, in *Encyclopedia of Condensed Matter Physics*, edited by T. Chakraborty (Academic Press, 2024), 2nd ed., p. 1003.
[4] B. M. Terhal, Quantum error correction for quantum memories, Rev. Mod. Phys. **87**, 307 (2015).
[5] L. Viola, E. Knill, and S. Lloyd, Dynamical decoupling of open quantum systems, Phys. Rev. Lett. **82**, 2417 (1999).
[6] P. Szańkowski, G. Ramon, J. Krzywda, D. Kwiatkowski, and Ł. Cywiński, Environmental noise spectroscopy with qubits subjected to dynamical decoupling, J. Phys.: Condens. Matter **29**, 333001 (2017).
[7] S. Foletti, H. Bluhm, D. Mahalu, V. Umansky, and A. Yacoby, Universal quantum control of two-electron spin quantum bits using dynamic nuclear polarization, Nat. Phys. **5**, 903 (2009).
[8] P. London, J. Scheuer, J.-M. Cai, I. Schwarz, A. Retzker, M. B. Plenio, M. Katagiri, T. Teraji, S. Koizumi, J. Isoya, R. Fischer, L. P. McGuinness, B. Naydenov, and F. Jelezko, Detecting and polarizing nuclear spins with double resonance on a single electron spin, Phys. Rev. Lett. **111**, 067601 (2013).
[9] J. Scheuer, I. Schwartz, Q. Chen, D. Schulze-Sünninghausen, P. Carl, P. Höfer, A. Retzker, H. Sumiya, J. Isoya, B. Luy, M. B. Plenio, B. Naydenov, and F. Jelezko, Optically induced dynamic nuclear spin polarisation in diamond, New J. Phys. **18**, 013040 (2016).
[10] M. D. Shulman, S. P. Harvey, J. M. Nichol, S. D. Bartlett, A. C. Doherty, V. Umansky, and A. Yacoby, Suppressing qubit dephasing using real-time Hamiltonian estimation, Nat. Commun. **5**, 5156 (2014).
[11] T. Nakajima, A. Noiri, K. Kawasaki, J. Yoneda, P. Stano, S. Amaha, T. Otsuka, K. Takeda, M. R. Delbecq, G. Allison, A. Ludwig, A. D. Wieck, D. Loss, and S. Tarucha, Coherence of a driven electron spin qubit actively decoupled from quasistatic noise, Phys. Rev. X **10**, 011060 (2020).
[12] J. Kim, J. Yun, W. Jang, H. Jang, J. Park, Y. Song, M.-K. Cho, S. Sim, H. Sohn, H. Jung, V. Umansky, and D. Kim, Approaching ideal visibility in singlet-triplet qubit operations using energy-selective tunneling-based Hamiltonian estimation, Phys. Rev. Lett. **129**, 040501 (2022).
[13] A. Vepsäläinen, R. Winik, A. H. Karamlou, J. Braumüller, A. D. Paolo, Y. Sung, B. Kannan, M. Kjaergaard, D. K. Kim, A. J. Melville, B. M. Niedzielski, J. L. Yoder, S. Gustavsson, and W. D. Oliver, Improving qubit coherence using closed-loop feedback, Nat. Commun. **13**, 1932 (2022).
[14] J. Yun, J. Park, H. Jang, J. Kim, W. Jang, Y. Song, M.-K. Cho, H. Sohn, H. Jung, V. Umansky, and D. Kim, Probing two-qubit capacitive interactions beyond bilinear regime using dual Hamiltonian parameter estimations, npj Quantum Inf. **9**, 30 (2023).
[15] 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).
[16] 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, arXiv:2403.02666.
[17] 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).
[18] P. Cappellaro, Spin-bath narrowing with adaptive parameter estimation, Phys. Rev. A **85**, 030301(R) (2012).
[19] 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).
[20] C. Bonato and D. W. Berry, Adaptive tracking of a time-varying field with a quantum sensor, Phys. Rev. A **95**, 052348 (2017).
[21] E. Scerri, E. M. Gauger, and C. Bonato, Extending qubit coherence by adaptive quantum environment learning, New J. Phys. **22**, 035002 (2020).
[22] A. Gutierrez-Rubio, P. Stano, and D. Loss, Optimal frequency estimation and its application to quantum dots, arXiv:2004.12049.
[23] K. Craigie, E. M. Gauger, Y. Altmann, and C. Bonato, Resource-efficient adaptive Bayesian tracking of magnetic fields with a quantum sensor, J. Phys.: Condens. Matter **33**, 195801 (2021).
[24] J. Benestad, J. A. Krzywda, E. van Nieuwenburg, and J. Danon, Efficient adaptive Bayesian estimation of a slowly fluctuating overhauser field gradient, arXiv:2309.15014.
[25] C. Bonato, M. S. Blok, H. T. Dinani, D. W. Berry, M. L. Markham, D. J. Twitchen, and R. Hanson, Optimized quantum sensing with a single electron spin using real-time adaptive measurements, Nat. Nanotechnol. **11**, 247 (2015).
[26] 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).
[27] R. Schirhagl, K. Chang, M. Loretz, and C. L. Degen, Nitrogen-vacancy centers in diamond: Nanoscale sensors for physics and biology, Annu. Rev. Phys. Chem. **65**, 83 (2014).
[28] G. Burkard, T. D. Ladd, A. Pan, J. M. Nichol, and J. R. Petta, Semiconductor spin qubits, Rev. Mod. Phys. **95**, 025003 (2023).
[29] G. Feher, J. P. Gordon, E. Buehler, E. A. Gere, and C. D. Thurmond, Spontaneous emission of radiation from an electron spin system, Phys. Rev. **109**, 221 (1958).

[30] J. P. Gordon and K. D. Bowers, Microwave spin echoes from donor electrons in silicon, Phys. Rev. Lett. **1**, 368 (1958).
[31] J. M. Taylor, J. R. Petta, A. C. Johnson, A. Yacoby, C. M. Marcus, and M. D. Lukin, Relaxation, dephasing, and quantum control of electron spins in double quantum dots, Phys. Rev. B **76**, 035315 (2007).
[32] F. K. Malinowski, F. Martins, Ł. Cywiński, M. S. Rudner, P. D. Nissen, S. Fallahi, G. C. Gardner, M. J. Manfra, C. M. Marcus, and F. Kuemmeth, Spectrum of the nuclear environment for GaAs spin qubits, Phys. Rev. Lett. **118**, 177702 (2017).
[33] S. Asmussen, *Applied Probability and Queues*, 2nd ed., Stochastic Modelling and Applied Probability, vol. 51, Springer, New York, 2003.
[34] R. S. Gupta and M. J. Biercuk, Machine learning for predictive estimation of qubit dynamics subject to dephasing, Phys. Rev. Appl. **9**, 064042 (2018).
[35] F. Fedele, A. Chatterjee, S. Fallahi, G. C. Gardner, M. J. Manfra, and F. Kuemmeth, Simultaneous operations in a two-dimensional array of singlet-triplet qubits, PRX Quantum **2**, 040306 (2021).
[36] QDevil model QDAC-II, www.quantum-machines.co.
[37] Quantum Machines model OPX+, www.quantum-machines.co.
[38] F. Vigneau, F. Fedele, A. Chatterjee, D. Reilly, F. Kuemmeth, M. F. Gonzalez-Zalba, E. Laird, and N. Ares, Probing quantum devices with radio-frequency reflectometry, Appl. Phys. Rev. **10**, 021305 (2023).
[39] N. G. Van Kampen, *Stochastic Processes in Physics and Chemistry* (North-Holland, Amsterdam, 1990).
[40] Since we only work with positive frequencies f_B, one should be aware that this becomes inaccurate for distributions that have significant weight close to $f_B = 0$, i.e., that have a variance larger than the square of the mean.
[41] We choose our prior distribution to be nonzero for positive frequencies only, resulting in a unimodal final distribution. As the sign of the Overhauser gradient is unknown, the true final distribution would always be symmetric around zero.
[42] See Supplemental Material at http://link.aps.org/supplemental/10.1103/PhysRevApplied.22.014033, which includes Ref. [15], for numerical simulations of different Bayesian estimation schemes, and a detailed analysis of the experimental setup limitations.
[43] We programmed the frequency resolution on the quantum controller to be 1 MHz. The associated minimum standard deviation of $P_{\text{final}}(f_B)$ calculated on the FPGA is approximately 0.8 MHz.
[44] Due to the numerical precision of the quantum controller and the discreteness of the t_i that can be implemented, the actual ratio between $1/\sigma_{i-1}$ and t_i varies slightly between FID probe cycles.
[45] To increase the estimation bandwidth within the FPGA memory constraints, the quantum controller overwrites Bayesian updates $P_i(f_B)$ and only records $P_{\text{final}}(f_B)$ at the end of each sequence, as well as all N measurement outcomes.
[46] H. T. Dinani, D. W. Berry, R. Gonzalez, J. R. Maze, and C. Bonato, Bayesian estimation for quantum sensing in the absence of single-shot detection, Phys. Rev. B **99**, 125413 (2019).