Two-Level-System Dynamics in a Superconducting Qubit Due to Background Ionizing Radiation

Ted Thorbeck,[1,*] Andrew Eddins,[2] Isaac Lauer,[1] Douglas T. McClure,[1] and Malcolm Carroll[1]

[1] *IBM Quantum, IBM T. J. Watson Research Center, Yorktown Heights, New York 10598, USA*

[2] *IBM Quantum, MIT-IBM Watson AI Laboratory, Cambridge, Massachusetts 02142, USA*

(Received 8 November 2022; accepted 16 February 2023; published 29 June 2023)

Superconducting qubit lifetimes must be both long and stable to provide an adequate foundation for quantum computing. This stability is imperiled by two-level systems (TLSs), currently a dominant loss mechanism, which exhibit slow spectral dynamics that destabilize qubit lifetimes on hour time scales. Stability is also threatened at millisecond time scales, where ionizing radiation has recently been found to cause bursts of correlated multiqubit decays, complicating quantum error correction. Here, we study both ionizing radiation and TLS dynamics on a 27-qubit processor, repurposing the standard transmon qubits as sensors of both radiation impacts and TLS dynamics. Unlike prior literature, we observe resilience of the qubit lifetimes to the transient quasiparticles generated by the impact of radiation. However, we also observe a new interaction between these two processes, "TLS scrambling," in which a radiation impact causes multiple TLSs to jump in frequency, which we suggest is due to the same charge rearrangement sensed by qubits near a radiation impact. As TLS scrambling brings TLSs out of or into resonance with the qubit, the lifetime of the qubit increases or decreases. Our findings thus identify radiation as a new contribution to fluctuations in qubit lifetimes, with implications for efforts to characterize and improve device stability.

DOI: 10.1103/PRXQuantum.4.020356

I. INTRODUCTION

The drive to build a superconducting quantum computer has led to rapid increases in both the number of qubits in a device and the lifetimes, T_1, of those qubits. The increase in lifetimes has been remarkable given the sensitivity of the qubits to environmental noise [1]. This sensitivity can be harnessed by using the qubits as sensors to better understand the noise. For example, two-level systems (TLSs) are currently a dominant loss mechanism in superconducting qubits but superconducting qubits are also useful to study TLSs [2]. A key diagnostic tool has been TLS spectroscopy, in which T_1 is measured as the frequency of the qubit is swept [3–7]. An individual TLS can be resolved as a dip in T_1 as the qubit is tuned to the TLS frequency. Spectroscopy has revealed that TLSs can drift, appear, and disappear over time [6–8]. When a TLS moves into resonance with a qubit, T_1 can be suppressed by an order of magnitude in a modern device [7,8]. This instability in T_1 is a threat to quantum computers: when lifetimes decrease, either the computation suffers or the device must be taken off line and retuned, potentially by changing the qubit frequency or by relearning a noise model for error mitigation [9]. Therefore understanding TLS dynamics is important for improving superconducting quantum processors. Prior work on TLS dynamics has focused on the interactions between TLSs [10,11]. In this model, each high-frequency TLS (i.e., TLSs that are near resonant with the qubit frequency) is coupled to many low-frequency TLSs that occasionally flip states due to the ambient thermal energy, perturbing the high-frequency TLS. Evidence for this model has been observed in both superconducting qubits [6,7,12–14] and resonators [11,15–17]. However, TLS-TLS interactions, via either electric or elastic dipole interactions, are very short range [12,18] compared to the large size of the transmon capacitor paddles (approximately 100 μm) or the Josephson junction (approximately 100 nm), so it is unlikely that multiple high-frequency TLSs would interact with the same set of low-frequency TLSs. Therefore, the interacting-defect model would struggle to explain simultaneous dynamics in multiple high-frequency TLSs.

Improvements in T_1 have helped to reveal previously unobservable decoherence mechanisms, such as the impact of ionizing radiation [19–21]. Ionizing radiation generates

*ted.thorbeck@ibm.com

quasiparticles (QPs) that poison the device, creating a transient dip in T_1, but in the literature [19–23], the magnitude and duration of the dip vary widely. On multiqubit devices, the QPs can poison many qubits at the same time, potentially the entire chip, generating correlated errors that quantum error-correction algorithms struggle to correct [22–25]. Unfortunately, background radiation is ubiquitous. γ rays and cosmic ray muons cannot be easily shielded and small amounts of α and β radiation sources could even lurk inside the device packaging. In one recent experiment, background radiation, which hit the chip on average every 10 s, suppressed T_1 to less than 1 μs, and rendered the entire chip unusable for about 100 ms [23,26]. Therefore drastic steps have been proposed to mitigate these potentially catastrophic events, such as moving the experiments deep underground [21], adding potentially lossy QP traps [24,27–32], and distributing the quantum information across multiple chips [33,34].

In this paper, we repurpose the standard qubits of a 27-qubit IBM Quantum Falcon R6 processor as multiphysics sensors—electrometers to detect the impact of radiation and spectrometers to monitor TLS dynamics—to study the effect of radiation not only on the qubits but also on the TLS. In Sec. II, we report localized multiqubit offset-charge jumps consistent with the impact of radiation. In contrast to previous reports [22–24], in Sec. III we observe minimal reduction in T_1 during impact, showing the potential for superconducting qubits to be robust against ionizing radiation. No special measures are taken to shield the device or to mitigate QPs, suggesting that a combination of materials, packaging, and architecture determines the susceptibility of the qubit to radiation. However, we sometimes observe an unexpected long-lasting change in T_1 after the radiation impact. In Sec. IV, we show that an impact can scramble the TLS spectrum, by which we mean that several TLSs that are near resonant with the qubit appear, disappear, or change frequency at the same time. Because scrambling involves multiple high-frequency TLSs, the interacting-defect model is insufficient to explain these dynamics. In Sec. V, we suggest that the offset-charge jumps and TLS scrambling can both be explained by charge rearrangement after the radiation impact event. We thus identify a new mechanism contributing to fluctuations in T_1 over time, defining a new focus for research to improve stability of superconducting quantum processors.

II. CHARGE DETECTION OF RADIATION IMPACT EVENTS

First, we briefly review the dynamics of radiation impacting a superconducting qubit chip as described in the literature [19,22–24,35]. We do not deliberately introduce any radioactive sources [19,21], so we consider only background radiation. Multiple layers of shielding protect our device from low-energy photons such as thermal radiation from the higher-temperature stages [36–38] and also attenuate any α or β radiation coming from outside the qubit packaging. However, γ rays and cosmic ray muons or neutrons cannot be easily shielded and small amounts of α and β sources can exist in the qubit packaging. Here, we discuss the impact of a γ ray but the impact of other forms of radiation would be similar. The likely radioisotopes in the vicinity, such as ^{40}K, ^{232}Th, and ^{238}U in building materials such as concrete, emit γ rays with an energy of order 1 MeV [19,21,22]. At these energies, Compton scattering is the most likely interaction of the γ ray with the substrate, as shown in Fig. 1(a). During Compton scattering, the γ ray ionizes a silicon atom, ejecting an electron with an energy of the order of 100 keV [22]. This high-energy electron initiates a cascade process, scattering off of atomic electrons and generating a large number of electron-hole pairs [Fig. 1(b)]. As the

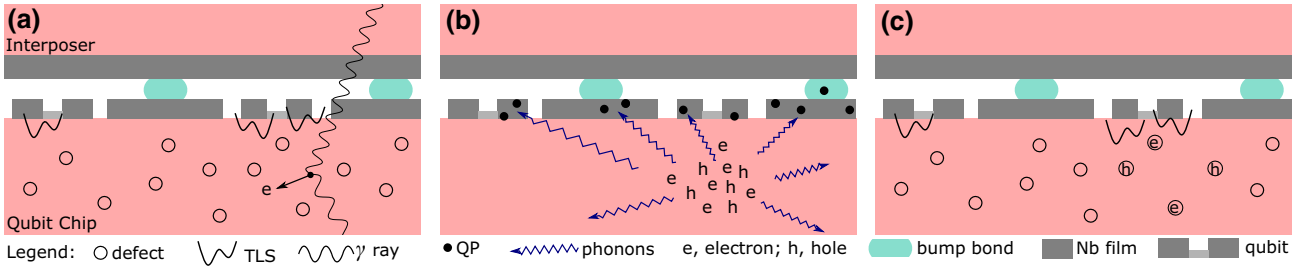

FIG. 1. A schematic of a typical γ-ray impact. The device consists of a qubit chip bump bonded onto an interposer. Each chip includes a thin patterned niobium film (gray) on a silicon substrate (pink). (a) Impact: a γ-ray Compton scatters off of a silicon atom, ejecting an electron. (b) Transient: the transient response of the device after the impact. [Defects and TLSs are not shown in (b) to avoid clutter.] The electron scatters off other atoms in the substrate, generating electron-hole pairs. Relaxation and recombination of the electrons and holes creates phonons that propagate outward, potentially for several millimeters, creating QPs when they hit either the ground plane or a qubit. The bump bonds provide both a lower-gap material, which may trap QPs, and a thermalization path for the phonons. (c) Quasistatic: after the transient response, most of the charges recombine but some of the electrons and holes get trapped at defects, resulting in a long-lasting charge rearrangement, which is sensed by both the local TLS and by the offset charge on the qubits.

electrons and holes relax and recombine, they emit photons and phonons. The phonons downconvert in energy until they travel ballistically through the silicon, potentially for several millimeters [29]. When a phonon reaches the superconductor at the surface of the silicon, it can break Cooper pairs, generating QPs, which are a source of loss for the qubits. Because the phonons can travel for millimeters, many qubits may be simultaneously poisoned by the QPs, resulting in correlated energy-relaxation events [22–24]. As shown in Fig. 1(c), the electrons and holes that do not quickly recombine will diffuse with a characteristic trapping length of a few hundred microns until becoming trapped in the substrate or at the surface, resulting in a long-lasting charge rearrangement that can be sensed by nearby qubits.

Radiation impact events can be detected using different techniques, so we choose on the basis of the capabilities of our quantum processor. The low-Q ($Q \sim 1200$) resonators used for fast readout rule out impact-detection methods based on changes in kinetic inductance [20,21,39]. Another method, monitoring for bursts of correlated qubit decays indicative of QP poisoning, has been demonstrated in a similar-scale quantum processor [23]; however, when we attempt this method, we are not able to resolve any events. Instead, following Ref. [22], we use the qubits to sense the abrupt changes in the local charge environment caused by an impact. The energy levels of the transmon, which retains the charge-qubit Hamiltonian, weakly depend on the local charge environment, quantified by the offset charge, n_{g0}. Therefore, a change in offset charge causes a small change in the qubit frequency, which can be detected by a Ramsey measurement. We do not make offset-charge-sensitive transmons, so our qubits do not have enough charge dispersion to easily detect an offset-charge jump [22]. However, the higher energy levels have larger charge dispersion [Fig. 2(b)], so we use the ef qubit subspace instead [40].

The experiment to detect an offset-charge jump is shown in Fig. 2(a). First, measurement $M0$ followed by a conditional X pulse if the qubit is in $|g\rangle$ prepares the qubit in $|e\rangle$. An $X_{ef}/2$ pulse rotates the qubit from $|e\rangle$ to the equator of the ef Bloch sphere, after which a fixed-delay Ramsey experiment is performed. Driving the $X_{ef}/2$ pulses at the mean transition frequency $\bar{\omega}_{ef}$ causes the transmon to evolve around the equator of the Bloch sphere during T_{Ramsey} at the detuning $\omega_{ef}(p, n_{g0}) - \bar{\omega}_{ef}$, where p is the QP parity. Because we only want to detect offset-charge jumps and not QP parity flips, here the two QP parities are symmetrically detuned above and below $\bar{\omega}_{ef}$, so they yield the same measurement result [Figs. 2(b) and 2(c)]. Similar sequences, replacing the second $X_{ef}/2$ with $Y_{ef}/2$, have been used to measure the QP parity [40,41]. At the end of the experiment, we map the state back to the ge subspace and measure the qubit ($M1$), where $P(M1 = 1) = (1 + \cos(\epsilon_{ef}\cos(2\pi n_{g0})T_{\text{Ramsey}}/2))/2$, in the absence of

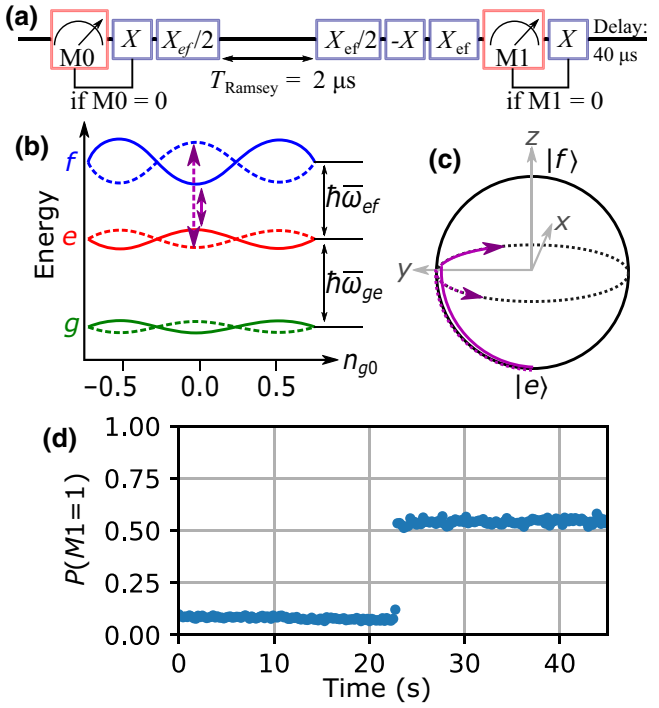

FIG. 2. Ramsey-based offset-charge jump measurement. (a) Measurement conditional X prepares the qubit in $|e\rangle$. $X_{ef}/2$ − idle − $X_{ef}/2$ performs a fixed-delay Ramsey sequence on the ef transition, with an idle time of 2 μs. The sequence $-X - X_{ef}$ maps the ef subspace to the ge subspace for Ramsey measurement $M1$. A conditional X then returns the qubit to $|e\rangle$ and a 40-μs delay between repetitions provides a fixed-delay T_1 measurement using the outcome of $M0$. (b) An exaggerated transmon energy-level diagram showing the dependence on the unitless and periodic offset charge, n_{g0}. For any n_{g0}, there are two ef-transition frequencies (purple), one for each QP parity, symmetrically detuned about the mean transition frequency $\bar{\omega}_{ef}$. (c) A Bloch-sphere illustration of the qubit state evolution during the early part of the Ramsey sequence, showing rotation along the equator at equal rates but in opposite directions depending on the QP parity. (d) A single-qubit trace of $P(M1 = 1)$ as a function of the experiment time, showing a jump 23 s into the experiment. Here, the 1×10^6 repetitions are grouped into 200 time bins to compute probabilities.

decoherence. Our transmons have ef charge dispersions $\epsilon_{ef}/2\pi \sim 800$ kHz. For fixed T_{Ramsey}, offset-charge jumps appear as abrupt jumps in $P(M1 = 1)$, as shown in Fig. 2(d). Because the relationship between $P(M1 = 1)$ and n_{g0} is nonlinear and periodic, the magnitude of the change in $P(M1 = 1)$ during a jump does not reliably indicate the magnitude of the change in n_{g0}. Therefore, we cannot use the magnitude of the change in $P(M1 = 1)$ on different qubits to triangulate the location of the impact.

We run the jump detector simultaneously on qubits across the chip to identify radiation impact events. Figure 3(a) shows the connectivity of half of the qubit chip;

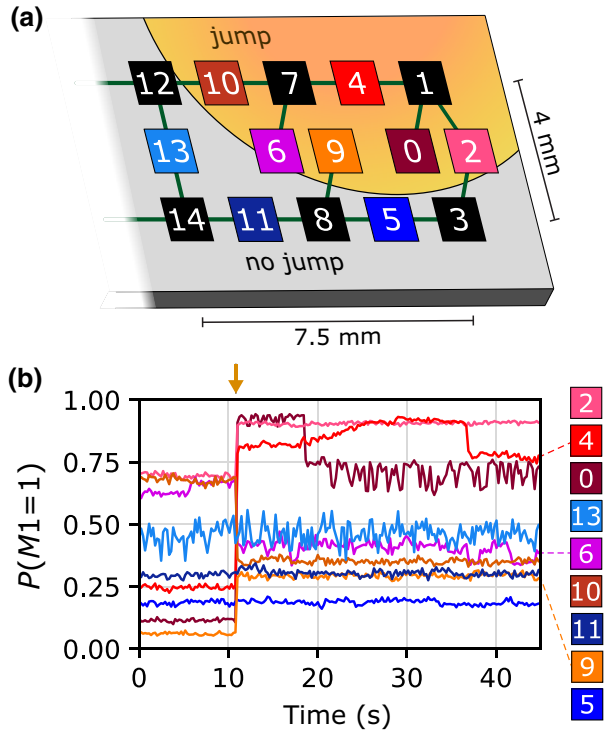

FIG. 3. Simultaneous offset-charge jumps. (a) A map of the physical layout of half of the device. The Ramsey-based jump detector is run on a set of next nearest neighbors (colored squares), while the other qubits are idled (black squares). (b) The results of a jump-detection run for nine qubits on the right half of the device. Simultaneous jumps are observed on six out of the nine qubits but not on any on the left half of the device (not shown). As seen on the map, the qubits that jumped are localized in the upper right portion of the device (colored gradient), which we attribute to a radiation event impacting that region.

the other half is similar. Dark lines between qubits indicate couplings via bus resonators that mediate two-qubit gates. These couplings induce a small shift in the qubit frequency that depends on the state of its neighbors. To avoid contamination of the Ramsey phase by this interaction, we restrict the simultaneous jump detection to a set of 17 non-neighboring qubits (0, 2, 4, 5, 6, 9, 10, 11, 13, 15, 16, 17, 20, 21, 22, 24, and 26), while the remaining ten qubits are idled. The results from one run of the jump detector are shown in Fig. 3(b). About 11 s into the run, simultaneous jumps can be seen on qubits 0, 2, 4, 6, 9, and 10, while no jumps are observed on qubits 5, 11, and 13 or any of the qubits on the left half of the chip (not shown). The qubits that jumped are localized to a small area of the chip, with a radius of a few millimeters, as shown by the colored background in Fig. 3(a). Random coincidence cannot explain simultaneous jumps on so many qubits in a small portion of the device (Appendix A 1), so these qubits must be sensing a common change in the local environment. While models of offset-charge drift have traditionally focused on local charge rearrangement very close to an island such as a TLS dipole flip, metallic grain charging, or a fluctuating patch potential [42–46], these models fail to explain simultaneous discrete jumps on qubits that are millimeters apart and well isolated by ground planes (Fig. 1). We thus attribute these multiqubit jumps to radiation generating large charge rearrangements in the substrate, which are not as effectively screened by the ground plane. As further evidence that the jumps we observe are due to radiation, in Appendix A 2 we show that the rate of events we extract from our device is consistent with values from the literature. In our device, the typical spacing between physically adjacent qubits is 1–2 mm, which has been too distant to detect simultaneous jumps in prior experiments [22,47]; however, the effective sensing distance of a transmon will depend on the geometry and layout of the qubit, the substrate thickness, and the isolation of the qubit from the ground plane and so will vary from device to device. In our device, each qubit sits in a $(650 \times 650)\ \mu\text{m}^2$ pocket in the ground plane, with a 90-μm separation between its two niobium capacitor pads, each 500 μm long and 50 μm wide. The significantly larger physical size of the transmon pads in this work, as compared to Refs. [22,47], is the simplest explanation for the larger sensing distance but it is also possible that the charges diffuse further in the substrate after the impact in our device as compared to the literature.

III. TRANSIENT EFFECT ON T_1

Having demonstrated the detection of radiation impacts, we next look for a reduction in T_1 due to the QPs generated by the impact. The pulse sequence in Fig. 2(a) uses a 40-μs delay between repetitions as a fixed-delay T_1 measurement, acquired in $M0$, interleaved with the jump detection. This sequence is repeated 1×10^6 times per experiment, which lasts 44 s. We run this experiment on the above set of non-neighboring qubits 500 times, resulting in a total "detector time" of over 6 h. As detailed in Appendix A, we use a matched filter and thresholding to detect multiqubit jumps and calculate the most likely time step of each jump, t_{detect}. We consider two jumps simultaneous if the extracted t_{detect} on the two qubits are within 10 ms of each other. Using a selective threshold value, we detect 43 multiqubit jumps comprising a total of 108 single-qubit jumps. For those 108 events, we then look for excess decays during the fixed-delay T_1 measurement. In Fig. 4, we plot $\overline{P(M0=1)}$ for 50 ms before and after the jump using the t_{detect} that we have determined earlier as the reference for the time of the jump. The overline is to indicate that data from multiple qubits are averaged together. Excess loss due to transient QP generated during an event will show up as a dip in this plot, with a width set by the timing precision of the jump-detection algorithm and the duration of the reduction in T_1 [22].

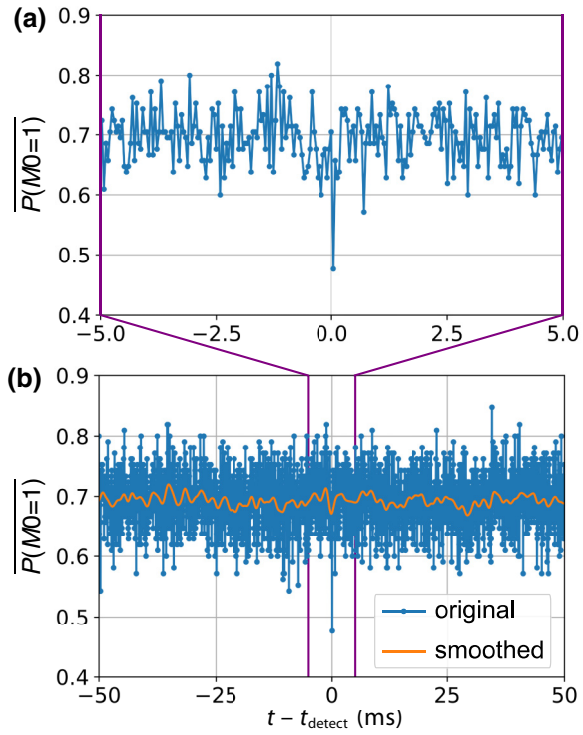

FIG. 4. The search for excess qubit loss associated with multiqubit offset-charge jumps. Collecting the 108 constituent single-qubit jumps, we average the probability of the qubit remaining in $|e\rangle$ after the 40-μs delay between shots. We show $\overline{P(M0=1)}$ (blue) for (a) 5 ms and (b) 50 ms before and after t_{detect}, the time step of the jump as determined using the jump-detection algorithm. Smoothing with a Gaussian filter with a sigma of ten data points (0.44 ms, orange) helps to show the absence of a trend. Excess qubit loss due to transient QPs corresponds to a dip in this plot. We see that $\overline{P(M0=1)}$ for the data point immediately after the extracted t_{detect} is suppressed by 4.9σ compared to the background. This is consistent with a small amount of qubit loss induced by a short-lived excess of QPs.

The averaged data (Fig. 4) are well described by a Gaussian distribution centered about a constant background level of $\overline{P(M0=1)} \approx 0.7$ (corresponding to an average $T_1 \sim 100$ μs), except for the point immediately after $t = t_{\text{detect}}$, which is 4.9σ below the mean. Since the nearby data points do not indicate a broader dip, either the data point is somehow spurious or it suggests that loss occurs only in a very short interval around the impact, e.g., QP poisoning that recovers within the 44-μs repetition time. To help assess the feasibility of our scheme resolving a real single-point feature, in Appendix A 4 we analyze simulated data and find that in some conditions the analysis algorithm can indeed be accurate to within a few time steps. Appendix A 3 also shows how the depth of the dip depends on the threshold used to declare a jump.

The absence of a sustained dip in Fig. 4 demonstrates the resilience of our quantum processor to the impact of radiation. We do not take any special measures to shield or reduce environmental radiation [19,21] or add any structures specifically to reduce the impact of transient QPs [24,27,28,31,32,47–49]. In Appendix A 1, we show that the impact rate that we measure is very similar to previous measurements. Therefore, the robustness of T_1 to radiation must derive from the packaging, design, or materials of the device. These qubits consist of niobium capacitor pads with aluminum leads for the Josephson junction. The bump bonds provide both a path for thermalization [24] and a lower-gap trap for QPs from the niobium ground plane. Going forward, additional reports will likely elucidate which variables determine the severity of the impact of radiation.

IV. EFFECT ON TLSs

While looking for a transient hit to T_1, we note that occasionally the T_1 of an individual qubit changes suddenly after a multiqubit jump. Unlike the transient dips previously discussed, here the T_1 is stable at the new value, which can be higher or lower. Next, we show that this is due to radiation "scrambling" the TLSs on a qubit, simultaneously changing the frequency of several TLSs near resonance with the qubit. When the scrambling brings a TLS into (out of) resonance with the qubit, the qubit T_1 decreases (increases). After the scrambling, the TLS spectrum is stable, so T_1 is stable at the new value.

We use TLS spectroscopy to monitor TLSs near resonance with the qubit. During TLS spectroscopy, the qubit frequency is swept and when the qubit is brought into resonance with a TLS, the two hybridize, decreasing the qubit T_1 [3–7]. Our transmons are not flux tunable but can be tuned via the Stark shift to probe TLSs using the pulse sequence shown in Fig. 5(a) [8,50]. A Stark tone amplitude Ω_s shifts the qubit frequency by $\Delta\omega_q = \alpha\Omega_s^2/2\Delta_s\,(\alpha + \Delta_s)$, where α is the qubit anharmonicity (typically, $\alpha/2\pi = -295$ MHZ) and Δ_s is the detuning of the Stark tone from the unshifted qubit frequency. With fixed detuning ($\Delta_s/2\pi = \pm 50$ MHz), we step Ω_s to sweep the Stark-shifted qubit frequency from $\Delta\omega_q/2\pi \approx -20$ MHz to $+20$ MHz. Figure 5(b) shows one example of $P(MS=1)$ as a function of $\Delta\omega_q$. We linearly ramp Ω_s, resulting in the quadratic spacing of $\Delta\omega_q$ in the horizontal axis. Dips in $P(MS=1)$ indicate excess loss attributable to TLSs. To look for changes in the TLS spectroscopy coincident with multiqubit jumps, we interleave the TLS spectroscopy with the jump detector [Fig. 5(a)]. Each iteration consists of two parts: TLS spectroscopy, in which the amplitude of the Stark pulse is swept in 251 steps, and 251 repetitions of the jump detector. A conditional X followed by a 20-μs delay prepares the qubit for the next repetition. This 24-ms sequence is repeated 10 000 times, taking a total of 4 min. We repeat the experiment 172 times for over 11 h of "detector time." A modified jump-detection

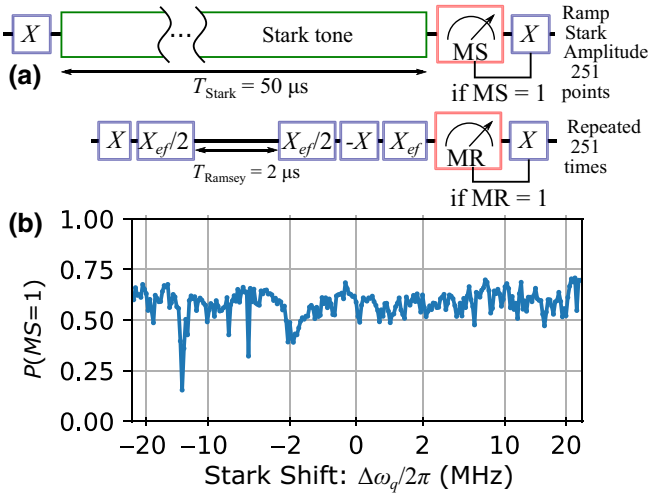

FIG. 5. Interleaved TLS spectroscopy and jump detection. (a) At each iteration, TLS spectroscopy (upper sequence) is executed by ramping the amplitude of the Stark pulse to sweep the qubit frequency, before being measured during *MS*; then the jump detector (lower sequence) is repeated 251 times, before being measured during *MR*. A conditional X to reset the qubit and a 20-μs delay follow each measurement. Each 4-min run consists of 10 000 iterations of this sequence. (b) A TLS spectrum, showing $P(MS = 1)$ as a function of the qubit Stark shift, averaged over all 10 000 iterations. Dips, such as that at $\Delta\omega_q/2\pi = -14$ MHz, correspond to excess loss when the qubit is in resonance with a TLS.

algorithm (Appendix B) run on the Ramsey-jump sequence identifies 656 multiqubit jumps. Interleaving the jump detector and the TLS spectroscopy allows us to look for changes in the TLS spectra associated with these multiqubit jumps and thus look for the impact of radiation on the TLSs in the device.

The three columns of Fig. 6 show three examples of TLS scrambling, each coincident with a multiqubit jump. In the top panel of Fig. 6(a), we plot the output of the jump detector for one run on qubits 0, 4, and 9, all of which experience a simultaneous jump 180 s into the experiment, indicating a radiation impact event. Interleaved TLS spectroscopy (middle panel) shows that the event coincides with abrupt changes in the spectral signatures of multiple TLSs (dark horizontal bands). The TLSs are stable both before and after the event. The dynamics during scrambling are too fast to resolve in this experiment. Note that prior to scrambling, when there is no Stark shift ($\Omega_S = 0$), T_1 suffers due to a TLS near resonance with the qubit. The scrambling brings that TLS out of resonance, with $P(MS = 1)$ increasing from 0.45 to 0.62, resulting in a long-lasting change to T_1. The other two columns in Fig. 6 show two additional examples of TLS-scrambling events. As in the event in Fig. 6(a), the event in Fig. 6(b) shows a long-lasting change at $\Omega_S = 0$, with $P(MS = 1)$ increasing from 0.44 to 0.60 but the event in Fig. 6(c) does not show a long-lasting change at $\Omega_S = 0$, with $P(MS = 1)$ remaining unchanged at 0.62 before and after the jump.

The jump-detector traces in the top row of Fig. 6 show additional jumps that do not correspond to TLS scrambling. In fact, most offset-charge jumps are not accompanied by TLS scrambling. To determine the fraction of multiqubit jumps associated with a TLS-scrambling event, we use the Pearson correlation coefficient r to quantify changes in the TLS spectra as shown in the bottom row of Fig. 6. Pearson's r measures the linear correlation between two ordered data sets, x and y, and is defined as the dot product $r = \hat{x} \cdot \hat{y}$ of the normalized zero-mean sets $\hat{x} = (x - \bar{x})/\sigma_x$, where σ_x is the standard deviation of x. For each time step, $x(y)$ is the average $P(MS = 1)$ for 200 data points, 4.8 s, before (after) the time step. Pearson's r is 1(0) if the two data sets are perfectly (un)correlated. In the bottom panels of Fig. 6, we see the TLS-scrambling events weaken the correlations between past and future spectra, causing r to dip below our selection threshold of 0.4 (black dashed line). However, crossing this threshold does not guarantee a scrambling event. For example, in the final minute of Fig. 6(b), a single pronounced TLS undergoing diffusion near $\Delta_S = 20$ MHz, presumably due to interaction with nearby thermal fluctuators, causes r to dip below 0.4 without TLS scrambling. Thus we observe both single TLS diffusion due to thermal fluctuators and TLS scrambling, which affects multiple TLSs simultaneously.

Of all r values computed for each qubit at each time point, r only falls below the 0.4 threshold 0.06% of the time. In contrast, one of the qubits participating in a multiqubit jump has an r below 0.4 during 34 of the 656 multiqubit jumps (5.2%). Therefore, because both multiqubit jumps and events where r dips below 0.4 are rare, it is unlikely that we will observe both simultaneously without a common origin. Having established that radiation is the most likely origin of the multiqubit jumps, we infer that radiation is also the likely origin of the TLS scrambling.

Most of the detected radiation impact events, as witnessed by multiqubit jumps, are not accompanied by detectable TLS scrambling. Also, in this data set we observe an absence of TLS scrambling on multiple qubits during the same event. This suggests a picture in which the offset charge on one of our qubits can sense a distant impact event, while the TLSs are less sensitive to distant events. We conjecture that some of the multiqubit jumps without TLS scrambling would be due to impacts that are not close enough to a qubit to induce TLS scrambling. Although this is a qualitatively plausible picture, we highlight that it is not the only possible explanation for this difference in response. It has been suggested in the literature that the offset charge is not capable of sensing distant charge rearrangements [47], in which case the impact would need to generate charge rearrangement over long distances within the chip. Long-distance charge rearrangement is possible, either by the direct diffusion

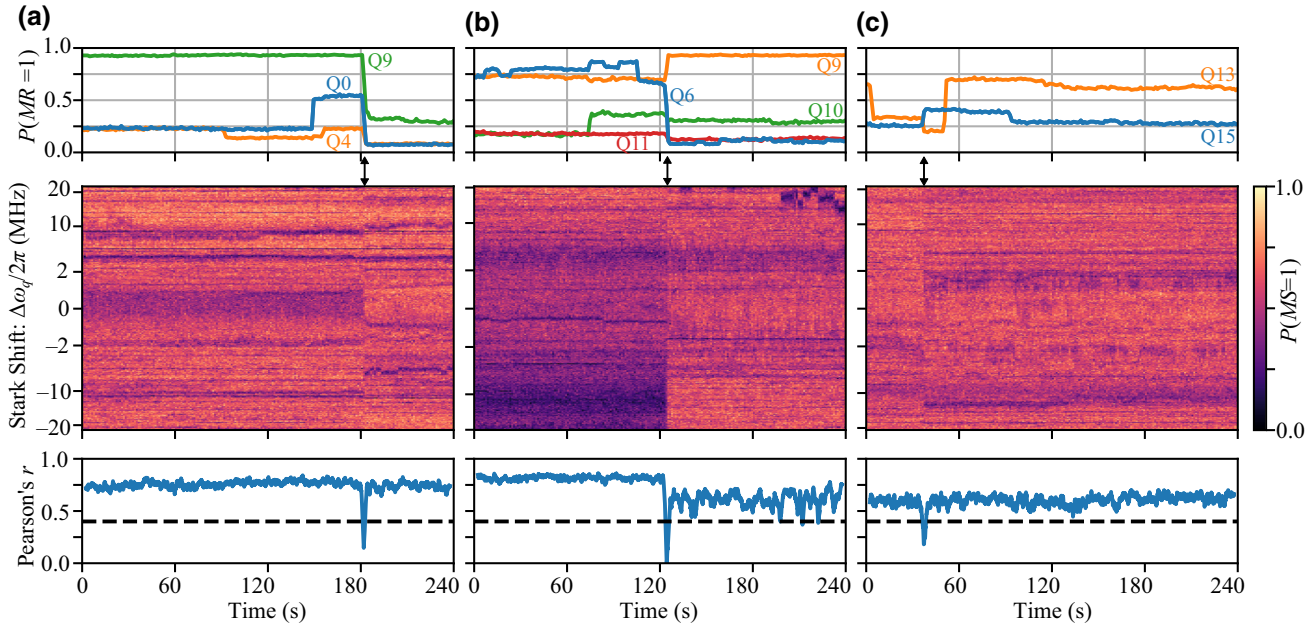

FIG. 6. TLS scrambling due to radiation. (a)–(c) Each column shows a TLS-scrambling event from a different run of the experiment: (a) run 18, qubit 0; (b) run 17, qubit 6; (c) run 21, qubit 15. The time axis is shared across all panels in a column. The values in the top and middle rows are averaged within 200 time bins to compute probabilities. The top panels show outcomes of Ramsey-jump detection (MR). The middle panels show the evolution of the TLS spectra over time. The frequency scale is nonlinear because of the quadratic dependence on the Stark tone amplitude. The color scale represents $P(MS = 1)$ at the end of the 50-μs Stark tone; dark horizontal lines correspond to excess loss when the qubit is resonant with a TLS. The bottom panels show Pearson's r for the TLS spectroscopy as defined in the main text. Dips correspond to sudden pronounced changes in spectral features in the middle panel. Dips crossing the threshold $r = 0.4$ (black dashed line) simultaneous with a multiqubit jump are declared TLS-scrambling events. Because we have previously established that multiqubit jumps are caused by the impact of radiation, we argue that the simultaneous TLS scrambling is also caused by the impact of radiation.

of charges generated by the impact or mediated by the diffusion of phonons and photons away from the initial impact site. There could be a higher energy threshold for the impact to scramble TLSs beyond the energy required to cause a multiqubit jump. Alternatively, TLS scrambling might be associated with only certain types of radiation.

V. DISCUSSION

What can we say about the interaction mechanism underlying TLS scrambling? Although the physical origin of TLSs remains the subject of much debate [2], TLSs are known to couple to the local electric field via a dipole moment. Given their simultaneity, it is natural to suppose that the TLS scrambling and offset-charge jumps are both responses to the same redistribution of charge. As shown in the insets of Fig. 1, some of the electrons and holes generated during the impact escape recombination and diffuse until becoming trapped at defects. This charge redistribution will change the electric field at the TLS, thus changing its frequency. The dipole moment of the TLS (approximately 1 eÅ [2]) is much smaller than the transmon dipole moment (approximately 100 e μm [51]), potentially explaining why qubits over a large area participate in a jump but the TLS scrambling is localized to the TLSs in the vicinity of at most one qubit. Because this picture predicts that some TLSs will scramble more than others, with TLSs at the metal-substrate or substrate-air interfaces more sensitive to charge rearrangements in the substrate than TLSs in the junction or at the metal-air interface, further insight could be gained by monitoring scrambling in a setup able to locate individual TLSs within a device [52].

Other mechanisms could also be at play. The scrambling could be mediated by the TLS elastic dipole coupling to local strains. Heat generated by an impact could cause brief expansion and thus local strain, shifting the TLS [53], akin to thermal cycling of the cryostat. Recently, stress-induced microfractures have been suggested as an alternative source of phonon and QP bursts [54] but in our architecture the electrical fields from microfractures due to the stress at the bump bonds or niobium-silicon interface are screened by the ground plane and thus unlikely to cause the offset-charge jumps. Repeating these measurements in different processors could reveal whether these dynamics depend on the materials and architecture, as in the case of the response to transient QPs. It would also

be enlightening to study other types of qubits or similarly susceptible devices such as single-electron transistors and SQUIDs [42,55].

Many proposed explanations for TLSs involve trapped charges: tunneling electrons [56–58], localized metal-induced gap states at the metal-insulator interface [46,59], or trapped QPs [60,61]. Because these charges can be created and destroyed during the impact, investigations of the response of TLSs to the impact of radiation may also tell us about the nature of the TLSs themselves.

VI. CONCLUSIONS

Using the standard transmons on a quantum processor as an array of electrometers, we detect radiation impacts and look for associated effects on T_1. We are barely able to resolve any transient effect, demonstrating a robustness to radiation in contrast with the severe hit to T_1 in the literature [22–24]. We do not take any steps to reduce the rate or severity of the impacts, which suggests that the resilience derives from the device materials and architecture and, moreover, can be obtained without compromising the baseline T_1. These results suggest that one may hope to realize traditional quantum error correction without the drastic interventions suggested in the literature [21,24,27–34]

Some radiation impacts coincide with long-lasting changes in T_1, both positive and negative, which we explain through radiation-induced TLS scrambling. These results reveal a new type of TLS dynamics beyond the interacting-defect model. We therefore identify radiation as one driver of fluctuations in T_1 over time, which is a central challenge to scalable superconducting quantum computing. Looking to the future, one can even imagine using an ionizing radiation source to controllably reset an uncooperative TLS spectrum plaguing a large quantum processor, with no need for a disruptive thermal cycle of the cryostat.

ACKNOWLEDGMENTS

The device was designed and fabricated internally at IBM. We acknowledge the use of IBM Quantum services for this work and these results were enabled by the work of the IBM Quantum software and hardware teams. This work was supported by the Intelligence Advanced Research Projects Activity (IARPA) under Logical Qubits (LogiQ) (Contract No. W911NF-16-1-0114). All statements of fact, opinion, or conclusions contained herein are those of the authors and should not be construed as representing the official views or policies of the U.S. Government. We thank Oliver Dial, Francesco Valenti, Youngseok Kim, Sami Rosenblatt, Joey Suttle, Jared Hertzberg, Chris Lirakis, Luke Govia, Ken Rodbell, Antonio Córcoles, Maika Takita, and Matthias Steffen for helpful conversations, device bring-up, and programmatic support.

APPENDIX A: AUTOMATIC DETECTION OF OFFSET-CHARGE JUMPS

Our experiments involve monitoring many qubits for hours at a time and the large amounts of generated data required efficient robust detection algorithms. The output of the Ramsey-jump detector from Fig. 2(d) is reproduced in Fig. 7(a). For plotting $P(M1 = 1)$, the data are binned into 200 bins. However, binning discards precise timing information about the jump. The inset of Fig. 7(a) shows the outcomes of the individual measurements immediately before and after the jump. By eye, we can see that we should be able to detect the time of this jump to within a few time steps. To detect jumps and extract the time of each jump, we use a matched filter with a step function with a duration of 50 ms as the template:

$$\text{Template}(t) = \begin{cases} -1, & -25\ \text{ms} < t < 0, \\ 0, & t = 0, \\ 1, & 0 < t < 25\ \text{ms}. \end{cases}$$

The cross correlation (SCIPY.SIGNAL.CORRELATE) of the raw output of the jump detector and the template is the result of sliding the template along the output of the jump

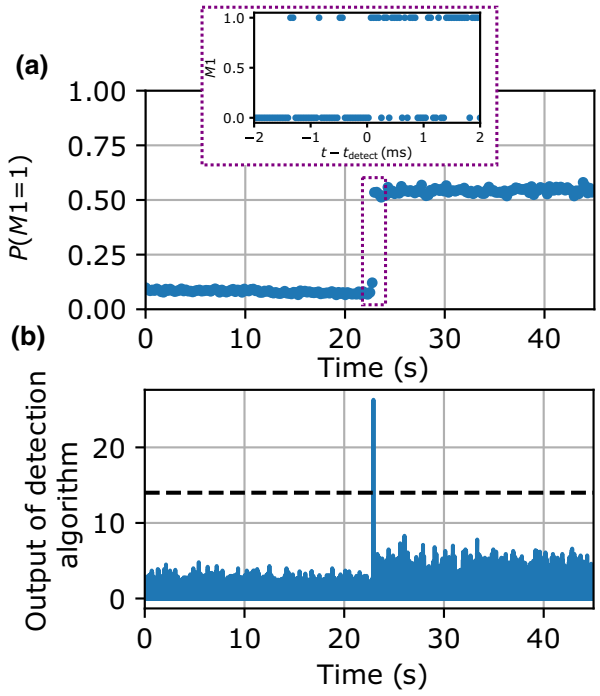

FIG. 7. The jump-detection algorithm. (a) The output of the Ramsey-based jump detector, reproduced from Fig. 2(d): the 1×10^6 shots of the experiment are binned into 200 bins to calculate $P(M1 = 1)$. The inset shows the raw output of M1 for 2 ms before and after the jump. (b) The output of the jump-detection algorithm after normalization by the median level, showing a peak (which sets t_{detect}) corresponding to the jump in (a). The dashed line is our default detection threshold of 14.

detector. Absent a jump in the data, the positive and negative halves of the step function null the correlation, while a jump causes a spike in the correlation. The problem of step detection thus reduces to peak detection in the jump signal, defined as the absolute value of this cross correlation, which we perform via thresholding. We choose this matched filter approach for simplicity; other techniques for offline change-point detection can likely improve sensitivity in future studies (e.g., information might be gleaned not only from the step in mean value when n_{g0} jumps but also from the step in variance). As the analyses of the 500 experiment runs are independent, time can be saved with parallel processing.

The choice of threshold contains two subtleties: how to normalize signals for cross-qubit comparisons and then how to set the overall threshold level. Instantaneous values of n_{g0} and $T^*_{2,ef}$ will vary over time and across qubits, causing the sensitivity of each qubit to fluctuate over time. We find it effective to normalize the jump signal trace from each qubit for each run by its median, such that a given threshold value roughly corresponds to a fixed confidence in (the SNR of) the jump signal, rather than to a uniform magnitude of the jump signal. Then, we choose an overall threshold of 14, which we empirically find provides good confidence in and time resolution of events while still accepting enough data for further statistical analysis (details in Appendix A 3). An example output of this algorithm with a large jump signal is shown in Fig. 7(b). A clear peak well above the threshold (black dashed line) indicates the single-qubit jump. The time step of this peak sets t_{detect}. Double counting of a noisy threshold crossing is avoided by setting a minimum time separation of 50 ms between detections on the same qubit, keeping the largest peak in such cases (SCIPY.SIGNAL.FIND_PEAKS).

Finally, the identified single-qubit jumps are further clustered into multiqubit jumps. Each jump is sorted by t_{detect} into a list. Multiqubit jumps are identified when the time between t_{detect} on the list is less than 10 ms. We caution that in principle this algorithm permits an edge case where a chain of < 10 ms delays generates an arbitrarily long single multiqubit jump but in our data all multiqubit jump durations are below 2 ms (median duration < 400 μs) for our choice of threshold = 14. With that threshold, we detect 43 multiqubit jumps, consisting of 108 single-qubit jumps.

1. Spatial and temporal jump correlations

How can we be sure that the events we identify as multiqubit jumps are true multiqubit jumps and not independent single-qubit jumps, that happen by coincidence to jump within 10 ms of each other? We support our interpretation of detected jumps as resulting from real charging events on the chip by looking for expected correlations in space and time.

First, multiqubit jumps with a common physical origin should preferentially involve qubits that are physically closer together. If the multiqubit jumps were random chance, distant qubits would be equally like to have simultaneous jumps as nearby qubits. For each pair of qubits used for jump detection, we calculate their separation and the rate at which the two qubits experience a simultaneous jump (in units of multiqubit jump events per hour of detector time). In Fig. 8(a), we show a histogram of the coincidence rate as a function of the distance between qubits. The distribution fits a Gaussian fall-off with a $\sigma = 1.9$ mm, indicating that the events are strongly localized. Correlated offset-charge jumps have not been previously reported at this length scale [22,47], though the characteristic length

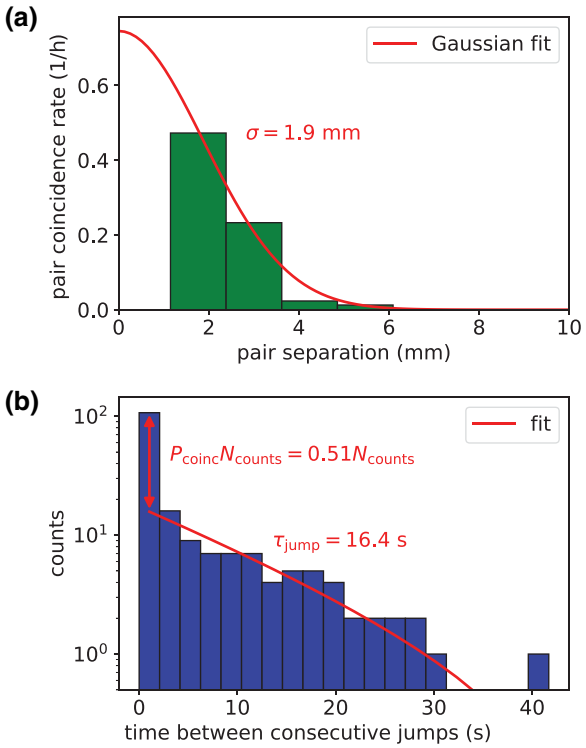

FIG. 8. Offset-charge jump correlations in space and time. (a) A histogram of the rate of simultaneous jumps for all pairs of physical qubits used for jump detection. The first bin is empty because there is a minimal distance of about 1 mm between pairs of qubits. In red is a Gaussian distribution, with $\sigma = 1.9$ mm, showing that qubits that experience simultaneous jumps are located within a few millimeters of each other. (b) A histogram of the time between consecutive jumps (including single-qubit jumps not associated with a multiqubit jump). The red line is a fit to a modified Poisson distribution, with a characteristic jump rate = 1/(16.4 s). The fit includes a correction to account for the finite duration of the experiment, 44 s. The Poisson distribution fails to account for the excess of events in the first bin, which accounts for 51% of total single-qubit jumps, indicating that the multiqubit jumps cannot be explained as random coincidences of uncorrelated single-qubit jumps.

scale is expected to vary from device to device, depending on the qubit design and chip layout.

Second, the time separation of independent random single-qubit jumps should be well described by a Poisson distribution but multiqubit jumps should instead exhibit more frequent coincidences where several qubits jump at once. For each 44 s run of the jump-detection experiment, we generate a single list of the times at which any individual qubit jumps and calculate the time delays between consecutive single-qubit jumps within that run. During multiqubit jumps, this time difference will be small and can be zero. Figure 8(b) shows a histogram of these time differences with 2-s bins. If each jump was an independent single-qubit event, the Poisson-distributed delays between events would follow an exponential probability distribution with some characteristic time scale t_{jump}. However, the histogram shows a clear excess of events in the first bin, which cannot be explained by individual qubits independently experiencing jumps.

While these results are qualitatively consistent with expectations, further analysis is complicated by the possibility of multiple mechanisms contributing to jump detections. In particular, single-qubit jumps can be explained by local processes not involving radiation [42]. For example, qubit 17 produces jumps in this data set at a higher rate than do the other qubits, likely due to a defect coupled to the *ef* transition, so this outlier qubit is excluded from this rate analysis. Some fraction of the single-qubit jumps from the included qubits are likely also produced by causes other than radiation. Nonetheless, we can grossly quantify the full distribution using the fit model

$$\text{PDF}(\Delta) = Ne^{-\Delta/\tau_{\text{jump}}}(1-\Delta/T)(1-P_{\text{coinc}}) + P_{\text{coinc}}\delta_{\Delta},$$

where PDF stands for probability density function, Δ is the time between events, $N = 1/\tau_{\text{jump}} - 1 - e^{-T/\tau_{\text{jump}}}/T$ is a normalization constant, and the factor $1-\Delta/T$ accounts for the finite duration $T = 44$ s of the experiment. The free parameters are τ_{jump} and P_{coinc}, representing the excess probability that a single-qubit jumps occurs in coincidence with the preceding single-qubit jump (here, the excess probability mass in the first histogram bin). To fit the histogram, the PDF must be converted to a probability mass function by integrating over each bin, where we define δ_{Δ} to integrate to 1 in the first bin and zero otherwise. The simplest interpretation of the fit result is that events producing jumps above threshold occur with a time scale of 16.4 s and many of those events cause more than one qubit to jump, increasing the number of jump counts in the first bin roughly tenfold from the expectation if qubits were jumping independently.

2. Comparison of rate to literature

If we assume that all of the jumps that we observe are due to radiation, then we can place a bound on the rate of the impact of radiation on the chip and we can compare the this rate to other reports for the rate of radiation impacts from the literature, as shown in Table I. The rate of impact from radiation will depend on many physical parameters, including the amount and type of local radioactivity, the amount of shielding, and the thickness of the substrate. Remarkably, dividing the average impact rates by the relevant chip sizes yields values that are all within roughly a factor of 2. This is a striking agreement considering that these experiments use both different hardware technologies and detection techniques, which further supports the interpretation of our jumps as resulting from radiation. It also suggests that other differences in effects that we see, such as the effect on T_1, cannot be explained by a difference in the rate of events.

3. Alternative thresholds

We deliberately choose a high threshold of 14 for the jump-detection algorithm to suppress false-positive detections, especially given that the noise properties of the device are likely to fluctuate appreciably on the multihour time scale of the experiment. However, we are free to reanalyze the same data with other thresholds to learn how various quantities systematically vary with threshold choice.

The top row of Figs. 9(a)–9(c) shows the equivalent of Fig. 4 for thresholds ranging from 8 to 16. Interestingly, the dip in T_1 becomes less prominent as the threshold is lowered, as shown in Fig. 9(d). For the smallest threshold, Fig. 9(a), the dip all but vanishes. There are two possible explanations for this disappearance. First, a higher threshold might select for events that generate more QPs,

TABLE I. A comparison of the impact rates of radiation from the literature, showing that the rate of events normalized by the area of the chip that we observe is within a factor of 2 of other reports in the literature.

Reference	Technology	Measurement	Time between impacts (s)	Size of chip (mm^2)	Normalized rate [$10^{-3}/(\text{s mm}^2)$]
This work	Transmons	Offset charge	16	150	0.4
Wilen [22]	Transmons	Offset charge	50	39	0.5
McEwen [23]	Transmons	Correlated decays	10	100	1
Cardani [21]	Granular-aluminum resonators	Resonator frequency shift	10	120	0.8
Grünhaupt [20]	Granular-aluminum resonators	Resonator frequency shift	20	120	0.4

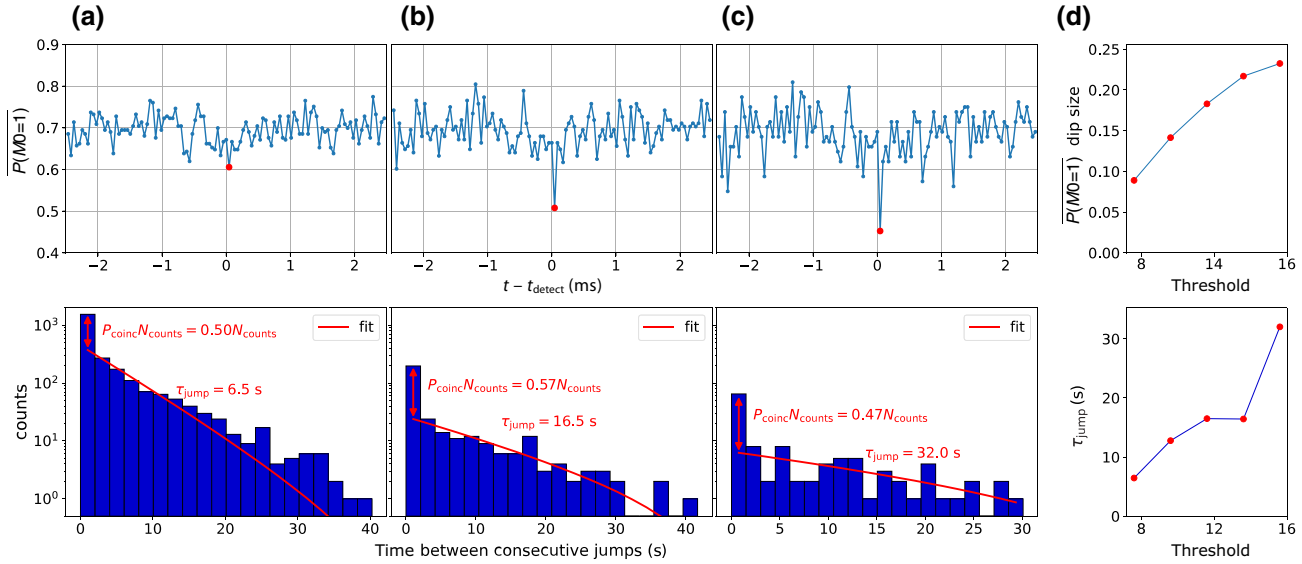

FIG. 9. Comparisons of various thresholds for the jump detection. (a)–(c) The top panels show the equivalent of Fig. 4(a) for increasing jump-detection thresholds of (a) 8, (b) 12, and (c) 16. (d) The size of the dip increases with threshold, perhaps due to tighter time resolution of the jump detection. The bottom panels of (a)–(c) show the equivalent of Fig. 8(b). Increasing the detection threshold decreases the characteristic rate at which events satisfying that threshold occur.

leading to a more prominent dip in T_1. Alternatively, keeping only jumps with a large clear change in $P(MR = 1)$ during the offset-charge jump-detector experiment should reduce uncertainty in the timing of the jump, improving the resolution of the dip. To assess the timing accuracy of the jump detector, we feed numerically simulated results of the experiment into the jump-detector algorithm (Appendix A 4).

The bottom rows of Figs. 9(a)–9(c) show the equivalent of Fig. 8(b) for the various thresholds. As expected, when we raise the threshold, we see the rate of events decline, from 6.5 s between impacts to 32 s between impacts. Deviation of the threshold-8 histogram from the modified Poisson fit may indicate an increasing significance of false-positive detections. At the other extreme, the threshold-16 histogram exhibits a scarcity of detections inhibiting statistical analysis. In Table I, we use the rate $\tau_{\text{jump}} = 16.4$ s, using the threshold-14 value from the main text. As we vary the threshold, τ_{jump} can change by a factor of 2 but this does not significantly affect the comparison.

4. Simulated timing accuracy of jump detector

Figures 4 and 9 show a reduction in T_1 after an impact that consists of a single data point. However, the Ramsey-based jump detector is probabilistic, so it is not obvious that the time of the jump can be determined to within a single time step even in principle. To probe the feasible timing accuracy of the jump-detector algorithm, we generate data from a model simulation of the experiment. In the simulation, we generate a list of impact times from a Poisson distribution of rate 1/(10 s), randomly select coordinates on the chip for each impact, and calculate the magnitude of Δn_{g0} for each impact based on the distance of each qubit to the impact site, assuming a maximum $\Delta n_{g0} = 0.1$ and a Gaussian fall-off using $\sigma = 1.5$ mm. For each qubit, n_{g0} is initialized to a uniformly random value and slowly diffused with a variance of 0.1 electrons per hour. Phase flips during Ramsey occur per a coherence time of the ef subspace $T^*_{2,ef} = 50$ μs and each qubit measurement has a symmetric error probability of 1.5 %. Finally, to permit comparison with Fig. 4, each simulated impact reduces T_1 of nearby qubits for a single time point, lowering T_1 to 1 μs at the epicenter with the same Gaussian spatial fall-off of $\sigma = 1.5$ mm.

We run the simulation 250 times, half as many as in the actual experiment, with the results shown in Fig. 10. Because this is a simulation, we can compare the true time of the jump, t_{true}, to the time of the jump extracted from the jump-detection algorithm, t_{detect}, in Fig. 10(a). Although the tails of the distribution are very long, far from either a Gaussian or a Lorentzian distribution, the extracted jump times are closely peaked around the true event time, such that very precisely resolving the time of the jump is possible. The resolution is better described by the median absolute deviation 132 μs (which is three time steps), than by the standard deviation 746 μs. Figure 10(b) shows the equivalent of Fig. 4, for the simulated data. The spread in $\overline{P(M0 = 1)}$ is reduced from the experiment, because the simulation does not account for all sources of variation, such as a spread in qubit T_1s. Notably, we do not observe single-time-step accuracy, as needed to explain Figs. 4

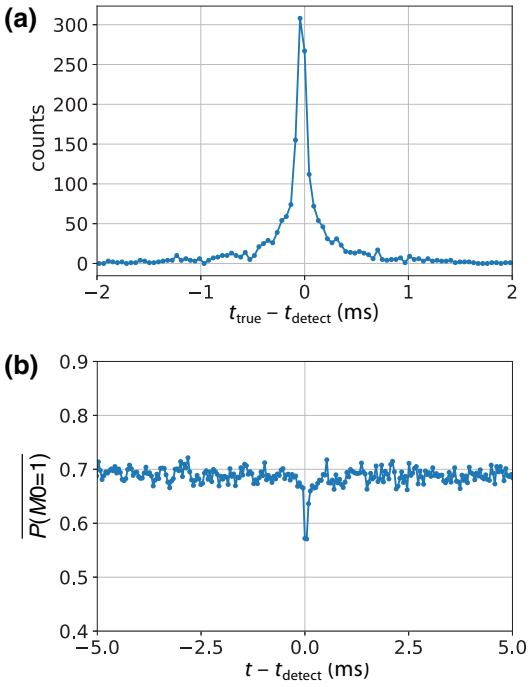

FIG. 10. The jump detection of numerically simulated data, for a threshold of 14. (a) Histograms of the time differences between the times at which the analysis detects a jump event and the true times of the corresponding simulated events. Each bin corresponds to a single time step (44 μs). (b) The dip resulting from a transient decrease in T_1, for comparison with Fig. 4. These results do not reproduce the single-point feature in Fig. 4 but do suggest that few-point resolution is possible.

and 9. This may still arise from discrepancies between simulated and experimental error processes, or perhaps the noise in Fig. 4 could be obscuring the full width of the dip, in which case future improvements in qubit coherence may help reveal the full duration of the transient dynamics.

APPENDIX B: AUTOMATING DETECTION OF TLS SCRAMBLING

For the TLS-scrambling experiment, we run the pulse sequence in Fig. 5(a), which uses a variation on the jump-detection sequence. The fundamental difference is that now the Ramsey-jump-detector experiment is run back to back 251 times in between sweeps of the TLS spectra. Because the Stark-shift measurements occupy three quarters of the experimental time, most of the radiation impact events will occur in between runs of the Ramsey-jump detector, so the very precise temporal resolution of the jump-detection algorithm from Appendix A 4 is irrelevant. Instead, we average those 251 shots to calculate $P(MR = 1)$ before running a slightly modified jump-detection algorithm. We use the same template as in Appendix A 1, except with a duration of the template of 200 sequences (5 s), to match the averaging time for the Pearson's r calculation in the main text. By feeding the averaged $P(MR = 1)$ into the jump-detection algorithm rather than the raw output of the Ramsey sequence, we greatly increase the sensitivity to small jumps, at the cost of decreasing the temporal resolution of the extracted jump time to the time to run a single sequence (about 24 ms). We are able to use a smaller threshold of 8 to detect a jump, because the signal is less noisy. For this reason, we detect more multiqubit jumps per hour than in Appendix A 1, as what would have been classified as single-qubit jumps are here classified as multiqubit jumps.

[1] I. Siddiqi, Engineering high-coherence superconducting qubits, Nat. Rev. Mater. **6**, 875 (2021).
[2] 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).
[3] 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).
[4] R. Barends, J. Kelly, A. Megrant, D. Sank, E. Jeffrey, Y. Chen, Y. Yin, B. Chiaro, J. Mutus, and C. Neill, *et al.*, Coherent Josephson Qubit Suitable for Scalable Quantum Integrated Circuits, Phys. Rev. Lett. **111**, 080502 (2013).
[5] A. Bilmes, S. Volosheniuk, J. D. Brehm, A. V. Ustinov, and J. Lisenfeld, Quantum sensors for microscopic tunneling systems, npj Quantum Inf. **7**, 1 (2021).
[6] 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).
[7] P. V. Klimov, J. Kelly, Z. Chen, M. Neeley, A. Megrant, B. Burkett, R. Barends, K. Arya, B. Chiaro, and Y. Chen, *et al.*, Fluctuations of Energy-Relaxation Times in Superconducting Qubits, Phys. Rev. Lett. **121**, 090502 (2018).
[8] M. Carroll, S. Rosenblatt, P. Jurcevic, I. Lauer, and A. Kandala, Dynamics of superconducting qubit relaxation times, npj Quantum Inf. **8**, 132 (2022).
[9] E. van den Berg, Z. K. Minev, A. Kandala, and K. Temme, Probabilistic error cancellation with sparse Pauli-Lindblad models on noisy quantum processors (2022), ArXiv:2201.09866.
[10] L. Faoro and L. B. Ioffe, Interacting tunneling model for two-level systems in amorphous materials and its predictions for their dephasing and noise in superconducting microresonators, Phys. Rev. B **91**, 014201 (2015).
[11] J. Burnett, L. Faoro, I. Wisby, V. Gurtovoi, A. Chernykh, G. Mikhailov, V. Tulin, R. Shaikhaidarov, V. Antonov, and P. Meeson, *et al.*, Evidence for interacting two-level systems from the $1/f$ noise of a superconducting resonator, Nat. Commun. **5**, 1 (2014).
[12] 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).

[13] 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).

[14] J. J. Burnett, A. Bengtsson, M. Scigliuzzo, D. Niepce, M. Kudra, P. Delsing, and J. Bylander, Decoherence benchmarking of superconducting qubits, npj Quantum Inf. **5**, 54 (2019).

[15] C. Neill, A. Megrant, R. Barends, Y. Chen, B. Chiaro, J. Kelly, J. Mutus, P. O'Malley, D. Sank, and J. Wenner, *et al.*, Fluctuations from edge defects in superconducting resonators, App. Phys. Lett. **103**, 072601 (2013).

[16] J. H. Béjanin, Y. Ayadi, X. Xu, C. Zhu, H. R. Mohebbi, and M. Mariantoni, Fluctuation Spectroscopy of Two-Level Systems in Superconducting Resonators, Phys. Rev. Appl. **18**, 034009 (2022).

[17] S. E. de Graaf, S. Mahashabde, S. E. Kubatkin, A. Y. Tzalenchuk, and A. V. Danilov, Quantifying dynamics and interactions of individual spurious low-energy fluctuators in superconducting circuits, Phys. Rev. B **103**, 174103 (2021).

[18] J. Lisenfeld, G. J. Grabovskij, C. Müller, J. H. Cole, G. Weiss, and A. V. Ustinov, Observation of directly interacting coherent two-level systems in an amorphous material, Nat. Commun. **6**, 1 (2015).

[19] A. P. Vepsäläinen, A. H. Karamlou, J. L. Orrell, A. S. Dogra, B. Loer, F. Vasconcelos, D. K. Kim, A. J. Melville, B. M. Niedzielski, J. L. Yoder, S. Gustavsson, J. A. Formaggio, B. A. VanDevender, and W. D. Oliver, Impact of ionizing radiation on superconducting qubit coherence, Nature **584**, 551 (2020).

[20] L. Grünhaupt, N. Maleeva, S. T. Skacel, M. Calvo, F. Levy-Bertrand, A. V. Ustinov, H. Rotzinger, A. Monfardini, G. Catelani, and I. M. Pop, Loss Mechanisms and Quasiparticle Dynamics in Superconducting Microwave Resonators Made of Thin-Film Granular Aluminum, Phys. Rev. Lett. **121**, 117001 (2018).

[21] L. Cardani, F. Valenti, N. Casali, G. Catelani, T. Charpentier, M. Clemenza, I. Colantoni, A. Cruciani, G. D'Imperio, and L. Gironi, *et al.*, Reducing the impact of radioactivity on quantum circuits in a deep-underground facility, Nat. Commun. **12**, 1 (2021).

[22] C. D. Wilen, S. Abdullah, N. A. Kurinsky, C. Stanford, L. Cardani, G. D'Imperio, C. Tomei, L. Faoro, L. B. Ioffe, C. H. Liu, A. Opremcak, B. G. Christensen, J. L. DuBois, and R. McDermott, Correlated charge noise and relaxation errors in superconducting qubits, Nature **594**, 369 (2021).

[23] M. McEwen, L. Faoro, K. Arya, A. Dunsworth, T. Huang, S. Kim, B. Burkett, A. Fowler, F. Arute, and J. C. Bardin, *et al.*, Resolving catastrophic error bursts from cosmic rays in large arrays of superconducting qubits, Nat. Phys. **18**, 107 (2022).

[24] J. M. Martinis, Saving superconducting quantum processors from decay and correlated errors generated by gamma and cosmic rays, npj Quantum Inf. **7**, 1 (2021).

[25] D. Aharonov, A. Kitaev, and J. Preskill, Fault-Tolerant Quantum Computation with Long-Range Correlated Noise, Phys. Rev. Lett. **96**, 050504 (2006).

[26] Z. Chen, K. J. Satzinger, J. Atalaya, A. N. Korotkov, A. Dunsworth, D. Sank, C. Quintana, M. McEwen, R. Barends, and P. V. Klimov, *et al.*, Exponential suppression of bit or phase errors with cyclic error correction, Nature **595**, 383 (2021).

[27] R.-P. Riwar, A. Hosseinkhani, L. D. Burkhart, Y. Y. Gao, R. J. Schoelkopf, L. I. Glazman, and G. Catelani, Normal-metal quasiparticle traps for superconducting qubits, Phys. Rev. B **94**, 104516 (2016).

[28] A. Hosseinkhani and G. Catelani, Proximity effect in normal-metal quasiparticle traps, Phys. Rev. B **97**, 054513 (2018).

[29] U. Patel, I. V. Pechenezhskiy, B. L. T. Plourde, M. G. Vavilov, and R. McDermott, Phonon-mediated quasiparticle poisoning of superconducting microwave resonators, Phys. Rev. B **96**, 220501(R) (2017).

[30] N. A. Court, A. J. Ferguson, R. Lutchyn, and R. G. Clark, Quantitative study of quasiparticle traps using the single-Cooper-pair transistor, Phys. Rev. B **77**, 100501(R) (2008).

[31] V. Iaia, J. Ku, A. Ballard, C. Larson, E. Yelton, C. Liu, S. Patel, R. McDermott, and B. Plourde, Phonon downconversion to suppress correlated errors in superconducting qubits, Nat. Commun. **13**, 6425 (2022).

[32] F. Henriques, F. Valenti, T. Charpentier, M. Lagoin, C. Gouriou, M. Martínez, L. Cardani, M. Vignati, L. Grünhaupt, and D. Gusenkova, *et al.*, Phonon traps reduce the quasiparticle density in superconducting circuits, App. Phys. Lett. **115**, 212601 (2019).

[33] J. L. Orrell and B. Loer, Sensor-assisted fault mitigation in quantum computation, Phys. Rev. App. **16**, 024025 (2021).

[34] Q. Xu, A. Seif, H. Yan, N. Mannucci, B. O. Sane, R. Van Meter, A. N. Cleland, and L. Jiang, Distributed Quantum Error Correction for Chip-Level Catastrophic Errors, Phys. Rev. Lett. **129**, 240502 (2022).

[35] S. W. Leman, Physics and Monte Carlo techniques as relevant to cryogenic, phonon, and ionization readout of cryogenic dark matter search radiation detectors, Rev. Sci. Instrum. **83**, 091101 (2012).

[36] R. Gordon, C. Murray, C. Kurter, M. Sandberg, S. Hall, K. Balakrishnan, R. Shelby, B. Wacaser, A. Stabile, and J. Sleight, *et al.*, Environmental radiation impact on lifetimes and quasiparticle tunneling rates of fixed-frequency transmon qubits, App. Phys. Lett. **120**, 074002 (2022).

[37] A. D. Córcoles, J. M. Chow, J. M. Gambetta, C. Rigetti, J. R. Rozen, G. A. Keefe, M. Beth Rothwell, M. B. Ketchen, and M. Steffen, Protecting superconducting qubits from radiation, App. Phys. Lett. **99**, 181906 (2011).

[38] R. Barends, J. Wenner, M. Lenander, Y. Chen, R. C. Bialczak, J. Kelly, E. Lucero, P. O'Malley, M. Mariantoni, and D. Sank, *et al.*, Minimizing quasiparticle generation from stray infrared light in superconducting quantum circuits, App. Phys. Lett. **99**, 113507 (2011).

[39] J. Baselmans, Kinetic inductance detectors, J. Low Temp. Phys. **167**, 292 (2012).

[40] D. M. Tennant, L. A. Martinez, K. M. Beck, S. R. O'Kelley, C. D. Wilen, R. McDermott, J. L. DuBois, and Y. J. Rosen, Low-Frequency Correlated Charge-Noise Measurements across Multiple Energy Transitions in a Tantalum Transmon, PRX Quantum **3**, 030307 (2022).

[41] D. Ristè, C. C. Bultink, M. J. Tiggelman, R. N. Schouten, K. W. Lehnert, and L. DiCarlo, Millisecond charge-parity

fluctuations and induced decoherence in a superconducting transmon qubit, Nat. Commun. **4**, 1913 (2013).

[42] M. D. Stewart and N. M. Zimmerman, Stability of single electron devices: Charge offset drift, App. Sci. **6**, 187 (2016).

[43] N. M. Zimmerman, W. H. Huber, B. Simonds, E. Hourdakis, A. Fujiwara, Y. Ono, Y. Takahashi, H. Inokawa, M. Furlan, and M. W. Keller, Why the long-term charge offset drift in Si single-electron tunneling transistors is much smaller (better) than in metal-based ones: Two-level fluctuator stability, J. App. Phys. **104**, 033710 (2008).

[44] S. Kafanov, H. Brenning, T. Duty, and P. Delsing, Charge noise in single-electron transistors and charge qubits may be caused by metallic grains, Phys. Rev. B **78**, 125411 (2008).

[45] B. G. Christensen, C. D. Wilen, A. Opremcak, J. Nelson, F. Schlenker, C. H. Zimonick, L. Faoro, L. B. Ioffe, Y. J. Rosen, J. L. DuBois, B. L. T. Plourde, and R. McDermott, Anomalous charge noise in superconducting qubits, Phys. Rev. B **100**, 140503(R) (2019).

[46] A. Pourkabirian, M. V. Gustafsson, G. Johansson, J. Clarke, and P. Delsing, Nonequilibrium Probing of Two-Level Charge Fluctuators Using the Step Response of a Single-Electron Transistor, Phys. Rev. Lett. **113**, 256801 (2014).

[47] X. Pan, H. Yuan, Y. Zhou, L. Zhang, J. Li, S. Liu, Z. H. Jiang, G. Catelani, L. Hu, and F. Yan, Engineering superconducting qubits to reduce quasiparticles and charge noise, Nat. Commun. **13**, 7196 (2022).

[48] K. Karatsu, A. Endo, J. Bueno, P. De Visser, R. Barends, D. Thoen, V. Murugesan, N. Tomita, and J. Baselmans, Mitigation of cosmic ray effect on microwave kinetic inductance detector arrays, App. Phys. Lett. **114**, 032601 (2019).

[49] R.-P. Riwar and G. Catelani, Efficient quasiparticle traps with low dissipation through gap engineering, Phys. Rev. B **100**, 144514 (2019).

[50] P. Zhao, T. Ma, Y. Jin, and H. Yu, Combating fluctuations in relaxation times of fixed-frequency transmon qubits with microwave-dressed states, Phys. Rev. A **105**, 062605 (2022).

[51] H. Paik, D. I. Schuster, L. S. Bishop, G. Kirchmair, G. Catelani, A. P. Sears, B. R. Johnson, M. J. Reagor, L. Frunzio, and L. I. Glazman, *et al.*, Observation of High Coherence in Josephson Junction Qubits Measured in a Three-Dimensional Circuit QED Architecture, Phys. Rev. Lett. **107**, 240501 (2011).

[52] 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**, 1 (2020).

[53] E. Nazaretski, R. D. Merithew, V. O. Kostroun, A. T. Zehnder, R. O. Pohl, and J. M. Parpia, Effect of Low-Level Radiation on the Low Temperature Acoustic Behavior of a-SiO_2, Phys. Rev. Lett. **92**, 245502 (2004).

[54] R. Anthony-Petersen, A. Biekert, R. Bunker, C. L. Chang, Y.-Y. Chang, L. Chaplinsky, E. Fascione, C. W. Fink, M. Garcia-Sciveres, and R. Germond, *et al.*, A stress induced source of phonon bursts and quasiparticle poisoning (2022), ArXiv:2208.02790.

[55] D. Gusenkova, F. Valenti, M. Spiecker, S. Günzler, P. Paluch, D. Rieger, L.-M. Pioraş-Ţimbolmaş, L. P. Zârbo, N. Casali, and I. Colantoni, *et al.*, Operating in a deep underground facility improves the locking of gradiometric fluxonium qubits at the sweet spots, App. Phys. Lett. **120**, 054001 (2022).

[56] L. Faoro and L. B. Ioffe, Quantum Two Level Systems and Kondo-Like Traps as Possible Sources of Decoherence in Superconducting Qubits, Phys. Rev. Lett. **96**, 047001 (2006).

[57] K. Agarwal, I. Martin, M. D. Lukin, and E. Demler, Polaronic model of two-level systems in amorphous solids, Phys. Rev. B **87**, 144201 (2013).

[58] R. M. Lutchyn, Ł. Cywiński, C. P. Nave, and S. DasSarma, Quantum decoherence of a charge qubit in a spin-fermion model, Phys. Rev. B **78**, 024508 (2008).

[59] S. K. Choi, D.-H. Lee, S. G. Louie, and J. Clarke, Localization of Metal-Induced Gap States at the Metal-Insulator Interface: Origin of Flux Noise in SQUIDs and Superconducting Qubits, Phys. Rev. Lett. **103**, 197001 (2009).

[60] A. Bespalov, M. Houzet, J. S. Meyer, and Y. V. Nazarov, Theoretical Model to Explain Excess of Quasiparticles in Superconductors, Phys. Rev. Lett. **117**, 117002 (2016).

[61] S. de Graaf, L. Faoro, L. Ioffe, S. Mahashabde, J. Burnett, T. Lindström, S. Kubatkin, A. Danilov, and A. Y. Tzalenchuk, Two-level systems in superconducting quantum devices due to trapped quasiparticles, Sci. Adv. **6**, eabc5055 (2020).