QUASIPARTICLE POISONING IN SUPERCONDUCTING QUANTUM COMPUTERS

José Aumentado, Gianluigi Catelani, and Kyle Serniak

Recent research has uncovered new insights into how some errors in superconducting qubits are generated and the best ways to mitigate them.

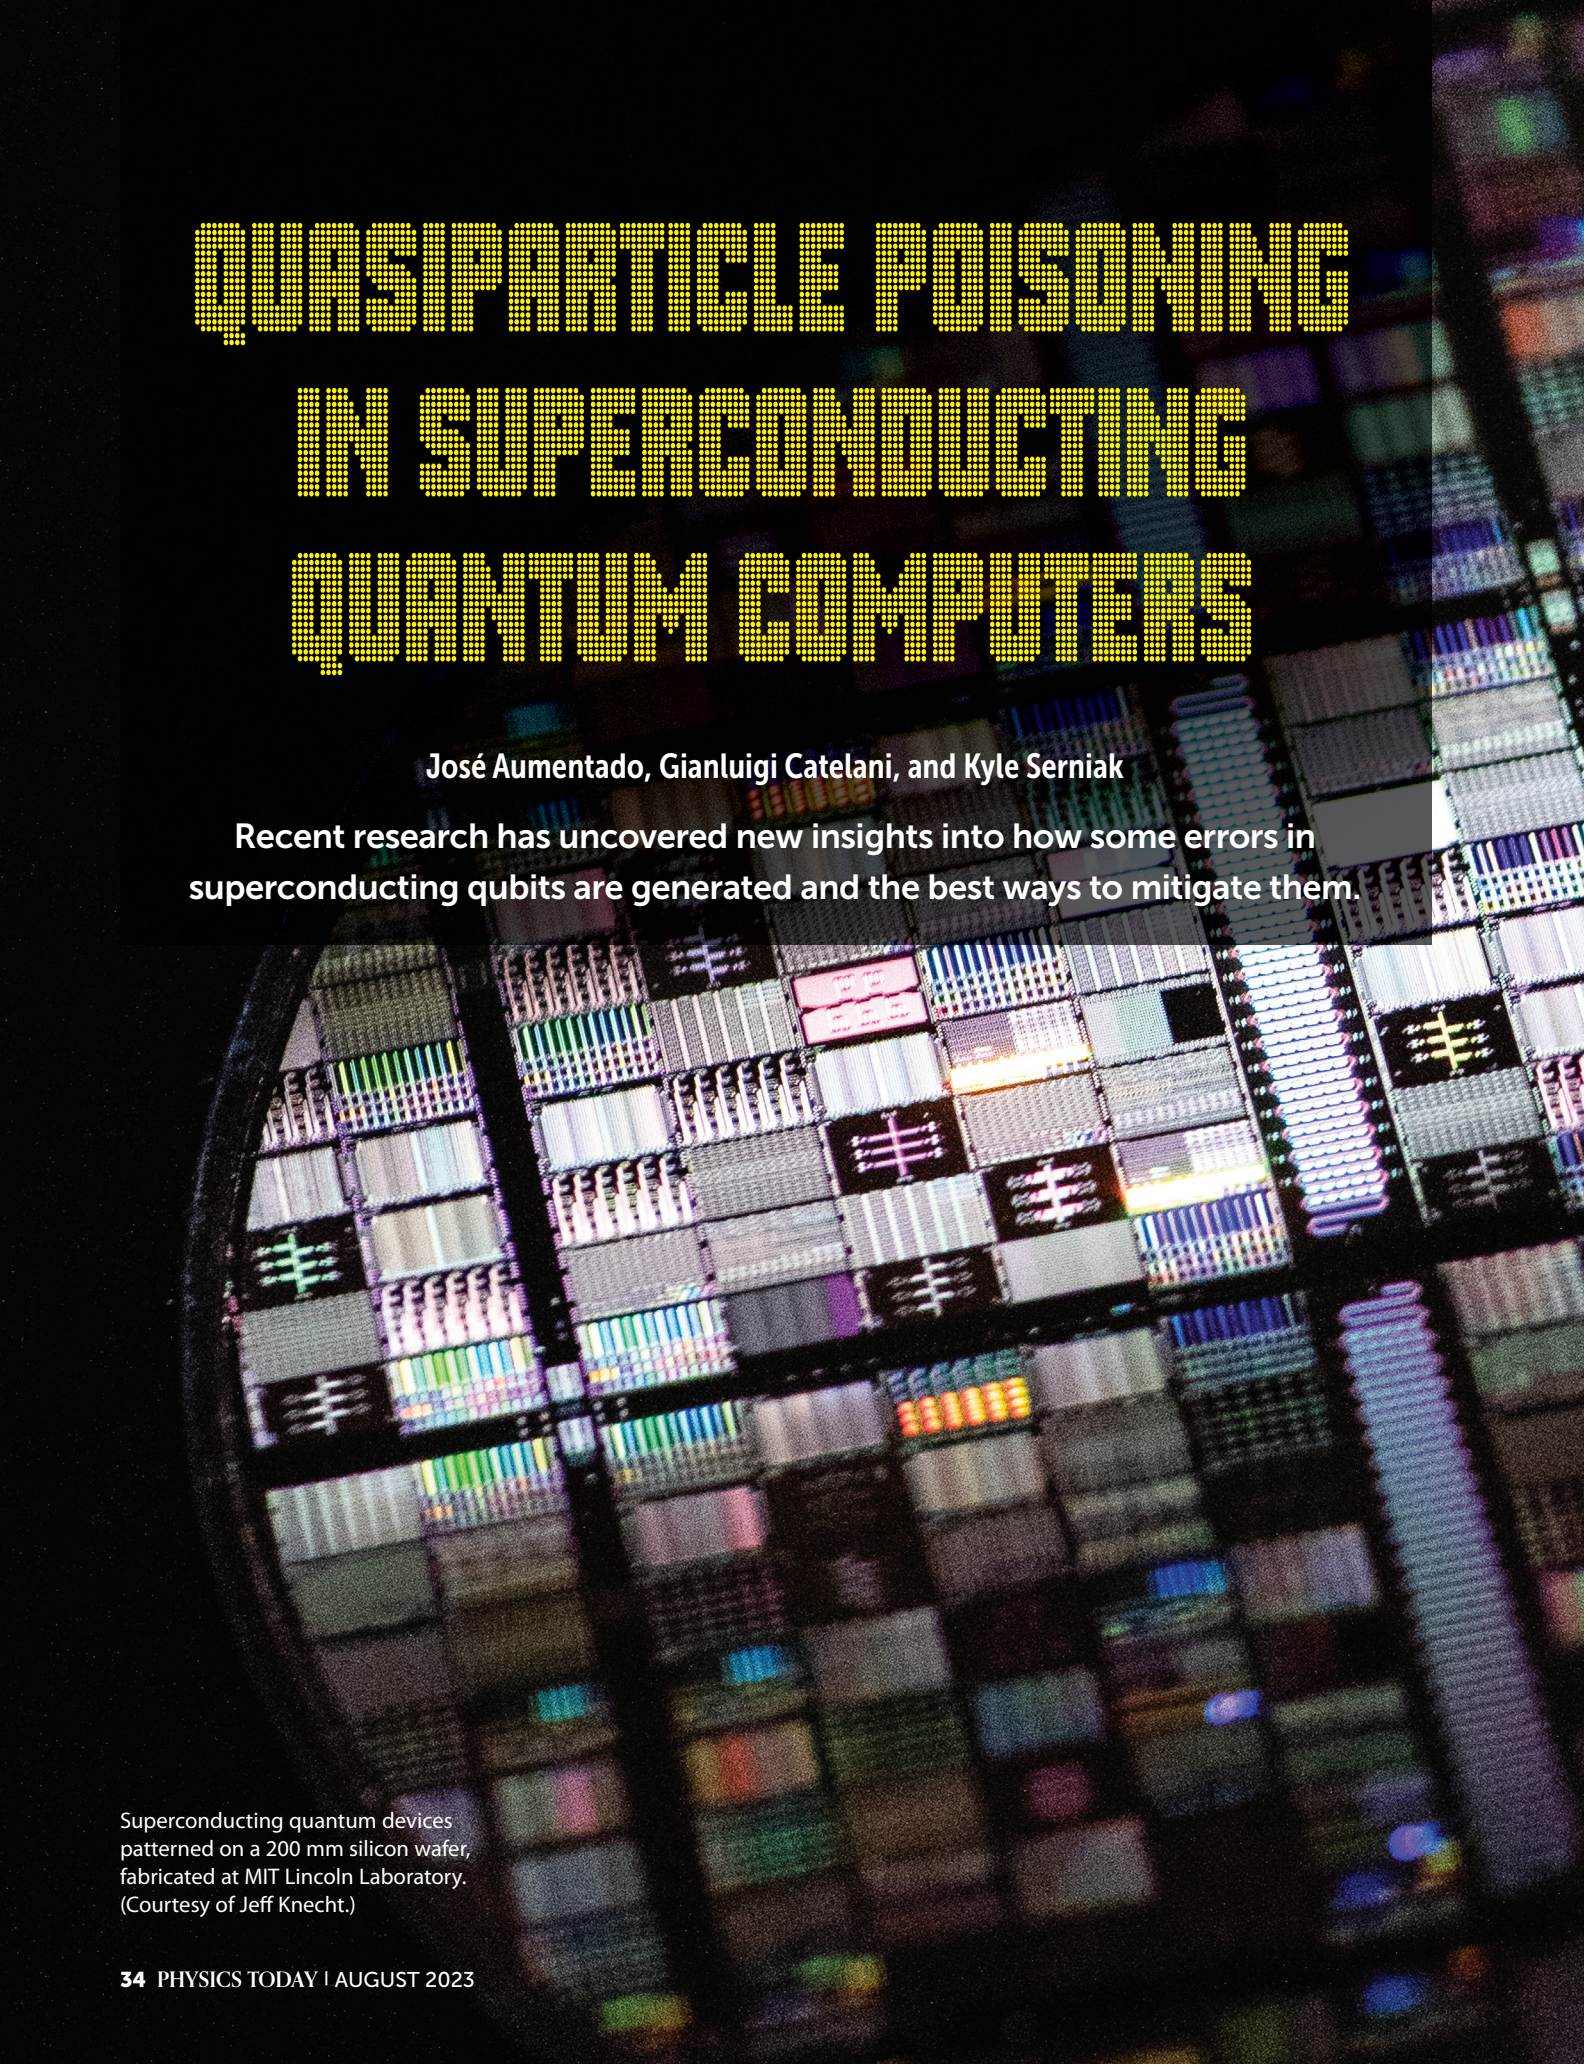

Superconducting quantum devices patterned on a 200 mm silicon wafer, fabricated at MIT Lincoln Laboratory. (Courtesy of Jeff Knecht.)

José Aumentado is a staff scientist in the Advanced Microwave Photonics Group at NIST in Boulder, Colorado, and is a senior fellow at Quantum Circuits Inc in New Haven, Connecticut. **Gianluigi Catelani** is a research scientist in the Peter Grünberg Institute at the Jülich Research Center in Germany and lead researcher in the Quantum Research Center at the Technology Innovation Institute in Abu Dhabi, United Arab Emirates. **Kyle Serniak** is a technical staff scientist at MIT Lincoln Laboratory in Lexington, Massachusetts, and at the MIT Research Laboratory of Electronics in Cambridge.


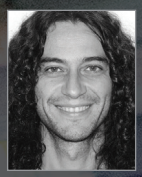


Although quantum computing is still in its infancy relative to the "classical" computing technology that we've come to know, love, and rely on, rapid advances over the past decade have taken it from the realm of science fiction to a probable reality of the not-so-distant future. Instead of manipulating bits of information by operating millions of transistors, a quantum computer relies on the precise control of many quantum subsystems—individual quantum bits, or qubits—along with an accurate readout of their quantum states. Many promising physical qubit platforms, such as trapped ions, neutral atoms, and solid-state defects (see the article by Christopher Anderson and David Awschalom on page 26), are based on building blocks that are typically thought of as archetypes of quantum behavior.

One of the leading candidate platforms for a useful quantum processor, however, is constructed from components that don't evoke a picture of tiny, microscopic particles with exotic properties. Instead, it consists of superconducting wires, capacitors, and inductors patterned on chips akin to existing semiconductor technologies. Those electronic circuits, which make up the superconducting qubit platform, embody many of the desirable properties of their atomic counterparts and have become the focus of several high-profile quantum computing efforts—led by both large companies, such as IBM, Google, and Alibaba, and startups, including Rigetti Computing, IQM, Alice & Bob, Oxford Quantum Circuits, and Quantware.[1] Those companies are leveraging modern clean-room fabrication tools to more easily engineer complex circuits with fast control.

In developing any quantum computing platform, a fundamental challenge arises from the tension between preserving quantum information and manipulating it: The former requires that qubits be isolated from their environment, while the latter demands that they have precise interactions with it. In fact, the key metrics for any platform can be summarized by the probability that an error will occur during a calculation and the time it will take to complete that calculation.

Currently, researchers looking at superconducting qubits are focusing on the error probability, which can be thought of as the ratio of how fast the qubit can be controlled to the rate at which it loses information to its environment. Of the primary mechanisms that are currently limiting superconducting qubit performance, one of the most intriguing and difficult to control is quasiparticle poisoning—the presence of charge carriers that do not participate in the superconducting condensate.

Superconducting qubits in a nutshell

Quantum effects are often weak and hard to observe in objects visible to the human eye. (For example, see Physics Today, July 2023, page 16.) So how is it that superconducting devices that are constructed from such circuit elements as inductors and capacitors and contain on the order of 10^{15} atoms behave quantum mechanically? As first shown by John Martinis, Michel Devoret, and John Clarke in 1987, a macroscopic degree of freedom can exhibit quantum behavior provided that energy dissipation is negligible and that the temperature of the system is low.

Thus the first ingredient to build a quantum circuit is to avoid energy dissipation, which leads to information loss. That's why circuit components are fabricated with superconducting materials. They can carry direct current without any resistance because the relevant charge carriers—electrons and holes near the Fermi energy—partner into Cooper pairs and condense into a macroscopic coherent state, as explained by the Bardeen-Cooper-Schrieffer (BCS) theory of superconductivity (see the article by Warren Pickett and Mikhail Eremets, Physics Today, May 2019, page 52). The condensate can be described by a complex-valued order parameter, the phase of which is critical to describe the physics of superconducting qubits.

Dissipationless transport is possible not only within bulk superconductors but also between two connected superconductors separated by what's called a weak link. The most widely used type of weak link is a tunnel barrier—a thin oxide layer separating two superconducting electrodes to form a Josephson junction. Importantly, Josephson junctions behave as nonlinear inductors: They lie at the heart of superconducting qubits, and the difference in the phase of the order parameter between the superconductors they connect is exactly the macroscopic degree of freedom that was shown to exhibit quantum behavior. In practice, aluminum is the superconductor of choice for Josephson junctions because it's compatible with relatively standard nanofabrication techniques and has a self-limiting few-nanometers-thick oxide at its surface, which is used for the junction barrier.

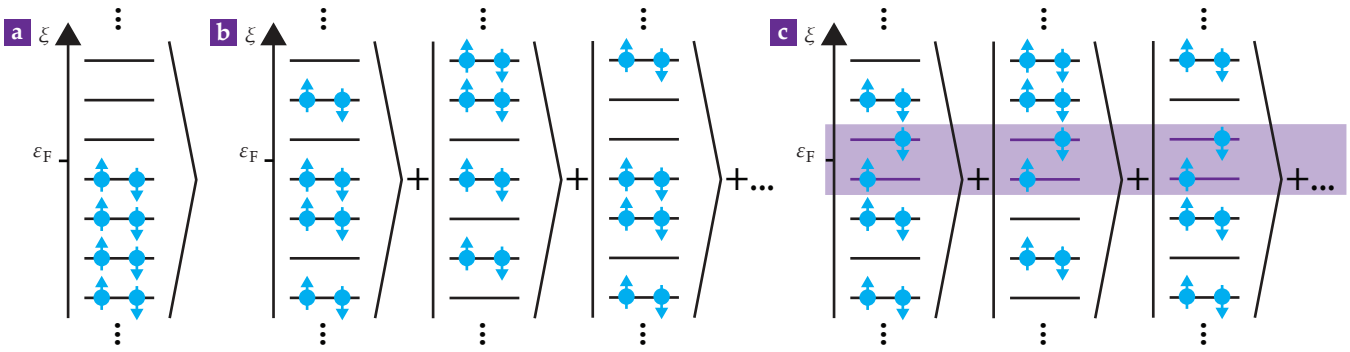

FIGURE 1. QUASIPARTICLE EXCITATIONS in superconductors. **(a)** In the ground state of a normal metal, spin-degenerate electrons (blue) occupy states with energy ξ up to the Fermi level ε_F. **(b)** The ground state of a Bardeen-Cooper-Schrieffer superconductor consists of a coherent superposition of all possible configurations of states, which have pair-correlated electron occupation in an energy window $\pm\Delta$ around the Fermi level. For simplicity, panels a and b neglect degeneracy or correlation in the momentum of the electrons. **(c)** When a phonon or photon with energy greater than 2Δ couples to the superconductor, the generated pair of quasiparticles poisons the superconductor: The two states the quasiparticles occupy (purple region) are fixed and don't participate in the coherent superposition of the superconducting condensate. (Adapted from ref. 9.)

The design flexibility of superconducting circuits originates from the many possible ways of combining the three basic circuit elements—capacitors, linear inductors, and nonlinear inductors (Josephson junctions), which all have parameters that can be tuned over a wide range. Is there a price to pay for such flexibility? Depending on how the components are arranged, quantum information can be encoded into the charge or the phase difference between superconducting condensates or as a combination of the two. The encoding methods hint at what can go wrong: The charge, the phase, or even the superconducting condensate itself can be disrupted.

Broadly speaking, the environmental effects acting on the charge or phase are known as charge noise and flux noise, respectively. They arise from materials defects and imperfections on the surface of the superconductor, at the interface with the substrate, in the oxide forming the Josephson junction, and in the substrate itself. At the microscopic scale, the sources of charge noise and flux noise arise from random changes in the configurations of charges and electron or nuclear spins.[2]

Another decoherence mechanism affecting charge and phase arises from the interaction of the superconductor with the electromagnetic environment: Like any other resonant electric circuit, a superconducting qubit can lose energy by emitting a photon. That's easy to visualize for the simplest superconducting qubit, called a transmon. Consisting of a Josephson junction in parallel with a capacitor, a transmon can be thought of as a nonlinear dipole antenna, which absorbs and emits photons at some characteristic frequency.

In contrast to the decoherence mechanism described above, the superconducting condensate can be directly disturbed by the environment via the breaking of Cooper pairs, a process that generates quasiparticle excitations in the superconductor itself. Cooper pairs comprise two electrons with opposite spin and momentum, and superconductivity results from the coherent superposition of the underlying many-body momentum states, which are either pair-occupied (electrons) or pair-unoccupied (holes), as illustrated in figure 1.

Picturing quasiparticles as broken Cooper pairs gives an idea of what they actually are. In a normal metal, electrons occupy various energy levels in a so-called Fermi sea, and when an electron is removed, what's left is a hole excitation. When removing an electron that was part of a Cooper pair, what's left is a coherent superposition of an electron and a hole, known as a Bogoliubov quasiparticle.

Whereas any small amount of energy is sufficient to generate an electron and a hole in a normal metal, it takes a finite energy, denoted as 2Δ, to break a Cooper pair. That energy, known as the superconducting gap, is proportional to the critical temperature T_c at which the superconductivity disappears: $\Delta \simeq 1.76\, k_B T_c$ for well-behaved BCS superconductors, such as aluminum. Because of the energy gap, at low temperature the thermally activated number of quasiparticles, which can be quantified as the fraction x_{QP} of broken Cooper pairs, should be exponentially small, $x_{QP} \sim \exp(-\Delta/k_B T)$. For aluminum at about 20 mK—the temperature at which aluminum-based superconducting qubits are typically operated—x_{QP} is expected to be about 10^{-46}, which is so small that in an Earth-sized block of superconducting aluminum, one would expect to find only two thermally excited quasiparticles. Unfortunately, as we will describe later, observed values of x_{QP} are much larger than expected.

So what happens if quasiparticles are present in a superconducting circuit? In bulk superconductors, they're responsible for finite AC dissipation proportional to x_{QP}. In qubit circuits comprising Josephson junctions, the situation is more complex. When a quasiparticle tunnels from one side of a junction to the other, its coupling to the phase difference across that junction makes it possible for the quasiparticle to absorb energy from the qubit, causing the qubit to decay. Similar to the dissipative response of bulk superconductors, the decay rate is proportional to x_{QP}. Even if the quasiparticle does not absorb energy, when it tunnels it can make the qubit frequency fluctuate, which leads to dephasing and a reduction of the qubit's coherence time. Both energy decay and dephasing originate from the dependence of the tunneling amplitude on the phase difference and have been investigated in a number of theoretical and experimental works (see references 3 and 4 and references therein).

The decoherence mechanisms are generic to any super-

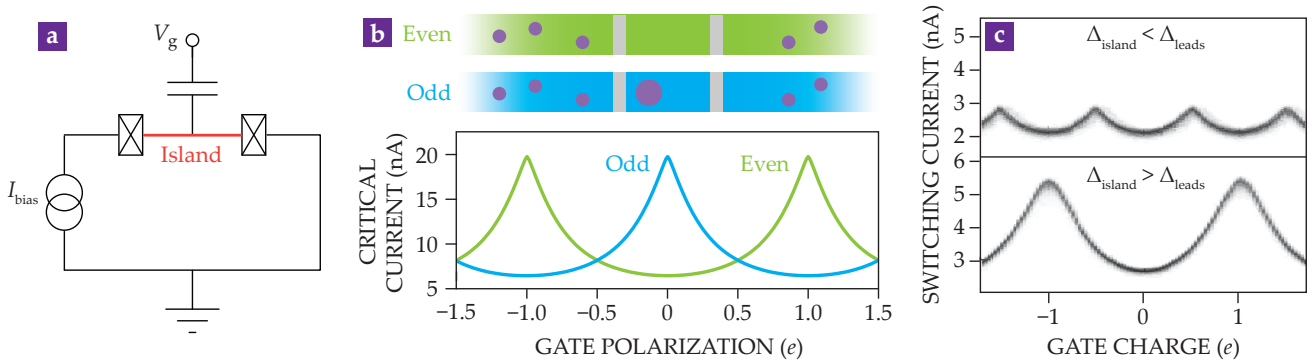

FIGURE 2. SUPERCONDUCTING CIRCUIT. (a) A Cooper-pair transistor circuit features two small Josephson junctions that isolate a submicron-scale superconducting island (red). **(b)** An odd parity state (blue) corresponds to an excess electron on the island, and an even parity state (green), to no excess electron. The effective critical current through the island modulates with an applied gate voltage V_g that corresponds to a change in the energy cost of placing additional Cooper pairs on the island. **(c)** The switching current, which is closely related to the critical current, has a value at a given gate voltage that reflects the presence or absence of quasiparticles poisoning the island charge state. The dips at $\pm 1e$ indicate that single quasiparticles occupy the island more often than not (top). The opposite (bottom) is true when the relative gap energy Δ of the superconducting island and superconducting lead is inverted. (Adapted from ref. 5.)

conducting qubit made with junctions, but different qubit designs have different sensitivities. In fact, qubits with junctions embedded in a superconducting loop can be tuned by threading a magnetic flux through that loop, and the sensitivity to quasiparticles can be suppressed at particular flux values known as sweet spots. The suppression is an interference effect that manifests the nature of quasiparticles as a coherent superposition of electron- and hole-like excitations. At the sweet spots, the sensitivity to flux noise is also minimized, making them by far the preferred operating point for such qubits.

The quasiparticle mystery

As mentioned above, no thermally excited quasiparticles should be present at temperatures sufficiently below T_c. Aluminum circuits with T_c = 1.2 K and at dilution refrigerator temperatures of 10 mK should be completely free of quasiparticles. So why worry about them at all?

In the 1990s several groups studied a class of superconducting charge-sensitive circuits that leveraged the so-called Coulomb blockade effect. In those devices, one or more submicrometer-scale superconducting islands were weakly coupled to connected electrodes by Josephson junctions. Importantly, the small size of the islands and junctions—typically no larger than 100 nm × 100 nm—fixed the islands' total capacitance C_Σ to less than a femtofarad. At that level, the corresponding charging energy for adding a single Cooper pair, $E_C = 2e^2/C_\Sigma$, where e is the electron charge, could easily exceed 10^{-23} J, or 1 K in temperature units.

In that parameter regime, the critical current and other electronic properties were sensitive to the addition or subtraction of single Cooper pairs and quasiparticles. Although quasiparticles do not have definite charge, when they tunnel on or off a superconducting island, the total charge on that island is shifted by the discrete value $\pm e$.

One of the simplest Coulomb blockade circuits is the single Cooper-pair transistor.[5] As shown in figure 2, the device has two small Josephson junctions that isolate a single superconducting island from superconducting leads, and a capacitively coupled gate electrode is placed nearby. In that configuration, the two junctions behave effectively as a single Josephson junction. Its critical current—the maximum current that the junction can carry while keeping the voltage across the junction close to zero—modulates with an applied gate voltage. Ideally, the modulation is a $2e$-periodic function of the gate charge $q_g = C_g V_g$ (where C_g is the gate capacitance to the island, and V_g is the gate voltage) and reflects the size of the Cooper-pair charge itself. As noted above, the presence of quasiparticles in the leads provides a source for single electrons to tunnel onto the island and offset the island's charge by an electron, which concomitantly shifts the current modulation by $1e$.

Many experimentalists therefore regarded a $1e$-periodic modulation to be indicative of the presence of quasiparticles. Indeed, one could turn a $2e$-periodic modulation into a $1e$-periodic modulation just by heating up the device to a few hundred millikelvin to create an abundance of thermally generated quasiparticles. It was common, however, to see $1e$-periodic modulation at much lower temperatures, even when controlling for other known causes of the behavior. It's known as quasiparticle poisoning, and its sporadic presence in some, but not all, devices was one of the first indications that the physics of quasiparticles was not fully understood.

Using a higher-speed DC measurement technique in the early 2000s, one of us (Aumentado) found evidence for quasiparticles at dilution-refrigerator temperatures, even in $2e$-periodic devices. The results showed that the tunneling of nonequilibrium quasiparticles on and off the island was sensitive to both gate voltage and the relative gap energies of the island and leads. Single Cooper-pair transistors share many things in common with today's superconducting qubit circuits, including the junction sizes and material choice of aluminum, and perhaps that's why it's not surprising that the basic phenomenon of nonequilibrium quasiparticle poisoning has persisted to the present day.

To probe the dynamics of nonequilibrium quasiparticles in superconducting qubits and test our understanding of quasiparticle poisoning, researchers have used many approaches over the years. For example, one can purposely add quasiparticles by increasing the system's temperature and then measuring such properties as the relaxation time T_1 (typically tens to hundreds of microseconds) and the qubit frequency ω_{10} (a few

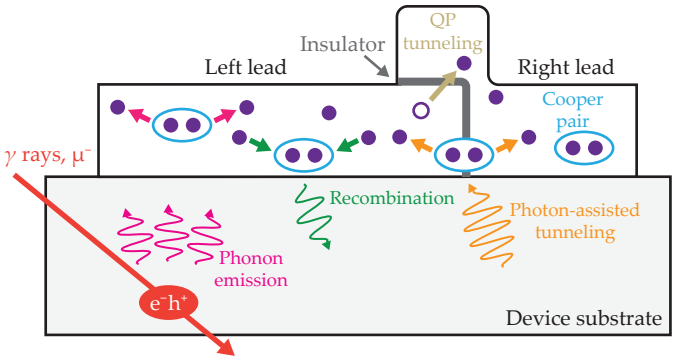

FIGURE 3. QUASIPARTICLE DYNAMICS. A Josephson junction, formed by a superconductor-insulator-superconductor heterostructure, is shown in cross section. Quasiparticles (purple) can undergo various inelastic processes. Some tunnel across the Josephson junction (yellow) and others are generated during photon-assisted tunneling of Cooper pairs (orange). Both processes can cause energy exchange between the quasiparticles and a qubit formed in part from the junction. Ionizing radiation can create in the substrate electron–hole pairs (red), which emit showers of phonons (pink) as they relax. Phonons with an energy of 2Δ or greater are sufficiently energetic to break Cooper pairs; freshly created quasiparticles in the device then lead to spatiotemporally correlated errors. Quasiparticles can also recombine and emit a phonon with energy greater than 2Δ (dark green).

gigahertz). Both those properties decrease when quasiparticles are present.[6]

Alternatively, nonequilibrium quasiparticles can be injected directly without raising the system temperature, and the expected relation between changes in T_1 and ω_{10} can be checked.[7] In fact, researchers have exploited the proportionality between $1/T_1$ and the quasiparticle density x_{QP} to monitor the dynamics of x_{QP}, and they have assessed to what extent quasiparticles were trapped by supercurrent vortices.[8] Such experiments also make it possible to place bounds on the density of nonequilibrium quasiparticles and to estimate their generation rate.

A more direct measure of quasiparticle effects in qubits is similar to the initial observations of $1e$ periodicity in single Cooper-pair transistors.[5] By explicitly reintroducing some charge sensitivity into a transmon circuit, researchers detected quasiparticle-induced errors via a correlated change in the odd–even "charge parity" of the circuit over a time τ_{QP} (see reference 9 and references therein). From those experiments, it's clear that modern-day superconducting qubits are still plagued by nonequilibrium quasiparticle poisoning.

Sources of nonequilibrium quasiparticles

Once physicists accepted that nonequilibrium quasiparticles were present in their superconducting devices, a simple question remained: Why? The answer boils down to the erroneous assumption that everything a qubit "sees" is perfectly isolated from the outside world and well-thermalized to the coldest stage of the cryostat. For low-noise experiments with superconducting qubits, researchers take a lot of care to filter and shield any unwanted noise. But qubits aren't ever completely sheltered. All it takes to produce a pair of quasiparticles in an otherwise isolated superconductor is an excitation with an energy greater than 2Δ, which for commonly used thin aluminum films corresponds to approximately 100 GHz, 5 K, or 400 μeV, depending on the preferred choice of units. That's not a lot of energy!

The many years that researchers have spent developing superconducting detectors have led to valuable insights into the dynamics of nonequilibrium quasiparticles. Figure 3 summarizes how all sorts of bad actors, including stray IR photons, mechanical vibrations of the device, and—most troubling of all—ionizing radiation from radioactive decay products and cosmic-ray secondary particles generate quasiparticles in qubits.

IR photons can leak into the experimental region of a cryostat, despite the best attempts to block or shield from them. Many popular cryogenic systems, including dilution refrigerators, consist of multiple temperature stages. Similar to a set of nested matryoshka dolls, a metal shield at each stage protects the next from the surrounding, hotter stage (see figure 4). The innermost shield should be thermalized to the lowest-temperature stage of the cryostat. Experiments with superconducting resonators, however, indicated that more shielding was needed: Some photons from higher-temperature stages can get through and reduce device performance.[10] Coating the experiment with IR-absorbing material is one remedy. It's the same principle that's used when painting stealth aircraft.

Researchers recently discovered that the qubit itself can act as an antenna that enhances the production of quasiparticles via absorption of IR radiation.[11] The absorption process is localized at the Josephson junctions of a qubit circuit; in addition to qubit relaxation, the process can explain recent observations of especially large qubit excitation rates.[9] Experiments have since demonstrated that the process does indeed contribute to quasiparticle generation and qubit excitation and that the process can be suppressed by improved filtering of the microwave lines feeding signals to the qubits and by proper design of the qubit and its surroundings.[9,12] Those improvements can lengthen by several orders of magnitude the time between quasiparticle tunneling events, from shorter than a millisecond to longer than a second.

Ionizing radiation is known to also produce quasiparticles in superconducting devices, and in many cases that's the desired effect. So-called pair-breaking detectors, such as microwave kinetic inductance detectors and transition-edge sensors, operate on the principle that ionizing radiation and other excitations deposit large amounts of energy into the crystalline device substrate in the form of ionized charge carriers and showers of high-energy phonons. In superconducting detectors, the phonons can produce quasiparticles, whose presence is inferred from a change in an observable parameter, such as kinetic inductance or critical current.

Although superconducting qubits are similar in construction to those types of detectors, it was only in hindsight that researchers realized that superconducting qubits could also act as detectors of ionizing radiation, with detection events translating into computational errors. Ionizing radiation reduces the performance of qubits.[13] Some of it, primarily γ rays, can be shielded by lead, but to cut down on the flux of pesky cosmic-ray muons, one needs to use the overburden of Earth's crust or to go deep underwater.[14]

The mechanism of quasiparticle production via cosmic-ray muons is particularly worrisome because about every 10 seconds a muon can generate bursts of quasiparticles throughout a device and knock out many nearby qubits simultaneously.[15]

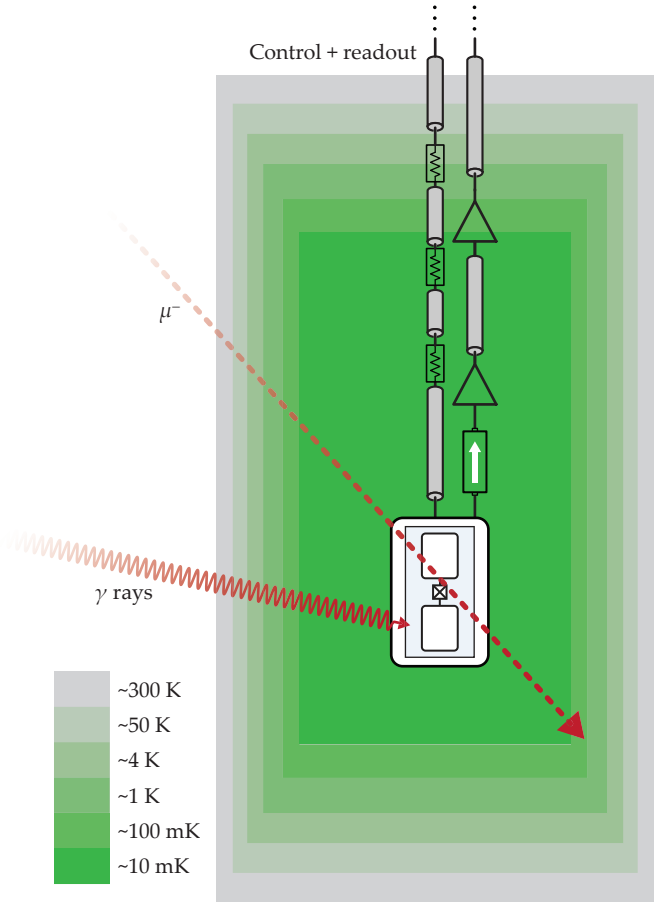

FIGURE 4. SUPERCONDUCTING QUBIT EXPERIMENTS often use dilution refrigerators with nested temperature stages. Each stage includes a metallic shield that blocks blackbody radiation from higher-temperature stages. Gamma rays and cosmic-ray muons, however, can penetrate through that shielding, sometimes hitting the superconducting quantum processor and creating spatiotemporally correlated, quasiparticle-induced errors.

Similar bursts were recently linked to mechanical relaxation of superconducting devices over the time scale of days. The link could explain an earlier observation of a slow decay in the generation rate over the course of an experiment. Those types of quasiparticle-induced spatiotemporally correlated errors are difficult to deal with in many quantum error-correction schemes, although they can be addressed if they're detected independently and if qubits likely to have been affected by errors can be excluded from further computation.[16]

Toward robust quantum computing

Qubit performance has improved by several orders of magnitude in the 25 years since the first demonstration of coherence in a superconducting qubit, but there is still a long road ahead. The consensus in the research community is that quantum error-correction techniques will be necessary to maintain complex multiqubit-state information for the duration of a useful computation. In such schemes, logical qubits are encoded in the combined state of many error-prone qubits, and higher error rates translate into stricter requirements on the total number of physical qubits.

An underlying assumption typical of quantum error-correction schemes is that physical errors are random. Using that thinking, researchers have steadily chipped away at the background population of nonequilibrium quasiparticles and suppressed their steady-state contribution to qubit errors to a sufficient level over time. But that assumption is violated by the aforementioned error bursts that arise from quasiparticles generated by ionizing radiation.

Luckily, there are many proposed—and some demonstrated—paths toward mitigating catastrophic error bursts. Having quasiparticles around is ok, so long as they don't tunnel across a qubit's Josephson junction. That could be achieved by using a superconductor for the ground plane with a smaller energy gap than the qubit superconductor or by adding normal-metal islands to the back of the chip.[17] Those design changes bring the energy of the phonons generated by radiation hits to below the gap of the qubit material, so that they cannot break Cooper pairs anymore. The few quasiparticles that are still generated in the qubit bulk can be kept away from the qubit's junctions by employing quasiparticle traps[9,18] or blocked from tunneling at the junctions via gap engineering.[5]

While those "on-chip" techniques are effective for many sources of quasiparticles, pesky cosmic-ray secondary particles such as muons are not effectively shielded except by massive amounts of material, which has led some scientists to suggest that underground facilities are critical to avoiding spatiotemporally correlated error bursts. Luckily for experimentalists who enjoy sunlight, there is hope that on-chip mitigation strategies could be combined with tungsten or lead shielding to provide sufficient protection. But such radiation-hardened superconducting qubits have yet to be fully demonstrated.

Nonequilibrium quasiparticles might sound like a bogeyman lurking in the shadows of superconducting quantum computing efforts, but they are just another item in the list of engineering and scientific challenges that must be met to make quantum computing a robust reality. There are many reasons to be optimistic: Recent research efforts have given more insight into quasiparticles' generation mechanisms and have provided a clear direction for future mitigation efforts.

REFERENCES

1. P. Krantz et al., *Appl. Phys. Rev.* **6**, 021318 (2019).
2. I. Siddiqi, *Nat. Rev. Mater.* **6**, 875 (2021).
3. L. I. Glazman, G. Catelani, *SciPost Phys. Lect. Notes* (2021), doi:10.21468/SciPostPhysLectNotes.31.
4. G. Catelani, J. P. Pekola, *Mater. Quantum Technol.* **2**, 013001 (2022).
5. J. Aumentado et al., *Phys. Rev. Lett.* **92**, 066802 (2004).
6. G. Catelani et al., *Phys. Rev. Lett.* **106**, 077002 (2011).
7. M. Lenander et al., *Phys. Rev. B* **84**, 024501 (2011).
8. C. Wang et al., *Nat. Commun.* **5**, 5836 (2014).
9. K. Serniak, "Nonequilibrium quasiparticles in superconducting qubits," PhD thesis, Yale U. (2019).
10. R. T. Gordon et al., *Appl. Phys. Lett.* **120**, 074002 (2022).
11. M. Houzet et al., *Phys. Rev. Lett.* **123**, 107704 (2019).
12. T. Connolly et al., https://arxiv.org/abs/2302.12330.
13. A. P. Vepsäläinen et al., *Nature* **584**, 551 (2020).
14. L. Cardani et al., *Nat. Commun.* **12**, 2733 (2021).
15. M. McEwen et al., *Nat. Phys.* **18**, 107 (2022).
16. Q. Xu et al., *Phys. Rev. Lett.* **129**, 240502 (2022).
17. V. Iaia et al., *Nat. Commun.* **13**, 6425 (2022).
18. R.-P. Riwar et al., *Phys. Rev. B* **94**, 104516 (2016).

PT