

Coherent suppression of electromagnetic dissipation due to superconducting quasiparticles

Ioan M. Pop^{1*}, Kurtis Geerlings^{1*}, Gianluigi Catelani^{1,2}, Robert J. Schoelkopf¹, Leonid I. Glazman¹ & Michel H. Devoret¹

Owing to the low-loss propagation of electromagnetic signals in superconductors, Josephson junctions constitute ideal building blocks for quantum memories, amplifiers, detectors and high-speed processing units, operating over a wide band of microwave frequencies. Nevertheless, although transport in superconducting wires is perfectly lossless for direct current, transport of radio-frequency signals can be dissipative in the presence of quasiparticle excitations above the superconducting gap¹. Moreover, the exact mechanism of this dissipation in Josephson junctions has never been fully resolved experimentally. In particular, Josephson's key theoretical prediction that quasiparticle dissipation should vanish in transport through a junction when the phase difference across the junction is π (ref. 2) has never been observed³. This subtle effect can be understood as resulting from the destructive interference of two separate dissipative channels involving electron-like and hole-like quasiparticles. Here we report the experimental observation of this quantum coherent suppression of quasiparticle dissipation across a Josephson junction. As the average phase bias across the junction is swept through π , we measure an increase of more than one order of magnitude in the energy relaxation time of a superconducting artificial atom. This striking suppression of dissipation, despite the presence of lossy quasiparticle excitations above the superconducting gap, provides a powerful tool for minimizing decoherence in quantum electronic systems and could be directly exploited in quantum information experiments with superconducting quantum bits.

Despite the success of Josephson's theoretical predictions, the tunnelling current across a Josephson junction in the presence of quasiparticles is still a subject of controversy³. It is customary to distinguish four contributions to the current through a superconducting tunnel junction, associated with Cooper pairs and quasiparticles^{2,3}: (1) the current of Cooper pairs, (2) the dispersive quasiparticle term, (3) the phase-independent dissipative term, and (4) the dissipative term proportional to the cosine of the phase difference. Contributions (1) and (2) comprise the supercurrent (that is, the Josephson current), which was measured shortly after its prediction and is nowadays extensively used in applications⁴. The cosine term, also known as quasiparticle-pair interference⁵, has eluded conclusive experimental observation³. Only very recently, thermal transport measurements showed the effect of quasiparticle-pair interference on the heat conductivity of a superconducting quantum interference device⁶. In a phase-biased Josephson junction, we expect the cosine and dissipative quasiparticle terms in the tunnelling current to add coherently and display a periodic dependence on flux^{7–9}.

The physical significance of the contributions (1)–(4) to the tunnelling current is best illustrated by the linear response of the junction to an a.c. voltage bias of frequency ω , $V(\omega)$, at a given d.c. phase bias, φ . This junction admittance, relating current and voltage by $I(\omega, \varphi) = Y(\omega, \varphi)V(\omega)$, takes the form

$$Y(\omega, \varphi) = Y_{\text{qp}}(\omega) \frac{1 + \varepsilon \cos \varphi}{2} - i \frac{\cos \varphi}{L_J \omega} \left(1 - 2x_{\text{qp}}^{\text{A}} \right) \quad (1)$$

The purely imaginary last term in equation (1) comes from the dissipationless supercurrent through a Josephson junction (contributions (1) and (2)), where L_J is known as the Josephson inductance. The factor x_{qp}^{A} accounts for the occupation of the 'Andreev' bound states, which carry the supercurrent^{10,11}; writing contribution (2) in this form allows for a generalization to a non-equilibrium distribution of quasiparticles (Methods). The complex function $Y_{\text{qp}}(\omega)$ depends on the bulk quasiparticle population through the density x_{qp} (normalized by the Cooper pair density)⁷, and its real part accounts for dissipation in the junction.

According to standard Bardeen–Cooper–Schrieffer theory, for junctions with identical electrode materials, the $\varepsilon \cos \varphi$ term (corresponding to contributions (3) and (4)) approaches unity at temperatures well below the critical temperature for superconductivity¹² and we expect the quasiparticle dissipation to be suppressed at $\varphi = \pi$ (ref. 8). There have been several attempts to measure the sign and amplitude of ε (refs 13, 14); however, owing to the intrinsic dissipative nature of the d.c. experimental methods available in the past, the results were difficult to interpret and proved inconclusive³. In this Letter we present direct measurements of dissipation as a function of the phase bias, φ , over a Josephson junction. We observe a sharp decrease in dissipation at $\varphi = \pi$, thus putting stringent bounds on the value of ε .

In principle, for commonly used superconductors such as aluminium and niobium, the quasiparticle fraction, x_{qp} , should ideally reach its thermal equilibrium value, which is practically zero ($x_{\text{qp}}^{\text{th}} \ll 10^{-20}$) at typical dilution refrigerator temperatures of $\lesssim 40$ mK. However, in recent years it has become increasingly clear that non-thermal quasiparticles can persist in superconducting circuits even at very low temperatures, inducing relaxation and dephasing^{15–20}. Using a fluxonium artificial atom²¹, which consists of a Josephson junction shunted by a superinductor²² (Fig. 1b), we directly measure the dissipation across a flux-biased Josephson junction due to these non-thermal quasiparticles. Superconducting quantum bits (qubits) are ideal testing systems for different dissipation mechanisms. They can be manipulated and measured at the single-photon level and their susceptibility to different loss mechanisms can be tuned *in situ*. We measure the dissipation rate by preparing the qubit in an excited state and observing its energy relaxation. In Fig. 1a, we present the schematic mechanism of quasiparticle dissipation across a phase-biased Josephson junction, the main element of a fluxonium artificial atom, whose first two energy levels constitute a qubit. A quasiparticle from the left electrode can tunnel through the junction while absorbing the qubit energy, resulting in qubit relaxation to the ground state. The quasiparticle can tunnel either as an electron or as a hole, with probability amplitude phases $+\varphi/2$ or $-\varphi/2$, respectively⁷. Because these two events are indistinguishable, the probability amplitudes add coherently to produce the ' $1 + \cos \varphi$ ' term in the total probability: $|e^{i\varphi/2} + e^{-i\varphi/2}|^2 = 2(1 + \cos \varphi)$.

To identify and quantify the multiple decay mechanisms acting on the qubit, we need to be able to decompose the total relaxation rate into its constituents. The fluxonium qubit (Fig. 1b) is a particularly attractive tool because its susceptibility to each loss mechanism has a distinctive functional dependence on applied magnetic flux, Φ_{ext} (ref. 23),

¹Department of Applied Physics, Yale University, 15 Prospect Street, New Haven, Connecticut 06511, USA. ²Forschungszentrum Jülich, Peter Grünberg Institut (PGI-2), 52425 Jülich, Germany.

*These authors contributed equally to this work.

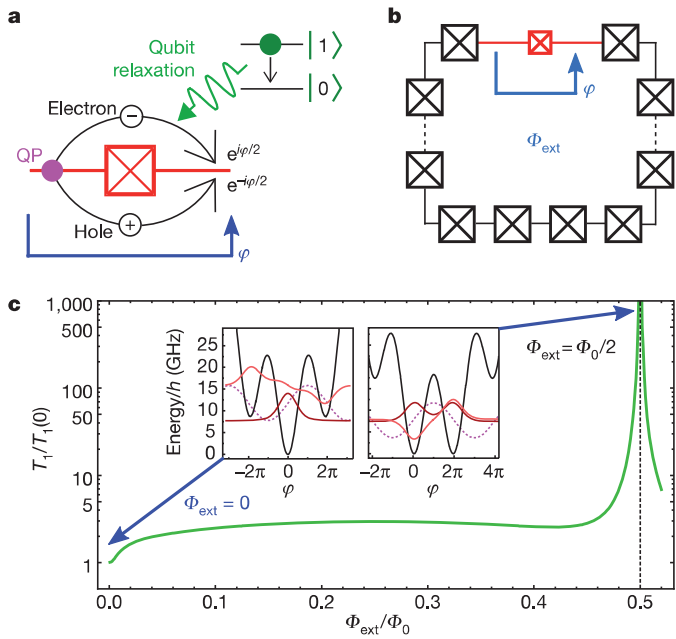


Figure 1 | Phase-dependent quasiparticle dissipation in a Josephson junction. **a**, In the case of a Josephson junction biased at phase φ , quasiparticles (QPs) receiving a quanta of excitation from the qubit or environment can tunnel across the junction either as electrons, acquiring a phase $+\varphi/2$, or as holes, acquiring a phase $-\varphi/2$. The φ dependence of quasiparticle dissipation results from the interference of these two indistinguishable paths. **b**, Schematic representation of the fluxonium qubit: the phase-slip junction (red) is shunted by an array of bigger junctions (black) that form the superinductor. In the fluxonium qubit, the external magnetic flux, Φ_{ext} , sets the average phase difference, φ , across the phase-slip junction. **c**, Flux dependence of the fluxonium relaxation time, T_1 , due to quasiparticle loss across the phase-slip junction. Quasiparticle loss is suppressed at $\Phi_{\text{ext}}/\Phi_0 = 0.5$, where the average phase difference over the small junction is π . The insets show the potential landscape (black), the qubit ground-state ($|0\rangle$; dark red) and excited-state ($|1\rangle$; light red) wavefunctions, and the quasiparticle dissipation operator⁷, $\sin(\varphi/2)$ (dashed magenta). The T_1 divergence at $\Phi_{\text{ext}}/\Phi_0 = 0.5$ is explained by the symmetry of $\sin(\varphi/2)$ around $\varphi = \pi$, together with the respective even and odd parities of $|0\rangle$ and $|1\rangle$, which lead to the vanishing of the quasiparticle matrix element $\langle 0 | \sin(\varphi/2) | 1 \rangle = 0$ (compare with equation (2)). Here the symbol φ represents the dynamical coordinate associated with the operator $\hat{\varphi}$.

effectively providing a fingerprint for the dominant dissipation. To illustrate this idea, in Fig. 1c we represent the expected magnetic flux dependence of the energy relaxation time of the fluxonium qubit, T_1 , assuming only quasiparticle dissipation. The T_1 peak around $\Phi_{\text{ext}} = 0.5\Phi_0$, where Φ_0 is the flux quantum, is predicted by the coherent cancellation of the dissipative part of equation (1), and is the signature of quasiparticle dissipation in the phase-slip junction⁸, so called because its phase difference fluctuates by an amount of the order of 2π .

In Fig. 2a, we present an electron-beam image of an actual device. The fluxonium artificial atom consists of a small Josephson junction




Figure 2 | Experimental set-up. **a**, Electron-beam image of fluxonium sample. The schematic diagram in the upper left corner is a three-dimensional model of the junction array. The Al/AIO_x/Al junctions are made using bridge-free, double-angle evaporation³⁰. SQUID, superconducting quantum interference device. **b**, Optical image of the antenna with the fluxonium qubit in the middle. **c**, Photograph of the sapphire chip inside the copper cavity (sample holder). **d**, Electrical scheme for the fluxonium qubit coupling to the input-output microwave lines.

providing the nonlinearity, in parallel with the ‘superinductor’, which shunts charge noise²¹. To suppress charge fluctuations effectively, the superinductor should provide, in the gigahertz domain an impedance larger than the resistance quantum $R_Q = h/(2e)^2 = 6.5 \text{ k}\Omega$, where h is Planck’s constant and e is the electron charge. Recent studies^{24,25} have demonstrated that superinductors can be reliably implemented using arrays of Josephson junctions. In our design (Fig. 2a), we use an array of 95 junctions, which yields a total inductance of 330 nH, or an equivalent impedance of 20 k Ω at 10 GHz. The relationship between the applied flux and the phase bias is $E_J \sin \varphi + E_L(\varphi - 2\pi\Phi_{\text{ext}}/\Phi_0) = 0$, where E_J is the Josephson energy of the phase-slip junction and E_L is the inductive energy of the superinductor.

In general, a superconducting qubit is simultaneously subjected to multiple dissipation channels, including capacitive, inductive, radiative and quasiparticle channels. To observe and identify one particular loss mechanism, in our case quasiparticle loss, we need to increase the characteristic times of all other relaxation channels so as to exceed the quasiparticle-induced T_1 . Recent experiments on transmon qubits suggest that this time is expected to be in the millisecond regime²⁰. To achieve this goal, we adopted a design strategy similar to the three-dimensional transmon set-up²⁶. The sapphire substrate of the fluxonium qubit is clamped in the middle of a copper waveguide cavity (Fig. 2c), which constitutes our sample holder. The coupling to the qubit is mediated by an antenna (Fig. 2b) that couples inductively to the qubit through the shared junctions, and capacitively to the waveguide cavity.

We perform microwave transmission measurements in a circuit quantum electrodynamics set-up²⁷ (Fig. 2d). The sample holder is thermally anchored to the mixing chamber of a dilution refrigerator with a base temperature of 17 mK. The cavity is housed inside a copper shield coated with infrared-absorbing material, and all microwave lines are filtered above 12 GHz using commercial low-pass filters, circulators and infrared absorbers similar to that in ref. 28. The output line is pre-amplified by a cryogenic high-electron-mobility transistor anchored to the 4 K stage of the cryostat. The copper cavity frequency is $f_c = 8.894 \text{ GHz}$, with a bandwidth of $\kappa/2\pi = 4 \text{ MHz}$, and the antenna mode frequency at zero flux is $f_a(0) = 10.6 \text{ GHz}$. The maximum qubit-cavity dispersive shift is $\chi/2\pi = 0.5 \text{ MHz}$.

To quantify the junction dissipation, we perform standard time-domain measurements of the energy relaxation time, T_1 : after exciting the qubit using a ‘saturation’ pulse, which divides its population equally between the ground and excited states, we measure the energy decay as a function of time. For any decay channel degree of freedom \hat{X} , coupled to the qubit by an operator \hat{C} , the relaxation rate can be written, following Fermi’s golden rule, as

$$\frac{1}{T_{1X}} = \frac{1}{\hbar^2} |\langle 0 | \hat{C} | 1 \rangle|^2 S_H^X(\omega_{01}) \quad (2)$$

where $|0\rangle$ and $|1\rangle$ are respectively the ground- and excited-state wavefunctions, \hbar is Planck’s constant divided by 2π , S_H^X is the current spectral density of noise for the decay channel²⁹ and ω_{01} is the transition frequency of the qubit:

$$S_{II}^X = \text{Re}[Y_X(\omega_{01})] \hbar \omega_{01} \left[1 + \coth\left(\frac{\hbar \omega_{01}}{2k_B T}\right) \right] \quad (3)$$

Here k_B is Boltzmann's constant.

When the qubit is coupled to dielectric, inductive and radiative loss channels, then $\hat{C} \propto \hat{\varphi}$. Inductive losses in the array would mainly be due to quasiparticles in the array junctions. When the qubit is coupled to quasiparticle dissipation in the phase-slip junction, then $\hat{C} \propto \sin(\hat{\varphi}/2)$. The dashed lines in Fig. 3a represent numerical evaluations of equation (2) for dielectric loss (green), radiative loss (or 'Purcell' loss; magenta) and quasiparticle loss (red). The continuous black line shows the expected T_1 as a function of applied flux, Φ_{ext} , from the sum of all relaxation rates. Although the qubit is limited by radiative and dielectric loss around $\Phi_{\text{ext}} = 0$, quasiparticle dissipation becomes dominant as Φ_{ext} increases above $\sim 0.25\Phi_0$. The remarkable, order-of-magnitude, increase in the measured value of T_1 (grey circles) in the vicinity of $\Phi_{\text{ext}}/\Phi_0 = 0.5$ is the unambiguous signature of coherent suppression of quasiparticle dissipation.

The sharpness of the measured T_1 peak at $\Phi_{\text{ext}}/\Phi_0 = 0.5$, which increases by more than an order of magnitude within a flux interval of only $\sim 2\%$ of Φ_0 , allows us to place a stringent bound (Methods) on the value of the $\cos \varphi$ prefactor, $\varepsilon > 0.99$.

In Fig. 3b, we plot typical energy relaxation curves for different flux biases in the vicinity of the half-flux quantum symmetry point, measured after a saturation pulse. The averaging time for one relaxation curve is 30 min.

A close-up view of the T_1 peak at $\Phi_{\text{ext}}/\Phi_0 = 0.5$ (Fig. 3c) reveals quantitative information about the non-equilibrium quasiparticle density, x_{qp} . The coloured lines correspond to calculations of T_1 for different x_{qp} values. Considering the volume of the superconducting 'islands' on each side of the phase-slip junction ($0.05 \mu\text{m}^3$), we estimate that $x_{\text{qp}} = 10^{-6}$ corresponds to approximately one quasiparticle per island.

Using the x_{qp} bounds from Fig. 3c, we deduce that the average number of non-equilibrium quasiparticles on the phase-slip junction islands is greater than 0.1 but less than 3. These bounds correspond to measurements taken over a period longer than two months.

As an independent check, we can estimate the average quasiparticle number by analysing individual energy decay curves measured after excitation with a calibrated π -pulse. For a small number of quasiparticles, fluctuations can become important, because the quasiparticle number can change from shot to shot when repeating the measurement. Assuming that the probability, $p_\lambda(n)$, of having n quasiparticles is a Poisson distribution with average λ , $p_\lambda(n) = \lambda^n e^{-\lambda}/n!$, we can then calculate the average time-domain relaxation of the qubit polarization:

$$\langle P(t) \rangle = e^{\lambda(\exp(-t/\tilde{T}_{1\text{qp}}) - 1)} e^{-t/T_{1\text{R}}} \quad (4)$$

Here $\tilde{T}_{1\text{qp}}$ is the relaxation time induced by one quasiparticle, and $T_{1\text{R}}$ is the relaxation time associated with the remaining dissipation channel, which in our case is likely to be dielectric loss. We note that, owing to fluctuations in the quasiparticle number, the first factor on the right-hand side of equation (4) is not an exponential decay. Indeed, as we show in Fig. 4a, the fast part of the measured decay is not linear when expressed on a logarithmic scale. The second exponential factor contributes to the linear dependence observed at longer times. Remarkably, the fitted value of the average quasiparticle number, $\lambda = 1.1$, falls within the bounds obtained from the above analysis of the data in Fig. 3c.

Because the π -pulse calibration unambiguously provides the signal levels corresponding to the offset and the amplitude of the qubit excitation, plotting the measured data on a logarithmic scale allows us to discriminate between different quasiparticle dynamics. At random moments in time, we measure qubit relaxations compatible with quasi-exponential decays, as shown by the green dashed and magenta solid lines in Fig. 4b. This behaviour can be explained by a reduction, during the measurement

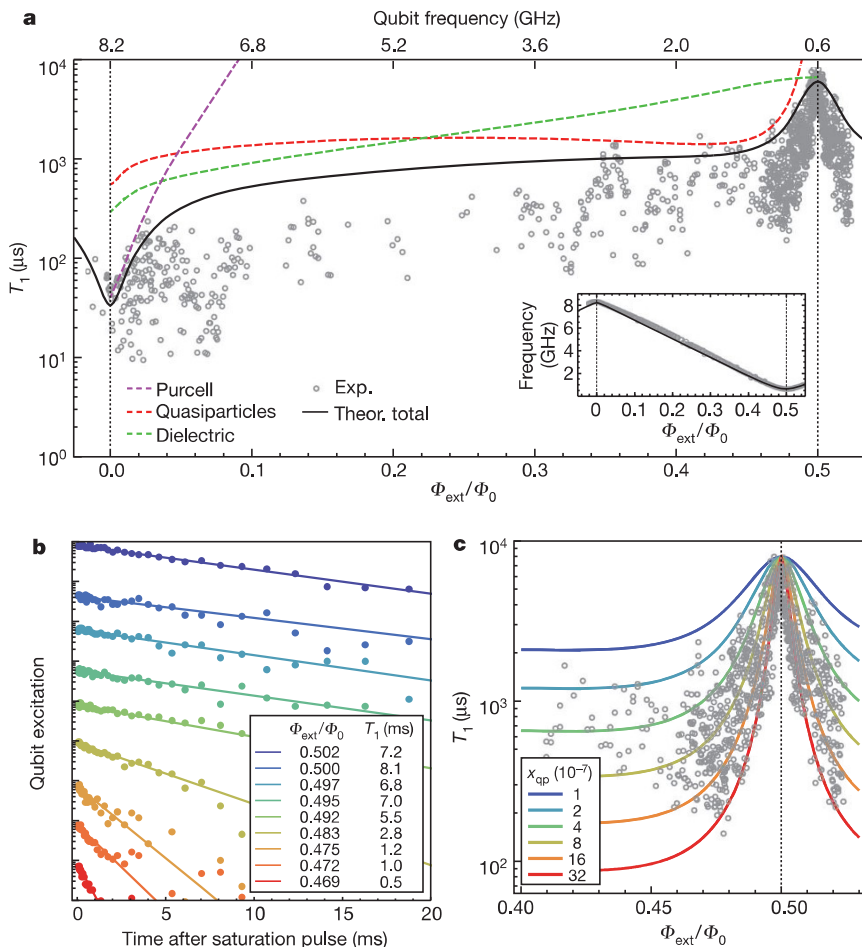


Figure 3 | Flux dependence of the fluxonium qubit T_1 measured by polarization saturation pulses. **a**, The measured values are represented by the grey circles, and the theoretical fit including all sources of dissipation is represented by the black line. The different contributions, T_{1X} , to the total relaxation rate, $1/T_1 = \sum 1/T_{1X}$, are plotted as coloured dashed lines. Inductive loss does not significantly contribute to the total decay rate, and the fit lies outside the figure scale (Methods). The inset shows the measured (grey circles) and the theoretical (black line) flux dependence of the $|0\rangle$ - $|1\rangle$ qubit transition frequency. **b**, Time-domain measurements of qubit free decay at different flux bias points. The y axis is shown on a logarithmic scale with the traces offset for clarity. The lines represent exponential fits and the inset table lists the fitted T_1 values. **c**, Detailed view in the vicinity of $\Phi_{\text{ext}}/\Phi_0 = 0.5$. The coloured lines show the calculated upper bound on T_1 resulting from the sum of quasiparticle and dielectric dissipation. For this narrow flux interval, we assumed that dielectric loss limits T_1 to a constant value of 8 ms.




Figure 4 | Time-domain measurements of qubit relaxation after π -pulse excitation. Typical measured traces of non-exponential (a) and quasi-exponential (b) decays. The magenta curves represent theoretical fits using equation (4). For comparison, the green dashed line in b is a fit with a simple exponential decay and time constant of 720 μ s. The flux bias point at $\Phi_{\text{ext}}/\Phi_0 = 0.48$ is chosen slightly off the maximum of the T_1 peak to allow faster repetition rates for the pulse calibration sequences. The data in a and b are taken several days apart; the acquisition time for data in each panel is 30 min. We observe non-exponential and quasi-exponential decays with similar occurrence probabilities. The offset and amplitude of the qubit excitation are calibrated using measurements of Rabi oscillations (insets). Error bars, 1 s.d. of the qubit excitation signal.

interval, of the average number of quasiparticles on the central islands ($\lambda = 0.3$), together with a downward fluctuation in T_{1R} . Such fluctuations, illustrated by the scatter of T_1 in Fig. 3a, are evident at any flux.

We have shown that the suppression of quasiparticle dissipation in the vicinity of π -phase bias across the junction confirms Josephson's original prediction². We establish the low-temperature value $\varepsilon = 1$ (expected from Bardeen–Cooper–Schrieffer theory) within less than 1% error for the prefactor in the $1 + \varepsilon \cos \phi$ interference term. The magnitude of the dissipation indicates the presence of a non-equilibrium population of quasiparticles. Although their origin remains unknown and further characterization of their dynamics goes beyond the scope of this work, the engineering of the qubit susceptibility to quasiparticle loss, as presented here, is a powerful tool for characterizing this ubiquitous dissipation mechanism and ultimately suppressing it altogether. The immediate implications for quantum information processing with superconducting circuits are evident in the achievement of relaxation times well above 1 ms in artificial atoms, an increase by two orders of magnitude.

METHODS SUMMARY

We have measured two nominally identical samples, fabricated in the same lithography cycle. In the main text, we focus on results obtained on sample A, which was measured extensively and exhaustively. In Methods, we show that our results and conclusions were also confirmed by measurements on sample B. Additionally, we describe in detail the microfabrication process, measurement set-up, sample parameters, coherence time measurements and derivations of theoretical expressions used in the main text.

Online Content Any additional Methods, Extended Data display items and Source Data are available in the online version of the paper; references unique to these sections appear only in the online paper.

Received 26 August 2013; accepted 13 January 2014.

- Mattis, D. C. & Bardeen, J. Theory of the anomalous skin effect in normal and superconducting metals. *Phys. Rev.* **111**, 412–417 (1958).
- Josephson, B. D. Possible new effects in superconductive tunnelling. *Phys. Lett.* **1**, 251–253 (1962).
- Barone, A. & Paternò, G. *Physics and Applications of the Josephson Effect* Ch. 2 (Wiley-VCH, 2005).
- Clarke, J. & Braginski, A. I. *The SQUID Handbook: Fundamentals and Technology of SQUIDs and SQUID Systems* (Wiley, 2004).
- Langenberg, D. Physical interpretation of the $\cos \phi$ term and implications for detectors. *Rev. Phys. Appl.* **9**, 35–40 (1974).
- Giazzotto, F. & Martinez-Perez, M. J. The Josephson heat interferometer. *Nature* **492**, 401–405 (2012).
- Catelani, G. *et al.* Quasiparticle relaxation of superconducting qubits in the presence of flux. *Phys. Rev. Lett.* **106**, 077002 (2011).
- Catelani, G., Schoelkopf, R. J., Devoret, M. H. & Glazman, L. I. Relaxation and frequency shifts induced by quasiparticles in superconducting qubits. *Phys. Rev. B* **84**, 064517 (2011).
- Leppäkangas, J., Marthaler, M. & Schön, G. Phase-dependent quasiparticle tunneling in Josephson junctions: Measuring the \cos term with a superconducting charge qubit. *Phys. Rev. B* **84**, 060505 (2011).
- Beenakker, C. W. J. Universal limit of critical-current fluctuations in mesoscopic Josephson junctions. *Phys. Rev. Lett.* **67**, 3836–3839 (1991).
- Brethreau, L., Girit, C. O., Pothier, H., Esteve, D. & Urbina, C. Exciting Andreev pairs in a superconducting atomic contact. *Nature* **499**, 312–315 (2013).
- Harris, R. E. Cosine and other terms in the Josephson tunneling current. *Phys. Rev. B* **10**, 84–94 (1974).
- Pedersen, N. F., Finnegan, T. F. & Langenberg, D. N. Magnetic field dependence and Q of the Josephson plasma resonance. *Phys. Rev. B* **6**, 4151–4159 (1972).
- Soerensen, O. H., Mygind, J. & Pedersen, N. F. Measured temperature dependence of the $\cos \phi$ conductance in Josephson tunnel junctions. *Phys. Rev. Lett.* **39**, 1018–1021 (1977).
- Aumentado, J., Keller, M. W., Martinis, J. M. & Devoret, M. H. Nonequilibrium quasiparticles and $2e$ periodicity in single-Cooper-pair transistors. *Phys. Rev. Lett.* **92**, 066802 (2004).
- Ferguson, A. J., Court, N. A., Hudson, F. E. & Clark, R. G. Microsecond resolution of quasiparticle tunneling in the single-Cooper-pair transistor. *Phys. Rev. Lett.* **97**, 106603 (2006).
- Shaw, M. D., Lutchyn, R. M., Delsing, P. & Ehternach, P. M. Kinetics of nonequilibrium quasiparticle tunneling in superconducting charge qubits. *Phys. Rev. B* **78**, 024503 (2008).
- Martinis, J. M., Ansmann, M. & Aumentado, J. Energy decay in superconducting Josephson-junction qubits from nonequilibrium quasiparticle excitations. *Phys. Rev. Lett.* **103**, 097002 (2009).
- Sun, L. *et al.* Measurements of quasiparticle tunneling dynamics in a band-gap-engineered transmon qubit. *Phys. Rev. Lett.* **108**, 230509 (2012).
- Ristè, D. *et al.* Millisecond charge-parity fluctuations and induced decoherence in a superconducting transmon qubit. *Nature Commun.* **4**, 1913 (2013).
- Manucharyan, V. E., Koch, J., Glazman, L. I. & Devoret, M. H. Fluxonium: single Cooper-pair circuit free of charge offsets. *Science* **326**, 113–116 (2009).
- Brooks, P., Kitaev, A. & Preskill, J. Protected gates for superconducting qubits. *Phys. Rev. A* **87**, 052306 (2013).
- Masluk, N. A. *Reducing the Losses of the Fluxonium Artificial Atom*, PhD thesis, Yale Univ. (2012).
- Masluk, N. A., Pop, I. M., Kamal, A., Mineev, Z. K. & Devoret, M. H. Microwave characterization of Josephson junction arrays: implementing a low loss superinductance. *Phys. Rev. Lett.* **109**, 137002 (2012).
- Bell, M. T., Sadovskyy, I. A., Ioffe, L. B., Kitaev, A. Y. & Gershenson, M. E. Quantum superinductor with tunable nonlinearity. *Phys. Rev. Lett.* **109**, 137003 (2012).
- Paik, H. *et al.* Observation of high coherence in Josephson junction qubits measured in a three-dimensional circuit QED architecture. *Phys. Rev. Lett.* **107**, 240501 (2011).
- Wallraff, A. *et al.* Strong coupling of a single photon to a superconducting qubit using circuit quantum electrodynamics. *Nature* **431**, 162–167 (2004).
- Santavica, D. F. & Prober, D. E. Impedance-matched low-pass stripline filters. *Meas. Sci. Technol.* **19**, 087001 (2008).
- Schoelkopf, R. J., Clerk, A. A., Girvin, S. M., Lehnert, K. W. & Devoret, M. H. in *Quantum Noise in Mesoscopic Physics* (eds Nazarov, Yu. V. & Blanter, Ya. M.) 175–204 (Kluwer Academic, 2002).
- Lecocq, F. *et al.* Junction fabrication by shadow evaporation without a suspended bridge. *Nanotechnology* **22**, 315302 (2011).

Acknowledgements We acknowledge discussions with L. Frunzio, A. Kamal, N. Masluk and U. Vool. Facilities use was supported by YINQE and NSF MRSEC DMR 1119826. This research was supported by IARPA under grant no. W911NF-09-1-0369, ARO under grant no. W911NF-09-1-0514, the NSF under grants nos DMR-1006060 and DMR-0653377, DOE contract no. DE-FG02-08ER46482 (L.I.G.), and the EU under REA grant agreement CIG-618258 (G.C.).

Author Contributions I.M.P. and K.G. performed the experiment and analysed the data, under the guidance of M.H.D. Theoretical support was provided by G.C. and L.I.G. The experimental design was proposed by I.M.P., K.G., R.J.S. and M.H.D. I.M.P. and M.H.D. led the writing of the manuscript. All authors provided suggestions for the experiment, discussed the results and contributed to the manuscript.

Author Information Reprints and permissions information is available at www.nature.com/reprints. The authors declare no competing financial interests. Readers are welcome to comment on the online version of the paper. Correspondence and requests for materials should be addressed to I.M.P. (ioan.pop@yale.edu).

METHODS

Microfabrication. The fluxonium qubit fabrication consists of a single step of double-angle evaporation of aluminium with an oxidation step in between. The main improvement compared with previous junction designs is the use of bridge-free fabrication³⁰ (BFF), which relaxes the constraints imposed on the junction geometry by the traditional Dolan bridge technique³¹. The use of BFF also allows a more aggressive cleaning of the substrate surface before the junction deposition, which has been proven to produce stable and reproducible tunnel barriers³².

The first step of the fabrication procedure is to clean the 430- μm -thick *c*-plane sapphire substrate via sonication in acetone for 1 min and then 5 min in an oxygen plasma at 300 mbar and a power of 300 W. The wafer is then soaked in heated *N*-methyl-2-pyrrolidone (NMP) at 90 °C for 10 min, sonicated in NMP and rinsed with acetone and methanol. Microchem EL-13 copolymer is spun onto the wafer at 2,000 r.p.m. for 100 s and baked at 200 °C for 5 min. The second resist layer, Microchem A-4 PMMA, is spun onto the wafer at 2,000 r.p.m. for 100 s and baked at 200 °C for 5 min. An antistatic gold coating is then deposited on the surface of the resist, in preparation for the electron-beam writing. The gold layer is deposited using a Cressington Sputter Coater 108 for 45 s with Ar flow adjusted to 0.08 mbar and a current of 30 mA, resulting in a 10 ± 1 nm gold layer.

The device is then written using the Vistec Electron Beam Pattern Generator (EBPG 5000+) using a 100-keV electron beam. The development is started by soaking in a potassium-iodide/iodine solution for 10 s to remove the gold layer. After being rinsed in water, the resist is developed in a 3:1 IPA:water mixture at 6 °C for 1 min. Ultrasound is then turned on for 15 s, and the wafer is left in the mixture for 15 s after that.

Metal deposition is then performed in the Plassys UMS300 UHV multichamber electron-beam evaporation system. First, the exposed surface is cleaned with an oxygen–argon plasma for 30 s to prepare a good surface for metal deposition. During this step, which is quasi-homogeneous, 10 ± 2 nm of resist is removed. A titanium sweep is then performed to absorb residual gases in the evaporation chamber, and the first layer of aluminium (20 nm) is deposited onto the substrate at an angle of 19°. This layer is then oxidized (in a separate oxidation chamber) with a 3:17 oxygen:argon mixture for 10 min at 100 torr. Immediately after this, the second layer of aluminium (30 nm) is deposited at an angle of -19° . A final oxidation for 10 min at 50 torr is then completed to passivate the aluminium surface, preparing it for removal from the vacuum system.

Finally, the resist layers are removed (lifted off) in a heated bath of NMP at 90 °C over 1 h, followed by 1 min of sonication and methanol rinsing. The 2-inch sapphire wafer is then diced under protective optical resist (SC1827 spun at 1,500 r.p.m. for 120 s and then baked at 90 °C for 5 min). Before mounting the fluxonium chips in the sample holder, they are rinsed in NMP and methanol.

The BFF relies on the ability to robustly produce undercuts in a PMMA/PMMA-MAA bilayer, while minimizing exposure due to proximity effects and forward scattering³⁰. The recent development of 100-keV EBPGs has made this task much easier. The dose calibration structures consist of long narrow openings (trenches) in the resist bilayer, with a designed undercut on the left side. In Extended Data Fig. 1, we present structures fabricated using the optimized doses of $1,200 \mu\text{C cm}^{-2}$ for the trenches and $450 \mu\text{C cm}^{-2}$ for the undercuts, which are the doses used for the fluxonium fabrication.

Measurement set-up. To measure the qubit state, we use a dispersive readout protocol that essentially consists of measuring the microwave transmission properties of the joint system formed by the sample holder cavity and the fluxonium qubit. When we send a measurement tone at the cavity resonance frequency, owing to the dispersive interaction between the cavity and the qubit, the phase and amplitude of the transmitted signal depend on the state of the qubit.

The schematic of a heterodyne interferometric experiment is shown in Extended Data Fig. 2. Three microwave waveform generators are shown, the qubit (ω_Q), cavity (ω_C) and LO ($\omega_{LO} = \omega_C + \omega_{IF}$) generators, which are commercial microwave generators. The qubit generator produces microwaves at the qubit frequency, mixed with pulse envelopes from an arbitrary waveform generator to create pulses to perform qubit rotations. The cavity and LO generator are used for the readout, with the cavity generator set to the interrogation frequency and the LO generator shifted by the IF frequency, $\omega_{IF} = 20$ MHz. The cavity signal (either continuous wave or pulsed) is split into two paths, making the interferometer. One of the paths goes through the device, and the other is used as reference; each path ends in the RF input of a microwave mixer. The LO generator is split to feed the LO signal of both RF mixers. This mixing operation produces signals at the sum and difference frequencies, $\omega_C + \omega_{LO}$ and $\omega_C - \omega_{LO}$. Additional filters remove the high-frequency component, and the ω_{IF} component is sent to a digitizer. The digitizer operates at 1 GS s^{-1} , and can effectively measure amplitude and phase of a signal at 20 MHz.

The sample holder and the low-temperature electronics are anchored to the dilution stage of a Kelvinox 400 system with a cooling power of approximately $400 \mu\text{W}$ at 100 mK and a base temperature of 15 mK, measured on the mixing chamber

plate. An additional thermometer mounted directly on the sample holder indicates a base temperature of 40 mK.

In Extended Data Fig. 3, we present a schematic of the low-temperature microwave circuitry. The signal is injected through -60 dB attenuated lines to the input of the cavity. The coupling quality factor at the input port is 10^2 . The cavity is over-coupled to the output port with a coupling quality factor of 2,000. The output signal travels through low-pass filters (see the following section for details on filtering) and cryogenic isolators (Quinstar 8–12 GHz) to the commercial CITCRYO1-12A Caltech high-electron-mobility transistor amplifier, which is coupled to the 4 K bath and has a typical gain of 35 dB within the 1–12-GHz band.

In the past few years, increasing evidence showed that infrared light leaking into the experimental set-up accounts for significant additional relaxation and decoherence^{33,34}. The goal of our filtering strategies is to attenuate any radiation at frequencies starting from 10 GHz and extending to infrared frequencies. For this purpose, we use a combination of commercial filters (K&L 12-GHz low-pass filters) and custom-made Eccosorb filters (which are similar to the ones used in ref. 28). Eccosorb CR-110, a proprietary material produced by Emerson & Cuming, was chosen for this filter because it has the least attenuation below 18 GHz. This material has a roughly linear attenuation-versus-frequency curve; higher frequencies are attenuated more than lower frequencies. The Eccosorb filter consists of a short section of microstrip placed in a copper box that serves as the ground plane. A thin metallic strip is used for the microstrip trace. Two SMA connectors are placed in the two opposing holes of the box, and then a metallic strip is cut to the size of a 50- Ω microstrip trace. For the box size used here, the microstrip trace is cut out of 0.01-inch thick copper, aiming for a width of 0.090 inch. The microstrip is then soldered in place onto the centerpins of the two SMA connectors. The Eccosorb CR-110 is then poured into the box, the lid is secured and the whole filter is cured at 70–90 °C, depending on cure time, for several hours.

In addition to filtering the noise inside the microwave lines, it is important to filter out infrared photons from external sources. These sources are likely to be hot components at other stages of the dilution refrigerator, or 300 K photons directly from room temperature components. Studied extensively in ref. 33, this stray infrared light can be avoided by building a sufficiently ‘light-tight’ shield around the sample box area. The infrared shield detailed here was based on this work. Additional work at IBM³⁵ showed that this infrared light can significantly affect qubit performance and is thus critical, especially for qubit experiments.

The infrared shield is an additional copper box surrounding the experimental area. The infrared shield is composed of two pieces, a plate and a can. The plate is mounted rigidly onto the mixing chamber and holds connectors through which all input and output microwave signals will travel. An indium seal is used between the plate and the can, and the can is screwed shut rigidly to achieve a hermetic seal. The entire structure can be observed in Extended Data Fig. 4.

In addition to preventing infrared radiation from reaching the sample box, this structure serves another purpose. A sheet of absorbing material is placed inside the copper can in an effort to further reduce infrared radiation from reaching the sample box. Even if infrared photons reach inside the can, they will quickly be absorbed on the walls. Extended Data Fig. 4b shows that this absorbent coating (a mixture of Stycast 2850 and carbon powder, containing 7% carbon powder by weight) is applied to a thin copper sheet that is then placed along the walls and bottom of the can.

Parameters of measured samples. The relaxation and coherence properties of two fluxonium qubits were measured. The qubits are nearly identical, with the only difference being in the coupling antennas. Fluxonium A had three coupling SQUIDs and a total antenna length of 1 mm, whereas fluxonium B had four coupling SQUIDs and a total antenna length of 2 mm. Fluxonium A was measured extensively and exhaustively over the course of 5 months; the results described in the main text were obtained from this sample. Fluxonium B was used to check the coherence repeatability, but was not measured in such detail.

Both qubits had visible spectroscopy peaks over the entire range of applied flux. Qubit frequency is plotted as a function of applied flux in Extended Data Fig. 5 for both fluxonium samples. This frequency dependence is fitted to that predicted from theory²¹. The fit parameters are the three energies that entirely characterize a fluxonium: E_J , the Josephson energy; E_C , the capacitive energy; and E_L , the inductive energy. From design and room-temperature measurements, the parameters were expected to be $E_J \approx 12$ GHz, $E_C \approx 2.6$ GHz and $E_L \approx 0.5$ GHz. Overall, the parameters are close between the two samples, with only slight discrepancies between expected and fit parameters.

Review of measured coherence times. In Extended Data Table 1, we list the measured coherence times for samples A and B at characteristic bias points. The measurements are taken using calibrated π -pulses. From measurement to measurement, the coherence times can vary by a factor of two. The minimum frequency (one-half flux quantum) and the maximum frequency (zero flux) are flux ‘sweet spots’. Therefore, it is not surprising that T_2 is much larger at these points than at intermediate fluxes. However, we note that T_2 is never limited by T_1 , which suggests the existence

of additional sources of dephasing. Near $\Phi_{\text{ext}}/\Phi_0 = 0.5$, the T_1 values of both fluxonium A and fluxonium B are observed to be roughly 1 ms (Extended Data Fig. 6). As discussed in the main text (Fig. 4), this value is subjected to fluctuations in time, owing to changes in the quasiparticle population.

Calculation of Purcell, capacitive, inductive and quasiparticle losses. As we mentioned above equation (2), for any decay channel, X , coupled to the qubit by an operator $\hat{C}X$, the relaxation rate is given, following Fermi's golden rule, by $1/T_{1X} = (1/\hbar^2)\langle 0|\hat{C}|1\rangle 2S_{II}^X(\omega_{01})$, where $|0\rangle$ and $|1\rangle$ are respectively the ground- and excited-state wavefunctions and S_{II}^X is the current spectral density of noise for the decay channel (equation (3)):

$$S_{II}^X = \text{Re}[Y_X(\omega_{01})]\hbar\omega_{01} \left[1 + \coth\left(\frac{\hbar\omega_{01}}{2k_B T}\right) \right]$$

In Extended Data Table 2, we list the respective expressions for the matrix element, the real part of the environment admittance at the qubit frequency, and the fitted bound for the quality factor of capacitive, inductive, quasiparticle and radiative loss channels.

In Extended Data Fig. 7, we plot the calculated bounds on T_1 for the four loss mechanisms listed in Extended Data Table 2, for different quality factors. Because the theoretical curves are upper bounds for T_1 , a good fit has to lie above the measured points. The quality factors that result in theoretical bounds compatible with the measured data are listed in the last column of Extended Data Table 2. The value for the capacitive quality factor is quoted at 5 GHz because it is expected to have a weak frequency dependence³⁶: $Q_{\text{cap}} \propto (f_0)^{-0.7}$.

Whereas radiative, capacitive and quasiparticle relaxation each dominate significant flux bias intervals, owing to its markedly different dependence on flux, inductive loss can only contribute over a narrow interval around $\Phi_{\text{ext}}/\Phi_0 = 0.5$ (Extended Data Fig. 7), where it adds to the quasiparticle and dielectric contributions. If the inductive loss is due to quasiparticles in the array⁸, we only need to introduce a minor modification to the flux dependence of the calculated loss, which can be accounted for by redefining $Q_{\text{ind}} \rightarrow Q_{\text{ind}}(1/2)\sqrt{\omega_{10}(f)/\omega_{10}(1/2)}$. The square-root dependence on the frequency originates in the quasiparticle current spectral density (equation (7)). This additional square-root factor weakens the flux dependence of the inductive loss, but does not qualitatively alter it compared with that shown in Extended Data Fig. 7b.

When analysing the energy relaxation of nonlinear oscillators, it is useful to introduce the notion of transition efficiency (also known as the reduced oscillator strength), defined as the ratio between the quality factor of one of the lossy components and the quality factor of the oscillator: $\eta = Q_x/Q_{1\rightarrow 0}$. For a harmonic oscillator, $\eta = 1$. For a qubit, depending on the value of the matrix element $\langle 0|\hat{C}|1\rangle$ and the real part of the environment impedance, η can be smaller or larger than unity. In Extended Data Fig. 8, we plot the transition efficiencies for capacitive, inductive and quasiparticle loss. The sharp suppression of η for quasiparticle loss in the vicinity of $\Phi_{\text{ext}}/\Phi_0 = 0.5$ is the signature of quasiparticle-hole interference. Also, notice that inductive η is larger than unity for $\Phi_{\text{ext}}/\Phi_0 > 0.42$. This implies that the transition efficiency for the qubit is larger than the efficiency of a harmonic oscillator at the same frequency, and the bound that we place on Q_{ind} is larger than the intrinsic quality factor of the transition, $Q_{1\rightarrow 0}$.

Relationship between relaxation time and admittance. Here we detail the relation between the admittance of the Josephson junction (equation (1)) and the qubit relaxation rate driven by quasiparticles (equations (2) and (3) and the definition of the operator \hat{C}). Most of what we summarize here is presented in greater detail in refs 7, 8.

The admittance in the form of equation (1) can be obtained from the standard microscopic theory of the Josephson current derived for an arbitrary time dependence of the phase bias $\varphi(t)$ (ref. 3). To find the linear response to a small alternating potential difference, we take $\varphi(t) = \varphi + (2eV/\omega)\sin(\omega t)$, where V and ω are respectively the amplitude and frequency of the alternating voltage, and φ is the constant component of the phase bias.

The first term on the right-hand side of equation (1) contains the dissipative part of admittance. For $\hbar\omega < 2\Delta$, where Δ is the superconducting gap, the a.c. perturbation does not break Cooper pairs. Therefore, the amplitude, $\text{Re}[Y_{\text{qp}}(\omega)]$, of the dissipative part is proportional to the dimensionless density of quasiparticles, x_{qp} . An extension^{7,8} of the standard theory³ relates $Y_{\text{qp}}(\omega)$ to x_{qp} even if the energy distribution of the quasiparticles deviates from equilibrium. Like in equilibrium, the parameter ε equals unity as long as the characteristic energy of quasiparticles (measured from Δ) is small compared with Δ . The $\cos\varphi$ functional form of the dissipative part of the admittance was noticed first in ref. 2.

The second term on the right-hand side of equation (1) is the non-dissipative, dispersive part of the admittance, which can be obtained by expansion of the conventional Josephson current formula, $j = j_c \sin\varphi(t)$, to linear order in V . The only generalization here is the replacement of the factor $\tanh(\Delta/2k_B T)$ in j_c (ref. 3) by the factor $1 - 2n_{\text{qp}}^A$. To explain it³⁷, we notice that $\tanh(\Delta/2k_B T) = 1 - 2n_{\text{qp}}(A)$, where

$2n_{\text{qp}}(A)$ is the value of the equilibrium Fermi occupation factor, $2n_{\text{qp}}(E) = 2[\exp(E/k_B T) + 1]^{-1}$, of an energy level E coinciding with the gap edge (the factor of 2 accounts for spin degeneracy). In fact, the equilibrium Josephson current evaluated beyond the lowest order in perturbation theory in tunnelling³⁸ contains the factor $1 - 2n_{\text{qp}}(E_A)$. The energy of an Andreev level, E_A , asymptotically approaches the value $E_A = \Delta$ in the limit of weak tunnelling, yielding the familiar result for j_c . In the generalization to non-equilibrium quasiparticles, the occupation factor of the Andreev levels, $n(E_A) = x_{\text{qp}}^A$, may differ from $n(A)$ even in the limit $E_A \rightarrow \Delta$; there is no set relation between the density, x_{qp} , of above-the-gap quasiparticles and x_{qp}^A .

The quasiparticle contribution, $1/T_{1\text{qp}}$, to the qubit relaxation rate can be written^{7,8} in the form of equation (2). In more detail, it reads

$$\frac{1}{T_{1\text{qp}}} = |\langle 0|\sin(\hat{\varphi}/2)|1\rangle|^2 S_{\text{qp}}(\omega_{01}) \quad (5)$$

where $S_{\text{qp}}(\omega)$ is related to the amplitude of the dissipative part of equation (1):

$$S_{\text{qp}}(\omega) \approx \frac{\omega}{\pi g_K} \text{Re}[Y_{\text{qp}}(\omega)] \approx x_{\text{qp}} \frac{8E_J}{\pi\hbar} \sqrt{\frac{2\Delta}{\hbar\omega}} \quad (6)$$

Here $g_K = e^2/\hbar$ is the conductance quantum. The derivations of equation (5) and the second part of equation (6) assume the characteristic quasiparticle energy to be small compared with the gap energy, Δ , and the energy of transition, $\hbar\omega_{01}$, respectively.

The fact that the transition rate $|1\rangle \rightarrow |0\rangle$ vanishes at $f = 1/2$ is the manifestation of symmetry between the electron-like and hole-like processes for the low-energy quasiparticles. This fact remains unchanged even if the quantum uncertainty of phase is not small^{7,8}. In the range $|f - 1/2| \lesssim \omega_{01}(f = 1/2)/(2\pi^2 E_L)$, the transition matrix element can be written as

$$\begin{aligned} & \langle 0|\sin(\hat{\varphi}/2)|1\rangle \\ & \approx \frac{\pi}{2} \left(f - \frac{1}{2}\right) \left[\frac{\omega_{01}(1/2)}{\omega_{01}(f)}\right]^{1/2} F\left(\frac{E_C}{E_J}, \frac{E_L}{E_J}\right) \end{aligned} \quad (7)$$

where the function F can be evaluated analytically in certain limiting cases⁸ or, in a generic setting, can be evaluated numerically using the approach described in appendix B of ref. 39. Combining equations (5)–(7), we find

$$\frac{1}{T_{1\text{qp}}} = \frac{\pi^2 F^2}{4} \left(f - \frac{1}{2}\right)^2 \frac{\omega_{01}(1/2)}{\pi g_K} \text{Re}[Y_{\text{qp}}(\omega_{01}(1/2))] \quad (8)$$

Now we compare the right-hand side of equation (8) with the dissipative part of the admittance of a classical junction biased by the flux $\varphi = 2\pi f$, with $f \approx 1/2$. For that, we expand the real part of the admittance equation (1) around $f = 1/2$:

$$\begin{aligned} & \text{Re} Y(\omega_{01}(1/2), 2\pi f) \\ & = \left[\frac{1-\varepsilon}{2} + \varepsilon\pi^2 \left(f - \frac{1}{2}\right)^2\right] \text{Re}[Y_{\text{qp}}(\omega_{01}(1/2))] \end{aligned} \quad (9)$$

For $\varepsilon = 1$, the right-hand sides of equations (8) and (9) are proportional to each other, and so for $f \approx 1/2$ there is a direct relation,

$$\frac{1}{T_{1\text{qp}}} \approx \frac{F^2 \omega_{01}(1/2)}{4 \pi g_K} \text{Re}[Y(\omega_{01}(1/2), 2\pi f)] \quad (10)$$

between the qubit relaxation rate and classical admittance, Y , of equation (1). We stress that equation (10) is based only on the assumption of small deviations from $f = 1/2$ ($|f - 1/2| < \omega_{01}(1/2)/4\pi^2 E_L$) and small quasiparticle energy compared with the qubit frequency.

Being a consequence of the fluctuation-dissipation relation, equation (10) should hold for any value of ε (not only for the value $\varepsilon = 1$, which corresponds to low characteristic energy of quasiparticles). In the next section, we use equation (10) to establish an experimental bound on ε .

Placing a lower bound on ε . By expanding equation (10) around $f = 1/2$, within the interval $|f - 1/2| < \omega_{01}(1/2)/4\pi^2 E_L$, we find that

$$\frac{1}{T_{1\text{qp}}} \approx A(1 - \varepsilon) + 2\pi^2 A\varepsilon \left(f - \frac{1}{2}\right)^2 \quad (11)$$

The parameter $A(1 - \varepsilon)$ determines the relaxation time at $f = 1/2$, and $2\pi^2 A\varepsilon$ gives the width of the $T_{1\text{qp}}$ peak; both terms are due to quasiparticle loss in the phase-slip junction. The relationship between A and the parameters in equation (10) can be worked out, but it is not needed.

We use equation (11) to fit the measured data in the vicinity of $f = 1/2$, with A and ε as adjustable parameters. The model is additionally constrained to give a higher bound for all measured values of T_1 . This procedure yields $\varepsilon = 0.991$ with

the standard error of 0.002 for the lower bound of ε . For illustration, in Extended Data Fig. 9 we plot $T_1(f)$ for a set of values of ε (keeping the value of $A(1-\varepsilon)$ unchanged), overlaid on the measured T_1 values. Comparing the theoretical curves with the measured data, it is clear that the bound $\varepsilon \geq 0.99$ is quite conservative.

Derivation of equation (4). As discussed in the previous section, it was shown in refs 7, 8 that the quasiparticle-induced relaxation rate is, in general, a function of reduced flux, f , and is proportional to the normalized quasiparticle density, x_{qp} :

$$\frac{1}{T_{1\text{qp}}} = x_{\text{qp}} \mathcal{F}(f) \quad (12)$$

where the function \mathcal{F} depends implicitly on the qubit parameters. This form for the relaxation rate was obtained in the limit of a large number of quasiparticles by averaging over their possible microscopic states. To investigate the effects of quasiparticles number (n_{qp}) fluctuations on the measured time evolution of the qubit population, we explicitly factorize n_{qp} by rewriting the relaxation rate in the equivalent form

$$\frac{1}{T_{1\text{qp}}} \equiv n_{\text{qp}} \tilde{T}_{1\text{qp}}(f) = n_{\text{qp}} \frac{1}{\tilde{T}_{1\text{qp}}} \quad (13)$$

where $\tilde{T}_{1\text{qp}}$ represents the qubit relaxation rate in the presence of a single quasiparticle and is inversely proportional to the electrodes' volume, V (to recover the density dependence in equation (12)).

For a given number of quasiparticles, n , by Matthiessen's rule the qubit polarization P behaves as

$$P(t; n) = e^{-t/T_{1\text{R}} - n\tilde{T}_{1\text{qp}}t}$$

where $1/T_{1\text{R}}$ is the total relaxation rate due to all non-quasiparticle processes. For a small number of quasiparticles, fluctuations can have important effects because the quasiparticle number can change from shot to shot when repeating the measurement many times. We assume that for such repeated attempts the probability, $p_\lambda(n)$, of having n quasiparticles is a Poisson distribution with average λ :

$$p_\lambda(n) = \frac{\lambda^n e^{-\lambda}}{n!}$$

We can then calculate the average relaxation probability, and find that

$$\langle P(t) \rangle = \sum_{n=0}^{\infty} P(t; n) p_\lambda(n) = e^{-t/T_{1\text{R}}} e^{\lambda(\exp(-\tilde{T}_{1\text{qp}}t) - 1)} \quad (14)$$


which coincides with equation (4) in the main text.

As a check, we consider the thermodynamic limit in which $\lambda, V \rightarrow \infty$ with $\lambda/V \propto x_{\text{qp}}$ constant; then $\tilde{T}_{1\text{qp}} \rightarrow 0$ and, keeping only the leading term inside the brackets on the right-hand side of equation (14), we have

$$\langle P(t) \rangle \approx e^{-t/T_{1\text{R}}} e^{-\lambda \tilde{T}_{1\text{qp}} t}$$

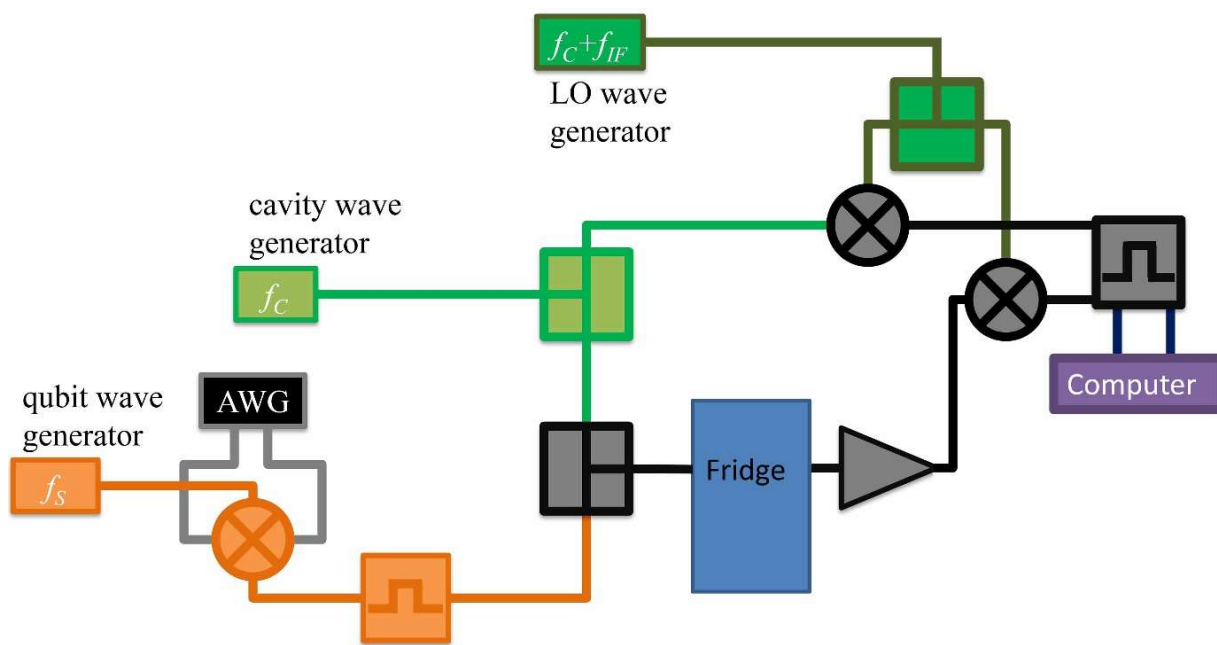
That is, we correctly recover the formula originally derived in limit of large quasiparticle number (because $\lambda \tilde{T}_{1\text{qp}} = 1/T_{1\text{qp}}$; compare with equation (13)). Note that in the limit $t \rightarrow \infty$ the last factor in equation (14) becomes $e^{-\lambda}$ and does not vanish: this accounts for the possibility that, with probability $e^{-\lambda}$, there are no quasiparticles in the electrodes, and so the quasiparticle mechanism does not contribute to the qubit relaxation.

31. Dolan, G. J. Offset masks for lift-off photoprocessing. *Appl. Phys. Lett.* **31**, 337–339 (1977).
32. Pop, I. M. *et al.* Fabrication of stable and reproducible submicron tunnel junctions. *J. Vac. Sci. Technol. B* **30**, 010607 (2012).
33. Barends, R. *et al.* Minimizing quasiparticle generation from stray infrared light in superconducting quantum circuits. *Appl. Phys. Lett.* **99**, 113507 (2011).
34. Rigetti, C. *et al.* Superconducting qubit in a waveguide cavity with a coherence time approaching 0.1 ms. *Phys. Rev. B* **86**, 100506 (2012).
35. Corcoles, A. D. *et al.* Protecting superconducting qubits from radiation. *Appl. Phys. Lett.* **99**, 181906 (2011).
36. Braginsky, V. B., Ilchenko, V. S. & Bagdassarov, K. S. Experimental-observation of fundamental microwave-absorption in high-quality dielectric crystals. *Phys. Lett. A* **120**, 300–305 (1987).
37. Kos, F., Nigg, S. E. & Glazman, L. I. Frequency-dependent admittance of a short superconducting weak link. *Phys. Rev. B* **87**, 174521 (2013).
38. Beenakker, C. W. J. *Transport Phenomena in Mesoscopic Systems* (Springer, 1992).
39. Zhu, G., Ferguson, D. G., Manucharyan, V. E. & Koch, J. Circuit QED with fluxonium qubits: theory of the dispersive regime. *Phys. Rev. B* **87**, 024510 (2013).




Extended Data Figure 1 | Scanning electron microscope imaging of controlled undercuts. a, 500-nm-deep cut; b, 900-nm-deep cut. For the purpose of scanning electron microscope imaging, the entire structure is

covered with a 10-nm layer of sputtered gold. Notice that the residual undercut on the right-hand side of the trenches is at least an order of magnitude smaller than the designed undercut.




Extended Data Figure 2 | Heterodyne measurement experimental set-up. Schematic diagram of experimental set-up to perform heterodyne measurement, involving an interferometric measurement, which compares a microwave signal going through the device under test with a signal bypassing the device. Two microwave generators (cavity and LO) are mixed together to

produce a lower frequency tone at the difference frequency, ω_{IF} , that can be digitized in the computer. The additional (qubit) microwave generator can be used to stimulate the device and the effect on the cavity transmission can be measured.

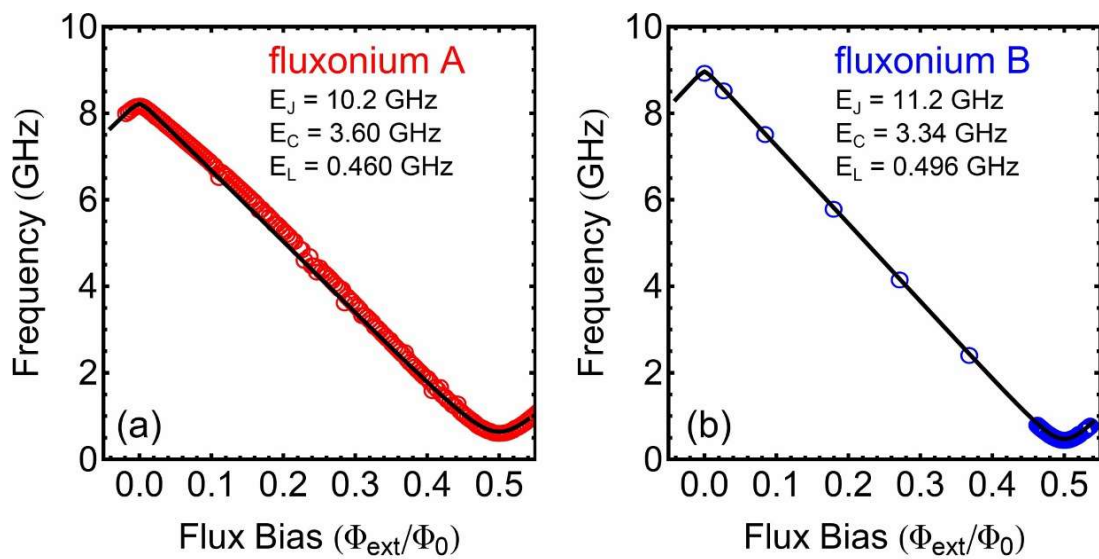


Extended Data Figure 3 | Microwave cryogenic measurement set-up.



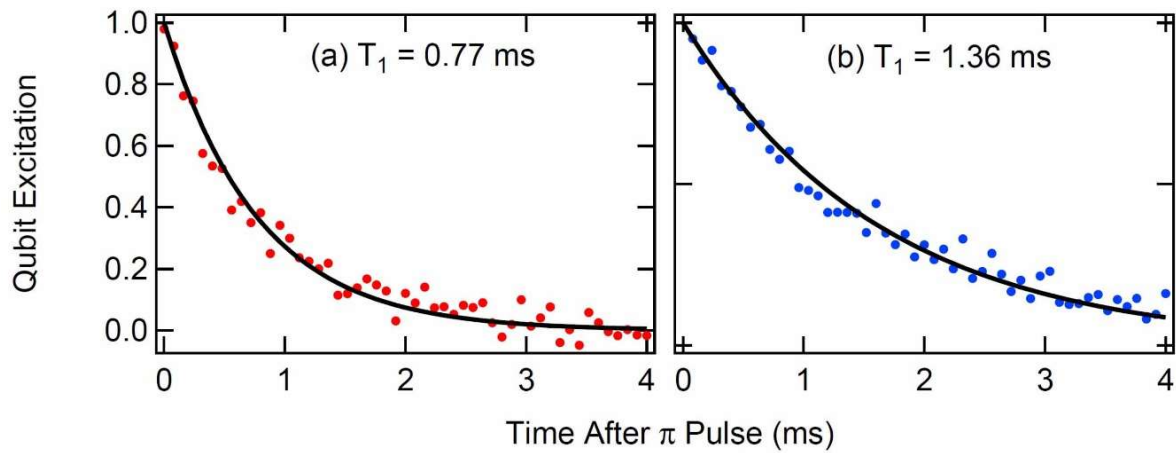
Extended Data Figure 4 | Infrared shielding. **a**, Top plate of the infrared shield and the attached samples connected by microwave coaxial lines. The hermetic seal is on the top side of the plate. The top half of the cryoperm shield

is also visible. **b**, Inside of the infrared shield can. An infrared absorbent coating was applied to a thin copper sheet and placed on the walls and bottom of the can. **c**, Closed infrared shield, completely enclosing the experimental area.



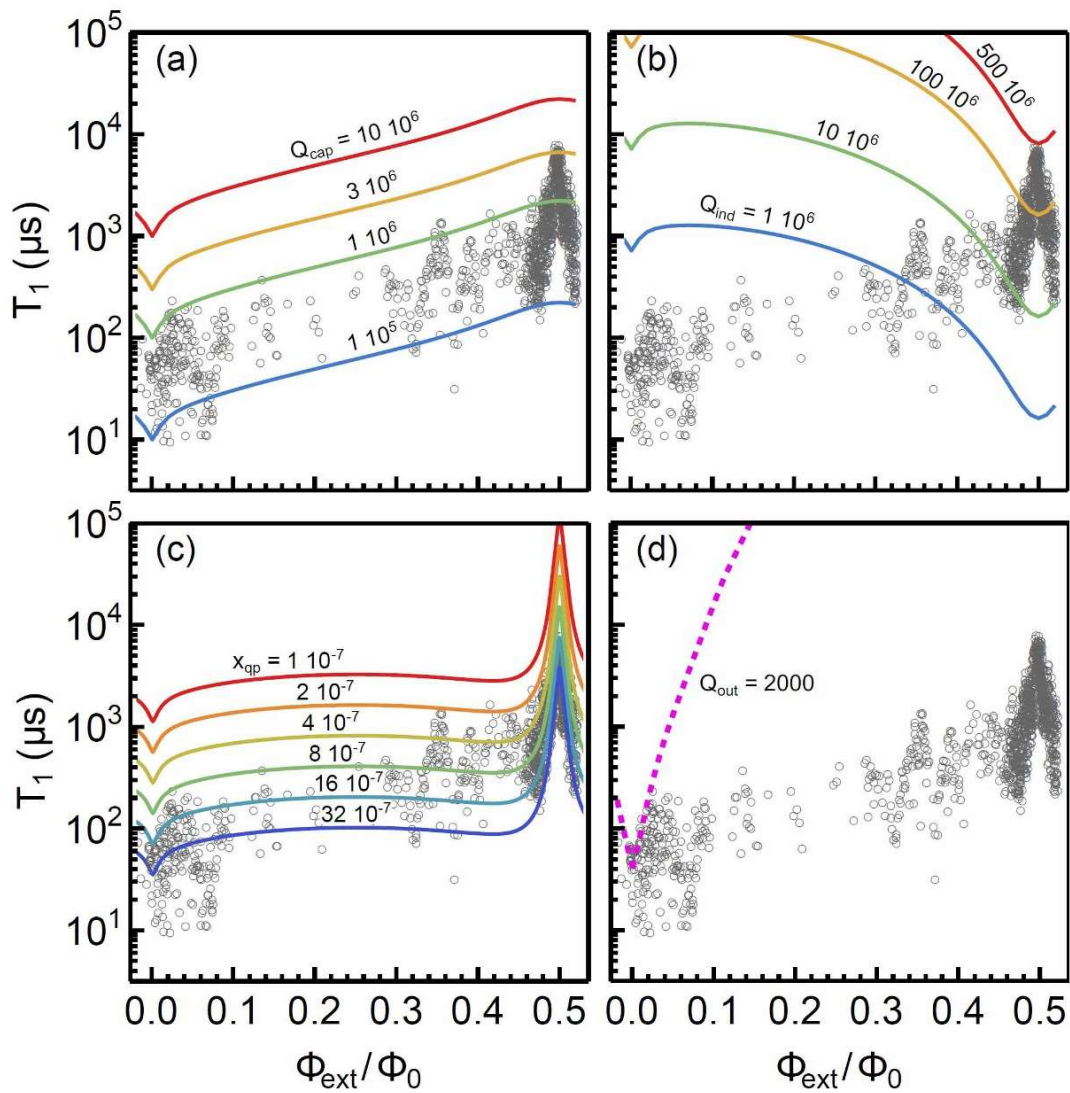
Extended Data Figure 5 | Measured qubit frequency as a function of applied flux over the entire tunable range. Fits of the expected frequency dependence from theory match well with the measured data and yield parameters as

listed for each fluxonium sample. The flux dependence of fluxonium B was sampled more sparsely than that of fluxonium A.




Extended Data Figure 6 | Measured relaxation times near $\Phi_{\text{ext}}/\Phi_0 = 0.5$. Data are fitted to a single exponential and reveal that lifetimes are ~ 1 ms for fluxonium sample A at $f_{01} = 640$ MHz (a) and fluxonium sample B at


$f_{01} = 750$ MHz (b). The presence of single exponentials as shown here fluctuates in time, as shown in Fig. 4.



Extended Data Figure 7 | Measured relaxation times. Measured T_1 values and theoretical bounds for capacitive (a), inductive (b), quasiparticle (c) and radiation (d) loss.



Extended Data Figure 8 | Transition efficiency of the fluxonium qubit. Shown as a function of applied magnetic flux for capacitive (red), inductive (blue) and quasiparticle (green) loss.



Extended Data Figure 9 | Placing a bound on ε . The lines represent T_1 values calculated from equation (11) for $\varepsilon = 0.9, 0.991$ (the fitted value), 0.999 and 0.9999 , respectively. The green line ($\varepsilon = 0.991$) bounds all measured points (grey circles), giving a conservative bound of $\varepsilon \geq 0.99$.

Extended Data Table 1 | Measured coherence times for fluxonium samples A and B at different bias points

Fluxonium	Φ_{ext}/Φ_0	f (GHz)	$T_1(\mu s)$	$T_2^R(\mu s)$	$T_2^E(\mu s)$
A	0.0	8.2	43	15	53
A	0.03	7.85	95	0.26	3.8
A	0.5	0.64	$\gtrsim 1000$	7.9	8.5
B	0.5	0.48	$\gtrsim 1000$	14	21

Extended Data Table 2 | Expressions used to calculate the qubit energy relaxation rate²³

loss mechanism	$\langle 0 \hat{C} 1 \rangle$	$\text{Re} [Y_X(\omega_{01})]$	quality factor
capacitive	$\Phi_0 \langle 0 \hat{\varphi} 1 \rangle$	$\frac{\omega_{01} C}{Q_{cap}}$	$Q_{cap} > 3 \cdot 10^6$
inductive	$\Phi_0 \langle 0 \hat{\varphi} 1 \rangle$	$\frac{1}{\omega_{01} L Q_{ind}}$	$Q_{ind} > 500 \cdot 10^6$
quasiparticle	$\langle 0 \sin(\hat{\varphi}/2) 1 \rangle$	$\frac{G_t}{2Q_{qp}} \left(\frac{2\Delta}{\hbar\omega_{01}} \right)^{3/2}$	$Q_{qp} \approx x_{qp}^{-1} > 0.3 \cdot 10^6$
radiative (Purcell)	$\Phi_0 \langle 0 \hat{\varphi} 1 \rangle$	HFSS numerical simulations	$Q_{out} = 2000$

The values of capacitance, C , inductance, L , and junction conductance, G_t , can be estimated knowing the parameters $E_C = e^2/2C$, $E_L = \hbar^2/4e^2L$ and $E_J = G_t A \hbar / 8e^2$, respectively (compare with Extended Data Fig. 5).