



This is the accepted manuscript made available via CHORUS. The article has been published as:

Precision Measurement of the Microwave Dielectric Loss of Sapphire in the Quantum Regime with Parts-per-Billion Sensitivity

Alexander P. Read, Benjamin J. Chapman, Chan U Lei, Jacob C. Curtis, Suhas Ganjam, Lev Krayzman, Luigi Frunzio, and Robert J. Schoelkopf

Phys. Rev. Applied **19**, 034064 — Published 20 March 2023

DOI: [10.1103/PhysRevApplied.19.034064](https://doi.org/10.1103/PhysRevApplied.19.034064)

Precision measurement of the microwave dielectric loss of sapphire in the quantum regime with parts-per-billion sensitivity

Alexander P. Read,^{1,2,*} Benjamin J. Chapman,^{1,2,*} Chan U Lei,^{1,2} Jacob C. Curtis,^{1,2} Suhas Ganjam,^{1,2} Lev Krayzman,^{1,2} Luigi Frunzio,^{1,2} and Robert J. Schoelkopf^{1,2}

¹*Departments of Physics and Applied Physics, Yale University, New Haven, Connecticut 06511, USA*

²*Yale Quantum Institute, Yale University, New Haven, Connecticut 06511, USA*

Dielectric loss is known to limit state-of-the-art superconducting qubit lifetimes. Recent experiments imply upper bounds on bulk dielectric loss tangents on the order of 100 parts-per-billion, but because these inferences are drawn from fully fabricated devices with many loss channels, these experiments do not positively identify dielectric loss, nor rule it out. To resolve this ambiguity, we have devised a measurement method capable of separating and resolving bulk dielectric loss with a sensitivity at the level of 5×10^{-9} . The method, which we call the dielectric dipper, involves the in-situ insertion of a dielectric sample into a high-quality microwave cavity mode. Smoothly varying the sample's participation in the cavity mode enables a differential measurement of the sample's dielectric loss tangent. The dielectric dipper can probe the low-power behavior of dielectrics at cryogenic temperatures, and does so without the need for any lithographic process, enabling controlled comparisons of substrate materials and processing techniques. We demonstrate the method with measurements of sapphire grown by edge-defined film-fed growth (EFG) in comparison to high-grade sapphire grown by the heat-exchanger method (HEMEX). For EFG sapphire we infer a bulk loss tangent of $63(8) \times 10^{-9}$ and a substrate-air interface loss tangent of $15(3) \times 10^{-4}$ (assuming a sample surface thickness of 3 nm). For a typical transmon, this bulk loss tangent would limit device quality factors to $Q \lesssim 20$ million, suggesting that bulk loss is likely the dominant loss mechanism in the longest-lived transmons on sapphire. We also demonstrate this method on HEMEX sapphire and bound its bulk loss tangent to be less than $19(6) \times 10^{-9}$. As this bound is about four times smaller than the bulk loss tangent of EFG sapphire, use of HEMEX sapphire as a substrate would lift the bulk dielectric coherence limit of a typical transmon qubit to several milliseconds.

I. INTRODUCTION

Superconducting circuits are a promising hardware platform for quantum information science, prized for their union of rapid gates and long coherence times. That success is partly due to a two-decade effort which has prolonged the nanosecond-scale coherence of early superconducting qubits [1] by six-orders of magnitude [2–5].

Some of these gains have come from engineering matrix-elements [6] and insensitivity to decoherence mechanisms like $1/f$ noise [7]. Other improvements have been made by directly minimizing noise-spectral densities [8], for example by improving the device's structure, materials, or fabrication process. For either approach, the first step in finding the next order-of-magnitude improvement is to determine the dominant loss mechanism.

Identifying the dominant source of loss is complicated by the many loss channels present in superconducting qubits [9]. These losses add together to limit the qubit's quality factor Q .

$$\frac{1}{Q} = \frac{1}{Q_{\text{conductor}}} + \frac{1}{Q_{\text{dielectric}}} + \frac{1}{Q_{\text{radiative}}} + \dots \quad (1)$$

While loss Q^{-1} is straightforward to measure, it is unclear how to best improve qubit coherence without know-

ing which of these mechanisms is the dominant source of loss. To overcome this challenge and distinguish between various sources of loss, one can measure Q^{-1} on either a suite of devices, each designed to be more or less sensitive to particular loss channels, or a single device with some in-situ tunable experimental parameter, which preferentially aggravates or alleviates one source of dissipation.

Studies on dielectric loss have principally employed the suite approach [10–18]. For transmons on sapphire with $Q \leq 4$ million, those studies measured Q^{-1} that scaled proportionally with the fraction of the qubit's energy stored in surface dielectrics [12, 13]. This directly implicated surface dielectric loss as the dominant loss mechanism in these devices, motivating the use of fabrication processes compatible with more aggressive surface cleaning. The result was some of the highest published transmon lifetimes to-date [2, 3].

Advances in coherence make it necessary to reassess which loss channels limit coherence. Quality factors of transmon qubits have now increased to $Q \approx 10$ million [2–4]. Does bulk dielectric loss limit the coherence of these devices?

A desirable technique to answer this question would uniquely distinguish bulk loss from other sources of loss in the system, and do so with high resolution. Specifically, it would have sensitivity to loss tangents much smaller than those which would limit transmon coherence, which in current devices would be a bulk loss tangent of $\tan\delta_{\text{bulk}} \approx 10^{-7}$. It would also complement previous methodologies [16, 17] by measuring both power-

* These authors contributed equally to this work.
alex.read@yale.edu
benjamin.chapman@yale.edu

dependent and power-independent losses without requiring lithography or other processing. Such flexibility would allow measurements of materials for which there is not yet an established process for depositing and patterning films and Josephson junctions. Similarly, standard substrates could be studied mid-process to determine the process stage at which losses manifest.

Here we introduce a suitable technique for this task, designed for the study of dielectric loss at microwave frequencies, cryogenic temperatures, and low powers, which allows dielectric loss to be measured independently from other qubit loss-channels. The technique, which we call the dielectric dipper, involves the in-situ insertion of a dielectric sample into a high-quality microwave cavity mode. Smooth variation of the sample's participation in the cavity mode enables a differential measurement of the sample's dielectric loss, with a sensitivity of 5×10^{-9} .

We use this technique to measure sapphire grown by edge-defined film-fed growth (EFG), and determine its bulk and substrate-air interface loss tangents to be $63(8) \times 10^{-9}$, and $15(3) \times 10^{-4}$, respectively [19]. We also measure high-grade sapphire grown by the heat-exchanger method (HEMEX) and bound the bulk loss tangent below $19(6) \times 10^{-9}$. These measurements suggest that current record-lifetime transmon qubits fabricated on EFG sapphire are approaching the bulk dielectric loss limit, and that the use of HEMEX sapphire as a substrate would suppress that loss, lifting the transmon coherence limit from bulk dielectric loss to several milliseconds.

II. MEASUREMENT TECHNIQUE

A. Overview

One way to improve the resolution of a dielectric loss measurement is to heavily suppress the non-dielectric loss channels. This consideration suggests the use of high- Q microwave cavities as sensitive probes; the properties that make them promising as quantum memories—relatively few physical loss mechanisms and lifetimes from milliseconds to seconds [20, 21]—also make them a well-suited platform for precision loss measurements.

To further increase sensitivity, one can distinguish dielectric loss from the non-dielectric “background” by making a separate measurement of that background, then subtracting it from the loss measured in the presence of the dielectric under study [22]. If those measurements are made in separate cooldowns, however, then comparison with the background may be corrupted by cooldown-to-cooldown variation of cavity properties [23]. This problem can be avoided by engineering an in-situ tunable coupling between the cavity-mode and the dielectric under study, which enables both measurements to be performed in the same cooldown.

We accomplish this using a coaxial stub cavity [25] with a high-quality (internal quality factor $Q \approx 100$ million) $\lambda/4$ mode, and a piezoelectric positioner which can insert

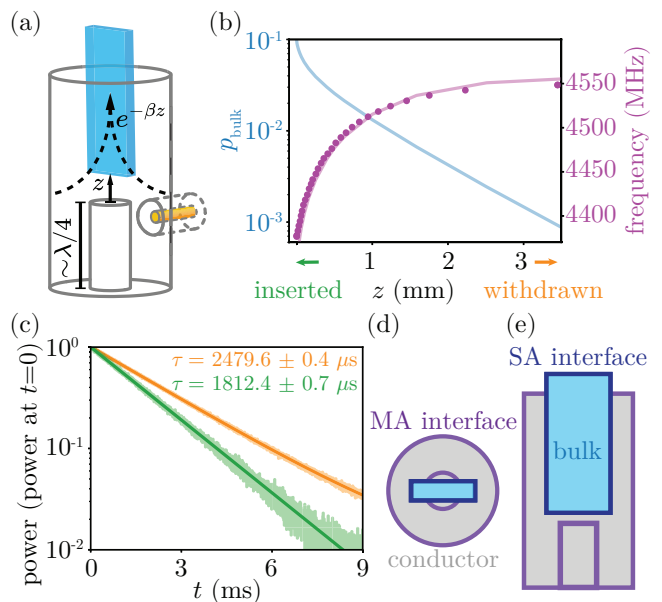


FIG. 1. **Measurement concept.** (a) A $\lambda/4$ coaxial stub cavity is machined from 4N aluminum, etched [24], and mounted to the base of a dilution refrigerator. A dielectric sample is inserted into the cavity waveguide to a position z , adjusted in-situ by a piezoelectric positioner (not pictured). (b) The fundamental cavity mode is far below the waveguide's cutoff, making the participation p_{bulk} of the cavity mode in the dielectric sample change exponentially with z (blue line). As p_{bulk} rises with insertion, the real part of the dielectric's permittivity causes the measured cavity frequency (purple circles) to drop, in agreement with electromagnetic simulations (purple line). (c) Decay of the cavity output power in a ring-down experiment which can distinguish between dephasing and decay (see App. A). The imaginary part of the dielectric's permittivity causes the energy decay rate of the cavity mode $1/\tau$ to change when the dielectric is withdrawn (orange) or inserted (green). (d) (Top view) In the withdrawn position, almost all of the measured loss comes from either the metal-air dielectric interface (aluminum oxide), or from losses in the conductor. (e). (Side view) As the sample is inserted, the loss channels of the substrate under study (the substrate-air interface and bulk of the dielectric) contribute to the measured loss.

a dielectric sample into the waveguide of the stub cavity (Fig. 1a). The fundamental cavity mode is designed to have a resonant frequency well below the cutoff of the waveguide. This causes its electric field to be attenuated exponentially along that waveguide, confining the cavity mode. Because of this attenuation, moving the sample along the waveguide changes the ratio between the electric energy stored in the volume of the dielectric sample U_{bulk} and the total energy stored in the cavity U_{tot} . This ratio is known as the bulk participation p_{bulk} :

$$p_{\text{bulk}} \equiv \frac{U_{\text{bulk}}}{U_{\text{tot}}} \equiv \frac{\frac{1}{2} \int_{V_{\text{sub}}} \mathbf{D} \cdot \mathbf{E} dV}{\frac{1}{2} \int_{V_{\text{tot}}} \mathbf{D} \cdot \mathbf{E} dV}. \quad (2)$$

This ratio p_{bulk} can be tuned in-situ (Fig. 1b), and with a large on-off ratio that is set by the stroke of the piezoelectric positioner and the attenuation length of the waveguide.

For materials with a permittivity different from vacuum, inserting the dielectric shifts the frequency of the cavity. While the positioner does feature a built-in resistive position indication, this frequency shift allows frequency to serve as a proxy for sample position (see Fig. 1b), bypassing the need for the position indication to determine system geometry. If the permittivity is not purely real, insertion also alters the energy decay rate of the cavity mode in a ringdown experiment, as shown in Fig. 1c (see App. A for details of the ringdown measurement). This in-situ tunability enables a differential measurement of sample loss to be performed in a single cooldown.

Ideally, all the additional loss observed after insertion of the sample could be attributed to the sample. However, as insertion of the sample perturbs the field distribution of the cavity mode, the cavity’s “background” loss also changes. This effect can be modeled with the participation formalism, in which the loss of an electromagnetic mode is the weighted-sum of material loss-factors q_j^{-1} :

$$Q^{-1} = \sum_j p_j q_j^{-1}, \quad (3)$$

where the weights $p_j \equiv U_j/U_{\text{tot}}$ given to each loss factor are known as the participations [24] (see App. B for details on the participations p_j). Implicit in this formalism are the assumptions that mechanisms for loss are linear and uniform within each material region. For comparison with other studies [12, 26], we assume that surface dielectric interfaces have a thickness of 3 nm (see App. B). We note the q_j^{-1} have natural physical interpretations: when j indexes a dielectric loss channel, q_j^{-1} is a loss tangent $\tan\delta_j$; and when j indexes a conductor loss channel, q_j^{-1} is the quotient of surface resistance and surface reactance: $q_{\text{cond}}^{-1} = R_s/(\omega\mu_0\lambda_L)$, where λ_L is the London penetration depth [27].

We consider a model with losses that originate in the cavity and losses that come from the dielectric sample. Fig. 1d indicates the two material regions of the cavity responsible for “background” losses: conductor loss in the superconducting walls of the cavity, and dielectric loss on the metal-air (MA) interface of those walls. While some energy does propagate along the waveguide and radiate out of the cavity, the waveguide length has been chosen to suppress this effect such that $Q_{\text{radiative}}^{-1} < 10^{-10}$, as verified by electromagnetics simulations. This is small compared to the sample losses intended to be resolved in this study, so we neglect radiative loss in our model.

Fig. 1e indicates the two regions associated with loss from the dielectric under study: the substrate-air (SA) interface [19] and the bulk of the dielectric. For a given dielectric sample, these two sources of dielectric loss depend sensitively on the sample position z , but their ratio

is predominantly set by the surface-to-volume ratio of the sample, which in the case of a thin sample, is largely insensitive to changes in z . This proportionality prevents the separate extraction of these individual loss rates from a measurement of a single sample. For this reason, we combine these two substrate loss channels into a single effective loss channel:

$$Q_{\text{sub}}^{-1} \equiv p_{\text{bulk}}q_{\text{bulk}}^{-1} + p_{\text{SA}}q_{\text{SA}}^{-1}. \quad (4)$$

From this composite loss channel we define an effective loss tangent of the sample: $q_{\text{sub}}^{-1} \equiv Q_{\text{sub}}^{-1}/p_{\text{bulk}}$.

Having chosen the loss channels under consideration, a set of measurements Q_i^{-1} , taken at a series of N different positions z_i , can be analyzed using a system of N equations:

$$\begin{pmatrix} Q_1^{-1} \\ Q_2^{-1} \\ \vdots \\ Q_N^{-1} \end{pmatrix} = \begin{pmatrix} \tilde{p}_{\text{cond},1} & p_{\text{MA},1} & p_{\text{bulk},1} \\ \tilde{p}_{\text{cond},2} & p_{\text{MA},2} & p_{\text{bulk},2} \\ \vdots & \vdots & \vdots \\ \tilde{p}_{\text{cond},N} & p_{\text{MA},N} & p_{\text{bulk},N} \end{pmatrix} \begin{pmatrix} \tilde{q}_{\text{cond}}^{-1} \\ q_{\text{MA}}^{-1} \\ q_{\text{sub}}^{-1} \end{pmatrix} \quad (5)$$

Here $\tilde{p}_{\text{cond},i} \equiv p_{\text{cond},i}/\omega_i$, and $\tilde{q}_{\text{cond}}^{-1} \equiv q_{\text{cond}}^{-1}\omega_i$. This accounts for the frequency dependence of q_{cond} . We refer to the matrix in Eq. 5 as the participation matrix \mathbf{P} [12, 24] and determine it with electromagnetic simulations (see App. B). The material loss factors q_j^{-1} are extracted from Eq. (5) using bounded (see App. E3) least-squares regression.

To decompose q_{sub}^{-1} into its components q_{bulk}^{-1} and q_{SA}^{-1} , a second measurement may be made on a sample of the same material but with different thickness. We demonstrate this process with EFG sapphire in Sec. III.

When analyzing various samples measured with the same cavity, we concatenate multiple instances of Eq. (5), providing each sample its own substrate loss factor and associated matrix column. Because this joint analysis permits only a single value of q_{MA}^{-1} and q_{cond}^{-1} , the number of free parameters in the model is reduced by four (compared to fitting three samples individually), ensuring consistency in the extracted substrate loss factors. Joint analysis is valid only if the cavity’s bare losses do not change between cooldowns, but this can be verified by measuring the cavity with each sample in the withdrawn position.

B. Sensitivity

When solving for the inferred material loss factors q_j^{-1} , the structure of \mathbf{P} determines how errors in the measured Q_i^{-1} propagate into loss factor errors $\sigma_{q_j^{-1}}$. If the propagated fractional error $\sigma_{q_j^{-1}}/q_j^{-1}$ is greater than 100%, then the loss factor is said to be unresolved as it is below the sensitivity of the measurement. One option for

predicting $\sigma_{q_j^{-1}}$ is to use the condition number [28] of \mathbf{P} to infer a bound on the severity of this error amplification. Such a perspective is useful in experimental design, and has been used to improve the conditioning of matrix problems like Eq. (5) [16, 17].

Another tool for predicting the error and sensitivity of an experiment is the covariance matrix. Given some participation matrix \mathbf{P} and an approximate expectation for the measurement errors $\sigma_{Q_i^{-1}}$, the covariance matrix \mathbf{C} allows the individual $\sigma_{q_j^{-1}}$ to be calculated (and not just bounded) from its diagonal elements $C_{jj} = \sigma_{q_j^{-1}}^2$. For the case of linear least-squares,

$$\mathbf{C} \equiv (\tilde{\mathbf{P}}^T \tilde{\mathbf{P}})^{-1}, \quad (6)$$

with $\tilde{P}_{ij} \equiv P_{ij}/\sigma_{Q_i^{-1}}$ [29]. The specificity of this information makes the covariance matrix a powerful tool for experimental design. We use it to predict the sensitivity of the dielectric dipper.

Accurately predicting an extracted loss factor error $\sigma_{q_j^{-1}}$ by solving Eq. 6 requires two things: the participation matrix \mathbf{P} (obtained from simulations) and an anticipated absolute measurement error, $\sigma_{Q_i^{-1}}$. The error in the measurement of a line-width often scales with the line-width, so we assume a fractional measurement error $\sigma_{Q_i^{-1}}/Q_i^{-1}$. Converting an anticipated fractional error into an anticipated absolute error requires an accurate prediction of Q_i^{-1} , which can be made using participations from simulation, and assumed values for q_j^{-1} as inputs to eq. (3). In this experiment, reasonable values for q_{cond}^{-1} and q_{MA}^{-1} can be chosen during the design phase from measurements of a cavity made of the same materials as the planned design (or the exact cavity to be used if it is available). However, during the design phase, the substrate material loss factors q_{bulk}^{-1} and q_{SA}^{-1} are unknown. We therefore calculate $\sigma_{q_{\text{sub}}^{-1}}$ as a function of these unknown substrate loss factors. The results of this error analysis are plotted in color in Fig. 2 in terms of the 95% confidence interval $C.I. = 2\sigma_{q_{\text{sub}}^{-1}}$. Also seen in Fig. 2 are white contours which indicate the results of converting these absolute errors into fractional errors, $2\sigma_{q_{\text{sub}}^{-1}}/q_{\text{sub}}^{-1}$, and provide a visualization of which loss tangent combinations would be resolvable by this measurement with 95% confidence.

When the losses from the sample are sufficiently small, the fractional error exceeds one and the substrate material quality factor q_{sub}^{-1} can no longer be distinguished from zero. The 100% fractional error contour indicates the boundary of this region. The calculated error near the boundary of this region, $\sigma_{q_{\text{sub}}^{-1}} < 5$ parts-per-billion, indicates the predicted sensitivity of the experiment. The sensitivity of our experiment is predominantly limited by our ability to know the composition of the cavity losses, and in turn our ability to distinguish substrate losses from losses in the metal-air interface of the cavity (see App. C for details). Thus, even though multiple participations change with sample insertion to some degree, we




FIG. 2. **Simulated sensitivity.** Estimated q_{sub}^{-1} 95% confidence interval ($C.I.$), plotted (in color) as a function of hypothetical dielectric material quality factors q_{SA}^{-1} and q_{bulk}^{-1} . White curves indicate the 3%, 10% and 100% fractional error contours. We assume fractional measurement error $\sigma_{Q_i^{-1}}/Q_i^{-1} = 1\%$ (limited by our ability to separate coupling loss from total loss), cavity material quality factors $q_{\text{MA}}^{-1} = 3 \times 10^{-2}$, and $q_{\text{cond}}^{-1} = 2 \times 10^{-5}$ and $N = 30$ measurement positions. Participations are as simulated for the 460 μm thick EFG sapphire sample.

can still measure dielectric losses with 5 part-per-billion sensitivity.

III. SAPPHIRE MEASUREMENTS

To demonstrate the utility of the dielectric dipper, we used it to determine whether bulk dielectric loss plays a role in the lifetimes of state-of-the-art superconducting qubits on sapphire [2, 3]. To do this, we measured samples from wafers of sapphire cut from sheets grown by the same edge-defined film-fed growth method (EFG) [30] as the substrates used in those studies. Samples of two different thicknesses (and the same length and width) were measured in order to localize substrate loss to either the bulk or the substrate-air interface using Eq. (4). Motivated by previous studies [13, 31] on low-loss sapphire substrates grown with a different process, we also measured a third sapphire sample which was cut from a boule grown by the heat-exchanger method (HEM) [30] and screened after production for low optical absorption and wavefront aberration (HEMEX) [32]. For all three samples, we performed power sweeps with the sample inserted to ensure that following measurements are made in the low power regime. Withdrawn power sweeps are also performed to infer bounds on the background losses of the cavity. We then performed a position sweep at low excitation power to extract the substrate loss tangents.




FIG. 3. **Power dependent losses.** Measured internal loss Q^{-1} with the sample withdrawn (orange) and inserted (green) as a function of the average initial cavity photon number n in the ringdown measurement, with 100 μm thick EFG sapphire (a), 460 μm thick EFG sapphire (b), and 440 μm thick HEMEX sapphire (c). Curves are fits of the TLS model in Eq. (7) and the shaded region reflects propagation of the standard error of the fit parameters q_j^{-1} to Q^{-1} . For the inserted measurements of EFG sapphire, fits are only to measurements with $n \lesssim 2 \times 10^{10}$, as above this photon number the loss deviates from the TLS model. The origin of this deviation is unclear but we conjecture it to be related to heating of the dielectric sample.

A. Power-dependence of loss

Due to the saturability of dielectric loss, measurements which aim to be representative of qubit performance must be taken in this same low-power regime where superconducting qubits operate. To identify this regime, we begin with the sample completely inserted, and measure the cavity's impulse response as a function of the excitation power in the ringdown experiment. We parameterize cavity excitation by the average number of photons n created by the impulse (see App. A 4). The green circles in Fig. 3 show the results for the 100 μm EFG sample (a), the 460 μm EFG sample (b), and the 440 μm HEMEX sample (c). In each case, the total internal loss of the cavity asymptotically approaches some high-power value Q_{hp}^{-1} , and some low-power value $Q_{\text{lp}}^{-1} = Q_{\text{hp}}^{-1} + Q_{\text{sat}}^{-1}$, with a smooth transition between these two regimes. This power-dependence can be captured by a two-level-system (TLS) model for dielectric loss [33] which we fit to the data:

$$Q^{-1} = Q_{\text{hp}}^{-1} + \frac{Q_{\text{sat}}^{-1}}{\sqrt{1 + \left(\frac{n}{n_c}\right)^\alpha}}. \quad (7)$$

Here, the critical photon number n_c sets the scale for saturation and α sets the width of this transition. Table I shows the extracted fit parameters for these sample-inserted power sweeps.

The inserted power sweeps show that cavity population $n = 10^5$ is sufficiently below n_c , such that measured loss is only weakly dependent on power and closely approximates the low-power behavior Q_{lp}^{-1} , as is appropriate for predicting their performance as superconducting qubit substrates. We therefore fix the excitation photon number for the remaining measurements to $n = 10^5$ (see App. E 1 for details). This results in a peak electric field

Sample	n_c (millions)	α
100 μm EFG sapphire	600(300)	0.42(5)
460 μm EFG sapphire	2000(1800)	0.29(4)
440 μm HEMEX sapphire	300(30)	0.40(1)

TABLE I. Fit results for the inserted power sweeps in Fig. (3) using the TLS model in Eq. (7).

in the sample of about 4×10^{-8} V/m. Note that the extracted values of n_c are a factor of 10^6 greater than those found in coplanar waveguide resonators [34, 35], reflecting the much larger mode volume of a 3D cavity [36].

To measure the power dependence of the cavity background alone and bound the low-power values of q_{cond} and q_{MA} , we repeat the power sweep with the sample completely withdrawn (orange circles in Fig. 3). Previous studies on conductor losses in aluminum cavities have found no power dependence up to 10^{10} photons [20]. If we assume this trend continues up to 10^{14} photons, the measured power dependence may be attributed entirely to saturation of dielectric losses at the cavity's metal-air interface. This measurement therefore provides two useful constraints on the cavity loss factors q_{cond}^{-1} and q_{MA}^{-1} : first, the conductor loss factor is bounded from above by the measurement with the smallest loss $\min(Q^{-1})$ (typically a high-power measurement, ideally described by Q_{hp}^{-1}); and second, the metal-air loss factor is bounded from below by the change in the cavity loss ΔQ^{-1} (ideally described by Q_{sat}^{-1}). More precisely,

$$q_{\text{cond}}^{-1} \leq \frac{\min(Q^{-1})}{p_{\text{cond}}}, \quad (8)$$

$$q_{\text{MA}}^{-1} \geq \frac{\Delta Q^{-1}}{p_{\text{MA}}}.$$

Table VI in App. E 3 lists the bounds obtained in this way.

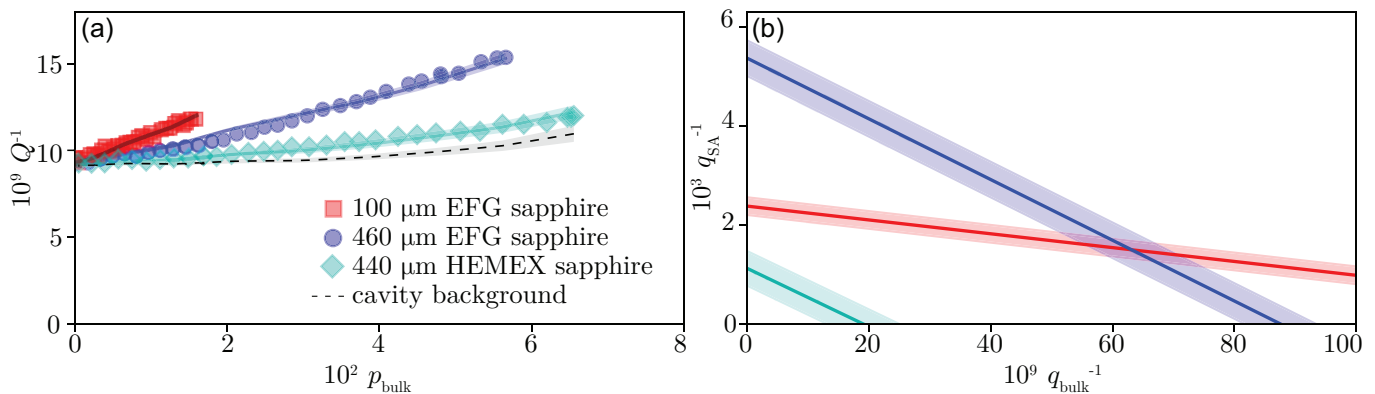


FIG. 4. **Measurements of bulk and surface loss tangents.** (a) Measured internal loss Q^{-1} plotted as a function of bulk participation p_{bulk} , for 100 μm thick EFG sapphire, 460 μm thick EFG sapphire, and 440 μm thick HEMEX sapphire. The excitation pulse for these measurements created approximately 10^5 photons (low-power regime), and the mixing chamber temperature was 20 mK (low-temperature regime). Solid lines are fits to the three-loss model in Eq. (5) consistent with bounds provided by the power sweeps (see App. E3). The shaded region shows propagation of the 95% confidence interval of q_{sub}^{-1} , and the dashed line is the estimated cavity background, comprised of conductor and MA dielectric losses. (b) The effective substrate loss factor q_{sub}^{-1} is a weighted average of substrate bulk and surface loss factors. By Eq. (4), a measurement of q_{sub}^{-1} defines a line in the 2-d parameter space of q_{bulk}^{-1} and q_{SA}^{-1} . The intersection of two such lines, each defined by measurements on the same material with different form factors, identifies that material’s bulk and SA loss factors. The single measurement of HEMEX sapphire does not separately resolve bulk and SA losses, but bounds both the bulk and SA loss factors.

B. Inserting the sample in-situ

With the cavity loss factors constrained and the edge of the low-power regime identified, we proceed with the position sweeps described in Sec. II. The results are shown in Fig. 4a. For each sample in this study, internal loss Q^{-1} of the cavity mode is plotted against participation of the substrate bulk in the cavity mode. When the substrate is withdrawn from the cavity, p_{bulk} approaches zero, and the measured loss is that of the empty cavity. As the sample is inserted, p_{bulk} rises, and the cavity mode inherits loss from the sample. All samples are inserted until $z \approx 30 \mu\text{m}$.

Examining the position sweeps for the two EFG samples, one can see that the total change in Q^{-1} is larger for the thicker sample. As the samples have almost identical surface areas, this indicates that surface losses do not account for all of the measured loss. To check that intuition and distinguish the bulk and surface contributions, we first extract the substrate loss tangent q_{sub}^{-1} for each sample by solving Eq. (5) (constrained by Eq. (8)). This separates substrate loss from the “background” cavity losses. As q_{sub}^{-1} is a weighted sum of the bulk and surface loss tangents of the sample material per Eq. (4), a measurement of q_{sub}^{-1} constrains those loss tangents, in the space of all possible ordered pairs $(q_{\text{bulk}}^{-1}, q_{\text{SA}}^{-1})$, to lie along a line (see Fig. 4b). A single constraining line has several useful geometric interpretations: its x and y -intercepts correspond to upper bounds on the bulk and SA interface loss tangents; its distance from the origin is proportional to the measured value of q_{sub}^{-1} ; and its slope is determined by the ratio of surface and bulk participations $p_{\text{SA}}/p_{\text{bulk}}$.

Because the two EFG sapphire samples have different surface-to-volume ratios, the constraint lines associated with these two measurements have different slopes. Their intersection marks the only viable $(q_{\text{bulk}}^{-1}, q_{\text{SA}}^{-1})$ pair which is consistent with the measured q_{sub}^{-1} of both samples. In this way, we distinguish between bulk loss and surface loss in EFG sapphire and extract $q_{\text{bulk}}^{-1} = 63(8) \times 10^{-9}$, and $q_{\text{SA}}^{-1} = 15(3) \times 10^{-4}$ [19].

Fig. 4a also displays the results of the position sweep on a HEMEX sapphire sample. The more gradual slope of this data set, compared to the EFG data sets, indicates that HEMEX sapphire has lower dielectric losses. Having measured one HEMEX sample, there is only a single contour to constrain possible $(q_{\text{bulk}}^{-1}, q_{\text{SA}}^{-1})$ pairs, and the bulk and surface contributions are not separately resolved. Nonetheless, the measurement bounds both bulk and surface losses. These bounds are represented graphically by the x and y -intercepts of the HEMEX constraint line in Fig 4b. The inferred bound on bulk loss is $q_{\text{bulk}}^{-1} < 19(6) \times 10^{-9}$, consistent with the findings of Ref. [31], and smaller than the bulk loss tangent of EFG sapphire by a factor of four (see Table II.)

Material	q_{bulk}^{-1}	q_{SA}^{-1}
EFG sapphire	$63(8) \times 10^{-9}$	$15(3) \times 10^{-4}$
HEMEX sapphire	$< 19(6) \times 10^{-9}$	$< 11(4) \times 10^{-4}$

TABLE II. Measurements and bounds of bulk and surface dielectric loss tangents for EFG and HEMEX sapphire [19].

IV. IMPLICATIONS FOR TRANSMON LIFETIMES

Having sensitively measured the loss tangents of these forms of sapphire, we now calculate the current and future coherence limits that they imply for typical transmon qubits. Assuming a qubit transition frequency of 4 GHz and a bulk participation of $p_{\text{bulk}} = 80\%$, these measurements imply that a transmon fabricated on EFG sapphire would be bulk limited to a quality factor of $Q \leq 20$ million, or $T_1 \leq 800 \mu\text{s}$. Several comments can be made on this coherence limit, which is illustrated in Fig. 5 alongside the coherence reported in several recent studies of transmons on sapphire.




FIG. 5. **Inferred bulk coherence limit and comparison to measured transmon lifetimes.** The coherence limit imposed by EFG sapphire and HEMEX sapphire assuming a bulk participation of $p_{\text{bulk}} = 0.8$, compared to quality factors of transmons on sapphire from select studies. Comparisons are made in terms of total loss Q^{-1} and in terms of T_1^{eff} ; a T_1 calculated from Q^{-1} assuming a 4 GHz qubit transition frequency.

First, this limit is consistent with previous studies [12] in which transmon lifetimes were found to be limited by surface losses, as those studies used devices with quality factors up to 4 million, and therefore would not have observed a strong effect from a coherence limit at the level of 20 million. This is illustrated in Fig. 5 by the EFG sapphire bulk loss limit being significantly below the total loss rates Q^{-1} observed in that study.

Second, this limit is consistent with the best-reported lifetimes of transmons made on EFG sapphire [2, 3, 37], as no such study has yet produced a transmon with quality factor $Q > 20$ million. Interestingly, the quality factors of the longest lived-transmons are approaching this bound, as seen in Fig. 5. Supposing that the fabrication process used to make those transmons did nothing to improve the bulk loss tangent of the substrate, this comparison suggests those transmons may be significantly bulk limited. For bulk-limited devices, replacing EFG with HEMEX sapphire will substantially increase transmon lifetimes. The increased coherence limit is also shown in Fig. 5. Further, the shared chemistry and lattice constant

of these materials implies that this substitution would not require any modification of the fabrication process.

A general insight gained from this comparative measurement is that not all sapphire is equal. As with high-power measurements made at several kelvin [38–41], growth method has a meaningful impact on the millikelvin and low-power loss tangent of sapphire, and in turn the lifetime of transmons built upon that sapphire.

Additionally, for completeness we use the measured value of q_{SA}^{-1} to calculate the SA coherence limit for a typical transmon. Assuming $p_{\text{SA}} = 1.18 \times 10^{-4}$ [12], substrate-air surface dielectric loss would limit the quality factor of a transmon to $Q \leq 7$ million. In real devices, however, the loss tangent q_{SA}^{-1} will likely depend on the details of the substrate surface preparation. More definite claims can be made about the SA interface coherence limit of real devices by leveraging the flexibility of the dipper technique and measuring dielectric samples which have undergone the same surface preparation sequence as the substrates used for those devices (For details on the preparation of samples in this study, see App. D).

Finally, the model used in our analysis has assumed that all sample loss is dielectric in origin rather than magnetic, but we can extend the model and allow sample loss to be a combination of dielectric and magnetic contributions. By considering our results for EFG sapphire in comparison to performance of striplines from a previous study [42], we can make two bounding statements about magnetic bulk loss in EFG sapphire. First, the EFG bulk loss tangent measured in this study must be at least 97% truly dielectric in origin, and no more than 3% could be from misattributed magnetic loss. This validates the choice to neglect magnetic sample loss in our analysis. Second, bulk magnetic loss is not a major source of loss in even the longest-lived transmons on EFG sapphire, accounting for less than 10% of the total loss. See App. H for details.

V. CONCLUSION

A precise knowledge of the decoherence mechanisms in superconducting qubits aids in all efforts to improve their coherence. Here, we present a method to precisely measure bulk dielectric loss with 5×10^{-9} sensitivity, and apply it in a study of EFG sapphire. At low powers and temperatures, we measure the bulk loss tangent at microwave frequencies to be $63(8) \times 10^{-9}$. This entails a bulk limit to the quality factor of a typical transmon on an EFG sapphire substrate of $Q \lesssim 20$ million, consistent with the longest lifetimes of transmons on sapphire, and suggests bulk dielectric loss is a major source of loss in those devices. We also measure HEMEX sapphire, and bound its bulk loss tangent to be less than $19(6) \times 10^{-9}$. Such a low-loss substrate would reduce one of the dominant sources of loss by a factor of four or more, and could enable longer transmon lifetimes.

Looking forward, measurements of other materials,

such as silicon, quartz, or lithium niobate, can inform efforts to improve coherence of a variety of devices, including qubits, on those substrates. We also expect that correlation of precision dielectric loss measurements with other physical probes (e.g. x-ray spectroscopy or scattering) will shed light on the underlying physical mechanisms behind this loss.

We acknowledge N. Ofek, Y. Liu, and P. Reinhold, for their work building the FPGA firmware and software used in this experiment, C. J. Axline and K. Li for their contributions in developing earlier versions of the experiment, and N. P. de Leon, A. Walter, A. Barbour, M. H. Devoret and S. M. Girvin for useful discussions. Use of facilities was supported by the Yale Institute for Nanoscience and Quantum Engineering (YINQE), Yong Sun, and the Yale SEAS cleanroom. We thank MIT Lincoln Lab and IARPA for providing a Josephson travelling wave parametric amplifier. This research, including funding for A.P.R., was supported by the U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers, and the Co-design Center for Quantum Advantage (C2QA) under contract number DE-SC0012704. Further support, including funding for B.J.C., comes from the Army Research Office (ARO), under Grant Number W911NF-18-1-0212. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official parties, either expressed or implied, of the Army Research Office (ARO) or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for government purposes notwithstanding any copyright notation herein. Contributions: B.J.C. and R.J.S. conceived the experiment; A.P.R. and B.J.C. designed and characterized the experimental apparatus, carried out the experiments, analyzed the data, and wrote the manuscript with input from all authors; L.F. and R.J.S. supervised the work. Disclosure: L.F. and R.J.S. are founders and shareholders of Quantum Circuits Inc.

Appendix A: Ringdown measurement

To measure the internal energy decay rate κ_{int} of a microwave cavity with (classical) mode amplitude a , the cavity is excited with an input field a_{in} (Fig. 6a) and the output field a_{out} is monitored to determine the rate at which its energy is lost.

During the pulse, the cavity mode amplitude a rises. After the pulse, a decays exponentially (Fig. 6b). While a portion of the energy from the excitation pulse makes its way into the cavity mode, the rest is reflected off the cavity port and propagates to the receiver (see App. J for the complete wiring diagram). At the same time, the cavity emits a field proportional to a . The reflected pulse and the cavity emission interfere to create the output field a_{out} , which is measured at the receiver (Fig. 6c).




FIG. 6. **Dynamics of a resonant ringdown measurement.** The incoming field a_{in} (a), cavity mode amplitude a , (b), and outgoing field a_{out} (c) as a function of time. Figure is not to scale.

We determine the total loss of the cavity mode κ_{tot} by fitting the decay rate of $|a_{\text{out}}|^2$ for times $t > t_p$, during which $a_{\text{out}} \propto a$. To distinguish between the effects of dephasing and actual energy decay, the emitted field must be sampled at a rate faster than the dephasing time of the cavity (see App. A 2), and data from multiple shots of the experiment must be processed with a phase-insensitive averaging scheme (see App. A 1).

The external coupling-rate of the cavity mode κ_{ext} is encoded in the overall amplitude of the emitted field in comparison to the amplitude and duration of the excitation pulse, and can be inferred from the outgoing field after the beginning ($t = 0^+$) and after the end ($t = t_p^+$) of the excitation pulse (see App. A 3).

Those two loss rates can then be subtracted to determine the internal energy decay rate $\kappa_{\text{int}} = \kappa_{\text{tot}} - \kappa_{\text{ext}}$, which can be divided by the cavity mode's resonant frequency ω_a to yield a dimensionless loss: $Q_{\text{int}}^{-1} = \kappa_{\text{int}}/\omega_a$.

1. Decay of a harmonic oscillator with time-varying resonant frequency

In the presence of a fluctuating cavity frequency, the cavity output field $|\langle a_{\text{out}} \rangle|$ (averaged over many shots) does not decay at the energy decay rate κ_{tot} . Rather, it decays at a rate which is also affected by cavity dephasing. For this reason, measurements of $|\langle a_{\text{out}} \rangle|$, such as those from a vector network analyzer, cannot be used to reliably extract κ_{tot} . Instead, one must measure the ensemble-average of the cavity output power, $\langle |a_{\text{out}}|^2 \rangle$, which is a phase-insensitive quantity and decays at the same rate, κ_{tot} , as the cavity energy. In many cases, this distinction is irrelevant, but in our experiment, mechanical vibrations of the high-permittivity dielectric sample cause the cavity frequency to jitter by as many as 100 line widths (see Fig. 7). When the dielectric is inserted, accurate measurement of κ_{tot} therefore requires fitting the decay of $\langle |a_{\text{out}}|^2 \rangle$, rather than that of $|\langle a_{\text{out}} \rangle|^2$. In this section we describe how the energy decay rate κ_{tot} can be inferred from ringdown measurements of an oscillator with a fluctuating resonance frequency.

Suppose an oscillator has a time-varying resonant frequency $\omega_a = \omega_0(1 + \epsilon(t))$. Assume the fluctuations of the resonance $\epsilon(t)$ are caused by a stochastic process with zero mean,

$$\langle \epsilon(t) \rangle = 0, \quad (\text{A1})$$

and characterized by a power spectral density $S(f)$:

$$\langle \epsilon(t)^2 \rangle = \int_0^\infty df S(f) \cos(2\pi ft). \quad (\text{A2})$$

where the brackets $\langle \cdot \rangle$ denote ensemble averaging.

The Langevin equation [43] dictates how a evolves in time. In the absence of an excitation pulse (here true for $t > t_p$), the measured output field a_{out} is proportional to a and governed by an equation of the same form:

$$\dot{a}_{\text{out}} = [i\omega_0(1 + \epsilon(t)) - \kappa_{\text{tot}}/2]a_{\text{out}}. \quad (\text{A3})$$

After separation of variables, integration yields

$$a_{\text{out}}(t) = e^{(i[\omega_0 t + \theta(t)] - \kappa_{\text{tot}} t/2)}, \quad (\text{A4})$$

with

$$\theta(t) \equiv \omega_0 \int_0^t \epsilon(t') dt'. \quad (\text{A5})$$

Here we have chosen the constant of integration to reflect an initial condition of unit amplitude in the oscillator mode.

The average of the output power decays as

$$\langle |a_{\text{out}}(t)|^2 \rangle = e^{-\kappa_{\text{tot}} t}. \quad (\text{A6})$$

We identify the decay time of the oscillator as $\tau \equiv 1/\kappa_{\text{tot}}$. $\langle |a_{\text{out}}(t)|^2 \rangle$ is independent of $\theta(t)$ and thus unaffected by the fluctuations in the resonant frequency.

This is in contrast to the output field,

$$\begin{aligned} |\langle a_{\text{out}}(t) \rangle| &= e^{-\kappa_{\text{tot}} t/2} |\langle e^{i\theta(t)} \rangle|, \\ &= e^{-\kappa_{\text{tot}} t/2} \left(1 - \frac{1}{2} \langle \theta(t)^2 \rangle + \mathcal{O}(\langle \theta(t)^3 \rangle) \right). \end{aligned} \quad (\text{A7})$$

By our assumption that the fluctuations are zero-mean, the linear term in the expansion vanishes. (If the noise is Gaussian, the higher moments vanish and this approximation is exact.) To evaluate the quadratic term, we follow the treatment given in Ref. [44]:

$$\begin{aligned} \langle \theta(t)^2 \rangle &= \omega_0^2 \int_0^t \int_0^t dt' dt'' \langle \epsilon(t') \epsilon(t'') \rangle, \\ &= \omega_0^2 \int_0^t \int_0^t dt' dt'' \int_0^\infty df S(f) \text{Re}\{e^{2\pi i f(t' - t'')}\}, \\ &= \omega_0^2 \int_0^\infty df S(f) W_0(f, t), \end{aligned} \quad (\text{A8})$$

where the spectral weight function $W_0(f, t)$

$$\begin{aligned} W_0(f, t) &= \left| \int_0^t dt' e^{2\pi i f t'} \right|^2, \\ &= \frac{\sin(\pi f t)^2}{(\pi f)^2}, \end{aligned} \quad (\text{A9})$$

low-pass filters the fluctuations.

We therefore have

$$\begin{aligned} |\langle a_{\text{out}} \rangle| &= e^{-\kappa_{\text{tot}} t/2} \times \\ &\left(1 - \frac{1}{2} \omega_0^2 \int_0^\infty df S(f) W_0(f, t) + \mathcal{O}(\langle \theta(t)^3 \rangle) \right). \end{aligned} \quad (\text{A10})$$

The exponential prefactor represents energy decay from the oscillator, and the rest of the expression accounts for the effect of pure dephasing. The form of this dephasing depends on the noise spectrum $S(f)$.

In summary, the exponential decay of the phase-insensitive quantity $\langle |a_{\text{out}}|^2 \rangle$ yields the time-scale for energy decay τ , and the phase-sensitive quantity $|\langle a_{\text{out}} \rangle|$ has a $1/e$ time which is affected by both energy decay and dephasing in the cavity. This is analogous to the distinction between T_1 and T_2 of a spin ensemble. Fluctuations in the cavity's resonant frequency can be detected by comparing τ to the $1/e$ time of $\langle |a_{\text{out}}|^2 \rangle$. By converting these decay rates to quality factors, they can be checked for consistency with quality factors extracted from circle fits of spectroscopic measurements [45] (See Fig. 7).

2. Finite measurement bandwidth

The above arguments can be generalized to account for the effects of finite measurement bandwidth, such as that imposed by a non-instantaneous measurement. In short, finite measurement bandwidth causes a reduction in the contrast of the measurement $\langle |a_{\text{out}}^m|^2 \rangle$, but does not change the time scale for decay. For completeness,




FIG. 7. Consistency between ringdown and spectroscopic measurement. When the sample is withdrawn from the cavity, Q_{tot} inferred from the decay rates of $\langle |a(t)|^2 \rangle$ (blue) and $|\langle a(t) \rangle|^2$ (red) are in close agreement with each other, and with Q_{tot} from circle fits [45] of measurements taken by a vector network analyzer $\langle S_{21}(f) \rangle$ (green). As the sample is inserted, energy decay and frequency jitter tend to reduce Q_{tot} extracted from $|\langle a(t) \rangle|^2$ or $\langle S_{21}(f) \rangle$, however $\langle |a(t)|^2 \rangle$ is not susceptible to frequency jitter (see Eq. A6) and thus only experiences reduction in Q_{tot} cause by energy decay. Q_{tot} can be, in general, non-monotonic in p_{bulk} due to frequency-dependent coupling loss, which is measured and subtracted off before further processing (see App. A 3).

we present those arguments below, as the suppression of contrast is an effect which must be controlled in order to accurately measure Q_{ext}^{-1} (a measurement discussed in the next subsection, App. A 3).

The result of a finite bandwidth measurement of the output field $a_{\text{out}}^m(t)$ is the convolution of the signal $a_{\text{out}}(t)$ with the impulse response $h(t)$ of the detector.

In practice, the detector's bandwidth is set by the inverse of the period t_m over which the signal is averaged:

$$a_{\text{out}}^m(t) \approx \frac{1}{t_m} \int_{t-t_m/2}^{t+t_m/2} dt' e^{i\theta(t') - \kappa_{\text{tot}} t'/2}. \quad (\text{A11})$$

The resonant frequency ω_0 doesn't appear in this expression because we assume that detection occurs after demodulation.

To evaluate the effect on the measured coherence time of the cavity, we focus on the measured average output field

$$|\langle a_{\text{out}}^m \rangle| = \left| \left\langle \frac{1}{t_m} \int_{t-t_m/2}^{t+t_m/2} dt' e^{i\theta(t') - \kappa_{\text{tot}} t'/2} \right\rangle \right|.$$

Expanding as in Ref. [44],

$$|\langle a_{\text{out}}^{\text{m}} \rangle| = \left| \frac{1}{t_{\text{m}}} \left[\frac{\sinh(\kappa_{\text{tot}} t_{\text{m}}/2)}{(\kappa_{\text{tot}}/2)} e^{-\kappa_{\text{tot}} t/2} - \frac{1}{2} \omega_0^2 \int_0^\infty df S(f) \times \int_{t-t_{\text{m}}/2}^{t+t_{\text{m}}/2} dt' e^{-\kappa_{\text{tot}} t'/2} W_0(f, t') + \mathcal{O}(\langle \theta(t')^3 \rangle) \right] \right|. \quad (\text{A12})$$

Comparing this result to eq. (A10), we can see that the time-averaging has two effects. The first effect is independent of the frequency fluctuations, as can be seen by setting $S(f) = 0$:

$$|\langle a_{\text{out}}^{\text{m}} \rangle| = \frac{\sinh(\kappa_{\text{tot}} t_{\text{m}}/2)}{(\kappa_{\text{tot}} t_{\text{m}}/2)} e^{-\kappa_{\text{tot}} t/2}. \quad (\text{A13})$$

In this limit, the time-dependence is an exponential decay characterized by a timescale $T_2 = 2/\kappa_{\text{tot}}$. The averaging, however, can reduce the initial measured amplitude. In the limit of fast measurements ($\kappa_{\text{tot}} t_{\text{m}} \ll 1$), unit amplitude is recovered.

The second effect of time averaging is a redefinition of the spectral weight function. The new function is a weighted average of W_0 taken over the acquisition window.

By a similar argument, we can evaluate the effect of measurement bandwidth on the measured energy decay rate. The average output power evaluates to

$$\langle |a_{\text{out}}^{\text{m}}|^2 \rangle = e^{-\kappa_{\text{tot}} t} \left(1 - \omega_0^2 \int_0^\infty df S(f) W_1(f) \right) + \mathcal{O}(\langle [\theta(t') - \theta(t'')]^3 \rangle) + \mathcal{O}(\kappa_{\text{tot}} t_{\text{m}})^2.$$

with

$$W_1(\omega) \equiv \frac{1 - \text{sinc}(\omega t_{\text{m}}/2)^2}{(\kappa_{\text{tot}}/2)^2 + \omega^2} + \mathcal{O}(\kappa_{\text{tot}} t_{\text{m}})^3, \quad (\text{A14})$$

and $\omega = 2\pi f$.

From Eq. (A14), we see that the fluctuations change the contrast, but do not change the decay rate. The amount by which they change the contrast is controlled by the spectrum of the fluctuations and the averaging time t_{m} because the spectral weight $W_1(f)$ acts as a low-pass filter with corner frequency $1/t_{\text{m}}$.

3. Measurement of the cavity coupling rate

After measuring the total loss κ_{tot} of the system by fitting $\langle |a_{\text{out}}|^2 \rangle$ to an exponential decay, measurement of the coupling loss κ_{ext} enables the isolation of internal loss $\kappa_{\text{int}} = \kappa_{\text{tot}} - \kappa_{\text{ext}}$. We now describe how the external coupling of the cavity mode can be inferred from a ringdown experiment. Consider a cavity with classical mode amplitude a and frequency ω_{a} and total loss rate κ_{tot} coupled to a single port with rate κ_{ext} . The cavity may be driven through this port by an incident field a_{in}

with its carrier frequency ω_{d} detuned from the cavity by $\Delta \equiv \omega_{\text{d}} - \omega_{\text{a}}$. In the rotating frame of the drive, the Langevin equation [43] is

$$\dot{a} = (i\Delta - \kappa_{\text{tot}}/2)a + \sqrt{\kappa_{\text{ext}}}a_{\text{in}}. \quad (\text{A15})$$

The general solution for the cavity field is

$$a(t) = -\sqrt{\kappa_{\text{ext}}} \int_0^t dt' a_{\text{in}}(t') e^{(i\Delta - \kappa_{\text{tot}}/2)(t-t')}. \quad (\text{A16})$$

Let the envelope of the incident field be that of a rectangular pulse,

$$a_{\text{in}} = \begin{cases} 0 & t < 0, \\ a_0 & 0 < t < t_{\text{p}}, \\ 0 & t_{\text{p}} < t. \end{cases}$$

The resulting cavity field is given by

$$a(t) = \begin{cases} 0 & t < 0, \\ \tilde{a}(t) & 0 < t < t_{\text{p}}, \\ \tilde{a}(t_{\text{p}}) e^{(i\Delta - \kappa_{\text{tot}}/2)(t-t_{\text{p}})} & t_{\text{p}} < t. \end{cases}$$

where

$$\tilde{a}(t) \equiv \frac{a_0 \sqrt{\kappa_{\text{ext}}}}{i\Delta - \kappa_{\text{tot}}/2} \left(1 - e^{(i\Delta - \kappa_{\text{tot}}/2)t} \right). \quad (\text{A17})$$

We probe the cavity dynamics by measuring the emitted field a_{out} , which input-output theory relates to the incident field and the cavity mode [43]

$$a_{\text{out}} = \sqrt{\kappa_{\text{ext}}}a + a_{\text{in}}, \quad (\text{A18})$$

Substituting in the expressions for the input and cavity fields,

$$a_{\text{out}}(t) = \begin{cases} 0 & t < 0, \\ \sqrt{\kappa_{\text{ext}}}\tilde{a}(t) + a_0 & 0 < t < t_{\text{p}}, \\ \sqrt{\kappa_{\text{ext}}}\tilde{a}(t_{\text{p}}) e^{(i\Delta - \kappa_{\text{tot}}/2)(t-t_{\text{p}})} & t_{\text{p}} < t. \end{cases} \quad (\text{A19})$$

The emitted fields slightly after the start of the pulse ($t = 0^+$) and slightly after the end of the pulse ($t = t_{\text{p}}^+$) are:

$$\begin{aligned} a_{\text{out}}(0^+) &= a_0, \\ a_{\text{out}}(t_{\text{p}}^+) &= \sqrt{\kappa_{\text{ext}}}\tilde{a}(t_{\text{p}}). \end{aligned} \quad (\text{A20})$$

These two points contain all the information required to extract the external coupling rate κ_{ext} . The output field $a_{\text{out}}(t)$ at these points in time can be converted to phase-insensitive output powers $\langle |a_{\text{out}}|^2 \rangle$. Assuming the cavity has been driven on-resonance ($\Delta = 0$),

$$\begin{aligned} \langle |a_{\text{out}}(0^+)|^2 \rangle &= a_0^2, \\ \langle |a_{\text{out}}(t_{\text{p}}^+)|^2 \rangle &= \frac{4a_0^2 \kappa_{\text{ext}}^2}{\kappa_{\text{tot}}^2} \left(1 - 2e^{-\kappa_{\text{tot}} t_{\text{p}}/2} + e^{-\kappa_{\text{tot}} t_{\text{p}}} \right). \end{aligned} \quad (\text{A21})$$

Solving for the external coupling κ_{ext} in terms of measured quantities and the experimental parameter pulse length t_p ,

$$\kappa_{\text{ext}} = \frac{\kappa_{\text{tot}}}{2(1 - e^{-\kappa_{\text{tot}}t_p/2})} \sqrt{\frac{\langle |a_{\text{out}}(t_p^+)|^2 \rangle}{\langle |a_{\text{out}}(0^+)|^2 \rangle}}. \quad (\text{A22})$$

As the pulse is typically very short compared to $1/\kappa_{\text{tot}}$, we can expand Eq. (A21) to second order in $\kappa_{\text{tot}}t_p/2$,

$$\begin{aligned} \langle |a_{\text{out}}(0^+)|^2 \rangle &= a_0^2, \\ \langle |a_{\text{out}}(t_p^+)|^2 \rangle &\approx a_0^2 \kappa_{\text{ext}}^2 t_p^2, \end{aligned} \quad (\text{A23})$$

and solve for κ_{ext} ,

$$\kappa_{\text{ext}} \approx \frac{1}{t_p} \sqrt{\frac{\langle |a_{\text{out}}(t_p^+)|^2 \rangle}{\langle |a_{\text{out}}(0^+)|^2 \rangle}}. \quad (\text{A24})$$

Though the output field at two precise moments in time is the only information needed for calculating κ_{ext} , a more precise determination of $\langle |a_{\text{out}}(t_p^+)|^2 \rangle$ can be had by fitting the entire ringdown portion of the data ($t < t_p$) to an exponential decay and extracting the ringdown amplitude.

A similar strategy could be employed to improve the precision of $\langle |a_{\text{out}}(0^+)|^2 \rangle$, but not as lucratively; because jitter in the system requires the pulse to be short (see App. A 4), there is not much pulse data ($0 < t < t_p$) for the fit. A workaround comes from noticing that $\langle |a_{\text{out}}(t)|^2 \rangle$, if not for interference from cavity emission, would have a constant amplitude $\langle |a_{\text{out}}(0^+)|^2 \rangle$, even for an arbitrarily long pulse. Assuming the transmission background of the entire measurement system $|S_{21}|$ is smooth and varies only slowly with frequency, $\langle |a_{\text{out}}(0^+)|^2 \rangle$ can be determined by measuring the transmitted power of a long pulse far above (perhaps hundreds of linewidths) and equally far below the cavity frequency ω_a , then averaging the two for an approximation of $\langle |a_{\text{out}}(0^+)|^2 \rangle$ at ω_a . Pulses played far off-resonant from the cavity will not excite the cavity, and there will be no emission signal to interfere with the reflected pulse signal.

If the cavity is frequency-tunable, as it is in the dipper by movement of the sample, the off-resonance condition can be met by changing the frequency of the cavity instead of the frequency of the pulse. If the cavity is moved away from the frequency of interest, then a single pulse at that frequency can be used to measure $\langle |a_{\text{out}}(0^+)|^2 \rangle$ at that frequency directly, rather than estimating it by linear interpolation between two measurements at other frequencies. In practice this is the approach we take, though both methods give similar results.

Finally, the short pulses ($t_p \approx 5 \mu\text{s}$) and weak coupling ($Q_{\text{ext}} \approx 2 \times 10^8$; coupling rate $\kappa_{\text{ext}}/2\pi \approx 100 \text{ Hz}$) used in this measurement mean that, for the same pulse amplitude, $\langle |a_{\text{out}}(0^+)|^2 \rangle$ will be larger than $\langle |a_{\text{out}}(t_p^+)|^2 \rangle$ by a few orders of magnitude. To avoid overdriving the

receiver, the transmission measurements of $\langle |a_{\text{out}}(0^+)|^2 \rangle$ can be performed with a pulse amplitude smaller by some reduction factor to make the $\langle |a_{\text{out}}(0^+)|^2 \rangle$ measurements more comparable in size to the $\langle |a_{\text{out}}(t_p^+)|^2 \rangle$ measurement. Before using this reduced-pulse measurement to calculate Q_{ext} , its result must be re-scaled such as to compensate for not only that reduction factor, but also for any compression on the way to the cavity. Even if this compression does not affect this smaller pulse for measuring $\langle |a_{\text{out}}(0^+)|^2 \rangle$, it may still affect the larger pulse used for measuring $\langle |a_{\text{out}}(t_p^+)|^2 \rangle$.

4. Calculating photon number

A cavity state with mode amplitude a has an average of n photons:

$$n(t) = |a(t)|^2. \quad (\text{A25})$$

To calculate n , we can use Eq. (A18) which relates the mode amplitude a to more directly accessible quantities $a_{\text{in}}(t)$ or $a_{\text{out}}(t)$. To infer cavity photon number from the field emitted by the cavity, we can start by solving Eq. (A18) for a after the pulse has stopped ($a_{\text{in}}(t) = 0$):

$$a(t) = \frac{a_{\text{out}}(t)}{\sqrt{\kappa_{\text{ext}}}}. \quad (\text{A26})$$

Substituting into eq. A25 gives an expression for $n(t)$ in terms of the output field:

$$n(t) = \frac{|a_{\text{out}}(t)|^2}{\kappa_{\text{ext}}}. \quad (\text{A27})$$

This output field equation for photon number can alternately be expressed in terms of an output power $P_{\text{out}}(t) = \hbar\omega_a |a_{\text{out}}(t)|^2$. Solving for $n(t)$, we have

$$n(t) = \frac{P_{\text{out}}(t)Q_{\text{ext}}}{\hbar\omega_a^2}. \quad (\text{A28})$$

Using eq. A21, this output field equation can be converted into an input field equation which predicts the number of photons injected by a rectangular excitation pulse of duration t_p and amplitude a_0 :

$$n(t_p) = \frac{4\kappa_{\text{ext}}a_0^2}{\kappa_{\text{tot}}^2} \left(1 - 2e^{-\kappa_{\text{tot}}t_p/2} + e^{-\kappa_{\text{tot}}t_p} \right). \quad (\text{A29})$$

In the limit of a short pulse,

$$n(t_p) = \kappa_{\text{ext}}a_0^2t_p^2 + \mathcal{O}(\kappa_{\text{tot}}t_p)^2. \quad (\text{A30})$$

This input field equation for photon number can alternately be expressed in terms of an input power $P_{\text{in}} = \hbar\omega_d |a_{\text{in}}|^2 = \hbar\omega_d a_0^2$, which allows Eq. (A30) to be expressed as

$$n(t_p) = \frac{P_{\text{in}}t_p^2}{\hbar Q_{\text{ext}}}. \quad (\text{A31})$$

This section so far has assumed no jitter in the cavity frequency. In addition to simplifying Eq. (A30), a short drive pulse is important for proper control of photon number in the presence of jitter. In general, the photon number will be dependent upon the inner product of the cavity spectrum and the pulse spectrum. If the pulse is short enough that its spectrum is approximately uniform over the jitter width of the cavity, then that inner product is approximated well by the above expressions (such considerations are also essential for accurate measurement of Q_{ext}^{-1} per App. A 3).

Appendix B: Participation simulations

To calculate the participation ratios p_j , we use finite-element analysis software (Ansys HFSS) to solve for the eigenmodes of Maxwell's equations in our cavity. The participation ratios are defined in terms of material parameters and the field solutions \mathbf{E} and \mathbf{H} for the fundamental cavity mode. We consider four participations in our model: the bulk dielectric participation p_{bulk} , the substrate-air surface dielectric participation p_{SA} , the metal-air surface dielectric participation p_{MA} , and the conductor participation p_{cond} (equal to the kinetic inductance ratio in bulk superconductors [46]). We define these participations as

$$\begin{aligned} p_{\text{bulk}} &\equiv \frac{U_{\text{bulk}}}{U_{\text{tot}}}, \\ p_{\text{SA}} &\equiv \frac{U_{\text{SA}}}{U_{\text{tot}}}, \\ p_{\text{MA}} &\equiv \frac{U_{\text{MA}}}{U_{\text{tot}}}, \\ p_{\text{cond}} &\equiv \frac{U_{\text{cond}}}{U_{\text{tot}}}, \end{aligned} \quad (\text{B1})$$

where the energies U_j in Eq. (B1) are defined as [12, 24]

$$\begin{aligned} U_{\text{tot}} &\equiv \frac{1}{2} \text{Re} \left\{ \int_{V_{\text{tot}}} \mathbf{E} \cdot \mathbf{D}^* dV \right\}, \\ &= \frac{1}{2} \text{Re} \left\{ \int_{V_{\text{tot}}} \mathbf{H} \cdot \mathbf{B}^* dV \right\}, \\ U_{\text{bulk}} &\equiv \frac{1}{2} \text{Re} \left\{ \int_{V_{\text{sub}}} \mathbf{E} \cdot \mathbf{D}^* dV \right\}, \\ U_{\text{SA}} &\equiv \frac{1}{2} t_{\text{SA}} \text{Re} \left\{ \int_{S_{\text{sub}}} \mathbf{E} \cdot \mathbf{D}^* dA \right\}, \\ U_{\text{MA}} &\equiv \frac{1}{2} t_{\text{MA}} \text{Re} \left\{ \frac{\epsilon_0}{\epsilon_{\text{MA}}} \int_{S_{\text{cav}}} \mathbf{E} \cdot \mathbf{D}^* dA \right\}, \\ U_{\text{cond}} &\equiv \frac{1}{2} \lambda_L \text{Re} \left\{ \int_{S_{\text{cav}}} \mathbf{H} \cdot \mathbf{B}^* dA \right\}. \end{aligned} \quad (\text{B2})$$

Here, V_{sub} and S_{sub} refer to the volume and surface of the dielectric substrate under study, S_{cav} is the surface of

the aluminum cavity, and V_{tot} is the volume of the entire system.

When solving for the eigenmodes of our system, we account for the birefringence of sapphire by assigning a tensorial relative permittivity $\overleftrightarrow{\epsilon}_{\text{bulk}}$ to the sample region. We assume a literature value for relative permittivity of 11.35 parallel to the c -axis, and 9.27 perpendicular to the c -axis [47].

For energies in material interface regions; U_{SA} , U_{MA} and U_{cond} , volume integrals have been reduced to surface integrals by assuming field uniformity over the depth of the interface and integrating over the depth to yield a corresponding length prefactor. This simplification allows these interface regions with distinct material properties to be represented by interfaces between regions defined in software, rather than as a regions of their own. This amounts to a significant savings in processing time, as accurate field solutions in high aspect ratio volumes requires fine meshing, which can be computationally intensive. In this sense, these interface regions are treated as perturbations on the system, as the assumed material properties and thickness of these interfaces do not feed back and affect the field solutions acquired by the simulation.

The participations quoted in this work assume the surface dielectric thickness parameters t_{SA} and t_{MA} to be 3 nm, but we emphasize that this is a choice made for ease of comparison with surface loss tangents on silicon [26], and is not based on any direct knowledge of the surface thicknesses in our system. The London penetration depth λ_L is set to 50 nm, which we confirm to be approximately accurate by fitting temperature sweeps of the cavity resonance frequency and quality factor to Mattis-Bardeen theory (see Table VIII in App. F).

For calculation of U_{SA} , it is crucial to perform the integral to one side of the substrate-air (SA) interface (rather than directly over it) to avoid ambiguity about the permittivity assigned by the software. The expression for U_{SA} found in eq. (B2) is able to take on its compact form because of two choices made in our analysis. For one, we chose to perform the integral over the substrate-side of the SA-interface. Additionally, we assume the permittivity of the sample surface ϵ_{SA} to be the same as that of the sample bulk ϵ_{bulk} , in which the integration was performed. The expression for U_{SA} can be generalized to relax this equal permittivity assumption. This is done by decomposing the integrand into its components normal and tangent to the SA interface and invoking continuity of \mathbf{E} and \mathbf{D} where appropriate:

$$\mathbf{E} \cdot \mathbf{D}^* \rightarrow E_{\parallel} \left(\frac{\epsilon_{\parallel}^{\text{SA}}}{\epsilon_{\parallel}^{\text{bulk}}} D_{\parallel} \right)^* + \left(\frac{\epsilon_{\perp}^{\text{bulk}}}{\epsilon_{\perp}^{\text{SA}}} E_{\perp} \right) D_{\perp}^*. \quad (\text{B3})$$

After this generalizing substitution, U_{SA} is still expressed in terms of the fields on the substrate-side of the SA interface. However, similar attention to boundary conditions

permits U_{SA} to be alternately expressed in terms of the fields in the vacuum immediately outside the sample, in case it is preferable to evaluate the surface integral over that side of the SA interface instead.

Unlike the SA interface, the MA interface exists at the outer boundary of the simulation. This means when calculating U_{MA} , there is no ambiguity in the value of permittivity assigned by the software, nor is there any choice about which side of the interface to perform the integration; the integral will be done in the vacuum region. Again invoking continuity of \mathbf{E} and \mathbf{D} , fields in the MA interface region, along with the resulting U_{MA} , can be expressed in terms of the fields in the vacuum, where the field solutions are available. Unlike for the SA interface, however, this transformation is simplified by the conductor at the MA interface boundary, which requires all electric fields to be normal to the surface. With no tangential term in the integrand, the permittivity factor multiplying the normal term can be factored out of the integral. The result is that U_{MA} uniformly scales inversely with ϵ_{MA} , which we assume to be equal to 10.

Appendix C: Measurement sensitivity in a polynomial basis

To gain intuition about what sets the sensitivities plotted in Fig. 2, we can use the mapping between z and $p_{\text{bulk}}(z)$ (Fig. 1b) to express $p_{\text{cond}}(z)$, $p_{MA}(z)$, and $Q^{-1}(z)$ as polynomials in p_{bulk} . In practice, 2nd-order polynomials closely approximate the system, as shown in Fig. 8:

$$\begin{aligned} Q^{-1} &= y_0 + y_1 p_{\text{bulk}} + y_2 p_{\text{bulk}}^2, \\ p_{\text{cond}} &= x_0^{\text{cond}}, \\ p_{\text{bulk}} &= p_{\text{bulk}}, \\ p_{MA} &= x_0^{\text{MA}} + x_1^{\text{MA}} p_{\text{bulk}} + x_2^{\text{MA}} p_{\text{bulk}}^2. \end{aligned} \quad (\text{C1})$$

Recasting the problem in this fashion is convenient because conductor and substrate losses are orthogonal in the p_{bulk} -polynomial basis; the resulting matrix equation,

$$\begin{pmatrix} y_0 \\ y_1 \\ y_2 \end{pmatrix} = \begin{pmatrix} x_0^{\text{cond}} & 0 & x_0^{\text{MA}} \\ 0 & 1 & x_1^{\text{MA}} \\ 0 & 0 & x_2^{\text{MA}} \end{pmatrix} \begin{pmatrix} q_{\text{cond}}^{-1} \\ q_{\text{sub}}^{-1} \\ q_{MA}^{-1} \end{pmatrix}, \quad (\text{C2})$$

is nearly diagonal.

Conveniently, it yields a single equation for q_{sub}^{-1} :

$$q_{\text{sub}}^{-1} = y_1 - x_1^{\text{MA}} q_{MA}^{-1} \quad (\text{C3})$$

Error propagation on Eq. C3 gives the desired expression for the sensitivity:

$$\sigma_{q_{\text{sub}}^{-1}} = \sigma_{y_1} + |x_1^{\text{MA}}| \sigma_{q_{MA}^{-1}} + \sigma_{x_1^{\text{MA}}} q_{MA}^{-1}. \quad (\text{C4})$$

The terms in Eq. C4 reveal what sets the measurement sensitivity, and how it can be improved. To estimate the




FIG. 8. **Polynomial representation.** Measured inverse internal quality factors (a), simulated MA participation (b), and simulated conductor participation (c) of the HEMEX sample, as a function of bulk participation. All lines are 2nd-order polynomial fits per Eq. (C1).

relative importance of the terms in Eq. C4, we calculate their magnitudes assuming a hypothetical lossless substrate with the same permittivity as sapphire [47]. Neglecting errors in alignment and simulation, the first term $\sigma_{y_1} \approx \sqrt{\frac{1}{N} \frac{\sigma_{Q^{-1}}}{\Delta p_{\text{bulk}}}} \approx 1$ ppb. The next two terms can be approximated using the coefficient x_1^{MA} listed in Table III. The two summands are then each approxi-

k	$10^9 y_k$	$10^6 x_k^{\text{cond}}$	x_k^{bulk}	$10^9 x_k^{\text{MA}}$
0	9.18(0.05)	43.92(0.01)	0	249(1)
1	26(3)	0	1	-300(60)
2	240(40)	0	0	15800(900)

TABLE III. Polynomial coefficients \mathbf{y} and \mathbf{x}^j extracted from the fits to Q^{-1} and p_j in Fig. 8.

mately 2×10^{-9} . If this analysis is repeated with polynomial order of 3 or 4, the x_1^{MA} term grows to approximately 5 ppb. Taken together, these considerations suggest that the sensitivity floor estimated from this polynomial analysis is several ppb, and is limited by our ability to separate MA losses from those caused by the substrate. This baseline is consistent with the sensitivity shown in Fig. 2b.

Comparing this sensitivity of several ppb with the assumed measurement error of approximately 0.1 ppb indicates that input errors are amplified by a factor of roughly 100 during the regression. This factor is much less than the condition number of the participation matrix, which is roughly 10^7 for the samples considered in this work. Such a discrepancy indicates the value of using the full covariance matrix.

Note that the above argument neglects the 1% fractional change in p_{cond} and the 3% fractional change in the cavity frequency that occur when the sample is inserted, which make conductive losses slightly dependent on insertion. Accounting for these effects, such as by replacing p_{cond} with \tilde{p}_{cond} , results in a small but non-vanishing x_1^{cond} . We have omitted these effects as they change the cavity loss Q^{-1} by less than 10^{-10} .

Appendix D: Sample preparation

All samples were sourced from double-side polished c -plane wafers (i.e., the extraordinary axis normal to the plane of the wafer). Samples were cut from the wafers using a LatticeGear FlexScribe to have dimensions of 6 mm x 45 mm. Samples were then cleaned with isopropanol and blown dry with nitrogen gas. Sample dimensions were verified by a wafer micrometer and Zeta-20 optical profilometer. EFG samples were sourced from Kyocera, and the HEMEX sample was sourced from GT Advanced Technologies (now called Crystal Systems).

Appendix E: Power sweeps

1. Defining the low-power regime

We define the low-power regime as the range of cavity photon number n which is sufficiently small, such that total loss in the system Q^{-1} is nearly desaturated, and approximately the same as it would be at single-photon powers Q_{lp}^{-1} . This deviation of Q^{-1} from single-photon behavior can be compared to the total saturable loss in the system to determine the saturation fraction F of loss in the system:

$$F \equiv \frac{Q_{\text{lp}}^{-1} - Q^{-1}}{Q_{\text{sat}}^{-1}} \quad (\text{E1})$$

This value is identically zero when system behavior matches Q_{lp}^{-1} . Using Eq. (7), the saturation fraction can be expressed in terms of n , n_c and α .

$$F = 1 - \frac{1}{\sqrt{1 + \left(\frac{n}{n_c}\right)^\alpha}} \quad (\text{E2})$$

Solving for n can tell you a maximum allowable photon number, given a maximum allowable saturation fraction and the parameters n_c and α from measurement of the system.

$$\frac{n}{n_c} < \left(\frac{1}{(1-F)^2} - 1 \right)^{1/\alpha} \quad (\text{E3})$$

We define the boundary of the low power regime as the n for which $F < 10^{-1}$. With $\alpha \approx 1/2$, this requires $n/n_c < 5 \times 10^{-2}$. (Smaller values of α put less stringent bounds on n). To be conservative we choose $n = 10^5$ such that $n/n_c < 10^{-3}$. Low-power measurements are made with a Josephson travelling wave parametric amplifier [48].

2. Power sweep bulk participations

Table IV shows the bulk participations for the inserted and withdrawn power sweeps on all samples.

Material	withdrawn p_{bulk}	inserted p_{bulk}
100 μm EFG	3.6×10^{-5}	1.7×10^{-2}
460 μm EFG	1.2×10^{-4}	5.66×10^{-2}
440 μm HEMEX	4.9×10^{-4}	7.1×10^{-2}

TABLE IV. Bulk participations for inserted and withdrawn power sweeps

For all of the EFG samples, p_{bulk} was chosen to match that of the endpoints of the position sweeps in Fig. 4. This was done for ease of cross-comparison of power sweeps and position sweeps (as both have measurements with equal p_{bulk} and $n = 10^5$).

For HEMEX (measured by an earlier version of our protocol), we measured the inserted power sweeps with the sample in contact with the post, rather than some small distance ($z \approx 30 \mu\text{m}$) away from the post. This was done with the intent of maximizing the sensitivity to the sample loss, but had the drawback of removing direct cross-comparability between the power sweep and position sweep of HEMEX. Based on the trend in Q^{-1} measured as a function of p_{bulk} (Fig. 4), the large participation of the inserted power sweep is insufficient to explain the large Q^{-1} measured at the lowest excitation powers (Fig. 3). We posit this discrepancy is caused by measuring the inserted HEMEX power sweep with the sample in contact with the top of the post, which could make the cavity mode more sensitive to surface losses associated with surface roughness or material non-homogeneity, neither of which are captured by our simulations.

The different condition under which HEMEX was measured does not invalidate the low-power verification done by the inserted power sweep, as further insertion of the sample tends to reduce the critical photon number n_c . This can be seen by comparing Table I to Table V. The

result of this effect is to give a stricter criterion for the boundary of the low-power regime, making the verification process conservative.

Sample	n_c (billions)	α
100 μm EFG	7(3)	0.5(0.1)
460 μm EFG	8(1)	0.4(0.16)
440 μm HEMEX	14(4)	0.32(0.03)

TABLE V. Fit results for the withdrawn (orange) power sweeps in Fig. (3) using the TLS model in Eq. (7).

3. Constraining the position sweep fit

When the sample is withdrawn, the total internal loss of the system Q^{-1} consists predominantly of two loss channels: conductivity loss Q_{cond}^{-1} and metal-air interface dielectric loss Q_{MA}^{-1} . Power sweeps of the cavity provide some information about these individual underlying loss mechanisms. This information can be translated into bounds on the low-power material loss factors q_j^{-1} associated with the cavity, and those bounds can be applied to the analysis of the position sweep described in II.

As with any resonator measurement, Q^{-1} can act as an upper bound on the loss of each underlying loss channel. For instance, the low-power value of Q_{MA}^{-1} can be bounded from above by the low-power total Q^{-1} , and that upper-bound can be converted into a bound on the low-power value of q_{MA}^{-1} . Additional information about Q_{MA}^{-1} comes from assuming that conductivity loss Q_{cond}^{-1} is power-independent up to $n = 10^{14}$ (an assumption consistent with observations up to $n = 10^{10}$ in prior studies of aluminum cylindrical cavities [20]). If Q_{cond}^{-1} is power-independent, then any power-dependence observed in Q^{-1} must be attributed to Q_{MA}^{-1} , and this attribution can be converted into a lower-bound on q_{MA}^{-1} (see eq. 8).

An upper-bound on Q_{cond}^{-1} can be inferred in the same way, but assuming that Q_{cond}^{-1} is power-independent implies a slightly stricter bound. If the low-power and high-power values of Q_{cond}^{-1} are assumed to be the same, then the high-power Q^{-1} (or any power Q^{-1}) can serve as an upper-bound on the low-power Q_{cond}^{-1} , and converted into an upper-bound on q_{cond}^{-1} (see eq. 8).

The precise bounds obtained for each sample are listed in Table VI. These bounds are obtained separately for each cooldown, to control for potential temporal variations in the bare quality factor of the cavity. Over the course of the measurements presented in this study, however, no such variations were observed. Consequently, the bounds in Table VI are relatively similar for all three measurements.

When performing the insertion analysis in Fig. 4, the three measurements are fit jointly, which has the advantage that it removes four free parameters from the fit,

Sample material	Cavity loss factor bounds
100 μm EFG	$8.62 \times 10^{-3} \leq q_{\text{MA}}^{-1} \leq 38.2 \times 10^{-3}$ $0 \leq q_{\text{cond}}^{-1} \leq 1.66 \times 10^{-4}$
460 μm EFG	$9.18 \times 10^{-3} \leq q_{\text{MA}}^{-1} \leq 37.6 \times 10^{-3}$ $0 \leq q_{\text{cond}}^{-1} \leq 1.60 \times 10^{-4}$
440 μm HEMEX	$9.33 \times 10^{-3} \leq q_{\text{MA}}^{-1} \leq 38.3 \times 10^{-3}$ $0 \leq q_{\text{cond}}^{-1} \leq 1.62 \times 10^{-4}$

TABLE VI. Constraints on cavity loss factors

reducing the risk of over-fitting. The “background” loss rates extracted in this way are shown in Table VII, and are consistent with other measurements of etched 4N aluminum cavities [20, 49]. (The quoted value for conductor loss is given at the withdrawn frequency of 4.55 GHz.) We could repeat the insertion analysis with the cavity loss factors constrained using the bounds from Table VI, but this has no effect on the fit, as the unconstrained fit already yields loss rates that fall within the bounds in Table VI. This demonstrates consistency between the insertion analysis from Sec. IIIA and the power analysis from Sec. IIIB.

Loss factor	Extracted value
q_{MA}^{-1}	$33(9) \times 10^{-3}$
q_{cond}^{-1}	$2(5) \times 10^{-5}$
R_s	$4(9) \times 10^{-8} \Omega$
q_{sub}^{-1} (thin EFG)	$170(13) \times 10^{-9}$
q_{sub}^{-1} (thick EFG)	$88(6) \times 10^{-9}$
q_{sub}^{-1} (HEMEX)	$19(6) \times 10^{-9}$
q_{bulk}^{-1} (EFG)	$63(8) \times 10^{-9}$
q_{SA}^{-1} (EFG)	$15(3) \times 10^{-4}$

TABLE VII. Summary of fit results.

Appendix F: Temperature dependence

Conductor losses in the aluminum walls of the cavity are temperature dependent. Near the vicinity of the superconducting transition temperature T_c , they are much greater than the metal-air (MA) surface dielectric losses of the cavity mode. We utilize this fact to experimentally confirm that insertion of the dielectric does not change the conductor participation of the cavity mode.

Fig. 9 shows measurements of cavity loss rate Q^{-1} as a function of the temperature T measured at the mixing chamber plate of the dilution refrigerator. Measurements of Q^{-1} are performed with the sample fully withdrawn (orange traces) and fully inserted (green traces).

Two distinct regimes are visible in the measurements. For temperatures at or near T_c , conductor losses dominate, and system behavior can be described by Mattis-Bardeen theory [50], which we fit to the temperature de-

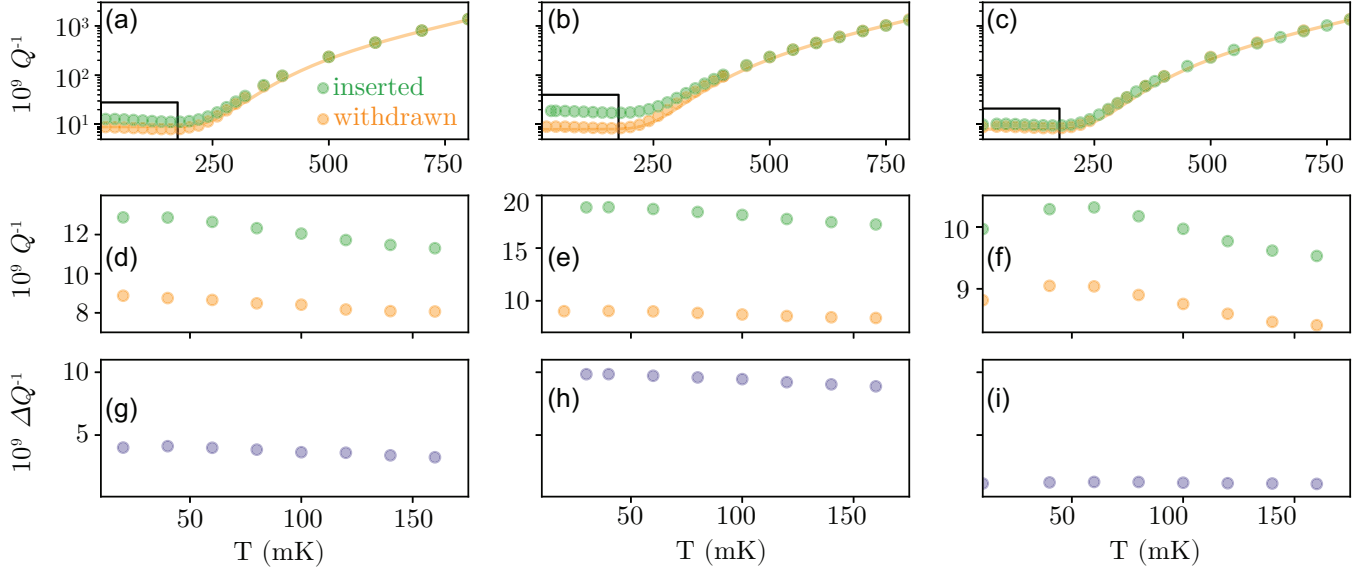


FIG. 9. **Temperature sweeps.** Measured inverse internal quality factor Q^{-1} with the sample withdrawn (orange) and inserted (green) as a function of the mixing-chamber plate temperature, for 100 μm thick EFG sapphire (a), 460 μm thick EFG sapphire (b), and 440 μm thick HEMEX sapphire (c). Orange lines are fits to Mattis-Bardeen theory [50], as described in the supplement of Ref. [36].

pendence of the withdrawn temperature sweeps (orange traces). The results are shown in Table VIII.

Material	Thickness	T_c (K)	λ_L (nm)
EFG sapphire	100 μm	1.13	135
EFG sapphire	460 μm	1.19	90
HEMEX sapphire	440 μm	1.13	90

TABLE VIII. Results from fitting the inverse quality factors measured during withdrawn (orange) temperature sweeps in Fig. (9) using Mattis-Bardeen theory [50].

Near T_c , conductivity loss grows to be orders of magnitude larger than the other loss channels, making total loss approximately proportional to participation in the conductor: $Q^{-1} \approx p_{\text{cond}} q_{\text{cond}}^{-1}$. A pair of high-temperature measurements with the sample withdrawn and inserted may therefore serve as a probe of the relative change in p_{cond} at the two measured positions.

$$\frac{1/Q^I - 1/Q^W}{1/Q^W} \approx \frac{p_{\text{cond}}^I - p_{\text{cond}}^W}{p_{\text{cond}}^W}. \quad (\text{F1})$$

(Here the superscripts indicate whether a measurement is made with the sample withdrawn or inserted). With this argument, high-temperature measurements with the sample withdrawn and inserted bound the relative change in conductor participation to less than 1%, in agreement with the numerical simulations shown in App. C.

For temperatures $T \ll T_c$, conductor losses are only weakly dependent upon temperature, so we attribute the

observed temperature dependence to the system's dielectrics. Fig. 9d-f shows this regime in greater detail. Weak temperature dependence is visible in both the inserted and withdrawn sweeps, indicating some saturability of the dielectric losses. However, subtracting the withdrawn and inserted measurements, as shown in Fig. 9g-i, reveals that inserting the sample does little to change this temperature dependence, indicating that the temperature dependence arises primarily from the cavity's metal-air interface, rather than from the sample under study. In the future, modifications to the experimental design could improve the contrast of such temperature-dependence comparisons by better thermalizing the dielectric substrate to the mixing chamber, or by constructing cavities from superconductors with higher transition temperatures, such as niobium.

Appendix G: Repeatability and cross-wafer variation

As a control test to check for measurement reproducibility and wafer-to-wafer variation in nominally identical samples, we repeated the position sweep of 460 μm thick EFG sapphire on a second control sample cut from a different wafer. The results are shown in Fig. 10. The two measurements are quite similar, suggesting that the dipper is a measurement technique with high reproducibility and that the dielectric properties of these two samples are the same, up to the resolution of the measurement.




FIG. 10. **Position sweep on nominally identical samples.** The position sweep in Fig. 4 was repeated with a second piece of 460 μm thick EFG sapphire denoted as a control. The control was prepared in nominally the same way, but cut from a different wafer of the same material.

Appendix H: Bounding magnetic bulk loss

Implicit in our analysis is the assumption that the origin of loss from the sample is purely dielectric and none of the loss is magnetic. In this section we relax that assumption, and modify our model to consider the possibility that some of the loss from the substrate could be magnetic in origin. We then compare our measurement results of EFG sapphire with quality factors of stripline resonators fabricated on the same material. From this comparison, we can make bounding statements about the composition of bulk loss in our sample, and about the prevalence of bulk magnetic loss in superconducting qubits on EFG sapphire.

Just as substrate loss can be decomposed into bulk and surface contributions (See Eq. 4 and Fig. 4b), we can further decompose loss from the bulk of our sample Q_{bulk}^{-1} into contributions from bulk dielectric loss $Q_{\text{bulk(E)}}^{-1}$ and bulk magnetic loss $Q_{\text{bulk(H)}}^{-1}$. These losses can be treated with the same participation formalism and attributed to an associated dielectric loss tangent $q_{\text{bulk(E)}}^{-1}$ and a magnetic loss tangent $q_{\text{bulk(H)}}^{-1}$, respectively:

$$Q_{\text{bulk}}^{-1} = p_{\text{bulk(E)}} q_{\text{bulk(E)}}^{-1} + p_{\text{bulk(H)}} q_{\text{bulk(H)}}^{-1}. \quad (\text{H1})$$

Because of the quasi-linear relationship between bulk dielectric participation $p_{\text{bulk(E)}}$ and bulk magnetic participation $p_{\text{bulk(H)}}$ in our geometry (see Fig. 11a), our measurement has limited ability to distinguish between these two loss channels, but it does constrain them to a contour in the space of all ordered pairs $(q_{\text{bulk(E)}}^{-1}, q_{\text{bulk(H)}}^{-1})$. In close analogy to the method of bulk-surface loss separation performed in Fig. 4b, an additional measurement with a different ratio of participations $p_{\text{bulk(E)}}/p_{\text{bulk(H)}}$ would suggest another constraint line with a different slope, and further constrain the ac-

cessible $(q_{\text{bulk(E)}}^{-1}, q_{\text{bulk(H)}}^{-1})$ pairs. Because our measurements probe loss in our sample with the fringe field of the cavity's charge antinode, our measurement is preferentially sensitive to dielectric loss, and achieves $p_{\text{bulk(E)}}/p_{\text{bulk(H)}} \approx 200$. To acquire another constraint line, we can consider the performance of a stripline fabricated on EFG sapphire from another study [42]. A stripline of this design can achieve $Q \approx 8 \times 10^6$ and has $p_{\text{bulk(E)}}/p_{\text{bulk(H)}} = 0.40/0.31 \approx 1.3$; producing a contour with a much more gradual slope which we pair with our steeper measurement contour (See Fig. 11b). Because the Q of the stripline is the result of a multitude of possible loss channels, the associated contour represents an inequality:

$$Q^{-1} \geq p_{\text{bulk(E)}} q_{\text{bulk(E)}}^{-1} + p_{\text{bulk(H)}} q_{\text{bulk(H)}}^{-1}. \quad (\text{H2})$$

This means that the intersection of the two contours cannot be used to extract the true value of $(q_{\text{bulk(E)}}^{-1}, q_{\text{bulk(H)}}^{-1})$ associated with our sample, however upper and lower bounds can be inferred on both these loss tangents using the region where our measurement contour overlaps the region bounded from above by the stripline contour.




FIG. 11. **Bounding magnetic bulk loss.** (a) Bulk magnetic participation $p_{\text{bulk(H)}}$ as a function of bulk electric participation $p_{\text{bulk(E)}}$. (b) Contours representing viable decompositions of bulk substrate loss into dielectric and magnetic components in our experiment (Eq. H1) and in a stripline resonator [42] (Eq. H2). After our model is generalized to allow losses from the substrate bulk to be a combination of dielectric and magnetic loss (purple contour), the relative contributions of these losses can be constrained by the performance of a stripline fabricated on the same substrate (orange bounding contour). To be consistent with stripline Q , it must be that at least 97% of the sample loss observed in our measurement must be dielectric in origin. These constraints also suggest that magnetic loss is at most a 10% contributor to the decoherence of transmons on EFG sapphire.

The contours bound the bulk dielectric loss tangent to $6.1 \times 10^{-8} \leq q_{\text{bulk(E)}}^{-1} \leq 6.3 \times 10^{-8}$, meaning that even considering the possibility that there may be magnetic losses associated with the sample bulk, this can account for no more than 3% of the sample loss detected in our measurement, and at least 97% of bulk loss we detect must be dielectric in origin. This validates the analysis in

the main text, which neglects magnetic losses. Though prior studies of HEMEX sapphire dielectric resonators have observed tunable loss due to coupling with electron spins of magnetic impurities brought into resonance by Tesla-scale magnetic fields [51–53], our result shows that dielectric loss is the dominant source of dissipation in EFG sapphire at low magnetic fields.

The contours also bound the bulk magnetic loss tangent to $q_{\text{bulk(H)}}^{-1} \leq 3.3 \times 10^{-7}$. This bound is much looser than the bounds for $q_{\text{bulk(E)}}^{-1}$, largely due to our measurement’s relative insensitivity to magnetic losses. Fortunately, because of its high kinetic inductance, a typical transmon qubit has $p_{\text{bulk(H)}} \approx 2.5\%$ (relatively small compared to the 31% of a typical stripline), so this bound on $q_{\text{bulk(H)}}^{-1}$ can still be used to infer a useful transmon coherence limit of $Q_{\text{bulk(H)}} > 1.2 \times 10^8$. This suggests that even in today’s longest-lived transmons on sapphire, bulk magnetic loss accounts for less than 10% of total loss, making bulk magnetic loss at most a minor contributor to their decoherence.

Appendix I: Comparison to other studies

Fig. 12 compares the loss tangents measured in this work to those reported in other studies on the microwave loss of sapphire at low temperatures and low excitation powers.




FIG. 12. Comparison of low-power and low-temperature loss-tangent measurements of sapphire.

Appendix J: Microwave measurement circuit

Fig. 13 shows the microwave wiring diagram of the measurement system.

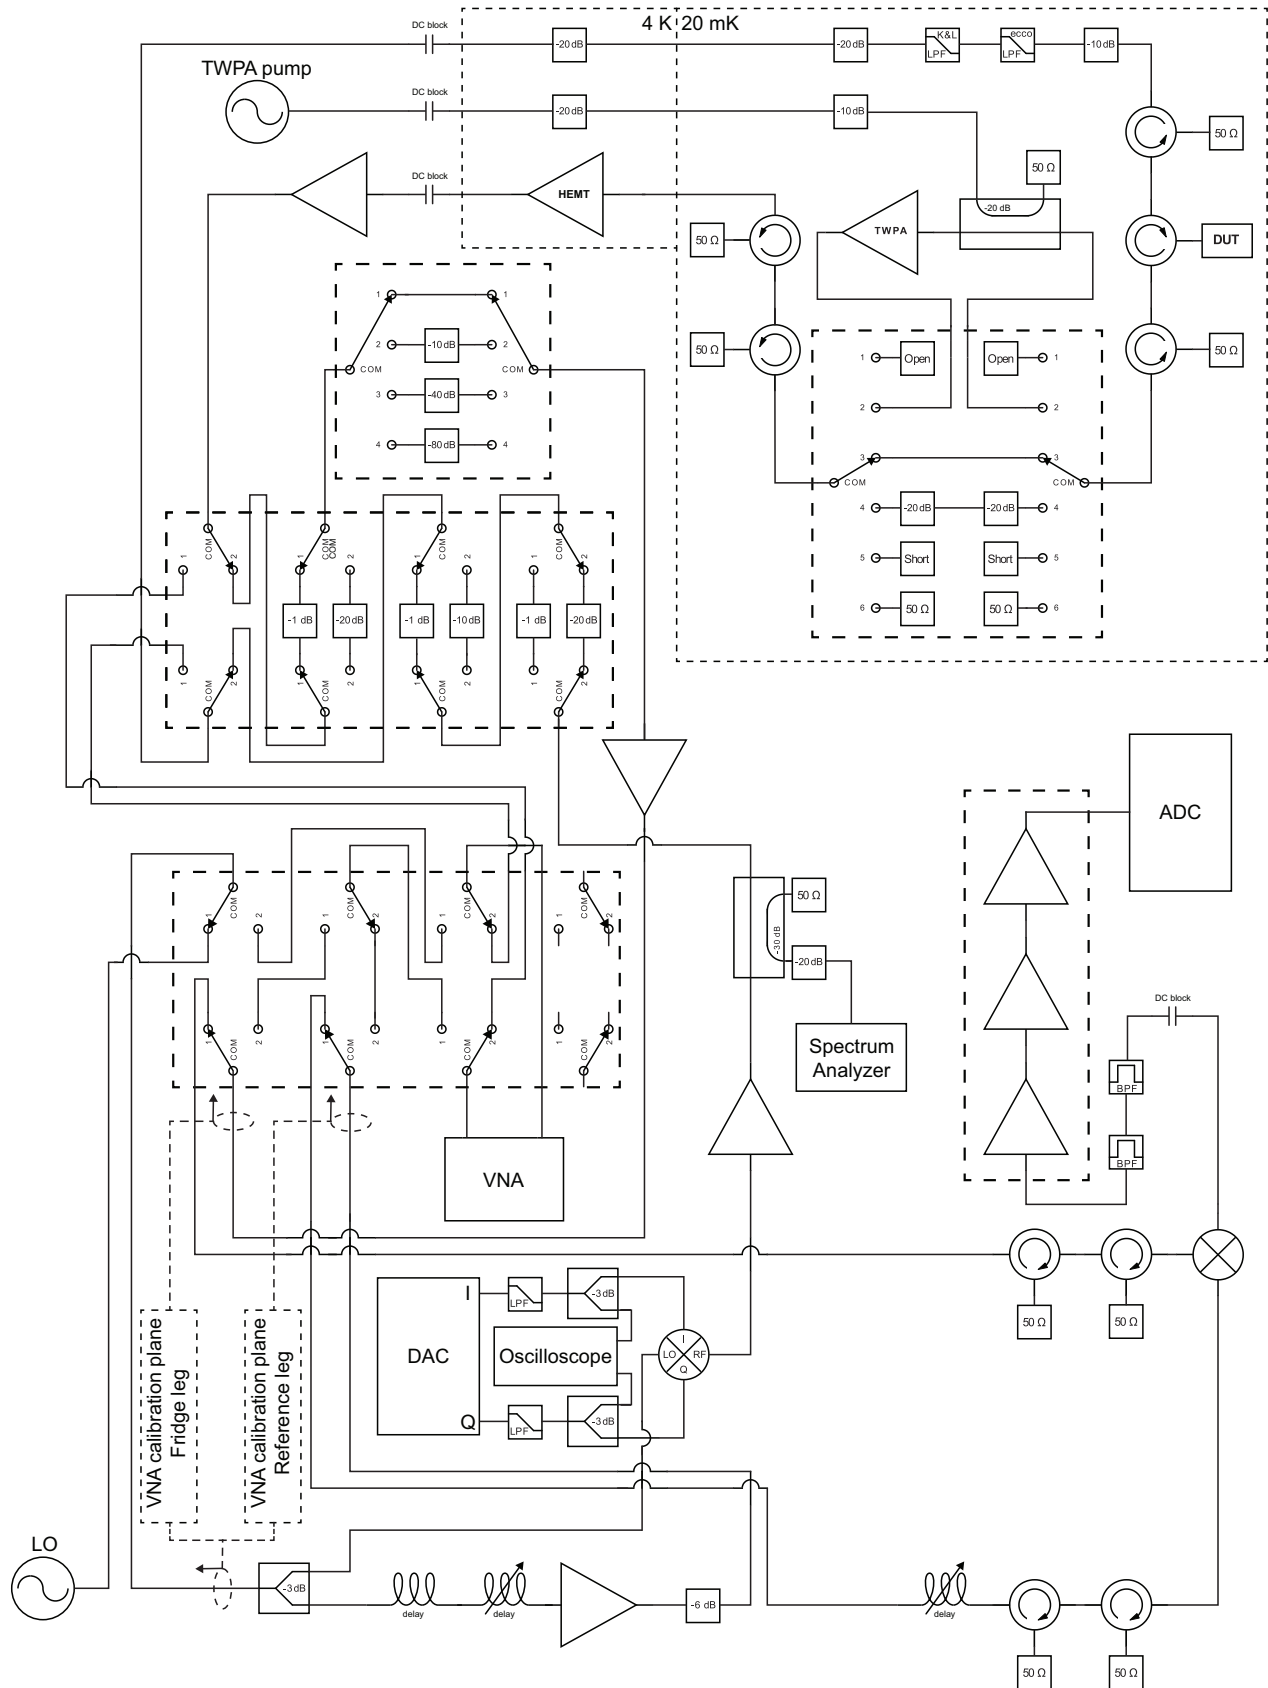


FIG. 13. **Microwave measurement circuit.** Microwave cavity (DUT) is cooled to 20 mK and measured with heterodyne detection, employing switched amplifier and attenuator banks for improved dynamic range.

-
- [1] Y. Nakamura, Y. Pashkin, and J. Tsai, Coherent control of macroscopic quantum states in a single-cooper-pair box, *nature* **398**, 786 (1999).
- [2] A. Place, L. Rodgers, P. Mundada, B. Smitham, M. Fitzpatrick, Z. Leng, A. Premkumar, J. Bryon, A. Vrajitoarea, S. Sussman, *et al.*, New material platform for superconducting transmon qubits with coherence times exceeding 0.3 milliseconds, *Nature communications* **12**, 1 (2021).
- [3] C. Wang, X. Li, H. Xu, Z. Li, J. Wang, Z. Yang, Z. Mi, X. Liang, T. Su, C. Yang, *et al.*, Transmon qubit with relaxation time exceeding 0.5 milliseconds, arXiv preprint arXiv:2105.09890 (2021).
- [4] R. Gordon, C. Murray, C. Kurter, M. Sandberg, S. Hall, K. Balakrishnan, R. Shelby, B. Wacaser, A. Stabile, J. Sleight, *et al.*, Environmental radiation impact on lifetimes and quasiparticle tunneling rates of fixed-frequency transmon qubits, arXiv preprint arXiv:2105.14003 (2021).
- [5] A. Somoroff, Q. Ficheux, R. Mencia, H. Xiong, R. Kuzmin, and V. Manucharyan, Millisecond coherence in a superconducting qubit, arXiv preprint arXiv:2103.08578 (2021).
- [6] V. Manucharyan, J. Koch, L. Glazman, and M. Devoret, Fluxonium: Single cooper-pair circuit free of charge offsets, *Science* **326**, 113 (2009).
- [7] J. Koch, T. Yu, J. Gambetta, A. Houck, D. Schuster, J. Majer, A. Blais, M. Devoret, S. Girvin, and R. Schoelkopf, Charge-insensitive qubit design derived from the cooper pair box, *Physical Review A* **76**, 042319 (2007).
- [8] H. Paik, D. Schuster, L. Bishop, G. Kirchmair, G. Catelani, A. Sears, B. Johnson, M. Reagor, L. Frunzio, L. Glazman, *et al.*, Observation of high coherence in josephson junction qubits measured in a three-dimensional circuit qed architecture, *Physical Review Letters* **107**, 240501 (2011).
- [9] C. McRae, H. Wang, J. Gao, M. Vissers, T. Brecht, A. Dunsworth, D. Pappas, and J. Mutus, Materials loss measurements using superconducting microwave resonators, *Review of Scientific Instruments* **91**, 091101 (2020).
- [10] J. Gao, *The physics of superconducting microwave resonators* (California Institute of Technology, 2008).
- [11] H. Wang, M. Hofheinz, J. Wenner, M. Ansmann, R. Bialczak, M. Lenander, E. Lucero, M. Neeley, A. O'Connell, D. Sank, *et al.*, Improving the coherence time of superconducting coplanar resonators, *Applied Physics Letters* **95**, 233508 (2009).
- [12] C. Wang, C. Axline, Y. Gao, T. Brecht, Y. Chu, L. Frunzio, M. Devoret, and R. Schoelkopf, Surface participation and dielectric loss in superconducting qubits, *Applied Physics Letters* **107**, 162601 (2015).
- [13] O. Dial, D. McClure, S. Poletto, G. Keefe, M. Rothwell, J. Gambetta, D. Abraham, J. Chow, and M. Steffen, Bulk and surface loss in superconducting transmon qubits, *Superconductor Science and Technology* **29**, 044001 (2016).
- [14] Y. Chu, C. Axline, C. Wang, T. Brecht, Y. Gao, L. Frunzio, and R. Schoelkopf, Suspending superconducting qubits by silicon micromachining, *Applied Physics Letters* **109**, 112601 (2016).
- [15] J. Gambetta, C. Murray, Y. Fung, D. McClure, O. Dial, W. Shanks, J. Sleight, and M. Steffen, Investigating surface loss effects in superconducting transmon qubits, *IEEE Transactions on Applied Superconductivity* **27**, 1 (2016).
- [16] G. Calusine, A. Melville, W. Woods, R. Das, C. Stull, V. Bolkhovskiy, D. Braje, D. Hover, D. Kim, X. Miloski, *et al.*, Analysis and mitigation of interface losses in trenched superconducting coplanar waveguide resonators, *Applied Physics Letters* **112**, 062601 (2018).
- [17] W. Woods, G. Calusine, A. Melville, A. Sevi, E. Golden, D. Kim, D. Rosenberg, J. Yoder, and W. Oliver, Determining interface dielectric losses in superconducting coplanar-waveguide resonators, *Physical Review Applied* **12**, 014012 (2019).
- [18] C. McRae, R. Lake, J. Long, M. Bal, X. Wu, B. Jugdersuren, T. Metcalf, X. Liu, and D. Pappas, Dielectric loss extraction for superconducting microwave resonators, *Applied Physics Letters* **116**, 194003 (2020).
- [19] We assume that surface dielectric interfaces have a thickness of 3 nm (see App. B).
- [20] M. Reagor, H. Paik, G. Catelani, L. Sun, C. Axline, E. Holland, I. Pop, N. Masluk, T. Brecht, L. Frunzio, *et al.*, Reaching 10 ms single photon lifetimes for superconducting aluminum cavities, *Applied Physics Letters* **102**, 192604 (2013).
- [21] A. Romanenko, R. Pilipenko, S. Zorzetti, D. Frolov, M. Awida, S. Belomestnykh, S. Posen, and A. Grassellino, Three-dimensional superconducting resonators at $t < 20$ mk with photon lifetimes up to $\tau = 2$ s, *Physical Review Applied* **13**, 034032 (2020).
- [22] M. Checchin, D. Frolov, A. Lunin, A. Grassellino, and A. Romanenko, Measurement of the low-temperature loss tangent of high-resistivity silicon wafers with high-q superconducting resonators, arXiv preprint arXiv:2108.08894 (2021).
- [23] J. Turneaure and I. Weissman, Microwave surface resistance of superconducting niobium, *Journal of Applied Physics* **39**, 4417 (1968).
- [24] M. Reagor, *Superconducting cavities for circuit quantum electrodynamics* (Yale University, 2016).
- [25] M. Reagor, W. Pfaff, C. Axline, R. Heeres, N. Ofek, K. Sliwa, E. Holland, C. Wang, J. Blumoff, K. Chou, *et al.*, Quantum memory with millisecond coherence in circuit qed, *Phys. Rev. B* **94**, 014506 (2016).
- [26] J. Wenner, R. Barends, R. Bialczak, Y. Chen, J. Kelly, E. Lucero, M. Mariantoni, A. Megrant, P. O'Malley, D. Sank, *et al.*, Surface loss simulations of superconducting coplanar waveguide resonators, *Applied Physics Letters* **99**, 113513 (2011).
- [27] J. Zmuidzinas, Superconducting microresonators: Physics and applications, *Annu. Rev. Condens. Matter Phys.* **3**, 169 (2012).
- [28] L. Trefethen and D. Bau III, *Numerical linear algebra*, Vol. 50 (Siam, 1997).
- [29] P. Richter, Estimating errors in least-squares fitting, *Telecommun. Data Acquisition Prog. Rep* **42**, 107 (1995).
- [30] D. Harris, A peek into the history of sapphire crystal growth, in *Window and Dome Technologies VIII*, Vol. 5078 (International Society for Optics and Photonics, 2003) pp. 1–11.

- [31] D. Creedon, Y. Reshitnyk, W. Farr, J. Martinis, T. Duty, and M. Tobar, High q-factor sapphire whispering gallery mode microwave resonator at single photon energies and millikelvin temperatures, *Applied Physics Letters* **98**, 222903 (2011).
- [32] C. Khattak, F. Schmid, and M. Smith, Correlation of sapphire quality with uniformity and optical properties, in *Window and Dome Technologies and Materials V*, Vol. 3060 (SPIE, 1997) pp. 250–257.
- [33] C. Müller, J. Cole, and J. Lisenfeld, Towards understanding two-level-systems in amorphous solids: insights from quantum circuits, *Reports on Progress in Physics* **82**, 124501 (2019).
- [34] C. Quintana, A. Megrant, Z. Chen, A. Dunsworth, B. Chiaro, R. Barends, B. Campbell, Y. Chen, I. Hoi, E. Jeffrey, *et al.*, Characterization and reduction of microfabrication-induced decoherence in superconducting quantum circuits, *Applied Physics Letters* **105**, 062601 (2014).
- [35] A. Dunsworth, A. Megrant, C. Quintana, Z. Chen, R. Barends, B. Burkett, B. Foxen, Y. Chen, B. Chiaro, A. Fowler, *et al.*, Characterization and reduction of capacitive loss induced by sub-micron josephson junction fabrication in superconducting qubits, *Applied Physics Letters* **111**, 022601 (2017).
- [36] C. Lei, L. Krayzman, S. Ganjam, L. Frunzio, and R. Schoelkopf, High coherence superconducting microwave cavities with indium bump bonding, *Applied Physics Letters* **116**, 154002 (2020).
- [37] K. Serniak, S. Diamond, M. Hays, V. Fatemi, S. Shankar, L. Frunzio, R. Schoelkopf, and M. Devoret, Direct dispersive monitoring of charge parity in offset-charge-sensitive transmons, *Physical Review Applied* **12**, 014052 (2019).
- [38] D. Strayer, G. Dick, and E. Tward, Superconductor-sapphire cavity for an all-cryogenic scsq, *IEEE Transactions on Magnetics* **19**, 512 (1983).
- [39] V. Braginsky, V. Ilchenko, and K. Bagdassarov, Experimental observation of fundamental microwave absorption in high-quality dielectric crystals, *Physics Letters A* **120**, 300 (1987).
- [40] A. Luiten, A. Mann, and D. Blair, Ultrahigh q-factor cryogenic sapphire resonator, *Electronics letters* **10**, 879 (1993).
- [41] J. Krupka, K. Derzakowski, M. Tobar, J. Hartnett, and R. Geyer, Complex permittivity of some ultralow loss dielectric crystals at cryogenic temperatures, *Measurement Science and Technology* **10**, 387 (1999).
- [42] C. Axline, M. Reagor, R. Heeres, P. Reinhold, C. Wang, K. Shain, W. Pfaff, Y. Chu, L. Frunzio, and R. Schoelkopf, An architecture for integrating planar and 3d cqed devices, *Applied Physics Letters* **109**, 042601 (2016).
- [43] M. Collett and C. Gardiner, Squeezing of intracavity and traveling-wave light fields produced in parametric amplification, *Phys. Rev. A* **30**, 1386 (1984).
- [44] J. Martinis, S. Nam, J. Aumentado, K. Lang, and C. Urbina, Decoherence of a superconducting qubit due to bias noise, *Phys. Rev. B* **67**, 094510 (2003).
- [45] P. Petersan and S. Anlage, Measurement of resonant frequency and quality factor of microwave resonators: Comparison of methods, *Journal of applied physics* **84**, 3392 (1998).
- [46] R. Meservey and P. Tedrow, Measurements of the kinetic inductance of superconducting linear structures, *Journal of Applied Physics* **40**, 2028 (1969).
- [47] J. Krupka, K. Derzakowski, A. Abramowicz, M. Tobar, and R. Geyer, Use of whispering-gallery modes for complex permittivity determinations of ultra-low-loss dielectric materials, *IEEE Transactions on Microwave Theory and Techniques* **47**, 752 (1999).
- [48] C. Macklin, K. O’Brien, D. Hover, M. Schwartz, V. Bolkhovskiy, X. Zhang, W. Oliver, and I. Siddiqi, A near-quantum-limited josephson traveling-wave parametric amplifier, *Science* **350**, 307 (2015).
- [49] C. Axline, *Building Blocks for Modular Circuit QED Quantum Computing* (Yale University, 2018) Chap. 5.
- [50] D. Mattis and J. Bardeen, Theory of the anomalous skin effect in normal and superconducting metals, *Physical Review* **111**, 412 (1958).
- [51] K. Benmessai, W. Farr, D. Creedon, Y. Reshitnyk, J.-M. Le Floch, T. Duty, and M. Tobar, Hybrid electron spin resonance and whispering gallery mode resonance spectroscopy of fe³⁺ in sapphire, *Physical Review B* **87**, 094412 (2013).
- [52] W. Farr, D. Creedon, M. Goryachev, K. Benmessai, and M. Tobar, Ultrasensitive microwave spectroscopy of paramagnetic impurities in sapphire crystals at millikelvin temperatures, *Physical Review B* **88**, 224426 (2013).
- [53] M. Goryachev, W. Farr, D. Creedon, and M. Tobar, Spin-photon interaction in a cavity with time-reversal symmetry breaking, *Physical Review B* **89**, 224407 (2014).