

Supplementary Information

I. MECHANICAL RESONATOR FABRICATION

Mechanical resonators were fabricated on 10 cm diameter, high-resistivity (>10 k Ω -cm) silicon wafers with a 150 nm thick dry thermal SiO₂ isolation layer. The lithographic steps used to generate classically-tested mechanical resonators were as follows:

1. Base electrode metallization was defined by sputter-depositing 130 nm of Al at a rate of ~ 0.4 Å/s. This layer was optically patterned by spin-coating with photoresist and exposing with an I-line UV stepper; after developing the resist, exposed Al was anisotropically etched in an inductively-coupled plasma reactive ion etcher (ICP RIE) using a 2:1 Cl₂/BCl₃ gas mixture (mass flow rate ratio). The photoresist was stripped in acetone with ultrasonic agitation.
2. A 330 nm thick piezoelectric AlN film was deposited in a radiofrequency reactive sputter system, using an Al target with a 3:10 N₂:Ar gas mixture. To define vias through the AlN to the Al base electrode, the AlN layer was lithographically patterned and etched using two steps of ICP RIE. The first step used a pure Cl₂ gas plasma that etched through most of the exposed AlN. The second step used an 8:1 Ar/CF₄ gas mixture to remove any remaining AlN, while effectively stopping on the underlying Al layer. The photoresist was then stripped.
3. The top electrode comprised a sputter-deposited 130 nm thick Al film. This metal layer was lithographically patterned and etched with a 2:1 Cl₂/BCl₃ gas mixture using the ICP RIE.
4. The mechanically suspended structure was defined by a second layer of lithography, using two steps of ICP RIE. The first step used a Cl₂ plasma to etch completely through the exposed Al/AlN/Al stack, stopping on the underlying SiO₂. The second etch used 8:1 Ar/CF₄ to remove the exposed SiO₂, exposing bare Si.
5. Prior to mechanical release, wafers were diced into individual chips containing one mechanical resonator per die. Mechanical release was performed using isotropic XeF₂ etching, by placing individual dies in a custom-built vacuum chamber that was flooded with XeF₂ gas, which selectively removes any exposed Si, releasing the structure.
6. Completed resonators were placed in a measurement mount, with electrical connections made with 25 μ m diameter Al wire bonds.

II. QUBIT-MECHANICAL RESONATOR FABRICATION

Coupled qubit-mechanical resonator samples were fabricated on the same type of Si wafers as above. The lithographic steps for these devices were as follows:

1. The mechanical resonator structure was defined first, as described above **with the exception that the AlN layer was reduced in thickness to 300 nm**. However, after depositing and etching the base electrode and AlN vias, steps 1 and 2 above, two process steps were added to remove as much of the mechanical resonator material from the Si wafer as possible. The first step created vertical sidewalls in a rectangular frame just outside the mechanical resonator structure and its lead wires (similar to a picture frame in shape). After patterning the photoresist, a pure Cl₂ ICP RIE plasma was used to etch through the exposed AlN and Al layers, stopping on the underlying SiO₂. The photoresist was stripped, and a second photoresist pattern defined to allow chemical removal of all Al and AlN past the rectangular frame region defined in the previous step. The wafer was immersed in heated ($\sim 60^\circ\text{C}$) Al etchant (“Type A” Al etchant, Transene Co., Danvers, Mass.) for 22 minutes, removing all exposed Al and AlN while leaving the underlying SiO₂ un-etched. After stripping the photoresist, the only features remaining on the wafer were the mechanical resonator structures.
2. The qubit base wiring and mechanical resonator top wiring was defined by sputter-depositing 154 nm of Al at ~ 0.4 Å/s. A photoresist pattern was generated and the wafer ICP RIE etched using a 2:1 Cl₂/BCl₃ gas mixture, etching through the exposed Al. The photoresist was then stripped.
3. Unwanted Al, AlN, and SiO₂ were removed from the area around the resonator structure, opening a via for the final XeF₂ release. This was done using a two-step ICP etch through a photoresist pattern, first using pure Cl₂ gas, followed by an 8:1 Ar/CF₄ mixture, etching to the Si substrate.
4. The remainder of the qubit fabrication was completed in a manner very similar to our standard procedure[?]; this process, together with the resonator fabrication, comprised 13 separate lithographic processing steps.
5. The wafer was diced into individual chips, each containing one qubit coupled to one mechanical resonator. The mechanical resonators were released by placing the dies in the XeF₂ etching chamber.

6. Dies were placed in our standard aluminum qubit mount for measurement, with electrical connections made with 25 μm diameter Al wire bonds.

III. VERIFICATION OF MECHANICAL NATURE OF RESONANCE

We performed extensive measurements on a variety of mechanical resonators prior to attempting the coupled qubit-resonator experiment. Measurements were made at room temperature with a network analyzer to measure the transmission S_{21} across the resonator, using the circuit shown in the inset to Fig. 1 of the main paper. These in part allowed us to validate that the resonances we observed were indeed mechanical in nature, rather than some sort of spurious, purely electrical resonance. Simple tests, such as breaking the suspended part of the resonator, or removing metal wiring leading to the resonator, would make the resonances disappear. More careful measurements included varying the thickness of the active part of the mechanical structure, and measuring resonators with different degrees of mechanical release.

In Fig. 1 we show the results of a series of measurements of different mechanical structures, all mechanically released, showing the dependence of the measured resonance frequency on the suspended structure thickness. We clearly observe that the resonance frequency f_r scales as expected with structure thickness t , $f_r \sim v/2t$, with fit effective sound speed $v \approx 9100$ m/s, in reasonable agreement with that expected for the multi-layer resonator structure. Note that as the composition of the resonators varies in this series of measurements, the average sound speed will vary with thickness, so this scaling is only approximate.

The frequency dependence on thickness shown in Fig. 1 is a strong indicator of mechanical response, and is not easily explained as an electrical resonance involving the capacitance of the resonator sandwich structure. For example, increasing the AlN dielectric thickness would correspond to a reduction in electrical capacitance and thus an increase in resonant frequency, the opposite of what is observed.

We performed a second series of experiments to further ensure the resonant response is indeed mechanical. Prior to the mechanical release using XeF_2 , the resonator structures are rigidly attached to the underlying Si substrate. When the structures are exposed to XeF_2 , the gas selectively and isotropically etches the exposed Si, gradually undercutting the resonator structure, starting from the sides. Our standard process is timed to allow complete release of the structure from the Si. For this experiment, we instead performed a series of short, timed XeF_2 etches on a single resonator, only partially undercutting the resonator in each step, leaving a pillar of Si connecting the mechanical resonator to the substrate, a pillar whose diameter was reduced in each subsequent etch step. This supporting pillar of silicon acts to damp

out the mechanical oscillations.

In Fig. 2 we show the mechanical response for the resonator, first measuring the resonator without any exposure to XeF_2 , for which it displays no resonant response. Without removing the resonator from the measurement mount, the resonator was then exposed to XeF_2 for a short period of time, releasing roughly 30% of the mechanically active part of the structure, and the device re-measured. A small resonant feature is detectable at the design resonance frequency. This process of exposing and measuring was repeated a total of three times, leaving the resonator undisturbed in the mount, with the resonant response becoming more pronounced at each step. This experiment provides further strong evidence that the resonance is indeed mechanical in nature.

We performed further tests on resonators that were co-fabricated with qubits. The wafer design layout for the coupled devices included a number of test dies, which included both the qubit and mechanical resonator, but fabricated without a coupling capacitor. With no connection between the mechanical resonator and qubit, we were able to perform the same standard network analysis measurement of the co-fabricated mechanical resonator. The resonant frequencies f_r of these test dies were found to be in the range 6.0 to 6.2 GHz, straddling the value of the resonator reported in the main paper; the spread in resonance frequencies is due to AlN thickness variations across the wafer.

IV. QUBIT MEASUREMENT

The qubit was measured using our standard single-shot state measurement procedure, in which we apply a flux bias current to the coupling inductor FB (see Fig. 2 in the main text) sufficient to selectively tunnel the qubit population in the excited qubit state $|e\rangle$, localized in the left well of the qubit potential, into the right well, while leaving the ground state $|g\rangle$ population trapped in the left well⁷. An on-chip superconducting quantum interference device (SQUID) completes this projective state readout by detecting the magnetic flux through the qubit loop, which differs by approximately one flux quantum for a qubit state in the left versus the right well. The result of the SQUID measurement is recorded, the system allowed to relax to its ground state, and the state preparation and measurement sequence repeated. Accumulating many of these preparation and measurement sequences allows us to calculate the probability P_e that the qubit was in the excited state prior to measurement. **The data in Figs. 2 and 3 of the main text were generated using 600 repetitions of the state measurement for each value of P_e shown. The data in Fig. 4 were generated using 6,000 repetitions. Each probability in Fig. 5b and c was taken with 1,500 measurements. The data in Fig. 6b and d were taken with 6,000 and 30,000 repetitions, respectively. Each probability in Fig. 7b was calculated from 600 measurements. Finally, the data in Sup. Fig.**

3a and b were taken using 6,000 and 3,000 repetitions, respectively, for each value of P_e .

V. QUBIT CHARACTERIZATION

The qubit energy relaxation time was measured with the qubit at its resting frequency of 5.44 GHz, well detuned from the resonator. We first allowed the qubit to relax to its ground state before exciting it with a variable-duration microwave tone at the resting frequency. The excited state probability P_e was then measured as a function of the pulse duration, as shown in Fig. 3. The qubit responds with the expected Rabi oscillations with a time decay, where the oscillation frequency is determined by the pulse amplitude and the decay rate is related to the qubit's energy relaxation time T_{1q} . From this measurement we determine the microwave pulse duration τ_π needed to swap the qubit ground and excited states (a π -pulse). We then directly measured the qubit relaxation time by exciting the qubit with a π -pulse and monitoring the qubit P_e as a function of delay τ after the pulse, as shown in Fig. 3b.

VI. SIMULATIONS

The simulations presented in Fig. 4b, 5b, and 7c of the main text were generated by a full time-dependent quantum simulation of the combined qubit and mechanical resonator Hamiltonian. Considering only two levels in the qubit, the Hamiltonian is

$$\frac{\hat{H}}{\hbar} = -\frac{f_q}{2}\sigma_z + f_r a^\dagger a - \frac{i\Omega}{2}(a\sigma_- - a^\dagger\sigma_+), \quad (1)$$

where the first two terms are the qubit and resonator Hamiltonians and the third term models the interaction in the rotating wave approximation, which we believe to be valid in our experiment[?]. Here σ_z , σ_\pm are the qubit z and lowering and raising operators, and a^\dagger and a are the resonator phonon raising and lowering operators; f_q and f_r are the qubit and resonator frequencies, and Ω is the coupling strength in frequency units. For the simulations in Fig. 4, we included five qubit levels, including the qubit's known level-dependent energy dispersion, so a straightforward extension of Eq. (1) was used. The resonator was modeled using six equally-spaced energy levels which was sufficient for those simulations. For the swap simulations (Fig. 6 in main text) we used two qubit levels, and five resonator levels. For the more energetic coherent state simulations (Fig. 7 in the main text), we used eight qubit levels with their known level-dependent energy dispersion, and 38 equally-spaced energy levels for the resonator, which provides sufficient energy range to accommodate relatively large resonator excitations.

Decoherence was incorporated using the Lindblad master equation[?],

$$\frac{d\rho}{dt} = -\frac{i}{\hbar}[H, \rho] + \frac{1}{2} \sum_k [L_k, \rho L_k^\dagger] + [L_k \rho, L_k^\dagger], \quad (2)$$

where standard Lindblad operators were used to include qubit energy decay, $L_{1q} = \sqrt{1/T_{1q}}a$, qubit pure dephasing $L_{\phi q} = \sqrt{2/T_{\phi q}}a^\dagger a$, resonator energy relaxation $L_{1r} = \sqrt{1/T_{1r}}a$ and resonator pure dephasing $L_{\phi r} = \sqrt{2/T_{\phi r}}a^\dagger a$. For all simulations we used the parameters $T_{1q} = 17$ ns and $T_{1r} = 6.1$ ns, as measured for both the qubit and resonator. The phase coherence times $T_{\phi r, q}$ were set to an arbitrary value of 100 ns, sufficiently large that the qubit and resonator T_2 lifetime, which satisfies $1/T_2 = 1/2T_1 + 1/T_\phi$, would exhibit $T_2 \approx 2T_1$ as measured. For the thermal simulations here, we included the thermal excitation operator $L_\uparrow = \sqrt{\Gamma_\uparrow}a^\dagger$. We set the thermal excitation and energy decay Lindblad coefficients to correspond to the Boltzmann factor, $\Gamma_\uparrow T_{1r} = e^{-\hbar f_r/k_B T}$, where T is the resonator temperature and T_{1r} the resonator energy relaxation time[?].

For the thermal simulations presented in Fig. 4b of the main text, the resonator was initially placed in a thermal state with temperature corresponding to a certain average number of phonons $\langle n \rangle$, the qubit placed in its ground state $|g\rangle$, and the Lindblad equation numerically integrated, for a given qubit-resonator detuning Δ , with a time step of 0.05 ns. We observed that the coupled system reached a steady state after roughly 150 ns, so the simulation was stopped at that time. The equilibrium qubit P_e was calculated for each simulation, as a function of the qubit-resonator detuning Δ and resonator thermal phonon occupation $\langle n \rangle$.

The upper bound we quote for the maximum average phonon number, $\langle n \rangle_{\max} < 0.07$, was obtained using a very conservative formulation: This phonon number yields a change in the qubit excited state probability P_e , between the calculated on-resonance maximum and off-resonance minimum values, corrected by the qubit visibility γ (the difference between the value of P_e measured for the qubit excited state, as prepared using a tuned qubit π -pulse, and that measured for the qubit ground state), equal to three times the standard deviation σ in the measured qubit P_e over the full range of detunings: $\langle n_{\max} \rangle \ni \gamma(P_{e, \max} - P_{e, \min}) \approx 3\sigma$.

For the swap simulations presented in Fig. 5b of the main text, the qubit was placed in its excited state $|e\rangle$ while at the resting frequency, and was then tuned to an interaction detuning Δ by applying a trapezoidal tuning pulse with a 1 ns rise time, a flat top at Δ detuning for a variable time, and a 1 ns fall time back to the resting frequency; the simulation used a 0.1 ns time step. The qubit state was evaluated immediately after returning it to the resting frequency. The time axis in Fig. 3 of the main text includes both the rise and fall times (i.e. $\tau = 0$ starts at the beginning of the tuning pulse).

The trapezoidal pulse was chosen as a reasonable ap-

proximation to the flux bias tuning pulse seen by the qubit, which includes an electronics-limited, ~ 1 ns rise and fall time. The exact shape of the frequency tuning pulse strongly effects the precise amplitude of the state transfer between qubit and mechanical resonator. For example, simulations show that when the rise and fall time is set to zero, no asymmetry in the response is observed. Using the trapezoidal tuning pulse as described, however, the pronounced asymmetry emerges, as in Fig. 5b of the main text. However, this pulse shape is only an approximation of the actual tuning pulse, so that some discrepancies remain between experiment and simulation.

In the coherent state simulation in Fig. 7c of the main text, the resonator was initially placed in a coherent state using a given microwave drive amplitude, quoted in terms of $\sqrt{\langle n \rangle}$, where $\langle n \rangle$ is the average phonon number of the resulting Poisson distribution. We numerically integrated the free evolution of the Lindblad equation with the qubit at the interaction frequency for a total time of 60 ns with 0.1 ns time steps. For a given value of $\sqrt{\langle n \rangle}$, the squared amplitude of the excited state of the qubit, equal to P_e , was recorded for each time step. The time dependence of P_e for a range of values of $\langle n \rangle$ was calculated.

VII. CLASSICAL CIRCUIT ANALYSIS

We estimate the coupling strength between the qubit and mechanical resonator by analyzing a linearized version of the combined qubit-resonator circuit shown in Fig. 2 of the main text. We approximate the qubit Josephson element as a variable inductor $L_J = \Phi_0 / (I_0 \cos \delta_0)$, where $\Phi_0 = h/2e$ is the flux quantum, I_0 is the Josephson critical current, and δ_0 is the value of the junction phase at the minimum of the left well in the qubit potential²⁷. The qubit is thus approximated as a parallel LC oscillator, with inductance L_{eff} given by the parallel combination of the shunt inductor L_q and the Josephson inductance, $L_{\text{eff}} = (1/L_q + 1/L_J)^{-1}$, and capacitance C_q . This linearized circuit will yield frequencies that are accurate to a few percent in the limit of small excitations, such that the circuit contains far less than one excitation quantum. The value of δ_0 , and thus the value of the effective inductance L_J , varies with the flux bias of the qubit, so the qubit LC oscillator frequency $f_q \approx 1/(2\pi\sqrt{L_{\text{eff}}C_q})$ is flux-tunable.

The circuit resonances are found by examining the elec-

trical admittance between the nodes separated by the qubit inductor L_{eff} , given by

$$Y = \frac{1}{Z_{L_{\text{eff}}}} + \frac{1}{Z_{C_q}} + \frac{1}{Z_{C_c} + \left(\frac{1}{Z_{C_0}} + \frac{1}{Z_{L_m} + Z_{C_m}}\right)^{-1}}, \quad (3)$$

where Z_x is the impedance of the element x . Solving for the frequencies where the admittance goes to zero yields the resonant modes of the circuit. The positive frequency solutions are given by

$$\omega_{\pm} = \sqrt{\frac{-B \pm \sqrt{B^2 - 4AC}}{2A}}, \quad (4)$$

where $A = C_m L_m L_{\text{eff}} (C_c C_q + C_0 (C_c + C_q))$, $B = -L_{\text{eff}} (C_c (C_0 + C_m) + C_q (C_0 + C_c + C_m)) - C_m L_m (C_0 + C_c)$, and $C = C_0 + C_c + C_m$. The calculated frequencies ω_{\pm} vary with the qubit inductance L_{eff} , matching the experimental data in Fig. 2 (main text): For the qubit well de-tuned from the resonator, the mode frequencies $\omega_{\pm}/2\pi$ correspond to the resonator frequency f_r and the flux bias-determined qubit frequency f_q , while for the qubit tuned close to the resonator, the level avoidance as shown in Fig. 2e of the main text is reproduced. The coupling strength Ω is then the minimum frequency difference between the two modes, $\Omega = |\omega_+ - \omega_-|_{\text{min}}/2\pi$.

The geometry of the mechanical resonator allows us to calculate the resonator capacitance $C_0 = 0.2$ pF, and the values for the piezoelectric coupling $k_{\text{eff}}^2 = 1.2\%$ (caption to Fig. 1 in the main text) and the frequency of the mechanical resonator $f_r = 6.17$ GHz (Fig. 3 in the main text) yield²⁷ $C_m = 0.70$ fF and $L_m = 949$ nH. Combined with the geometrically-determined qubit capacitance $C_q = 0.98$ pF and coupling capacitance $C_c = 0.5$ pF, we find the resonant frequencies as a function of L_{eff} , and then find the minimum frequency difference; this yields $L_{\text{eff, min}} = 580$ pH and $\Omega = 110$ MHz.

The experimental coupling strength is extracted by fitting the measured mode frequencies as a function of flux bias (shown in Fig. 2e) to this circuit model. We find an experimental coupling frequency $\Omega = 124$ MHz, differing by about 10% from that calculated above, and with the geometric values of C_0 , C_c and C_q , estimate the mechanical capacitance and inductance to be $C_m = 0.89$ fF and $L_m = 750$ nH, differing somewhat from the values given above but still quite reasonable. The value for $L_{\text{eff, min}}$ cannot be accurately extracted from this fit, but is found to be roughly 700 pH, again reasonably close to the value above.




FIG. 1: Classically-measured resonant frequency f_r of a series of mechanical resonators with varying total thickness t . Data (blue points) show the expected relationship between frequency and thickness (red line, $f = v/2t$), with thinner structures showing higher resonant frequency responses, and the data scaling with a sound speed $v = 9100 \text{ m/s}$. The spread in frequency values for a given quoted thickness is due to thickness variations across the wafer. The structure measured in Fig. 1 of the main paper is indicated by the red point.




FIG. 2: Resonant response of a single mechanical resonator, each measured after a brief exposure to XeF_2 gas, resulting in further partial release. The percentage undercuts are rough estimates of the area of the mechanical resonator that remained directly connected to the Si substrate after each exposure. No response is visible prior to undercut, and the amplitude of the response is seen to increase with degree of release; the response amplitude is found to scale with quality factor, thus showing that the quality factor for partially released structures is determined in part by mechanical coupling to the substrate.




FIG. 3: Qubit characterization. **a**, Qubit Rabi oscillation and pulse sequence. Inset: The qubit is tuned to the resting frequency $f_q = 5.44$ GHz, well out of resonance with the resonator, and allowed to relax to its ground state $|g\rangle$ before applying an inductively-coupled microwave tone, with fixed amplitude and variable duration τ . The resulting qubit excited state occupation P_e is then evaluated (see Supplementary Information). Main panel: Measured excited state probability P_e as a function of microwave tone duration τ . From this measurement we determine the duration τ_π needed for a microwave π -pulse, exchanging the qubit $|g\rangle \leftrightarrow |e\rangle$ states. **b**, Measurement of qubit relaxation time. Inset: Pulse sequence. While at its resting frequency, the qubit is excited from its ground to its excited state by a 5.0 ns microwave π -pulse (X_π). After a delay τ , we measure the qubit excited state probability P_e . Main panel: Qubit measured P_e vs. τ (blue), with a fit to a relaxation time $T_{1q} \approx 17$ ns (red).