Supplemental Material for Observation of Collisions between Two Ultracold Ground-State CaF Molecules

Lawrence W. Cheuk,[1,2,3,*] Loïc Anderegg,[1,2] Yicheng Bao,[1,2] Sean Burchesky,[1,2] Scarlett S. Yu,[1,2] Wolfgang Ketterle,[2,4] Kang-Kuen Ni,[1,2,5] and John M. Doyle[1,2]

[1]*Department of Physics, Harvard University, Cambridge, MA 02138, USA*
[2]*Harvard-MIT Center for Ultracold Atoms, Cambridge, MA 02138, USA*
[3]*Department of Physics, Princeton University, Princeton, NJ 08544, USA*
[4]*Department of Physics, Massachusetts Institute of Technology, Cambridge, MA 02139, USA*
[5]*Department of Chemistry and Chemical Biology, Harvard University, Cambridge, MA 02138, USA*

(Dated: February 8, 2021)

CAF ENERGY STRUCTURE

Laser cooling and imaging of CaF molecules is performed by addressing $|X^2\Sigma_{1/2}, N = 1, -\rangle \rightarrow |A^2\Pi_{1/2}, J = 1/2, +\rangle$ transitions. Here, $+/-$ indicate the parity of the state. This scheme allows photon-cycling, which is necessary for imaging and cooling.

The collisional loss rate measurements are conducted in hyperfine states from the first excited $(X^2\Sigma_{1/2}, N = 1)$ and ground rotational $(X^2\Sigma_{1/2}, N = 0)$ manifolds. Microwave transitions are used to transfer population between specific hyperfine states in the $N = 1$ and $N = 0$ rotational manifolds.

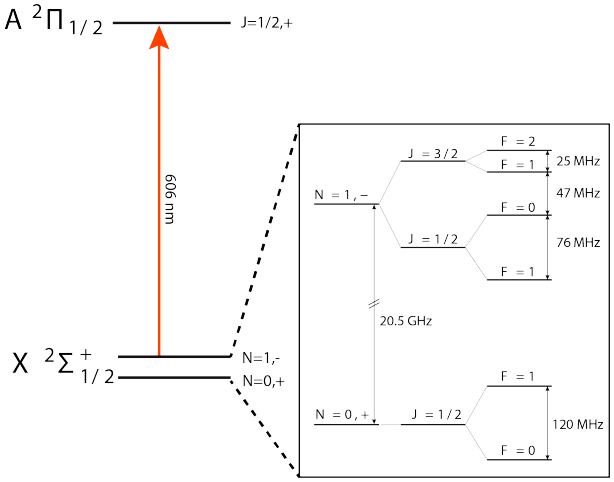

FIG. S1. CaF Energy Levels Relevant to This Work.

IMAGING SINGLE MOLECULES

The molecules are imaged using $\Lambda-$imaging [1]. Two hyperfine components ($|J = 1/2, F = 1\rangle$ and $|J = 3/2, F = 2\rangle$) remain on during Λ-imaging, with a power balance of approximately 5:1 and a total intensity of $\sim 40\,\mathrm{mW/cm^2}$. The single photon detuning is $\Delta = 2\pi \times 51\,\mathrm{MHz}$ with the 2-photon detuning set $\delta = 2\pi \times 200\,\mathrm{kHz}$ to the blue.

The Λ-imaging beams consist of 3 beams, one in each direction, with $1/e^2$ waists of about 800 μm to mitigate scatter in the chamber, which reduces the imaging fidelity. We estimate the imaging fidelity by looking at histograms with and without molecules. To isolate the contribution of the light from the molecules, we scale the background histogram and subtract it off from histogram with molecules loaded. There are two types of errors; unoccupied traps can be misidentified as traps containing molecules or loaded traps can be misidentified as empty. Misidentification of an unloaded tweezer as a loaded tweezer can contribute to the error in the number of surviving pairs. Loaded tweezers (1) misidentified as empty tweezers (0) are of no consequence, aside from a data rate reduction.

The probability of incorrectly detecting 0s as 1s is about 1% for the threshold used, and the probability of incorrectly detecting 1s as 0s is about 2%. The single shot imaging error is taken as the sum of the two possible errors, which yields a fidelity of 97%.

MERGING TWO TWEEZERS

The moving tweezer is generated by the first order diffraction peak of an acousto-optical deflector (AOD). The input waveform for the AOD is produced by a Universal Software Radio Peripheral (USRP).

During loading and imaging the tweezers are separated by $10\,\mu$m. After the first image, the AOD tweezer is moved via a linear frequency ramp to the position of the stationary tweezer over 1 ms. Once the two tweezers are overlapped, the light in the AOD tweezer is linearly ramped off in 1 ms. To separate the tweezers, the process proceeds in the reverse order.

BALANCING OF TWO TWEEZERS

Doubly occupied tweezers cannot be imaged due to light-assisted collisions, which cause 2-body loss of both

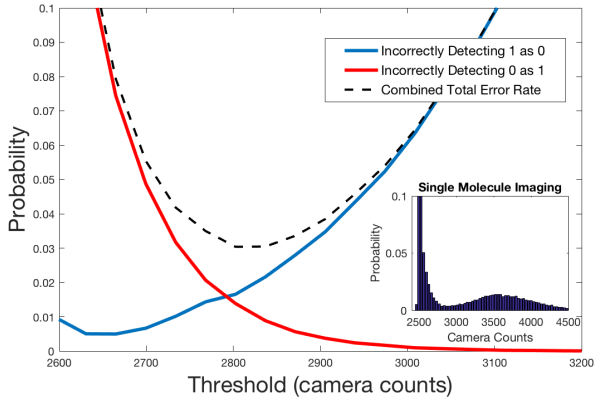

FIG. S2. Imaging Fidelity. The imaging fidelity is estimated by subtracting a scaled background light histogram from a histogram of images containing molecules. The blue curve shows the probability that a tweezer containing a molecule is detected as empty, plotted as a function of the threshold. The red curve shows the probability of incorrectly identifying an empty tweezer as occupied. The sum of the two errors is the total imaging error rate, shown in dotted black. The inset histogram shows a data set of about 50,000 images of tweezers loaded with molecules. The loading fraction can be estimated by the ratio of counts above and below the detection threshold value.

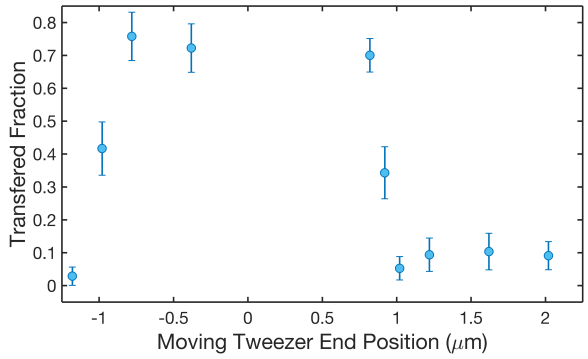

FIG. S3. Tweezer Spatial Overlap. The spatial overlap of the tweezers is scanned in 1D by sweeping the moving tweezer over the stationary tweezer and then ramping off the moving tweezer. For tweezers overlapped to within roughly the beam waist, the transfer fraction is insensitive to position.

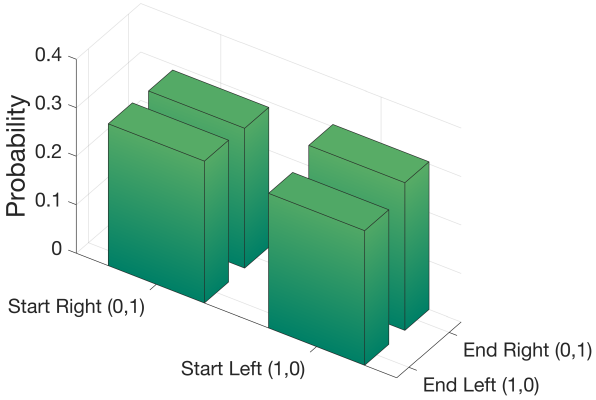

FIG. S4. Single Molecule Splitting Ratio. The data where one molecule is loaded is used to monitor the splitting ratio. Shown are the measured probabilities that a molecule loaded in either tweezer remains in the initial tweezer or switches tweezers after merging and splitting. When the tweezer trap depths are balanced and are spatially overlapped well, the probability of finding a single molecule in either tweezer after splitting is identical. Trap depth imbalances between the two tweezers bias the splitting towards the deeper tweezer. A mismatch in spatial overlap causes a molecule to remain in its initial tweezer with higher probability.

molecules before enough photons are collected. To detect survival of two molecules after merging, we turn on the second tweezer and separate the molecules before re-imaging. To maximize the probability of the two molecules separating in different tweezers, the spatial overlap and trap depth need to be tuned. To spatially overlap the tweezers, we scan the ending point of the AOD tweezer over the stationary trap and then ramp down the AOD trap. Transfer fraction is measured by selecting data in which the first image shows one molecule in the moving tweezer, and the second image one molecule in the stationary tweezer. The end position of the merge is then set between the shoulders of the scan in Fig. S3.

Power balancing is done by looking at the probability of a single molecule re-separating with equal probability into the two traps (Fig. S4). Uneven splitting does not affect the results of our measurement, but it reduces the data rate. The splitting ratio is available in the dataset, and is found to be stable throughout the measurements.

STATE PREPARATION

Prior to microwave transfer from $N = 0$ to $N = 1$, we optically pump the molecules into $|N = 1, F = 0, m_F = 0\rangle$ using a $100\,\mu$s optical pumping pulse. This consist of 2 mW of resonant $X - A$ light that addresses all hyperfine components in the $N = 1$ rotational level except for $F = 0$. After optical pumping, the tweezer depth is lowered and the molecules are transferred into the ground rotational manifold via microwaves. We populate the states $|N = 0, F = 1, m_F = \pm 1, 0\rangle$ by applying a π-pulse (Fig. S5).

Electric dipole transitions from $|N = 1, F = 0, m_F = 0\rangle$ to $|N = 0, F = 0, m_F = 0\rangle$ are strictly forbidden by angular momentum selection rules. In the presence of a magnetic field, the transition becomes weakly allowed due to mixing between $|N = 0, F = 0, m_F = 0\rangle$ and $|N = 0, F = 1, m_F = 0\rangle$. We estimate an effective $F = 0 \rightarrow F = 0$ Rabi frequency of $\Omega_r \sim$

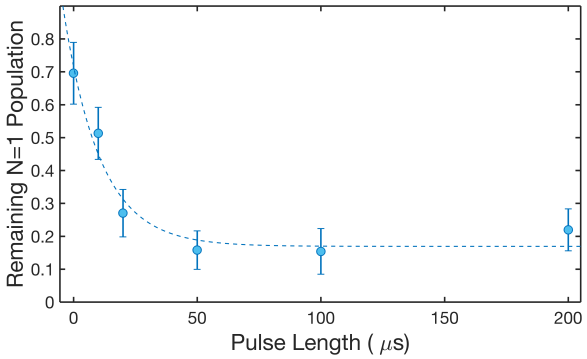

FIG. S5. Preparing Molecules in the Ground Rotational Manifold. Molecules are optically pumped to the $|N=1, F=0, m_F=0\rangle$ state. Subsequently, a microwave π-pulse is applied to transfer population in $|N=1, F=0, m_F=0\rangle$ to the $|N=0, F=1\rangle$hyperfine manifold in the ground rotational state. The molecules are then dark to the imaging light addressing $N=1$. Shown is the remaining number of molecules in $N=1$ as a function of the duration of the optical pumping pulse.

$\frac{\mu_b B}{\Delta_{\rm hf}}\Omega_{F=0\to F=1} \approx 2\pi\times 2\,$kHz, where hf is the ground rotational hyperfine splitting, and $\Omega_{F=0\to F=1}$ is the Rabi frequency for the allowed transition. Due to the small Rabi frequency, a Landau-Zener sweep was used to transfer population, since it does not require precise determination of the resonance nor timing of the pulse.

The microwave frequency was set to be on resonance with the transition at a field value of 3.5 G. To populate the state we swept the magnetic field down from 5 Gauss to 2.5 Gauss in 20 ms. Two effects lead to incomplete transfer. First, projection errors occur at finite detuning. These errors are given by Ω^2/Δ^2, where Ω is the Rabi frequency and Δ is the detuning. For the start and end points, this corresponds to 2×10^{-4} and 3×10^{-4} respectively. Note that the Rabi frequency is different for the start and end points, since the $|N=0, F=1, m_F=0\rangle$ varies linearly with magnetic field. The second type of error is due to Landau-Zener tunneling near the avoided crossing. This is given by $e^{-2\pi\frac{\Omega^2}{d\Delta/dt}}$. This can be improved with longer sweep times. However, in practice, no further gains are obtained once the sweep time approaches the coherence time. For the parameters used, we estimated that a fidelity of 87% can be obtained.

Experimentally, we measure the transfer efficiency by imaging $N=1$ molecules after successive sweeps (Fig. S6). Fitting the series of Landau-Zener sweeps gives a measured transfer fraction of about 74%. In addition to imperfect Landau-Zener sweeps, imperfect optical pumping also reduces the transfer efficiency.

To eliminate the residual population in $N=1$, we apply a 5 ms pulse of resonant $X-A$ light. This "cleanup" pulse heats the $N=1$ molecules out of the trap. The cleanup pulse contains all hyperfine components except $|N=1, J=3/2, F=1\rangle$. This is to avoid possible two-photon transfer between the $N=0$ hyperfine manifolds, since the $|N=1, J=3/2, F=1\rangle - |N=1, J=1/2, F=1, m_F\rangle$ splitting of 120 MHz is nearly resonant with the splitting between the $F=0$ and $F=1$ manifold in the $N=0$ ground rotational manifold.

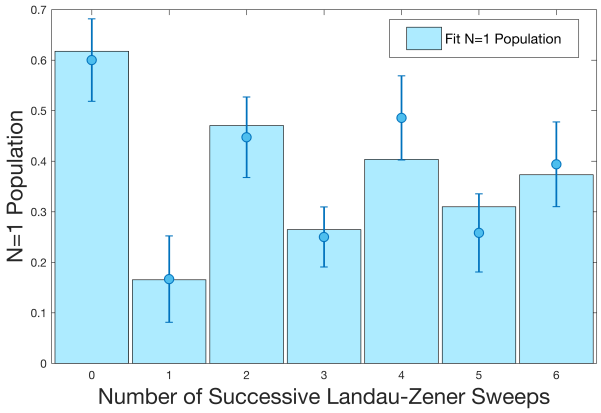

FIG. S6. Landau-Zener Transfer into the Absolute Ground State $|N=0, F=0\rangle$. Shown is the remaining $N=1$ population versus the number of Landau-Zener sweeps. By fitting the transfer efficiencies to a exponentially decaying sinusoidal curve with period set to two sweeps, we obtain the single sweep efficiency.

CHARACTERIZING THE TWEEZER TRAP

The tweezer trap is modeled by a Gaussian beam with astigmatism. Two parameters characterize the beam, the $1/e^2$ waist and the focal shift between two orthogonal directions in the radial plane perpendicular to the beam axis. These parameters are determined from trap frequency measurements and power measurements of the total power of the trapping beam, along with the theoretically calculated polarizability.

The trap frequencies are measured by modulating the trap depth and looking for parametric heating loss features at twice the trapping frequencies. We find a radial trap frequency of $\omega_r = 2\pi\times 76.5\,$kHz and an axial trap frequency $\omega_a = 2\pi\times 5.6\,$kHz at a power of 170 mW.

Next, various model trap potentials are constructed with beam waists near the diffraction limit, and focal shifts near the corresponding Rayleigh range. We then extract the trapping frequencies of each configuration using the local curvature at the trap center. We find that a $1/e^2$ beam waist of 1.53 μm and astigmatic focal shift of $\pm 6.2\,\mu$m matches the measured radial and axial trapping frequencies well. These parameters are used for subsequent simulations.

To simulate trajectories of single particle, we perform Monte-Carlo simulations using the model trap potential.

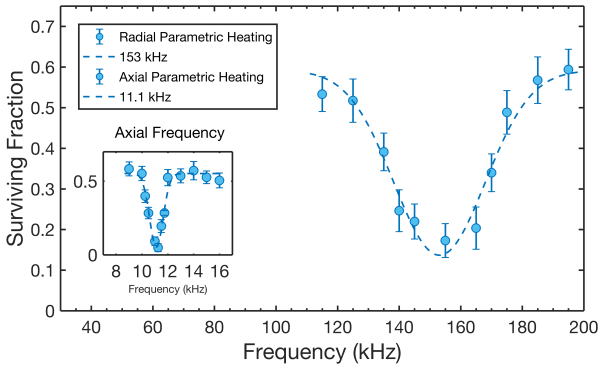

FIG. S7. Measuring trap Frequencies via Trap Modulation. The trap intensity is modulated, and the surviving fraction of molecules is recorded as a function of modulation frequency. Two loss features are seen (inset shows the axial trapping frequency), corresponding to the parametric heating resonances for the axial and radial directions. The heating resonances are at twice the trap frequencies.

This is done on a graphics processing unit (GPU), allowing a high degree of parallelization.

THERMOMETRY OF MOLECULES INSIDE OPTICAL TWEEZERS

Thermometry of molecules inside the optical tweezers is crucial for our work. It allows one to determine the collisional energy, and also extract the average in-trap density. We assume that the spatial and velocity distribution inside the tweezer trap is well-approximated by a thermal distribution after laser-cooling. Four different methods are used to measure the single molecule temperature. The time-of-flight expansion method is used since it requires minimal assumptions of trap geometry. Nevertheless, all four methods yield results that are similar.

Time-of-flight Expansion

Time-of-flight (ToF) is the most direct way to probe velocity distribution of an ultracold atomic or molecular cloud, and has the fewest number of assumptions. The tweezer is turned off for various durations before the molecules are imaged (Fig.S8(a)). To avoid affecting the spatial distribution of the expanded wavefunction due to photon recoil, we limit the number of photons scattered during fluorescent imaging. The resonance imaging beams provide a scattering rate around $10^6\,\mathrm{s}^{-1}$ and are kept on for $10\,\mu$s. On average, 100 photons are scattered and one photon is collected by the EMCCD camera. The distributions are then fitted to Gaussian curves, and the result is fitted a model of free expansion of a Gaussian sample. The finite imaging duration is taken into account. Using this method, we obtain a temperature of $41(12)\,\mu$K. The relatively large error bars are due to the low signal-to-noise of this method.

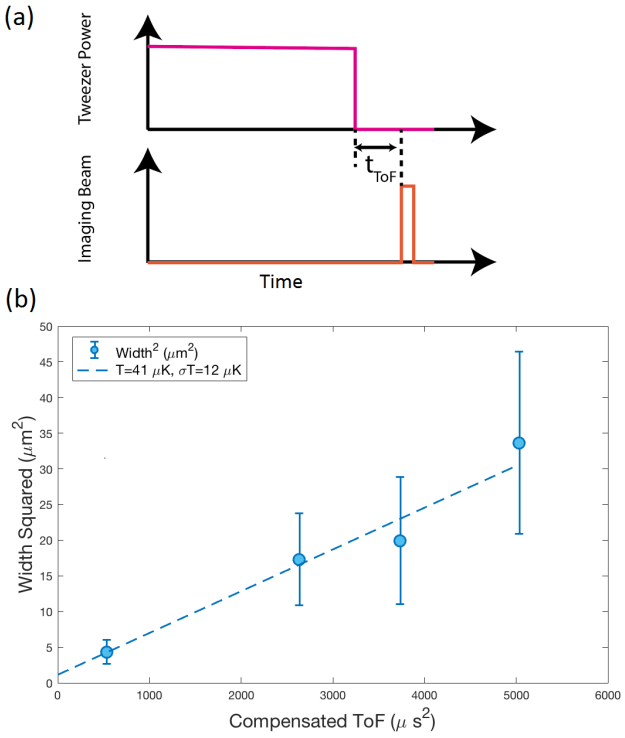

FIG. S8. Thermometry using Time-of-Flight Expansion. (a) Experimental sequence used for time-of-flight expansion. (b) Fit to time-of-flight data. The data is fit to a model of a freely expanding Gaussian sample. The finite imaging duration is taken into account by using the effective compensated time-of-flight.

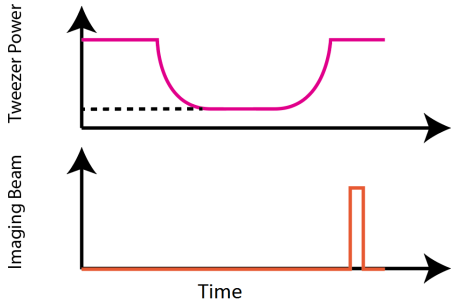

FIG. S9. Experimental Sequence for an Adiabatic Ramp.

Adiabatic Ramp

The tweezer is ramped down adiabatically following an exponential decay curve to various trap depths [2]. It is

then held at constant depth for 40 ms to allow molecules with higher energy to escape the trap (Fig.S9). Subsequently, the tweezer light is jumped back to the original trap depth. If the molecule is still inside the trap, it will be recaptured and imaged. During the ramp, the temperature of the molecule is adiabatically lowered since the phase space density is nearly conserved. If the adiabatic condition is met during the ramp, molecule temperature should scale linearly as the trap depth for a harmonic trap. This is also confirmed with Monte-Carlo simulations using the trap model mentioned in the previous section. Comparison to dynamic Monte-Carlo simulations gives a temperature of 24 μK. This method is sensitive to the details of the trapping potential away from the center, as the molecules invariably explore the edges of the trap.

Diabatic Ramp

The trap potential is instantaneously (in less than 100 ns) jumped to a lower level, and then stays constant for 40 ms in order for molecules with higher energy to escape (Fig.S10). This method can be understood as truncating the high energy tail of a Boltzmann distribution in the trap. The distribution of energy in the trap can then be mapped out by scanning the lowered trap depth. The survival probability as a function of final trap depth is then compared with dynamic Monte Carlo simulations, and yields a trap depth of 24(7) μK. Again, this method is sensitive to the shape of the trap away from the center, as molecules explore the edges of the trap.

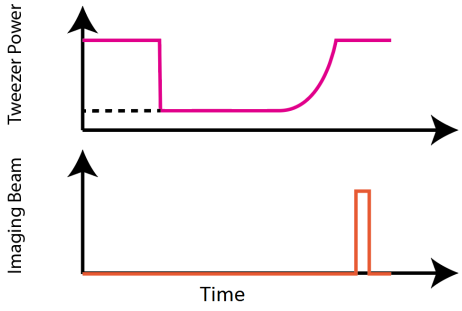

FIG. S10. Experimental Sequence for a Diabatic Ramp.

Release and Recapture

The trap is turned off for various amounts of time to allow the molecular distribution to expand. The trap is then turned back on to recapture molecules that have not traveled out of the trapping region (Fig.S11). The survival versus release time is then compared with Monte-Carlo simulations, and gives a temperature of 52(11) μK. This method is also sensitive to the details of the trap away from the trap center.

DETERMINING THE IN-TRAP DENSITY

To determine the in-trap density, we rely on Monte-Carlo simulation using the temperature determined via time-of-flight. One could worry that the anharmonicity of the tweezer trap could lead to systematic errors of the density. However, as we show below, in our temperature regime, the harmonic approximation suffices for determining the density. This means that one only needs the trapping frequencies, which are directly measured.

Using the trap parameters determined earlier, we obtain a trap depth of 1.38 mK. The temperature of the molecules is much lower, at 41(12) μK, giving a temperature to trap depth ratio of $\eta = 34$. For $\eta \gg 1$, the molecules only explore the center of the trap. By comparing the density obtained using a 3D harmonic trap to that obtained via Monte-Carlo simulations with the model potential, we confirm that both yield the same average density at these temperatures. We also find that the effect of gravity is negligible.

In Fig.S12, we show the trap-averaged single particle density $\frac{1}{N}\int d^3r\, n(r)^2$ for both the full Monte-Carlo simulation with the model potential (including astigmatism and gravity), and also for the 3D harmonic trap with corresponding trap frequencies. The overlap of the two curves confirm that at the temperature range explored in the experiment, the density can be estimated simply by using the harmonic approximation.

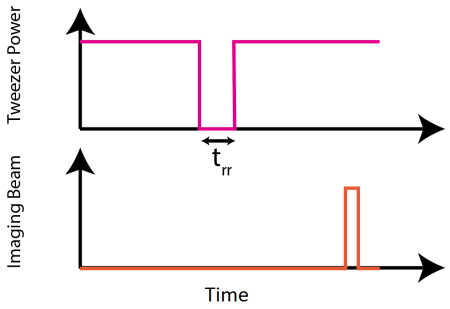

FIG. S11. Experimental Sequence for Release and Recapture.

DETERMINING COLLISIONAL LOSS RATE

The collisional loss rate β is given by $\beta = \frac{1}{\rho\tau}$, where ρ is the average single particle density in the trap and τ is the exponential time constant of the fit from the experimental decay curve.

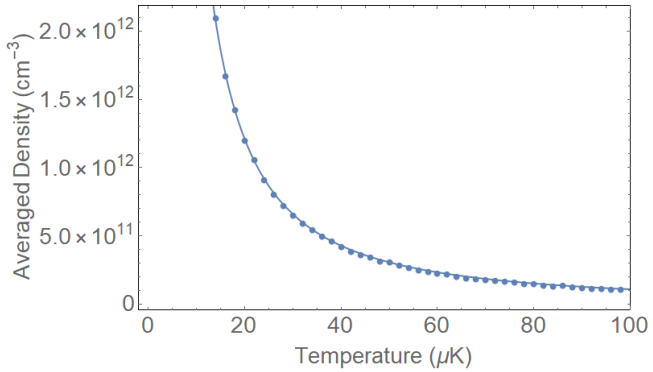

FIG. S12. Shown in dots is the average single particle density obtained from Monte-Carlo simulations. Shown in solid line is the calculated average single particle density for a 3D harmonic trap.

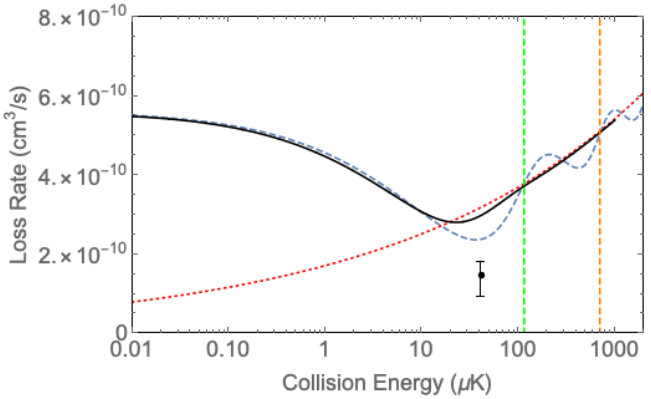

FIG. S13. Universal Loss Rates. The loss rate versus collisional energy is shown by the blue dashed curve. The thermally averaged loss rate is shown by the black solid curve. The classical Langevin limit is shown by the red dotted curve. The d-wave and g-wave barrier energies are indicated by the green and orange vertical dashed lines.

UNIVERSAL LOSS RATE

The universal loss rate is calculated using a single-channel model [3, 4] with unity loss at short range. The loss rate is first calculated as a function of collisional energy. Next, we perform a thermal average over the collisional energies in order to compare with the measured values [5]. Both loss rates are shown in Fig. S13.

* lcheuk@princeton.edu

[1] L. W. Cheuk, L. Anderegg, B. L. Augenbraun, Y. Bao, S. Burchesky, W. Ketterle, and J. M. Doyle, Physical Review Letters **121**, 083201 (2018).
[2] C. Tuchendler, A. M. Lance, A. Browaeys, Y. R. P. Sortais, and P. Grangier, Phys. Rev. A **78**, 033425 (2008).
[3] Z. Idziaszek and P. S. Julienne, Phys. Rev. Lett. **104**, 113202 (2010).
[4] M. D. Frye, P. S. Julienne, and J. M. Hutson, New Journal of Physics **17**, 045019 (2015).
[5] P. D. Gregory, M. D. Frye, J. A. Blackmore, E. M. Bridge, R. Sawant, J. M. Hutson, and S. L. Cornish, Nature Communications **10**, 3104 (2019).