Rb-87 Atom Source, Laser Cooling and Tweezer Loading

Oct 22, 2026 · D.F. Hakala

Overview

A neutral-atom array starts as a lump of rubidium metal and ends as single atoms at a few microkelvin, one per tweezer. Getting there is a chain of stages. Each one cools by roughly one to two orders of magnitude and narrows the atoms down to where they're needed. The cold cloud left after sub-Doppler cooling is the **reservoir** the tweezers load from.

Stage	What it does	Temperature	Atoms	Time
Vapor source	Releases Rb from metal or a salt	~300 K (hundreds of m/s)	—	continuous
2D⁺ MOT	Cools in two dimensions and sends a slow beam to the science cell	beam at ~10–20 m/s	~10^9–10^{10} /s	continuous
3D MOT	Captures and confines atoms at the tweezer plane	~100–200 μK	10^6–10^8	0.1–1 s
Optical molasses	Sub-Doppler (Sisyphus) cooling	~5–20 μK	same	5–20 ms
Gray molasses (optional)	Dark-state cooling on D1	~2–5 μK	same	ms
Tweezer loading	One atom per site via collisional blockade	~10–20 μK in-trap	0 or 1 per site, 50–90%	during molasses

The numbers are typical ranges for Rb-87 experiments, not specifications for a particular apparatus.

Background: how an optical tweezer traps an atom

An optical tweezer traps an atom with the light's electric field, not with radiation pressure. A tightly focused, far-detuned laser beam polarizes the atom. The induced dipole's energy is lowest where the light is brightest, and the focus is the brightest point in space, so the atom is pulled into it.

1. The light induces an oscillating dipole. The laser field $E(t) = E_0 \cos \omega t$ pushes the electron cloud one way and the nucleus the other, inducing a dipole $p = \alpha E$ oscillating at the laser frequency. $\alpha(\omega)$ is the atom's dynamic polarizability. The dipole's interaction energy with the field that created it, averaged over an optical cycle, is

$$U = -1/2 \langle \mathbf{p} \cdot \mathbf{E} \rangle = -\frac{\mathrm{Re}(\alpha)\, I}{2\varepsilon_0 c}$$

The potential is proportional to the local intensity I. This is the formula used in `tweezer_potential.py`.

2. Red detuning makes the atom seek high intensity.

- **Below resonance (red detuning).** The electron oscillates in phase with the driving field, like a mass on a spring driven below its natural frequency. Re(α) > 0, so U is negative and lowest where I is highest: the atom is pulled toward bright spots.
- **Above resonance (blue detuning).** The electron oscillates out of phase, Re(α) < 0, and the atom is pushed toward dark regions. That's used for "bottle beam" and other repulsive traps.

At 1064 nm, Rb is far red of both D lines (780 and 795 nm), so it's a bright-spot trap. The model uses α ≈ 687 atomic units for Rb at 1064 nm, a value quoted from memory that should be checked against a reference.

3. The quantum picture: the AC Stark (light) shift. The same effect is the light shift of the ground state. For a two-level atom with Rabi frequency Ω and detuning $\Delta = \omega - \omega_o$,

$$\Delta E_g = \frac{\hbar \Omega^2}{4\Delta}$$

Ω^2 is proportional to I, and $\Delta < 0$ for red light, so the ground state is pushed down in proportion to intensity. A ground-state atom's energy landscape follows the intensity pattern, and the trapping potential is that light shift. Real Rb has two D lines plus the far-detuned counter-rotating terms; summing them gives the polarizability in step 2. The excited states are shifted by a different amount, which is why the cooling and imaging resonances move inside a deep tweezer (Section 6).

4. The force is the gradient of the potential.

$$\mathbf{F} = -\nabla U = \frac{\mathrm{Re}(\alpha)}{2\varepsilon_0 c} \nabla I$$

This **dipole (gradient) force** is conservative: there's no friction, so it holds atoms but can't cool them. That's why atoms must be laser-cooled before or during loading. There's also a **scattering force** from absorbed and re-emitted photons, which pushes atoms along the beam and heats them. Far detuning keeps it negligible, because depth and scattering fall at different rates:

$$U \propto \frac{I}{\Delta}, \qquad \Gamma_{sc} \approx \frac{\Gamma}{\Delta}\frac{U}{\hbar} \propto \frac{I}{\Delta^2}$$

Detuning far and adding power keeps the depth while almost eliminating scattering. Ignoring counter-rotating terms, 1064 nm is about 100 THz from the D lines, so $\Gamma/\Delta \approx 6 \times 10^{-8}$. A 13 MHz-deep trap then scatters only about 5 photons per second, a heating rate of order μK per second, which is why 1064 nm traps hold atoms for seconds.

5. Focusing turns the beam into a 3D trap. An unfocused beam is bright along a line, so it would confine the atom sideways but not along the beam. A high-NA objective concentrates the light into a spot about a micron wide, making the intensity a true 3D maximum. Near the bottom the trap is harmonic:

$$\omega_r = \sqrt{\frac{4U_0}{mw_0^2}}, \qquad \omega_z = \sqrt{\frac{2U_0}{mz_R^2}}, \qquad z_R = \frac{\pi w_0^2}{\lambda}$$

- **Radially** the intensity falls off over the waist w_o, about 1 μm, which gives steep walls.
- **Along the beam** it falls over the Rayleigh range z_R, a few μm, so axial confinement is weaker.

From `tweezer_potential.py` (NA 0.5, 6.5 mW, x-polarized): depth 0.645 mK, radial frequencies 72 and 76 kHz (unequal because of polarization), axial frequency 12 kHz. High NA matters because z_R shrinks as $1/NA^2$, raising ω_z quickly, and the smaller trap volume is what lets collisional blockade leave one atom per site.

6. From one tweezer to an array. Each trap is its own intensity maximum, so an array needs many foci at chosen positions.

- **SLM (static grid).** The objective maps angle to position: a beam tilted by θ focuses at $x \approx f\theta$. The SLM hologram is a superposition of many tilts, one per site, and the objective turns each into its own spot. Gerchberg–Saxton or weighted GS computes that phase pattern and sets the spots' relative powers, which sets the depth uniformity.
- **AOD (moving traps).** Each RF tone at frequency f diffracts the beam at an angle proportional to f. Several tones give several spots, and sweeping a tone drags its spot, and the atom in it, smoothly across the plane.
- **Spacing.** Sites are a few μm apart and waists are about 1 μm, so the traps barely overlap. Each site's potential is essentially that of an isolated tweezer, apart from small interference ripples if neighbors are coherent.

7. Polarization adds a small state-dependent part. At high NA the focus has a longitudinal field component 90° out of phase with the main field, so the light off axis is slightly elliptical. That adds a vector light shift proportional to m_F on top of the scalar trap: about ±60 kHz for stretched states, and zero for the m_F = 0 clock states. With linear polarization and a bias field along the beam, it largely drops out.

In short: the light induces a dipole, red detuning makes its energy lowest at high intensity, far detuning makes that energy shift dominate photon scattering, and focusing to a micron spot turns the intensity maximum into a 3D potential well. The SLM and AODs multiply that well into an array.

1. Making the atomic vapor

Two common sources:

- **Alkali-metal dispenser.** A small metal boat holding a rubidium salt (Rb_2CrO_4) mixed with a reducing agent. Running a few amps through it (roughly 3–6 A) heats it, and the reaction releases pure Rb vapor. Dispensers are clean, compact and switch on and off quickly.
- **Oven or ampoule.** A sealed glass or steel reservoir of Rb metal, broken open under vacuum and warmed to a few tens of °C. Room-temperature Rb already has a useful vapor pressure, about 4×10^{-7} Torr at 25 °C, and heating raises it.

Many setups add **light-induced atomic desorption (LIAD)**: UV LEDs knock Rb atoms off the glass walls. That briefly raises the vapor pressure during loading without permanently degrading the vacuum.

The result is a thermal gas at room temperature with atoms moving at a few hundred m/s. Only the slow tail of the velocity distribution, below about 20–30 m/s, can be captured by laser cooling, so the next stages select and slow those atoms.

2. Source chamber and 2D⁺ MOT

Fast loading needs plenty of Rb vapor, but the qubits need ultra-high vacuum. Background-gas collisions knock atoms out of the tweezers, and trap lifetime scales inversely with pressure. So most systems use **two chambers joined by a narrow differential-pumping tube**: the source cell runs at a relatively high Rb pressure, and the science cell stays below about 10^{-11} mbar.

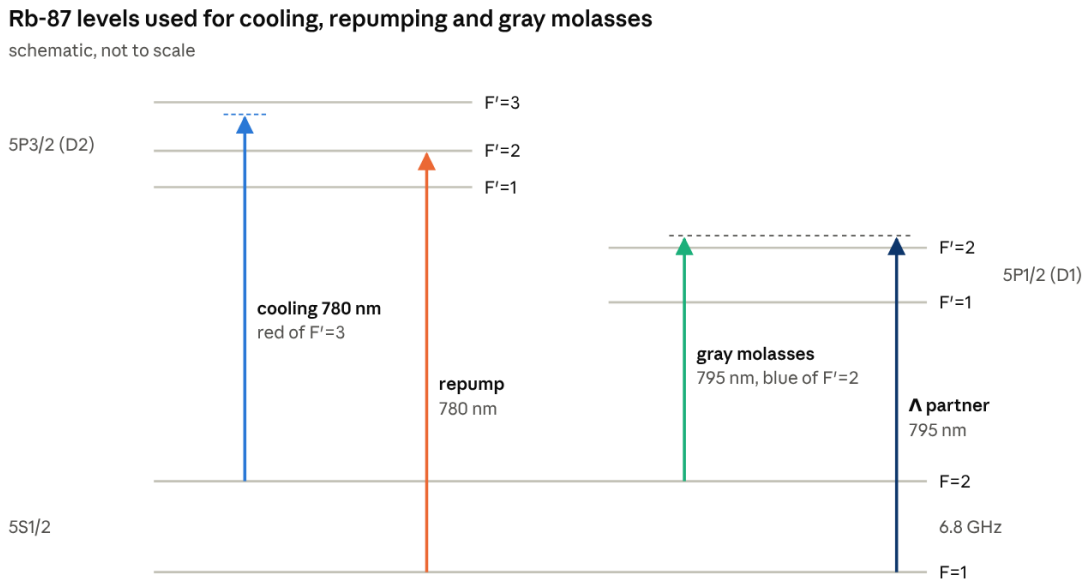
- **2D MOT.** Two pairs of red-detuned 780 nm beams plus a 2D quadrupole magnetic field cool atoms across two axes and squeeze them onto a line along the tube axis.
- **Push beam.** A weak beam along that axis shoves the cooled atoms through the tube and cools them along the beam too. That's the "+" in 2D⁺.
- **Output.** A cold atomic beam at about 10–20 m/s with a flux of roughly 10^9–10^{10} atoms/s lands in the science cell. The narrow tube keeps the gas flow out of the science cell small.

Atoms with narrow cooling lines, such as Sr and Yb, usually need a **Zeeman slower** from a hot oven instead, because their vapor pressure at room temperature is far too low. For Rb, a 2D⁺ MOT is the norm.

3. Capturing atoms in the 3D MOT

In the science cell a **magneto-optical trap** sits right where the tweezer array will be. It captures the slow atomic beam and holds a dense cloud.

- **Cooling light.** Three orthogonal pairs of counter-propagating beams at 780 nm, tuned to the $F = 2 \rightarrow F' = 3$ cycling transition and red-detuned by about 2–3 natural linewidths ($\Gamma = 2\pi \times 6.07$ MHz). A moving atom sees the beam heading into Doppler-shifted closer to resonance, so it scatters more photons from that beam and slows down.
- **Magnetic gradient.** Anti-Helmholtz coils give about 10–20 G/cm. With circularly polarized beams of the right handedness, the Zeeman shift makes the scattering force point toward the field zero. That confines atoms in position as well as slowing them.
- **Repump.** Occasionally an atom is excited to $F' = 2$ instead of $F' = 3$, decays to $F = 1$, and goes dark to the cooling light. A weak repump beam on $F = 1 \rightarrow F' = 2$ returns it to the cycle. Without the repump the MOT empties in milliseconds.
- **Result.** 10^6–10^8 atoms in about 0.1–1 s, at a density around 10^{10}–10^{11} cm^{-3} and a temperature near the Doppler limit, $\hbar\Gamma/2k_B \approx 146$ μK. Bias coils shift the cloud so it overlaps the tweezer focal plane.



Rb-87 cooling transitions · D2 for the MOT and molasses, D1 for gray molasses

The cooling light sits slightly below F′ = 3, shown by the dashed tick under the level. The D1 gray-molasses light sits slightly above F′ = 2, with its Λ partner tuned to the same two-photon resonance.

How Doppler cooling works, with the energy budget

Numbers are for Rb-87 on the D2 cooling transition (F = 2 → F′ = 3, 780.24 nm).

Step by step

1. **Tune the laser a little below resonance.** The laser frequency is $\omega_L = \omega_0 + \delta$ with $\delta < 0$, typically $\delta \approx -\Gamma/2$ to -3Γ, where $\Gamma = 2\pi \times 6.07$ MHz. An atom at rest scatters only weakly, because the light is off resonance.
2. **Use the Doppler shift to pick a beam.** An atom moving at velocity v sees the beam it's heading into shifted up by kv, toward resonance, and the beam behind it shifted down, further away. So it absorbs far more photons from the beam in front. Here $k = 2\pi/\lambda = 8.05 \times 10^6$ m^{-1}, so 1 m/s gives a 1.28 MHz shift.
3. **Absorb and get kicked backward.** Each absorbed photon delivers momentum $\hbar k$ against the motion, changing the velocity by the recoil velocity $v_r = \hbar k/m = 5.9$ mm/s.
4. **Re-emit in a random direction.** About 26 ns later the atom emits spontaneously. The direction is random, so on average the emission returns no momentum: the absorption kicks add up, and the emission kicks average to zero.
5. **Repeat.** At up to $\Gamma/2 \approx 1.9 \times 10^7$ scatterings per second, the kicks act like a friction force, $F \approx -\beta v$. The maximum deceleration is $v_r\Gamma/2 \approx 1.1 \times 10^5$ m/s², about 11,000 g. Three counter-propagating pairs damp motion along all three axes.
6. **Random kicks heat the atom.** Emission recoils point in random directions, and the number of photons absorbed from each beam fluctuates. The atom's momentum random-walks, which deposits energy.

7. **Balance gives the Doppler limit.** Cooling falls as v shrinks, because a slower atom sees less Doppler asymmetry, while heating stays roughly constant. They balance at

$$k_B T_D = \frac{\hbar \Gamma}{2} \quad (\delta = -\Gamma/2) \;\Rightarrow\; T_D \approx 146\ \mu\mathrm{K}$$

That corresponds to about 12 cm/s rms velocity per axis.

Energy budget for one scattering event

For an atom moving into the beam at speed v, energy is conserved across the absorb–emit cycle:

$$\hbar\omega_L + K_{before} = \hbar\omega_{emitted} + K_{after}$$

On average the atom's kinetic energy changes by −ħkv from the absorption kick, +E_r from the absorption recoil, and +E_r from the emission recoil (zero momentum on average, but not zero energy). Here $E_r = \hbar^2k^2/2m = h \times 3.77$ kHz $= 2.5 \times 10^{-30}$ J, about 0.18 μK × k_B. So, the emitted photon comes out bluer than the absorbed one, and the extra energy ħkv leaves with the scattered light:

$$\hbar\omega_{emitted} \approx \hbar\omega_L + \hbar k v - 2E_r$$

Worked example, an atom moving at 1 m/s:

Quantity	Value
Laser photon energy	1.589 eV (384.23 THz, about 3 MHz below resonance at δ = −Γ/2)
Atom's kinetic energy	7.2×10^{-26} J
Removed per scattering, ħkv	h × 1.28 MHz = 8.5×10^{-28} J (about 61 μK × k_B)
Added per scattering, 2E_r	h × 7.5 kHz = 5.0×10^{-30} J
Net removed per scattering	about 1.2% of its kinetic energy
Emitted photon's frequency shift	about +1.27 MHz relative to the laser

The atom loses only about a millionth of the photon's energy on each event, but it scatters millions of photons per second.

Budget for stopping an atom

- **Photons needed.** Each scattering removes about one recoil velocity, so stopping from speed v takes about v/v_r photons: about 170 from 1 m/s, about 3,400–5,100 from a typical MOT capture speed of 20–30 m/s, and about 51,000 from thermal speeds (about 300 m/s). That's why the MOT captures only the slow tail of the vapor.
- **Distance and time.** At maximum deceleration, an atom at 30 m/s stops in about 4 mm and 0.3 ms. Real scattering runs below the maximum, so a few mm to a cm and about a millisecond is typical, comparable to a MOT beam size.
- **Damping time near rest.** With one saturation intensity per beam and δ = −Γ/2, the friction coefficient is $\beta \approx 0.44\,\hbar k^2 \approx 3 \times 10^{-21}$ kg/s, so the damping time m/β is about 50 μs.

Budget at the Doppler limit

In steady state the cooling power βv^2 equals the heating power, which is about $2E_r$ per scattering event times the scattering rate. An atom at 12 cm/s has $\hbar kv \approx h \times 150$ kHz removed per photon, only about 20× the $h \times 7.5$ kHz added. At this speed the scattering asymmetry between the two beams is small, so the net cooling per photon shrinks until it matches the random heating. That balance is the 146 µK floor. Getting colder takes the sub-Doppler mechanisms in Sections 4 and 5, which keep cooling at speeds where the Doppler asymmetry has faded.

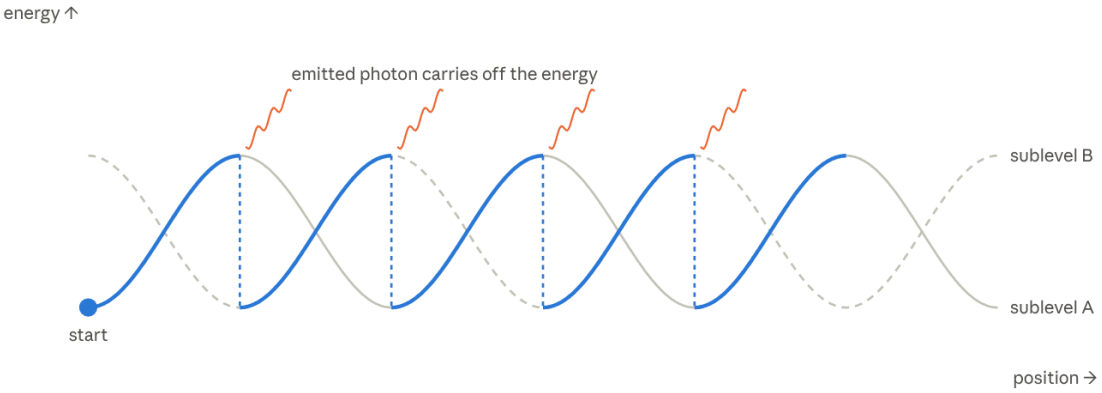
4. Optical molasses

What it is. Optical molasses is the MOT's six cooling beams with the magnetic gradient switched off. The red-detuned counter-propagating beams produce a friction force $F \approx -\beta v$, so any velocity gets damped, as if the atoms were moving through thick syrup. With no gradient there's no restoring force pulling atoms to a point, so the cloud diffuses slowly while it cools. That's fine for the few milliseconds it's needed. “Molasses light” means those beams used in this mode: the same 780 nm $F = 2 \rightarrow F' = 3$ light plus repump, with different settings than for the MOT.

Two cooling mechanisms work at once:

- **Doppler cooling.** The moving atom scatters more photons from the beam it's heading into. This stops at the Doppler limit, $\hbar\Gamma/2k_B \approx 146$ µK for Rb.
- **Polarization-gradient (Sisyphus) cooling.** This is the one that matters.
    a. Counter-propagating beams with different polarizations make a light field whose polarization changes over fractions of a wavelength. With crossed linear polarizations it cycles linear → σ^+ → linear → σ^-.
    b. Each ground-state sublevel (m_F) gets a different, position-dependent light shift. The atom moves through a corrugated landscape of hills and valleys about $\lambda/4$ apart, and the landscape differs by sublevel.
    c. Optical pumping is fastest where the atom's own sublevel sits high. So, the atom climbs a hill, losing kinetic energy, and near the top is pumped into a sublevel whose potential is a valley there. The emitted photon carries away the difference.
    d. Like Sisyphus pushing his boulder uphill forever, the atom keeps climbing, and each climb removes energy.

Sisyphus cooling: the atom keeps climbing, and each hilltop costs it energy



Climbing a hill turns the atom's kinetic energy into potential energy.
Near the hilltop, optical pumping drops it into the other sublevel's valley, and it climbs again.

polarization-gradient (Sisyphus) cooling · schematic

Settings and limits. The final temperature scales roughly as intensity divided by detuning, $T \propto I/|\Delta|$. The usual sequence:

1. Switch off the gradient.
2. Jump the detuning to about −10 to −20 Γ.
3. Ramp the intensity down over 5–20 ms.

That gives about 5–20 μK for Rb. The floor is set by photon recoil, about 0.36 μK per recoil for Rb, and in practice molasses stops at a few tens of recoils.

Imaging. The same light is used to image the atoms in the tweezers. They scatter plenty of photons for the camera while Sisyphus cooling keeps them from heating out of the traps.

How the Sisyphus cycle removes energy

The atom pays for each hill climb with its kinetic energy, and that energy then leaves with a scattered photon that has more energy than the photon the atom absorbed.

The setup. With red-detuned light, each ground sublevel is shifted down by an amount set by the local intensity and by how strongly that sublevel couples to the local polarization. In the textbook case of a ground state with m = ±½ in crossed linear beams:

- **σ^+ regions:** m = +½ couples about three times more strongly than m = −½, so it's shifted three times further down and sits in the valley.
- **σ^- regions:** the roles swap.
- **Optical pumping:** the same light pumps atoms toward the strongly coupled state. In a σ^+ region it pumps them into m = +½.

So, pumping always moves an atom into whichever sublevel is lowest at that spot. That is what makes the cycle work.

One cycle, as an energy ledger.

1. **Start in a valley.** The atom is in sublevel A at the bottom of its potential, moving with kinetic energy K.
2. **Climb.** Its potential energy rises toward the hilltop, a height U_o above the valley. Total energy is conserved on this leg, so its kinetic energy drops to $K - U_o$. Nothing irreversible has happened yet; the atom is simply slowing on the slope.
3. **Get pumped at the top.** Near the hilltop the local polarization favors sublevel B, which sits in its valley there. The atom absorbs a laser photon of energy $\hbar\omega_L$ and spontaneously emits one, ending in sublevel B. Its internal energy has dropped by U_o, so the emitted photon carries $\hbar\omega_L + U_o$: it is blue-shifted by U_o relative to the absorbed light. That's where the energy goes.
4. **Repeat.** The atom is back in a valley, now with kinetic energy $K - U_o$, and starts climbing the next hill on sublevel B.

Each cycle that ends with pumping near a hilltop removes about U_o. The emitted photon is only U_o/h higher in frequency, a few MHz on a 384 THz carrier, but the atom scatters thousands of photons per millisecond, so the losses add up quickly.

Why it removes energy instead of adding it. Pumping is random, with a rate γ_p that is highest where the atom's current sublevel is least favored, near its hilltop. Because pumping is most likely near hilltops and always lands the atom in the lower sublevel, energy is removed on average. This works best when the atom takes about one pumping time to cross a hill, moving about $\lambda/4$ in $1/\gamma_p$. Much faster atoms cross many hills between pumping events, the timing averages out, and they barely cool. That's why Sisyphus cooling acts on slow atoms (below a few tens of cm/s) and comes after Doppler cooling.

Why this lowers the temperature. For a gas, temperature measures the velocity spread: the mean kinetic energy per atom is $(3/2)\, k_B T$. Sisyphus cooling acts on each atom independently, with no collisions needed, so the velocity distribution narrows and the temperature falls. Spontaneous emission is the essential irreversible step: the scattered photons go off in random directions into many empty modes of the field, carrying away the energy and entropy. Without it, energy the light took out could flow back in, and there would be no net cooling.

What sets the final temperature. Cooling competes with heating from the same scattering. Each absorbed and emitted photon gives a random kick of $\hbar k$, and the dipole force fluctuates as the atom jumps between sublevels.

- **The equilibrium:** kinetic energy settles at roughly the hill height, $k_B T \sim U_o$. Atoms with less energy than that get trapped in valleys, and cooling becomes ineffective.
- **The scaling:** U_o is proportional to intensity divided by detuning, so $T \propto I/|\Delta|$. That's why the sequence jumps the detuning and ramps the intensity down.
- **The floor:** as U_o approaches the recoil energy, a single photon kick (about 0.36 μK for Rb) undoes a cycle's cooling. Molasses bottoms out at a few tens of recoils, the 5–20 μK range for Rb.

Comparison with Doppler cooling. Doppler cooling slows the atom by about one photon momentum per scattering event, because it absorbs preferentially from the beam it's moving into. The force needs the Doppler shift kv to be comparable to the linewidth Γ, so it fades for slow atoms and stops at about

146 μK. Sisyphus cooling relies instead on the light-shift landscape and the optical-pumping time, both of which can be much slower than 1/Γ. It keeps working where Doppler cooling has faded, giving roughly another factor of 10 in temperature.

5. Gray molasses

What it is. Gray molasses is a sub-Doppler cooling scheme that uses dark states. For Rb it runs on the D1 line at 795 nm, usually on F = 2 → F′ = 2, with the light **blue**-detuned instead of red-detuned.

How it works:

1. **Dark states.** On an F → F′ = F transition, for any light polarization there's a combination of ground sublevels that can't absorb the light at all.
2. **Blue detuning.** The light-shifted bright states sit *above* the dark state, and their energy rises and falls with position in the polarization pattern.
3. **Climb and drop.** An atom in a bright state climbs a hill and is optically pumped into the dark state near the top. That removes energy, as in ordinary Sisyphus cooling.
4. **Return at the bottom.** While it's dark, the atom's motion through the changing polarization slowly mixes it back into a bright state. That's most likely where the bright state comes closest in energy to the dark state: at the bottom of a valley. The atom re-enters there and climbs again.
5. **Velocity selection.** The mixing back is faster for faster atoms. Slow atoms stay dark longer, so they pile up cold and scatter almost no light. That's the "gray": mostly dark, but not completely.

Why it's useful:

- **Colder.** A few μK for Rb, below ordinary molasses, because the cold atoms barely scatter.
- **Works at high density.** Less scattering means less heating from photons re-absorbed inside a dense cloud.
- **Λ-enhanced version.** A second frequency on F = 1 → F′ = 2, tuned to exact two-photon (Raman) resonance with the first, creates extra dark states between the two hyperfine levels. That cools a little further and stops atoms leaking into F = 1.
- **Works inside tweezers.** Blue-detuned light is much less sensitive to the tweezer light shift than red molasses. It's used to cool atoms in the traps and for enhanced loading (Section 6).

Why D1 and not D2. D1's two excited hyperfine levels are about 812 MHz apart. On D2, the nearby F′ = 3 level, only about 267 MHz from F′ = 2, would keep scattering photons and spoil the dark states.

	Optical molasses (bright)	Gray molasses
Line	D2, 780 nm, F = 2 → F′ = 3	D1, 795 nm, F = 2 → F′ = 2 (+ F = 1 → F′ = 2 for Λ)
Detuning	Red, about −10 to −20 Γ	Blue, a few to about 10 Γ
Mechanism	Sisyphus between bright sublevels	Sisyphus plus velocity-selective dark states
Typical Rb temperature	5–20 μK	2–5 μK

	Optical molasses (bright)	Gray molasses
Scattering	High, good for imaging	Low, good for dense clouds and tweezers
Main uses	Post-MOT cooling, cooling during imaging	Deeper cooling, in-tweezer cooling, enhanced loading

6. Loading the tweezers from the cold reservoir

The cold molasses cloud is the reservoir. The tweezers are already on during the MOT and molasses stages, so an atom passing through a trap is cooled inside it and stays.

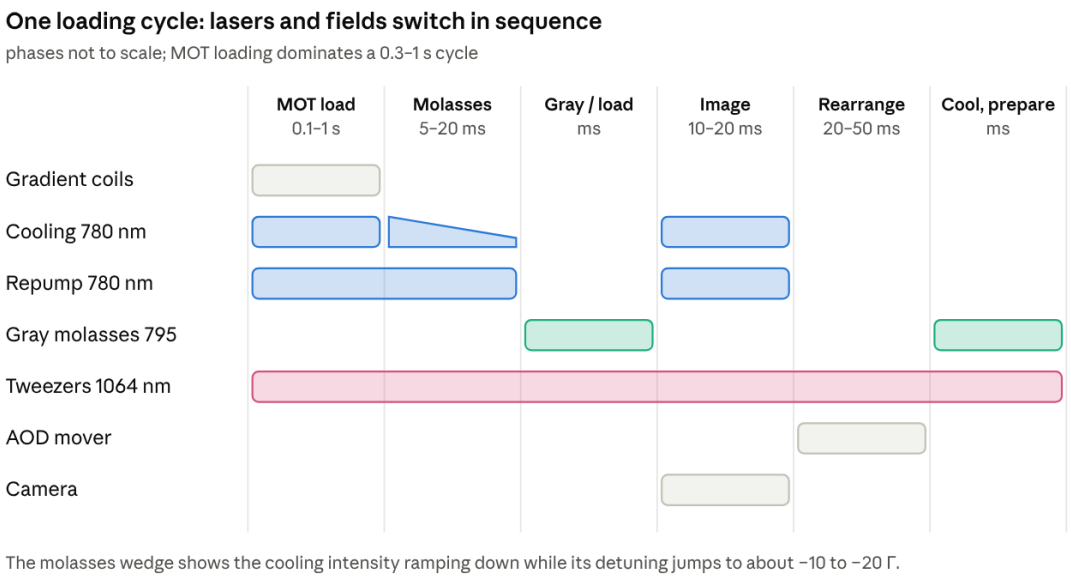
- **Collisional blockade.** Each tweezer is so small that a second atom triggers a light-assisted collision. The pair is excited together, gains kinetic energy, and both atoms leave. Each site ends up with 0 or 1 atom, at about 50–60% occupancy.
- **Enhanced loading.** Tuning the loading light blue of resonance (for example with gray molasses) makes these collisions release only enough energy to eject one atom. Occupancy rises to about 80–90%.
- **Light shift.** The deep tweezer light shifts the atomic resonance inside each trap. The loading and imaging detunings have to be chosen with the in-trap shift in mind, which is one reason trap-depth uniformity across the array matters.
- **Depth vs temperature.** Keep the trap depth at least about 10× the atom temperature. Atoms at 10–20 μK in a 0.65 mK trap give a ratio of about 30–60, which is comfortable.

7. After loading, and continuous reservoirs

Loading leaves a random half-to-mostly-filled array. The cycle then continues:

1. **Image.** Take a fluorescence image with molasses light on, so the atoms stay cold while scattering photons, to find which sites are filled.
2. **Rearrange.** Move atoms into a defect-free array with the AOD tweezers, then re-image to check.
3. **Cool to the motional ground state.** Raman sideband cooling or in-trap gray molasses, aiming for about 1–5 μK and mostly the ground state of the radial motion.
4. **Prepare the qubits.** Optically pump into the qubit state, such as the $|F = 1, m_F = 0\rangle$ clock state.

A full cycle takes about 0.3–1 s, and MOT loading is usually the longest part.



timing of one cycle · phases **not** to scale

Continuous reservoirs for large arrays. For 1,000+ atoms, newer systems don't discard the MOT each cycle. They keep atoms in a separate reservoir off to the side: a MOT, or a large optical lattice or dipole trap. Atoms are carried into the array region by a moving optical lattice (a “conveyor”) or by tweezers, so lost atoms can be replaced while the remaining qubits keep their states.

8. Laser requirements

Function	Wavelength	Transition	Notes
2D MOT, push beam, 3D MOT, molasses, imaging	780 nm	D2, F = 2 → F′ = 3	One laser, split and frequency-shifted with AOMs to set each detuning
Repump	780 nm	D2, F = 1 → F′ = 2	About 6.8 GHz from the cooling light; a separate laser or an EOM sideband
Gray molasses (optional)	795 nm	D1, F = 2 → F′ = 2	Blue-detuned; AOM for switching and detuning
Λ partner (optional)	795 nm	D1, F = 1 → F′ = 2	EOM sideband or second AOM at two-photon resonance
Tweezers	1064 nm	far off resonance	High-power fiber amplifier feeding the SLM, plus the AOD mover path

The frequency-setting lasers are locked to Rb saturated-absorption spectroscopy or to a ULE cavity. AOMs set the detunings and switch the beams on the timescales in the timing chart. Many tweezer labs consider the 795 nm D1 laser worth adding for the colder temperatures and higher loading probability.

Appendix: trap depth units and key numbers

Trap depth in kelvin. A trap depth quoted in mK is the potential energy divided by Boltzmann's constant, so the energy is

$$U_0 = k_B T_{depth}$$

with no factor of 3/2. For a 0.645 mK trap, U_o = (1.380649 × 10^{-23} J/K)(0.645 × 10^{-3} K) ≈ 8.9 × 10^{-27} J, the same as h × 13.4 MHz. Quoting depth as a temperature or as a frequency (U_o/h) is purely a unit convention.

The factor 3/2 belongs to the **atoms**: the mean kinetic energy of a gas at temperature T is (3/2) k_B T. In a 3D harmonic trap the mean total energy is 3 k_B T, half kinetic and half potential. That is the quantity to compare with the trap depth.

Key Rb-87 numbers used in this document

Quantity	Value
D2 line (cooling, repump)	780.24 nm
D1 line (gray molasses)	794.98 nm
Natural linewidth Γ (D2)	2π × 6.07 MHz
Ground hyperfine splitting	6.835 GHz
D2 excited splittings F′ = 3–2, 2–1	~267 MHz, ~157 MHz
D1 excited splitting F′ = 2–1	~812 MHz
Doppler limit ħΓ/2k_B	~146 μK
Recoil temperature	~0.36 μK
Vapor pressure at 25 °C	~4 × 10^{-7} Torr

These are standard reference values (for example, D. A. Steck, *Rubidium 87 D Line Data*), quoted here from memory. Check them against the reference before using them in calculations.