Atom Computing Alkaline-Earth System: Layout, Atoms and Nuclear-Spin Qubits

Dimensioned drawings, ^{87}Sr and ^{171}Yb atomic structure, and how the nuclear-spin qubit is initialized, gated and read out

Technical note · October 25, 2026 · D. Hakala

What this note shows. A companion to the Harvard–QuEra layout note, built from the Atom Computing whitepaper in your Neutral_atom_QC_companies folder and the five papers it cites. It gives (1) a side elevation of the two-chamber ^{171}Yb system with a laser key, (2) the atom-plane layout in micrometres, (3) level diagrams for ^{87}Sr and ^{171}Yb, and (4) the nuclear-spin qubit operations. Your guess is right: Atom Computing's first machine used strontium-87 and its second-generation machines use ytterbium-171. Numbers from papers are marked as published; chamber geometry, beam routing and cavity lengths are engineering estimates.

1. The company and its hardware lineage

Atom Computing was founded in 2018 by Ben Bloom and Jonathan King in Berkeley, California. Unlike QuEra, which encodes qubits in hyperfine states of the alkali atom ^{87}Rb, Atom Computing has always used **alkaline-earth-like atoms** and stores each qubit in the **nuclear spin** of the ground state. The first machine ("Phoenix", about 100 qubits) used ^{87}Sr; the second generation, with more than 1,200 trapping sites, uses ^{171}Yb and has been the basis of the joint logical-qubit work with Microsoft.

Paper	Atom	Traps	Key hardware result
Barnes et al., Nat. Commun. 13, 2779 (2022)	^{87}Sr	813.4 nm, LCoS SLM (WGS); 110 sites at 4 µm; NA 0.65	Nuclear-spin qubits with $T_2^* \approx 21$ s and T_2(echo) ≈ 42 s; 21-qubit register
Norcia et al., PRX 13, 041034 (2023)	^{171}Yb	483 nm tweezers, 532 nm transport lattice	Mid-circuit measurement with 459.6 nm "hiding" light; Ramsey contrast kept at 97.2 %
Norcia et al., PRX Quantum 5, 030316 (2024)	^{171}Yb	459.6 + 423.3 nm target array; 483 nm reservoir; 783.8 nm cavity lattices	1225-site array filled iteratively to 99 %; 99.92 % imaging fidelity
Muniz et al., arXiv:2411.11708 (2024)	^{171}Yb	460 nm science, 483 nm reservoir	Clock-shelved Rydberg CZ: 99.40 % raw, 99.72 % with post-selection; 1Q 99.963 %
Reichardt et al., arXiv:2411.11822 (2024)	^{171}Yb	register + interaction zone; AOD transport	24 logical qubits entangled ($[\![4,2,2]\!]$); 28 logical qubits in Bernstein–Vazirani; loss correction
Atom Computing whitepaper (2025)	^{171}Yb	multi-chamber, reloadable reservoir, cavities	Roadmap; tens-of-seconds coherence; 64-logical-qubit architecture

2. System layout

The whitepaper describes a multi-chamber design, and Norcia 2024 gives the details. Figure 1 draws it as a side elevation:

1. **Chamber 1 (source and MOT).** An oven and a 399 nm Zeeman slower feed a core–shell magneto-optical trap: a 556 nm core on the narrow $^1S_0$–$^3P_1$ line and a 399 nm shell on the broad $^1S_0$–$^1P_1$ line. Atoms are Doppler-cooled to about 10 µK.
2. **The optical elevator.** Two 532 nm beams form a vertical standing wave. Chirping the frequency of one beam makes the lattice slide, carrying atoms about 30 cm into the science chamber in 100 ms (75 µm beam diameter, about 30 µK). This keeps the hot oven and MOT gas away from the qubits, which is how the vacuum-limited lifetime reaches about 30 s.
3. **Chamber 2 (science).** Two cavity-enhanced 783.8 nm optical lattices hold the atoms for loading and imaging. The XY cavity lies in the atom plane (finesse 2900, waist 268 µm) and the Z cavity is tilted 15° out of it (finesse 3000, waist 183 µm). The cavity build-up reduces the required laser power by 255× and 132× compared with tweezers of the same depth, which makes deep (330 µK and 260 µK) imaging traps practical over the whole array.
4. **Tweezers.** Target tweezers at 459.6 nm and 423.3 nm (about 500 nm waist, about 50 µK, 100–200 µW each) form the computational array. A 483 nm reservoir array (about 250 µK) is shifted 170 µm in 10 ms by a galvo mirror, and an AOD-steered moving tweezer rearranges atoms. All of this light, and the 556 nm fluorescence going to the camera, passes through an NA 0.65 objective with a 0.5 mm field of view.
5. **Global beams.** 556 nm (imaging, optical pumping, Raman), 578 nm (clock shelving) and 302 nm (Rydberg UV), with a bias field of about 500 G (518 G in Muniz 2024).

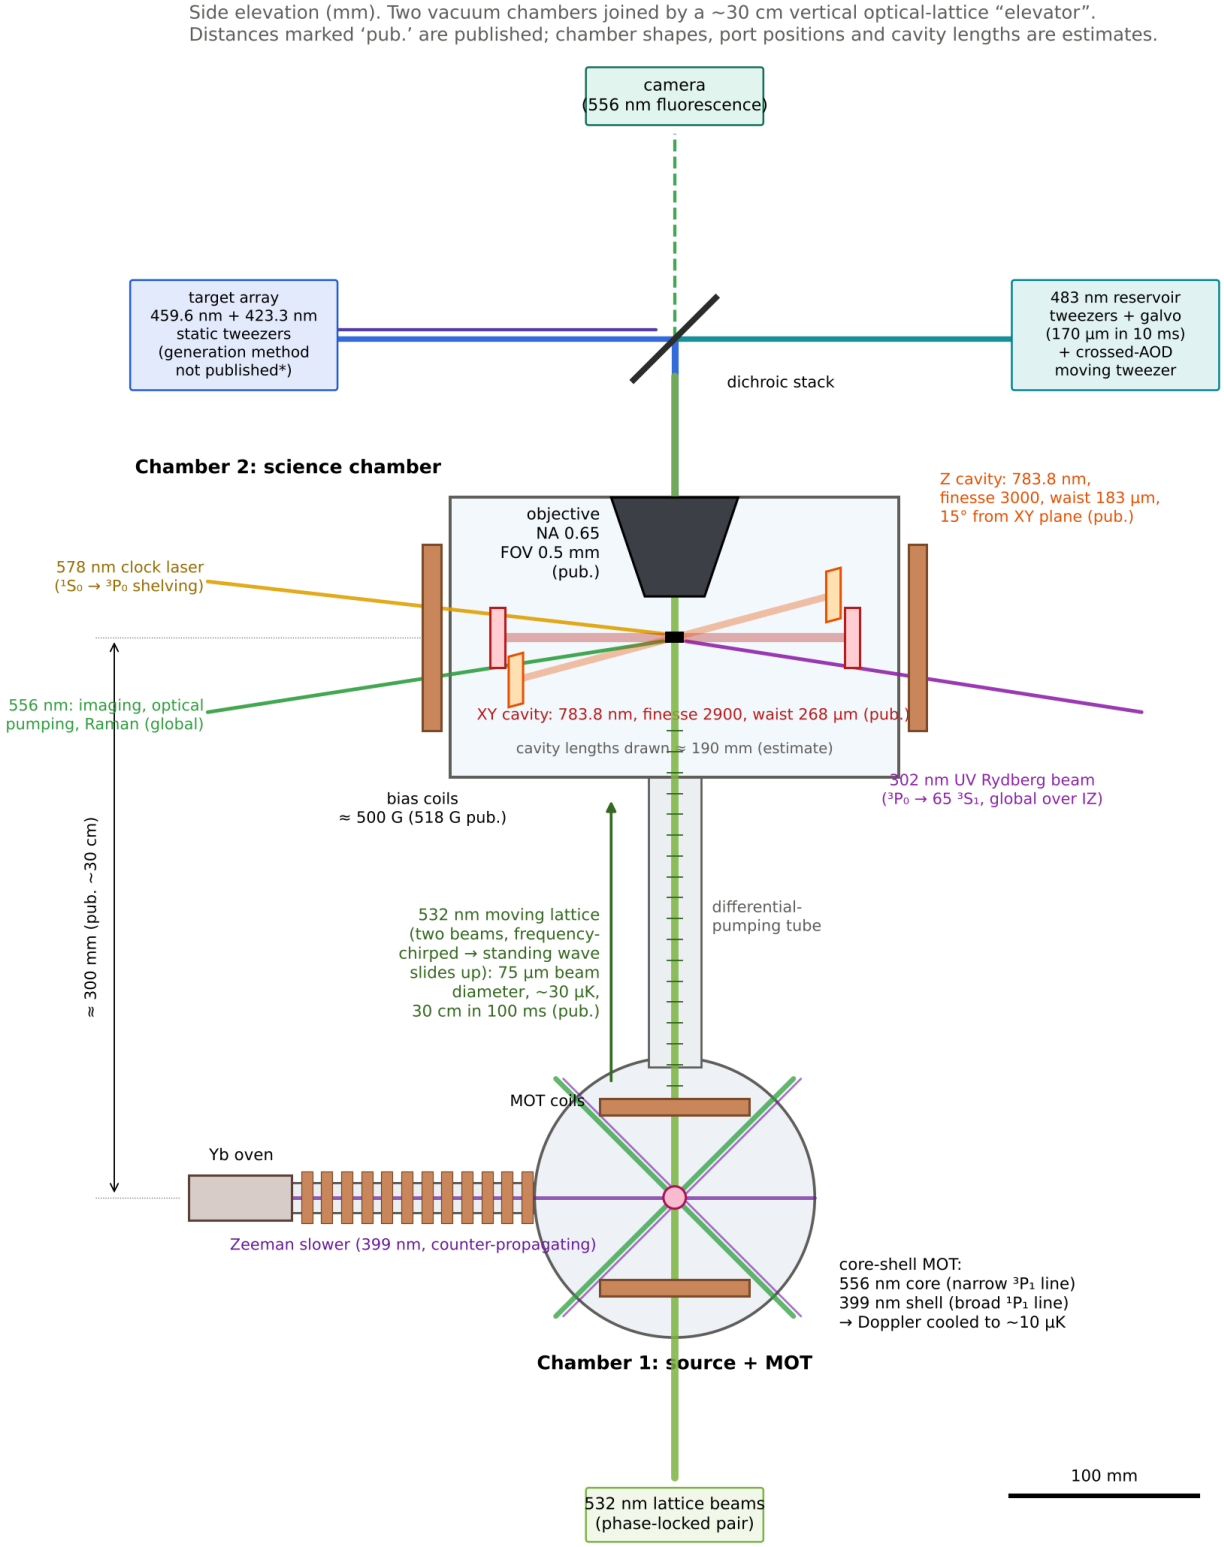


Figure 1. Side elevation of the second-generation ^{171}Yb system with laser key. The ~30 cm elevator, the cavity finesses, waists and tilt, the NA, the field of view and the bias field are published; chamber shapes, port positions and the drawn cavity length (~190 mm) are estimates.

Consistency check: why build cavities at all?

Each target tweezer needs about 100–200 µW for a 50 µK trap. Imaging needs traps several times deeper (hundreds of µK), so an array of 1,225 deep tweezers would need on the order of a watt of blue light at the atoms after the losses of an SLM or AOD path. A lattice whose intensity is built up 132–255× by a cavity provides deep traps over a 180–270 µm waist from a few watts of near-infrared light, so the tweezers only have to define positions and hold the atoms gently during gates.

3. Atom-plane layout

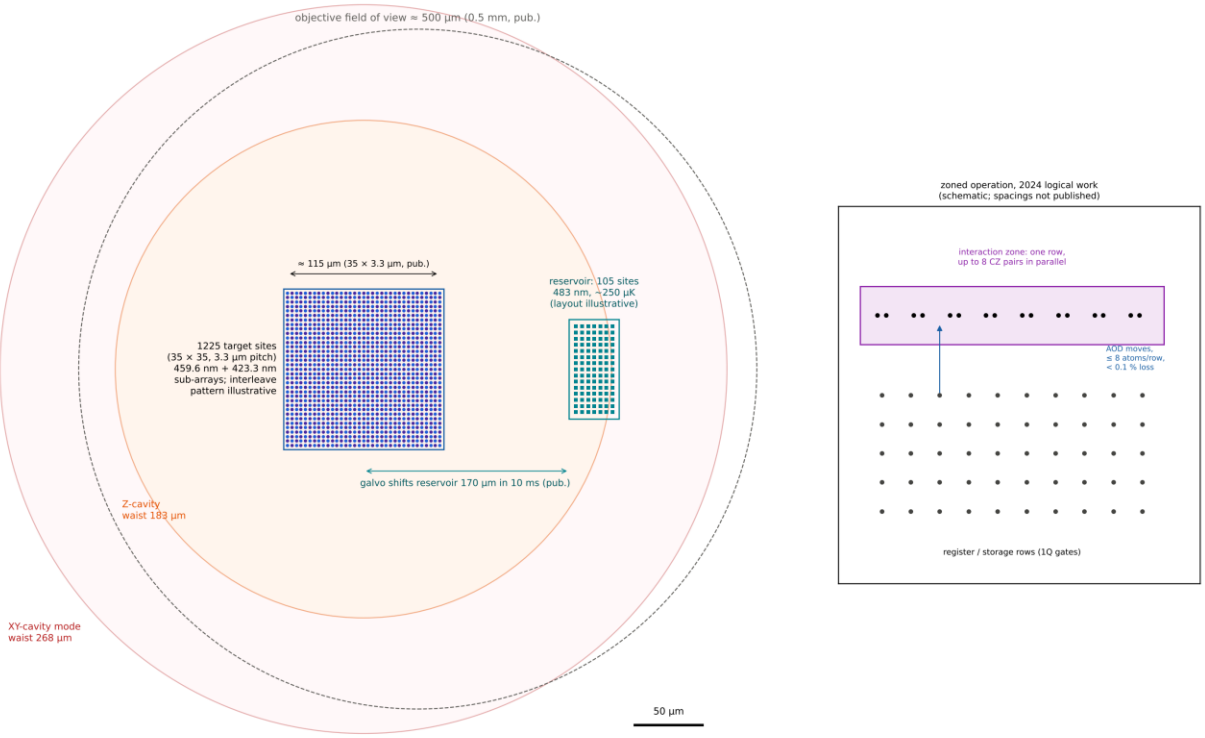


Figure 2. Atom plane in micrometers. The 35 × 35 target array at 3.3 µm pitch, the 105-site reservoir and its 170 µm galvo shift, the cavity waists and the field of view are published. The interleave pattern of the two target sub-arrays and the reservoir arrangement are illustrative. Right: zoned operation in the 2024 logical work (schematic).

Quantity	Published value	Source
Target array	1225 sites (35 × 35), 3.3 µm pitch, ≈ 115 µm square; two sub-arrays at 459.6 and 423.3 nm	Norcia 2024
Reservoir	105 sites at 483 nm, ≈ 250 µK; moved 170 µm in 10 ms by galvo	Norcia 2024
Filling	98 % before rearrangement, 99 % after; loading cycle ≈ 300 ms	Norcia 2024
Cavity modes	XY waist 268 µm; Z waist 183 µm, 15° from the atom plane	Norcia 2024
Field of view	0.5 mm, NA 0.65	Norcia 2024
Interaction zone	one row; up to 8 CZ pairs in parallel; pair spacing inside blockade radius	Muniz 2024, Reichardt 2024
Transport	AOD tweezers move up to 8 atoms of a row in parallel, paths may not cross; < 0.1 % loss per move	Reichardt 2024

Quantity	Published value	Source
Processor size used in logical work	up to 256 atoms available for computation	Reichardt 2024

4. The alkaline-earth atoms: ^{87}Sr and ^{171}Yb

Strontium (group 2) and ytterbium (a lanthanide with a filled 4f shell) both have **two valence electrons**. That gives them a spectrum unlike rubidium's and four properties Atom Computing builds on:

- **Singlet and triplet families.** The ground state is $^1S_0$ (total electronic angular momentum J = 0). There is a broad singlet transition ($^1S_0$–$^1P_1$, about 30 MHz wide) for fast slowing and trapping, and intercombination lines to the triplet ^3P states that are narrow because spin is nearly conserved.
- **A narrow cooling and imaging line.** $^1S_0$–$^3P_1$ is 7.4 kHz wide in Sr and 182 kHz in Yb. Doppler cooling on it reaches microkelvin temperatures, and in Yb it is broad enough to image atoms in a few milliseconds.
- **A metastable clock state.** $^3P_0$ lives for many seconds and is reached by an ultranarrow "clock" laser (698 nm in Sr, 578 nm in Yb). It acts as a shelf that the laser can move one qubit state into, and it is the launch point for single-photon UV Rydberg excitation.
- **A nuclear spin in a J = 0 ground state.** In the fermionic isotopes (^{87}Sr, I = 9/2; ^{171}Yb, I = 1/2) the nuclear spin is the only angular momentum in $^1S_0$. With no electron spin attached, it couples to magnetic fields thousands of times more weakly than an alkali hyperfine qubit and is nearly immune to the differential light shifts of the trapping light.

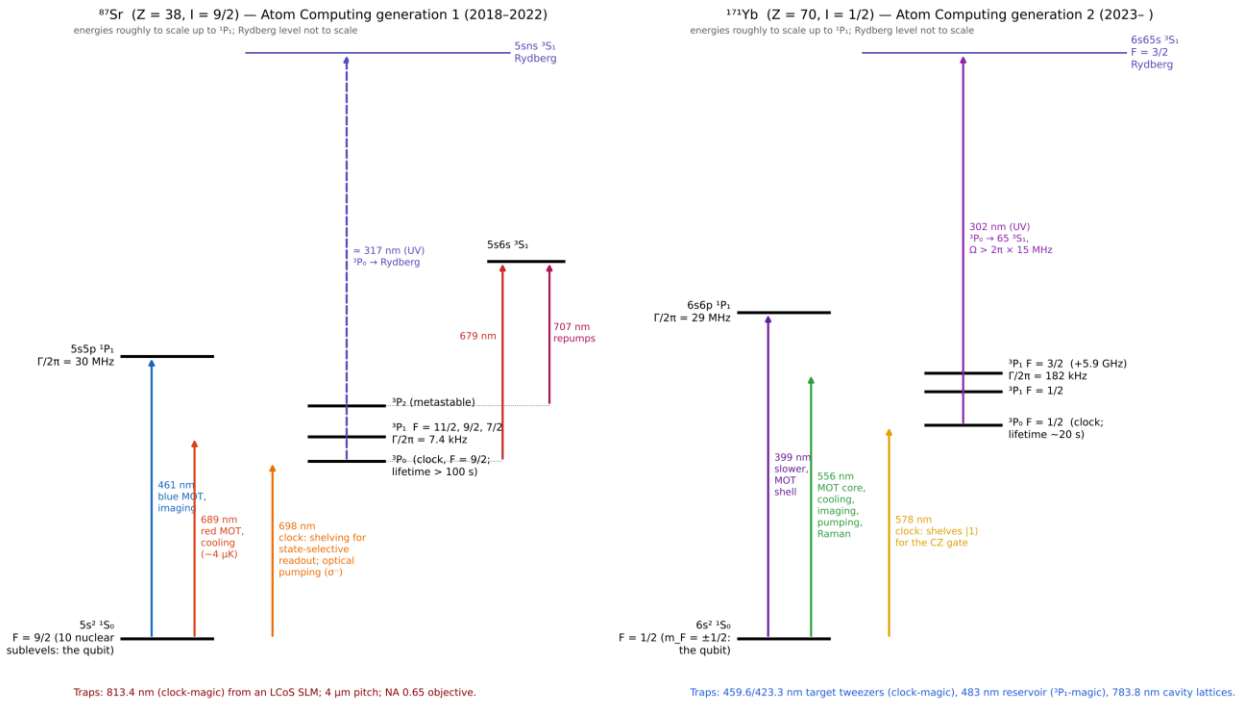


Figure 3. Level diagrams for ^{87}Sr (Atom Computing generation 1) and ^{171}Yb (generation 2), with the wavelengths each machine uses. Energies are roughly to scale up to $^1P_1$; the Rydberg levels are not.

Property	^{87}Sr	^{171}Yb
Natural abundance	7.0 %	14.3 %

Property	^{87}Sr	^{171}Yb
Nuclear spin I	9/2 (10 ground sublevels)	1/2 (exactly 2 ground sublevels)
Nuclear magnetic moment	−1.093 μ_N	+0.492 μ_N
Qubit splitting per gauss, eq. (1)	≈ 185 Hz/G between adjacent m_F	≈ 750 Hz/G
Broad line $^1S_0$–$^1P_1$	461 nm, 30 MHz	399 nm, 29 MHz
Narrow line $^1S_0$–$^3P_1$	689 nm, 7.4 kHz	556 nm, 182 kHz
Clock line $^1S_0$–$^3P_0$	698 nm	578 nm
Trap wavelengths used	813.4 nm (clock-magic)	459.6 / 423.3 nm (clock-magic), 483 nm ($^3P_1$-magic), 783.8 nm lattice
Rydberg route used	UV from $^3P_0$ (≈ 317 nm) in the wider Sr literature	302 nm from $^3P_0$ to 6s65s $^3S_1$, F = 3/2

The qubit splitting is set by the nuclear moment alone:

$$\Delta\nu_q \;=\; \frac{\mu_I}{I\,h}\,B \;=\; \gamma_I\,B \qquad (1)$$

where μ_I is the nuclear magnetic moment, I the nuclear spin and γ_I the resulting splitting per gauss. At 518 G this gives the published 388.9 kHz for ^{171}Yb. An alkali clock qubit, by contrast, sits on a GHz hyperfine splitting whose sensitivity is set by the electron.

Why the move from Sr to Yb?

The papers do not give the company's reasons, but the physics points to several. ^{171}Yb's I = 1/2 gives a ready-made two-level qubit with no leakage states, where ^{87}Sr needed an extra Stark-shift beam to isolate two of its ten sublevels. Yb's 182 kHz $^3P_1$ line images 25× faster than Sr's 7.4 kHz line, which matters for mid-circuit measurement. And Yb's 399 nm, 556 nm and 578 nm lines and its near-UV tweezer wavelengths allow traps that are magic for exactly the transition in use at each stage.

5. The nuclear-spin qubit: initialization and manipulation

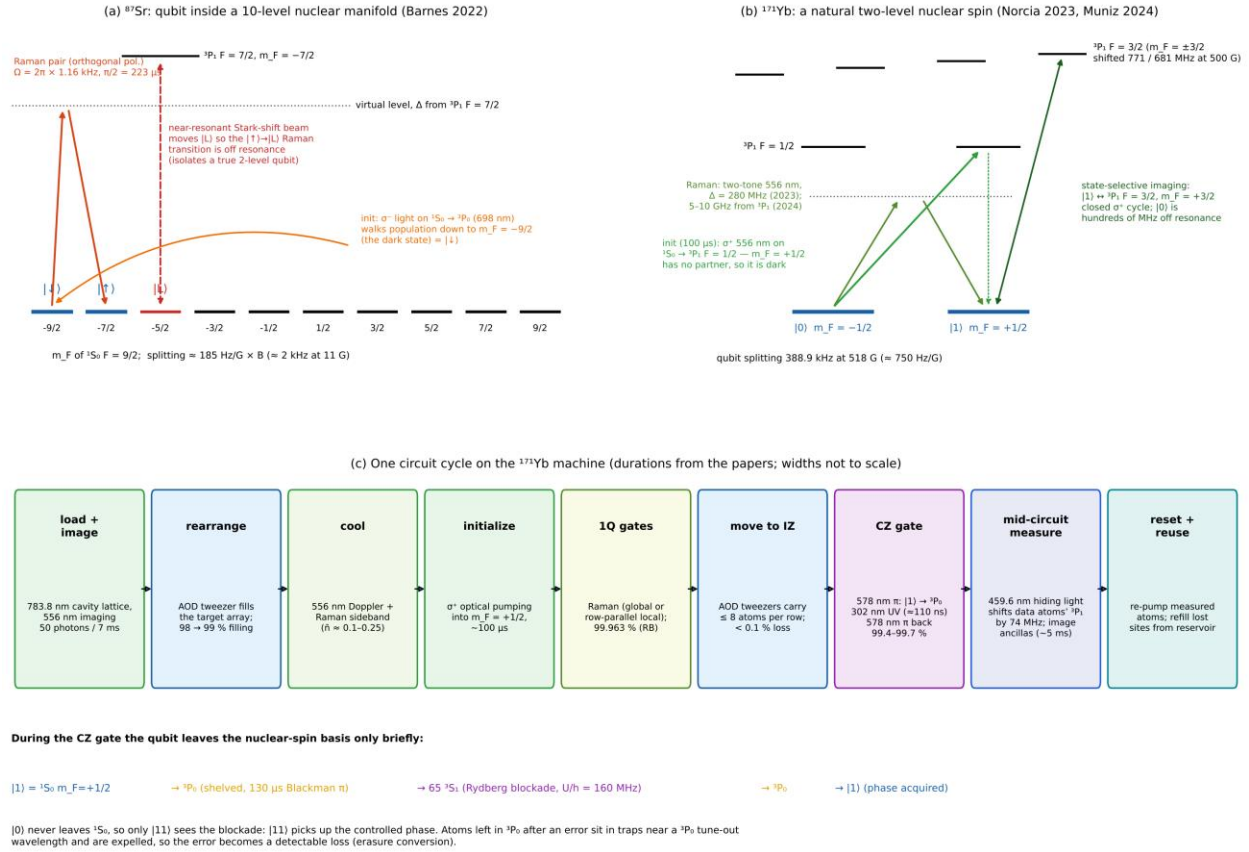


Figure 4. (a) ^{87}Sr qubit inside its ten-level ground manifold. (b) ^{171}Yb qubit with the transitions used for pumping, Raman gates and state-selective imaging. (c) One circuit cycle on the ^{171}Yb machine.

5.1 Generation 1: ^{87}Sr (Barnes 2022)

- **Encoding.** $|\downarrow\rangle = {}^1S_0\ |F = 9/2, m_F = -9/2\rangle$ and $|\uparrow\rangle = |m_F = -7/2\rangle$, at an 11 G quantization field.
- **Initialization.** σ^--polarized light near the ${}^1S_0 \rightarrow {}^3P_0$ transitions moves population down the manifold until it reaches $m_F = -9/2$, which has no σ^- partner and so is dark. Atoms are first cooled on 689 nm (red MOT, then Sisyphus cooling) to about 4 µK.
- **Isolating two of ten levels.** Because every adjacent pair of sublevels has the same Zeeman splitting, a Raman drive tuned to $|\downarrow\rangle \leftrightarrow |\uparrow\rangle$ also drives $|\uparrow\rangle \leftrightarrow |L\rangle = m_F = -5/2$. A near-resonant Stark-shift beam on $|L\rangle \rightarrow {}^3P_1\ F = 7/2, m_F = -7/2$ shifts $|L\rangle$ out of resonance while scattering only about 1 Hz from the qubit states.
- **Gates.** Two orthogonally polarized Raman beams detuned from ${}^3P_1\ F = 7/2$, shaped by EOMs and steered by crossed AODs: Rabi frequency $2\pi \times 1.16$ kHz ($\pi/2$ in 223 µs), up to 7 qubits per column driven at once.
- **Readout.** A clock π pulse shelves $|\downarrow\rangle$ in ${}^3P_0$; 461 nm fluorescence then shows only atoms left in $|\uparrow\rangle$, and a second image after repumping separates state from loss.
- **Coherence.** $T_1 \gg 5$ s, $T_2^* = 21 \pm 7$ s, $T_2(\text{echo}) = 42 \pm 6$ s, more than 10^5 times longer than a gate.

5.2 Generation 2: ^{171}Yb (Norcia 2023, Muniz 2024, Reichardt 2024)

Encoding. $|0\rangle = {}^1S_0$ m_F = −1/2 and $|1\rangle$ = m_F = +1/2, separated by 388.9 kHz at 518 G. There are no other ground-state sublevels.

Initialization. σ^+ light at 556 nm on ${}^1S_0 \rightarrow {}^3P_1$ F = 1/2 excites m_F = −1/2 to the excited m_F = +1/2 level. The ground m_F = +1/2 state would need an excited m_F = +3/2 level, which F = 1/2 lacks, so it is dark and population collects in $|1\rangle$ in about 100 µs. Before this, atoms are Doppler-cooled on 556 nm and Raman-sideband or gray-molasses cooled to $\bar{n} \approx 0.1$–0.25 motional quanta.

Single-qubit gates. A two-photon Raman transition through a virtual level detuned by Δ from ${}^3P_1$ drives $|0\rangle \leftrightarrow |1\rangle$:

$$\Omega_R = \frac{\Omega_1 \Omega_2}{2\Delta} \tag{2}$$

$$\frac{\Omega_R}{\Gamma_{sc}} \approx \frac{\Delta}{\Gamma} \qquad (\Omega_1 = \Omega_2) \tag{3}$$

Here Ω_1 and Ω_2 are the single-beam Rabi frequencies, Γ is the ${}^3P_1$ linewidth ($2\pi \times 182$ kHz) and Γ_sc the photon-scattering rate. Equation (3) says more detuning buys more gates per scattered photon. Atom Computing raised Δ from 280 MHz (2023, $\Omega_R \approx 2\pi \times 1.4$ kHz) to 5 GHz (Muniz 2024, $\Omega_R \approx 2\pi \times 7$ kHz) and about 10 GHz (Reichardt 2024), where $\Delta/\Gamma \approx 3 \times 10^4$. Z rotations are done virtually, as phase updates in the control software. Gates are applied globally or row-by-row to selected sites. Measured Clifford fidelity: 99.963(2) % (Muniz) and 99.85(2) % in the full logical-circuit setting (Reichardt).

Two-qubit CZ gate. The nuclear spin itself does not interact with neighbors, so the gate briefly converts it into an electronic state:

1. A 578 nm clock π pulse (130 µs Blackman envelope, $\Omega \approx 2\pi \times 7$ kHz) shelves $|1\rangle$ into ${}^3P_0$. $|0\rangle$ stays in ${}^1S_0$ because the pulse is selective for $|1\rangle$ (polarization and the Zeeman shift at 518 G). Shelving fidelity: 99.86 % per atom.
2. A 302 nm σ^- UV pulse (about 110 ns) drives ${}^3P_0 \rightarrow 6s65s\ {}^3S_1$ F = 3/2, m_F = −3/2 with $\Omega_{Ryd} > 2\pi \times 15$ MHz. Two atoms closer than the blockade radius cannot both be excited (U/h = 160 MHz), so only $|11\rangle$ acquires a different phase.
3. A second clock π pulse returns ${}^3P_0$ to $|1\rangle$.

A rough scale for the error caused by a finite blockade is

$$\epsilon_{\text{block}} \sim \left(\frac{\Omega_{\text{Ryd}}}{2U}\right)^2 \tag{4}$$

which is about 2×10^{-3} for 15 MHz and 160 MHz. Optimized pulse shapes compensate much of this, and the measured CZ fidelity is 99.40 % raw and 99.72 % after discarding runs with detected loss. The Rydberg lifetime is 65 µs, so the 110 ns pulse contributes little decay.

Erasure conversion. The storage tweezers sit near a ${}^3P_0$ tune-out wavelength, where the trap holds ${}^3P_0$ atoms only weakly. An atom left in ${}^3P_0$ by a failed shelving or Rydberg step is therefore lost rather than silently flipped. Its loss shows up in the next image, so the error is located (an erasure), which the ⟦4,2,2⟧ and ⟦9,1,3⟧ codes in Reichardt 2024 exploit.

Mid-circuit measurement and reset. To read ancillas without disturbing data qubits, 459.6 nm "hiding" light is applied to the data sites. It shifts their ${}^3P_1$ level by 74 MHz (spread 6 MHz), so the 556 nm imaging light no longer scatters from them. The ancillas are imaged state-selectively: $|1\rangle$ cycles on the closed σ^+ transition to

$^3P_1$ F = 3/2, m_F = +3/2, while |0⟩ is hundreds of MHz off resonance at 500 G (the m_F = ±3/2 levels shift by 771 and 681 MHz). About 30 photons are collected in 5 ms, and a 100 µs optical-pumping pulse then resets the measured atoms for reuse. Data-qubit Ramsey contrast falls only from 98.5 % to 97.2 % through a measurement. In the second generation, readout moves into the cavity-enhanced lattices, where state-resolved readout loses and errs at the few-parts-per-thousand level.

^{171}Yb performance figure	Value	Source
Single-qubit Clifford fidelity	99.963(2) %; 99.85(2) % in logical circuits	Muniz 2024; Reichardt 2024
CZ fidelity	99.40 % raw, 99.72 % post-selected; 99.6 % (0.4 % error) with moves	Muniz 2024; Reichardt 2024
Clock shelving	99.86 % per atom per pulse	Muniz 2024
Imaging	99.92 % fidelity, 50 photons in 7 ms	Norcia 2024
Mid-circuit measurement	~30 photons in 5 ms; data Ramsey contrast 98.5 → 97.2 %	Norcia 2023
Coherence	T_2^* = 430 ms (field-gradient limited), 3.3 s with echo; whitepaper: tens of seconds	Norcia 2023; whitepaper 2025
Atom loss per AOD move	< 0.1 %	Reichardt 2024

6. Comparison with the Harvard–QuEra architecture

Feature	Harvard / QuEra (^{87}Rb)	Atom Computing (^{171}Yb)
Qubit	Hyperfine clock states, 6.8 GHz	Nuclear spin m_F = ±1/2, ~389 kHz
Single-qubit gates	795 nm Raman, local 2D-AOD spot + global	556 nm Raman, global or row-parallel local
Rydberg excitation	Two-photon 420 + 1013 nm from ground	Clock shelve (578 nm) then single-photon 302 nm UV
Static traps	852 nm, LCoS SLM	459.6 / 423.3 nm target array; 783.8 nm cavity lattices
Moving traps	Crossed DTSX-400 AOD pair	AOD moving tweezer; galvo for reservoir
Atom delivery	2D$^+$ MOT into a glass cell	Second chamber fed by a 30 cm moving lattice
Readout	780 nm imaging; spin-to-position in a 795 nm lattice	556 nm state-selective imaging with hiding light; cavity lattices
Error-handling hook	Reservoir refill of lost atoms	Erasure conversion via $^3P_0$ tune-out traps + refill

7. Relevance to neutral-atom arrays and your LCoS design

- **Short-wavelength SLMs.** An LCoS hologram for a Yb target array runs at 459.6 nm or 423.3 nm, not 852 nm. That needs roughly half the liquid-crystal thickness for a 2π stroke, which helps speed and fringing, but blue light raises absorption, LC photo-degradation, and the phase error from a given flatness error (it scales as 1/λ). Coatings, cover glass and the backplane mirror must all be specified for the blue.
- **Two-color holograms.** The two target sub-arrays use different wavelengths. One LCoS cannot serve both with a single phase pattern, so the choice is two SLM paths combined by a dichroic, or a spatially split panel with per-region wavelength calibration and its own zero-order handling.

- **Magic wavelengths constrain the optics.** Every trap color is chosen to be magic for a specific transition ($^3P_0$ for gates, $^3P_1$ for imaging). A light shift from SLM intensity non-uniformity therefore matters less for the qubit frequency than in Rb, but uniformity still sets trap depth, loading and the tune-out-based erasure behavior.
- **Cavities take the load off the SLM.** Because the deep imaging traps come from a cavity lattice, the SLM only has to supply about 100–200 µW per site, which eases the damage-threshold and heat-load limits on a 1,000+ site hologram.
- **Pitch and field.** A 3.3 µm pitch over 115 µm, inside a 0.5 mm field at NA 0.65, is about 1.5× denser than the QuEra arrays. In your LCoS terms, that pushes the hologram toward the panel's highest usable spatial frequencies, where fringing-field crosstalk and pixel-edge effects reduce efficiency most.

Sources

Atom Computing, "Highly Scalable Quantum Computing with Neutral Atoms" (whitepaper, 2025) — copy in Neutral_atom_QC_companies

K. Barnes et al., "Assembly and coherent control of a register of nuclear spin qubits," Nat. Commun. 13, 2779 (2022); arXiv:2108.04790 — https://arxiv.org/abs/2108.04790

M. A. Norcia et al., "Midcircuit qubit measurement and rearrangement in a ^{171}Yb atomic array," Phys. Rev. X 13, 041034 (2023); arXiv:2305.19119 — https://arxiv.org/abs/2305.19119

M. A. Norcia et al., "Iterative assembly of ^{171}Yb atom arrays with cavity-enhanced optical lattices," PRX Quantum 5, 030316 (2024); arXiv:2401.16177 — https://arxiv.org/abs/2401.16177

J. A. Muniz et al., "High-fidelity universal gates in the ^{171}Yb ground-state nuclear-spin qubit," arXiv:2411.11708 (2024) — https://arxiv.org/abs/2411.11708

B. W. Reichardt et al., "Logical computation demonstrated with a neutral atom quantum processor," arXiv:2411.11822 (2024) — https://arxiv.org/abs/2411.11822

Atom Computing, "Highly coherent qubits and the path to error correction" — https://atom-computing.com/highly-coherent-qubits-and-the-path-to-error-correction/

Figures generated by atom_computing_layout_figs.py (saved alongside this note). Atomic constants (abundances, nuclear moments, linewidths) are standard reference values; the Sr Rydberg wavelength is from the wider Sr literature, not an Atom Computing paper.