Improving Rydberg Gate Fidelity

Oct 21, 2026 · D.F. Hakala

The baseline Rb atom design sits near 99.7% CZ fidelity. Most of the remaining error is 6P scattering and Rydberg decay, both reduced by more laser power, so laser power and phase-noise filtering come first; robust pulses, uniform illumination and colder atoms take it toward ~99.9%, and cryogenics or new gate types toward 99.95%.

Where the error comes from

Numbers come from the gate simulation built earlier (time-optimal CZ, $\Omega = 2\pi \times 4.6$ MHz, gate time 263 ns) plus standard ^{87}Rb values.

Error source	What was computed or assumed	Typical error per CZ
6P scattering	Error ≈ 3.8 Γ_{6p}/Δ; about 2.7×10^{-3} per driven atom at $\Delta = 2\pi \times 2$ GHz	~1.3×10^{-3} averaged over inputs
Rydberg decay	Average time in the Rydberg state $2.96/\Omega$ (simulated); lifetime ~150 µs at room temperature	6.8×10^{-4}
Ω error from intensity	1% Ω error → 3.5×10^{-4}; 2% → 1.4×10^{-3} (simulated)	10^{-4}–10^{-3}
Laser phase noise	1 kHz effective linewidth, no servo bumps	~1×10^{-4}, often 10× worse in practice
Doppler shift	At 10 µK the spread is 0.94% of Ω; a 0.01 Ω detuning costs 1.3×10^{-4} (simulated)	~1×10^{-4}

These sum to roughly 2.5–3.5×10^{-3}, about 99.65–99.75% fidelity, which is where the best labs are. The two largest terms both shrink with more laser power.

1. Reduce laser noise and add power

The effective two-photon Rabi frequency sets the trade-off:

$$\Omega = \frac{\Omega_{420}\,\Omega_{1013}}{2\Delta}$$

- **Larger detuning:** at fixed Ω, scattering error falls as 1/Δ, but Δ can only grow if $\Omega_{420} \cdot \Omega_{1013}$ grows. Quadrupling the power in both beams raises Δ by 4× and cuts scattering 4×.
- **Faster gate:** at fixed Δ, more power raises Ω, shortening the gate and reducing Rydberg-decay error proportionally.

How to get the power

- **Per-beam Rabi frequency:** Ω_{420} and Ω_{1013} each need ~$2\pi \times 100$–250 MHz to support $\Delta \approx 2\pi \times 2$–$7$ GHz at $\Omega \approx 2\pi \times 5$ MHz.
- **1013 nm:** a narrow ECDL seed plus a ytterbium fiber amplifier at several watts. Straightforward.
- **420 nm (the harder leg):** frequency-double an 840 nm diode + tapered amplifier in a resonant SHG cavity, or injection-lock direct 420 nm diodes to a stable seed. Expect hundreds of mW to ~1 W.
- **Beam size:** make it only as large as the active zone, or use the multi-spot approach in option 3.

Phase noise

- **Linewidth:** Pound–Drever–Hall lock both lasers to the ULE cavity for Hz–kHz linewidth.
- **Servo bumps:** the lock pushes noise to ~0.5–2 MHz, exactly where a gate at $\Omega \approx 2\pi \times 5$ MHz is most sensitive. This usually dominates the phase-noise error.
- **Fix for servo bumps:** use the light *transmitted* through the cavity, which is filtered, to injection-lock a slave diode or seed the amplifier, so the high-power beam inherits the clean spectrum. This is how the 99.5% results were achieved.
- **Measuring it:** self-heterodyne or beat-note against a second cavity; on the atoms, phase noise appears as Rydberg Rabi dephasing faster than decay alone explains.

Targets

- $\Delta \geq 2\pi \times 5$ GHz
- Phase-noise power spectral density below ~1 Hz^2/Hz in the 0.1–5 MHz band, with no visible servo bumps
- Intensity noise below 0.1% RMS during the gate

2. Use robust optimal-control pulses

The time-optimal pulse is fast but error-sensitive: in simulation a 1% Ω error costs 3.5×10^{-4} and a 0.01 Ω detuning costs 1.3×10^{-4}, both growing with the square of the error.

- **Optimize over a range of errors:** optimize the laser phase $\phi(t)$, and optionally the amplitude $\Omega(t)$, with the cost averaged over a spread of Ω errors (say ±2%) and detunings (±0.03 Ω). This is ensemble or robust GRAPE.
- **Or cancel first-order sensitivity:** add derivative terms that force the gate's sensitivity to zero at the nominal point.
- **Trade-off:** robust pulses are typically 10–30% longer, adding a little decay error; the right balance depends on the measured noise.
- **Smooth edges:** ~10–20 ns amplitude ramps so AOM and amplifier bandwidth don't distort the pulse.

Best way to achieve it

1. Extend `rydberg_cz_pulse_sim.py` with a robust cost function (a small change).
2. Include the measured laser noise spectrum and beam-intensity map in the model.

3. Calibrate on the machine with closed-loop optimization of a few pulse parameters (for example Bayesian optimization of the phase-profile coefficients) against measured CZ fidelity.

Target: error from ±1% Ω and ±0.02 Ω detuning below 10^{-4}.

3. Make the Rydberg illumination uniform

A Gaussian global beam 1% weaker at the edge sites adds ~3.5×10^{-4} error there. Because $\Omega \propto \sqrt{(I_{420} \cdot I_{1013})}$, each beam needs about 1% uniformity to hold the Ω error near 1%.

Approach	How	Strength	Limitation
Multi-spot Rydberg beams (preferred)	One 420 and 1013 nm spot per atom, from the same kind of hologram as the trap array	5–20× more intensity per atom at the same power (feeds option 1); uniformity fixed by the weighted Gerchberg–Saxton feedback used for traps	Spots must track the atom positions
Flat-top beam	Refractive beam shaper (πShaper type) or static diffractive element designed with the hologram code	Simple and stable	Wastes power outside the atoms
LCOS SLM directly at 420 nm	Same SLM optics at the blue wavelength	Reconfigurable	Organic alignment and LC degrade under blue light (oblique-SiO_2 alignment is the robust choice); dielectric mirror must suit 420 nm; check response and flicker

Many groups use the SLM to design and test the pattern, then fabricate a static fused-silica element for the production machine.

Calibration and specs

- Measure each site's Rydberg Rabi frequency (or each beam's light shift) on the atoms and feed it back into the hologram weights until the spread is under 0.5%.
- Pointing stability below 1% of the spot size.
- Pure polarization, so the coupling strength is identical at every site.

4. Use colder atoms

At 10 μK the Doppler spread is 0.94% of Ω (~1×10^{-4} error); thermal position spread adds a little more through intensity variation across each spot and variation in atom spacing.

1. **Λ-enhanced gray molasses on D1 (795 nm)** after loading reaches ~5 μK. Cheap: mostly AOM and EOM work on light already available.
2. **Adiabatic trap lowering** over ~1 ms just before the gate reduces temperature roughly in proportion to the trap frequency.
3. **Raman sideband cooling** to the motional ground state gives the lowest temperature. Needs a Raman beam pair and resolved trap sidebands; worthwhile if Raman beams already serve single-qubit gates.
4. **Counter-propagating 420 and 1013 nm beams** reduce the effective wavevector that drives the Doppler shift; keep that geometry.

Targets: ≤ 5 μK at gate time (halves the Doppler term), and pair spacing chosen so the blockade strength $V/\Omega > 50$ even at the edges of the thermal spread.

5. Convert errors into erasures

If most errors leave a detectable trace, the decoder knows where they happened, and a surface code tolerates several times more physical error: Wu, Kolkowitz, Puri and Thompson (2022) estimated the threshold rises from ~1% to ~4% when about 98% of errors are erasures.

How it works in ^{171}Yb

- The qubit sits in the metastable $^3P_0$ state (nuclear-spin states $m_I = \pm ½$).
- Most Rydberg decays end up outside the qubit subspace, mainly in the $^1S_0$ ground state.
- Imaging on a $^1S_0$ transition reveals those atoms without disturbing the $^3P_0$ qubits.
- The MPQ ^{88}Sr work converted about 91% of leakage into detected erasures this way.

What ^{87}Rb can do (part of the benefit)

- **Detect atom loss mid-circuit:** occupancy-only imaging, with missing atoms treated as erasures. Harvard's 2025 below-threshold result used this.
- **Pump out leftover Rydberg population** after each gate, so it becomes clean loss rather than a hidden error.
- **Detect leakage** to other hyperfine sublevels with state-selective imaging.

Decision point

Full erasure conversion means moving to Yb or Sr, with new lasers (399, 556 and 578 nm, plus Rydberg UV) and a magic trap wavelength that also removes trap-induced dephasing. It is the biggest long-term lever, but it is a species choice, not an upgrade.

6. Pushing past 99.9%

Option	How it works	Benefit	Cost or status
Cryogenic environment (~4 K)	Removes blackbody-driven Rydberg transitions; improves vacuum	Rydberg lifetime 2–4× longer: decay error 6.8×10^{-4} (150 μs) → 2.6×10^{-4} (400 μs), simulated; trap lifetime ~6,000	Cryostat with optical access, vibration control, trap-laser heat load

Option	How it works	Benefit	Cost or status
Choice of Rydberg level n	Lifetime grows as n^3, stray-field sensitivity as n^7	Cheap tuning knob; n ≈ 50–80 is the usual sweet spot	Needs electrodes to null stray fields and Rydberg spectroscopy to check stability
Förster-resonance gates	A small DC field makes two pair states degenerate; interaction switches from C_6/r^6 to C_3/r^3	Stronger, longer-range blockade, less sensitive to atom spacing	Precise field control
Dipolar iSWAP / exchange gates	Two different Rydberg states exchange energy through dipole coupling	Native iSWAP or parametrized exchange gates; proposed to reach 99.9% (Ildefonso et al., 2025)	Simulation only, not yet demonstrated

Recommended order

Step	Action	Expected CZ error after
Baseline	Current design	~3×10^{-3} (99.7%)
+ 1	$\Delta \geq 2\pi \times 5$ GHz, filtered phase noise	~1.5×10^{-3}
+ 2, 3	Robust pulses; multi-spot Rydberg beam with atom feedback	~1×10^{-3}
+ 4	≤ 5 μK at gate time	~8×10^{-4}
+ 6	Cryogenics or Förster gates	~$4–5 \times 10^{-4}$ (99.95%)
5, in parallel	Loss detection and Rydberg pump-out (Rb), or erasure conversion (Yb/Sr)	Raises what the error-correction code tolerates rather than the raw fidelity

These are estimates; real machines usually show extra error sources (stray fields, crosstalk, calibration drift) that only appear during commissioning.

- ☐ Extend the pulse simulation with a robust cost function and compare sensitivity curves with the time-optimal pulse

Sources

- Simulated values: `rydberg_cz_pulse_sim.py` (C:\LCoS\Neutral Atom QC)
- Universal gates for a metastable qubit in Sr-88 (arXiv 2506.10714)
- Native iSWAP and exchange gates from dipolar Rydberg interactions (arXiv 2512.05037)
- Architectural mechanisms of a universal fault-tolerant quantum computer (arXiv 2506.20661)

- Neutral Atom Quantum Computing: Principles, Routes, Progress, and Challenges (arXiv 2608.05010)
- Recalled, not re-checked: the 2.95/Ω time-in-Rydberg result (Jandura & Pupillo 2022; Evered et al. 2023), the erasure threshold (Wu et al., *Nat. Commun.* 13, 4657, 2022), and the ^{87}Rb $6P_{3/2}$ lifetime (~112 ns).