GS Holograms: Light Shift, Loading and Gate Detuning

How hologram quality in an SLM tweezer array sets trap light shifts, single-atom loading and gate detuning

Technical note · October 24, 2026 · D. F. Hakala

The Gerchberg–Saxton algorithm, or more often its weighted version (WGS), computes the phase pattern that turns one laser beam into a tweezer array. It only sets intensities, positions and phases in the focal plane; light shift, loading probability and gate detuning follow from how well the hologram and the optics behind it deliver them. Most of the chain comes down to one number per site, the trap depth $U_o \propto \mathrm{Re}[\alpha(\lambda)] \cdot P_{spot} / w_o^2$, so any sitewise error in spot power or waist feeds all three.

1. Light shift

A 1% spread in trap depth gives about a 1% spread in light shift, so hologram uniformity sets qubit and Rydberg detuning errors directly.

Physics

The tweezer lowers the ground-state energy by U_o. For Rb or Cs near 800–1064 nm, U_o/h is about 20.8 MHz per mK of depth. Two differential effects matter for qubits:

- **Hyperfine (qubit) differential shift:** $\delta_{hf} \approx \eta \cdot U_o/h$. The ratio η is roughly the hyperfine splitting divided by the effective detuning of the trap light. It is about 10^{-3} for red-detuned alkali traps, giving tens of kHz at 1 mK. The qubit frequency then depends on the local trap depth.
- **Ground–Rydberg differential shift:** a Rydberg electron is nearly free, with $\alpha_R \approx -e^2/(m_e\omega^2)$, which is small and negative. The ground–Rydberg transition is therefore shifted by about U_o/h, roughly 10–20 MHz in a typical trap. This is far larger than the qubit shift.

How the hologram and optics contribute

Contributor	Mechanism	Remedy
Spot-to-spot intensity nonuniformity	Plain GS leaves 10–30% variation; WGS reaches below 1% in simulation, 2–5% on hardware without feedback	Camera or atom-based feedback (Section 4)
SLM pixel envelope	Diffraction efficiency falls roughly as sinc² of steering angle, so spots far from the zero order come out dimmer	Pre-weight those spots in WGS
Waist variation across the field	Off-axis aberrations (field curvature, coma, astigmatism) enlarge w_o and lower peak intensity at constant power; a camera not conjugate to the atom plane can miss it	Field-dependent Zernike correction on the SLM
Interference between neighbors	GS leaves relative spot phases free; overlapping tails interfere at close spacing, so depth depends on those phases	Lock or constrain spot phases (phase-fixed WGS)
Ghost spots, zero order, fringing losses	Stray light near the array adds unwanted local light shifts	Tilt or defocus carrier to move the zero order out of the field, then block it
Temporal noise	SLM phase flicker at frame or sub-frame rate becomes intensity noise on every trap: light-shift noise and parametric heating near 2× the trap frequency	Faster addressing, analog drive, measure the flicker spectrum

Contributor	Mechanism	Remedy
Polarization	Tight focusing at high NA creates longitudinal fields and local ellipticity, giving vector light shifts (a fictitious magnetic field) that vary by site	Clean polarization after the SLM; keep it linear
Atom temperature	A thermal atom samples less than the peak intensity, spreading the shift	Colder atoms, deeper traps, spin echo

Rough scale

For Rb in 1 mK traps with $\eta \approx 10^{-3}$, δ_{hf} is about 20 kHz. A 1% depth spread gives about 200 Hz of sitewise detuning, which limits the free-evolution T_2^* to around a millisecond unless you use spin echo or lower the traps during gates.

2. Loading probability

Trap depth sets the effective cooling detuning at every site, so depth uniformity decides whether each site stays in the narrow window that high-efficiency loading needs.

Physics

In a tweezer below about 1 µm waist, light-assisted collisions eject atoms in pairs ("collisional blockade"). This leaves 0 or 1 atom with a single-atom probability of about 50–60%. Gray-molasses or blue-detuned "enhanced loading" reaches about 80–90%, but only in a narrow window of cooling detuning measured relative to the *light-shifted* transition inside each trap.

How the hologram contributes

- **Depth sets the effective cooling detuning at each site.** The trap light shifts the cooling transition by an amount set by U_o, and only the excited-state part differs. A depth error therefore detunes the molasses differently at different sites. Standard blockade loading is fairly tolerant. Enhanced loading is not: a 5–10% depth spread can move outer sites off the optimum and lower their loading by tens of percent.
- **The depth window.** Traps that are too shallow capture fewer atoms and lose them during imaging. Traps that are too deep shift the cooling transition so far that light-assisted collisions and cooling no longer behave as designed. Uniformity keeps every site inside the window.
- **Waist.** Collisional blockade needs a small trap volume. Sites with aberration-enlarged waists may load two atoms more often, show weaker blockade, or behave differently under the enhanced-loading process.
- **Ghost traps.** Ghost orders strong enough to trap atoms add stray atoms and background in imaging.
- **Imaging survival.** Shallow sites heat out during fluorescence imaging. That shows up as apparent loading failures and lowers rearrangement success.
- **Rearrangement.** If the SLM does the moving, the hologram must update every frame. Nematic LC switches over milliseconds, and during switching the intensity dips and the spots flicker. Phase-continuous hologram sequences, or a static SLM array with AOD-based moving tweezers, avoid this. Hand-off between the two needs SLM spot positions that match the AOD grid to a small fraction of w_o, which means calibrating GS spot positions against the actual aberrated optics.

- **Array statistics.** A defect-free N-site array needs a reservoir of about N/p sites, where p is the loading probability, plus high transport fidelity. Every point lost in p or in survival is paid for in extra SLM spots and laser power.

3. Gate detuning

With traps on during a Rydberg gate, a 1% depth spread costs roughly 10^{-3} infidelity per gate at $\Omega/2\pi = 5$ MHz; for SLM addressing beams, spot uniformity is the single-qubit gate calibration.

Two-qubit Rydberg gates (blockade CZ, time-optimal pulses)

These need the laser detuning Δ to match each atom's ground–Rydberg resonance well within the Rabi frequency Ω. The phase error and leakage scale as about $(\delta\Delta/\Omega)^2$.

- **Traps on during the gate:** the resonance shifts by about U_o/h. Taking 20 MHz at 1 mK, a 1% depth spread is about 200 kHz of sitewise detuning. With $\Omega/2\pi = 5$ MHz that is $\delta\Delta/\Omega \approx 0.04$, roughly 10^{-3} infidelity per gate from this alone. The trap light can also photo-ionize Rydberg atoms and anti-trap them.
- **Traps off or lowered during the gate (most common):** the light-shift term goes away, but release-and-recapture heating, Doppler shifts and atom loss appear instead. Dipping the traps briefly rather than switching them off trades between the two.
- **Rydberg-beam light shifts:** the two-photon excitation itself shifts the levels. In a ladder scheme the shift is about $(\Omega_1^2 - \Omega_2^2)/(4\Delta_{int})$, where Δ_{int} is the detuning from the intermediate state. Global Rydberg beams are Gaussian, so outer sites see different Ω and different shifts. An SLM running GS or a flat-top algorithm can homogenize these beams too.

Other contributors not set by the hologram

- Doppler shift, $k_{eff} \cdot v$ (thermal motion)
- Laser phase and frequency noise
- Stray electric fields (Rydberg states have large DC Stark shifts)
- Magnetic-field gradients across the array (Zeeman shifts)
- Blackbody-induced shifts and lifetime

Single-qubit gates

With Raman or microwave rotations, the differential light shift δ_{hf} at each site is a detuning error. It causes over- or under-rotation and phase drift. With **SLM-generated local addressing beams**, the light shift is deliberate: the addressed atoms are pushed off resonance, or a site-selective phase is applied. Here the GS spot uniformity *is* the gate calibration. A 2% spread in addressing intensity is a 2% spread in phase or detuning, which must be calibrated per site or suppressed with composite pulses.

4. Closing the loop and tightening the budget

Feedback measured on the atoms themselves is what takes depth uniformity from a few percent to a few tenths of a percent.

1. **Camera feedback.** Run WGS with a camera conjugate to the atom plane, using measured spot intensities rather than simulated ones. This typically reaches about 1–2% uniformity.

2. **Atom-based feedback.** Measure the light shift at each site directly, using microwave or Raman spectroscopy of the differential shift, or the shift of the push-out or imaging resonance. Feed those values back into the WGS weights. This reaches a few tenths of a percent and automatically includes waist and polarization effects that a camera misses.
3. **Fixed spot phases.** Keep the phases between spots fixed across recomputations, so that interference does not move the depths around.
4. **Aberration correction.** Apply field-dependent Zernike correction on the same SLM so that w_o is uniform, not just the power.
5. **Wavelength choice.** A trap wavelength farther detuned lowers η, and a magic wavelength cancels the differential shift where the species allows it. Ramping the traps down during gates removes most of the remaining light-shift detuning.

Target budget

Parameter	Target
Trap-depth uniformity after feedback	≤ 1% (stretch about 0.3%)
Spot position error	< 5% of w_o
Ghost and zero-order intensity near the array	< 10^{-3} of a trap
SLM flicker near 2× trap frequency	< 10^{-3} relative intensity

Numbers in this note are order-of-magnitude estimates for alkali atoms (Rb, Cs) in red-detuned tweezers; check them against your species, trap wavelength and gate scheme.