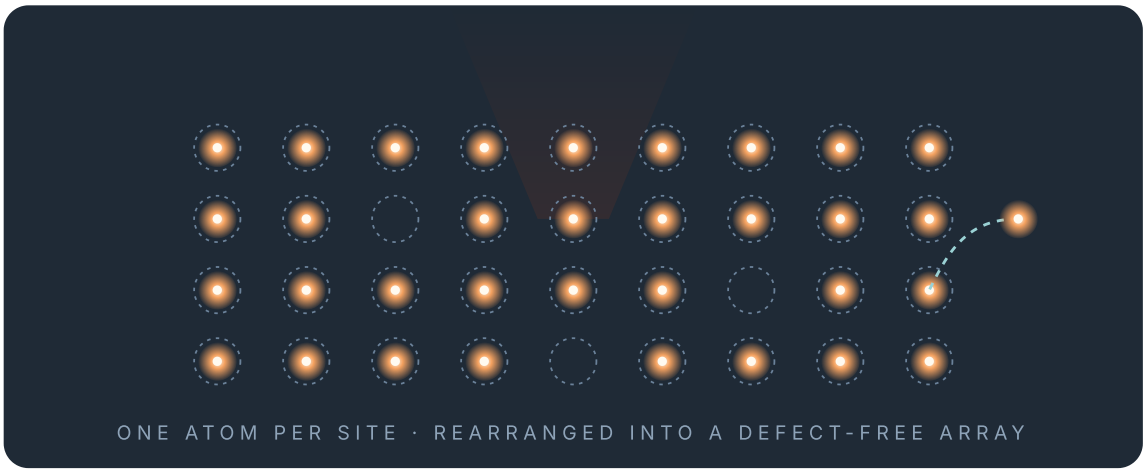
AMO TOOLKIT · TRAINING KIT

Building with Atoms

A training kit for experimental AMO and neutral-atom researchers



Saumitra Phatak

Edition 1 · November 2026 · compiled from amotoolkit.com
Interactive tools, animations and updates: amotoolkit.com

About this book

Building with Atoms collects everything on AMO Toolkit into one volume you can download, annotate and keep: the guides, the reference material, the derivations, the worked numbers and the figures of all 35 pages, arranged in the order a new researcher meets them. It was written for students joining atomic, molecular and optical physics groups, and especially for those building or running neutral-atom and optical-tweezer experiments.

How it relates to the website

Each chapter is one page of the website. The website stays the place to *use* the material: its calculators, simulators and animations respond to your inputs, and it is updated as the field moves. This book is a snapshot of it, made so you do not have to come back to look something up.

- **Interactive figures** appear here as snapshots at the page's default settings. A dashed ▶ box marks each place where the website has controls you can change, with a link straight to it.
- **Calculator readouts** are shown as small panels of numbers. They are the page's output for its default inputs, so they double as worked examples.
- **Python notebooks** reproduce 19 of the calculators and are linked at the start of their chapters and listed in Appendix A.
- **Problems** close most chapters; Part V lays the chapters out as a 14-week course.

How to read it

Part I follows the build of a lab, Part II the first experiment with single atoms, Part III the step from atoms to qubits, and Part IV the literature, the research groups and the industry. Read a part in order, or go straight to the chapter you need: every chapter starts with what it is for and links to its live page.

Boxes have consistent meanings. A **tinted box** is a key idea or a practical note; an **amber box** is a warning or a common mistake; a **bordered box** groups a worked example, a formula or a comparison. Tags such as PREPRINT flag the status of a source.

A word of care

The numbers in this book are estimates from simple models, checked against their sources when written. They are good for scale, design trade-offs and intuition; they are not a substitute for measuring your own system. Preprints are marked as such, and company and funding figures carry the date they were checked. If you find an error, the "Found a problem?" link on any page of the website reaches the author. Appendix D says more.

Contents

PART I · BUILD A LAB

1 Atomic Species & Level Structure 6
2 Laser Systems & Wavelength Architecture 16
3 Ultra-High Vacuum Engineering 19
4 Laser Frequency Stabilization 31
5 Polarimetry & Optical State Tomography 46
6 Magneto-Optical Traps 56
7 Experimental Control & Timing Systems 64
8 Optical Tweezer Arrays 74
9 Wavefront Aberrations & Zernike Modes 81
10 Laboratory Automation & Closed-Loop Optimization 88
11 Cavity Quantum Electrodynamics 101
12 AMO Laboratory Operations Handbook 109

PART II · TRAP, IMAGE AND COOL

13 Laser Cooling Forces & Sub-Doppler Mechanisms 129
14 Absorption & Dispersive Cloud Imaging 140
15 Time-of-Flight Thermometry 148
16 Single-Atom Fluorescence Imaging 154
17 Release-and-Recapture Thermometry 165
18 Ground-State Cooling of Trapped Atoms 171
19 End-to-End Tweezer Experimental Cycle 189
20 Everyday AMO Reference Calculations 194

PART III · FROM ATOMS TO A QUANTUM COMPUTER

21 Quantum Computing Fundamentals 203
22 Decoherence & Optical Bloch Equations 226
23 Dynamical Decoupling & Noise Spectroscopy 233
24 Rydberg Blockade Physics 244
25 Two-Qubit Gate Error Budgets 251
26 Randomized Benchmarking 257
27 Zoned Neutral-Atom Processors 267
28 Remote Entanglement & Quantum Networks 272

PART IV · INTO THE FIELD

29 Foundational AMO Literature 290
30 Global AMO Research Directory 303
31 Quantum Computing Industry Landscape 321
32 Neutral-Atom Qubits: ^{87}Rb vs. ^{171}Yb 341
33 Trapped-Ion Architectures: Quantinuum vs. IonQ 359
34 Superconducting Processors: Google vs. IBM 374
35 Photonic Architectures: PsiQuantum vs. Xanadu 390

PART V · TEACH WITH THIS BOOK

V The 14-week course 404

APPENDICES

A Python notebooks 407
B Constants used in this book 408
C Citing this book and the work behind it 409
D Using this material responsibly 410

PART I

Build a Lab

Twelve steps from an empty room to a working single-atom experiment: choose the atom, plan its lasers, pump down, lock, check polarisation, trap a cloud, wire the control system, build the tweezer array, correct the optics, automate the routine work, add a cavity, and run the lab day to day.

1	Atomic Species & Level Structure	Atomic Species Selector
2	Laser Systems & Wavelength Architecture	Laser System Planner
3	Ultra-High Vacuum Engineering	Vacuum Systems Guide
4	Laser Frequency Stabilization	Laser Locking Guide
5	Polarimetry & Optical State Tomography	Polarimetry Simulator
6	Magneto-Optical Traps	MOT Designer
7	Experimental Control & Timing Systems	Experiment Control Systems
8	Optical Tweezer Arrays	Tweezer Array Design Lab
9	Wavefront Aberrations & Zernike Modes	Zernike Wavefront Lab
10	Laboratory Automation & Closed-Loop Optimization	Automating the Experiment
11	Cavity Quantum Electrodynamics	Cavity QED Coupling Lab
12	AMO Laboratory Operations Handbook	AMO Lab Operations Handbook

CHAPTER 1 · BUILD STEP 01

1

Atomic Species & Level Structure

Atomic Species Selector · Reference

Compare laser-coolable atoms by the decisions an experimentalist actually faces: species choice, laser system, qubit platform, cooling strategy, and nearby papers or groups.

> *Every experimental AMO apparatus begins with a single choice that locks in every laser wavelength, vacuum source, and cooling scheme downstream: the atomic species.*

IN THIS CHAPTER No prerequisites

- Compare alkali (Rb, Cs, Li, Na, K) and alkaline-earth-like (Sr, Yb) level structures, D-lines, and hyperfine splittings.
- Read key atomic constants—natural linewidth Γ, saturation intensity I_sat, and Doppler/recoil temperatures—from standard data sheets.
- Match atomic properties (magic wavelengths, nuclear spin, Feshbach resonances) to quantum computing, clock, or simulation goals.

WEBSITE COMPANION Live interactive page ↗ · Week 1

Species choice workflow

1.1 Use the library to make a design decision, not just to look up a number.

The same atom can be attractive for one goal and painful for another. Start with your experimental goal, then use the tables and data sheets to check whether the lasers, cooling, and coherence story are compatible.

01 **Choose a species**

Start with mass, linewidth, recoil, hyperfine structure, and whether you need alkali or alkaline-earth physics.

02 **Compare platforms**

Map species to tweezers, clocks, molecules, Rydberg gates, quantum gases, and imaging constraints.

03 **Design lasers**

Use wavelengths and linewidths to plan cooling, repump, trapping, clock, and Rydberg beams.

04
Find groups and papers
Jump from species to canonical data sheets, representative labs, and literature trails.

1.2 Periodic Table Explorer

The 15 isotopes in this library come from nine elements, lit up below. Pick one to see its numbers, energy levels, data sheets and the groups that use it, or colour the whole table by a property to see the periodic trends behind choosing a species.

Colour by



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

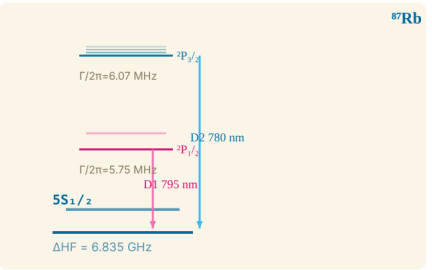
Alkali metals Two-electron (alkaline-earth-like) Magnetic lanthanides

Alkalis (one valence electron) give microwave hyperfine qubits and diode-laser D lines. The two-electron atoms Sr and Yb add narrow intercombination and mHz clock lines; Yb is a lanthanide, but its filled 4f shell leaves an alkaline-earth-like $6s^2$ outer shell. Dy and Er carry large magnetic moments from their open 4f shells.

^{87}Rb · Rubidium · Alkali metals

> Alkali atoms (Li, Na, K, Rb, Cs) have a single valence electron, giving a simple hydrogen-like level structure. Their D1 and D2 lines in the visible/near-IR are accessible with diode lasers. Hyperfine ground states form natural two-level qubits. **Rb-87** (6.8 GHz clock) and **Cs-133** (9.2 GHz, defines the SI second) are the two most widely used qubit atoms today.
>
> Alkaline-earth and alkaline-earth-like atoms (Ca, Sr, Yb) have two valence electrons, giving rich level structures including ultra-narrow *intercombination* and *clock* transitions. These enable narrow-line cooling near the recoil scale, exceptional coherence times, and optical-lattice clocks accurate to 1 part in 10^{18}. Fermionic isotopes (Sr-87, Yb-171) offer nuclear spin qubits decoupled from the electronic state, ideal for optical-clock quantum computing.
>
> Highly magnetic atoms (Cr, Dy, Er) have large magnetic moments (6–10 μ_B), making their inter-particle interactions strongly anisotropic and long-range. This opens the door to **dipolar quantum simulation**, enabling phenomena impossible with contact- interaction BECs. Dy has the largest magnetic moment of any element (10 μ_B).

^{87}Rb · energy levels



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

D2 line: 780.241 nm · $\Gamma/2\pi$ = 6.07 MHz
D1 line: 794.979 nm · $\Gamma/2\pi$ = 5.75 MHz
Ground HF splitting: 6.835 GHz
⚛ Clock qubit $|F{=}1\rangle \leftrightarrow |F{=}2\rangle$ at 6.835 GHz

Schematic: level spacings are not to scale.

^{6}Li / ^{7}Li, Lithium

Why it's used

Lightest alkali. Large recoil energy makes imaging and trapping more demanding, so D1 gray molasses or related sub-Doppler cooling is especially important. Strong Feshbach resonances allow tunable interactions. **Li-6** is the primary fermionic atom for strongly-correlated physics and BEC-BCS crossover experiments. Used in Li-Cs molecular assembly (Hood Lab, Purdue).

Key spectral lines

D1: 670.992 nm | D2: 670.977 nm | $\Gamma/2\pi$ = 5.87 MHz

Qubit transition

^{7}Li: $|F{=}1\rangle \leftrightarrow |F{=}2\rangle$ hyperfine qubit (803 MHz). ^{6}Li: fermionic spin states used for many-body qubits.

Data sheets

📄 ^{6}Li properties, Gehm (NCSU) 📚 Steck alkali data index 🔗 NIST Atomic Spectra Database

Further reading

Gehm (2003), Properties of ^{6}Li [NCSU tech doc]

^{23}Na, Sodium

Why it's used

One of the first two BEC atoms (Ketterle group, MIT, 1995, months after ^{87}Rb at JILA; Nobel Prize 2001). Yellow D-line at 589 nm. Larger scattering length suitable for BEC studies. Being revisited for molecule formation: NaLi, NaK, NaRb, NaCs.

Key spectral lines

D1: 589.756 nm | D2: 589.158 nm | $\Gamma/2\pi$ = 9.80 MHz

Qubit transition

$|F{=}1\rangle \leftrightarrow |F{=}2\rangle$ hyperfine qubit (1772 MHz)

Data sheets

📄 Na-23, Steck data sheet 📚 Steck alkali data index 🔗 NIST Atomic Spectra Database

Note

The Steck data sheet above is the canonical reference. Steck (2019, updated continuously).

^{39}K / ^{40}K / ^{41}K, Potassium

Why it's used

K-40 is, with ^{6}Li, one of the two fermionic alkali isotopes; a standard atom for Fermi-Hubbard model simulations in optical lattices. **K-39** has accessible Feshbach resonances for tuning interactions. All isotopes share the same 767/770 nm D-lines, diode laser accessible.

Key spectral lines

D1: 770.108 nm | D2: 766.701 nm | $\Gamma/2\pi$ = 6.04 MHz

Qubit transition

K-39: $|1,0\rangle \leftrightarrow |2,0\rangle$ clock transition (or $|1,-1\rangle \leftrightarrow |2,+1\rangle$ at its magic field)

Data sheets

📄 K properties, Tiecke (2010) — source offline 📚 Steck alkali data index 🔗 NIST Atomic Spectra Database

^{85}Rb / ^{87}Rb, Rubidium

Why it's used

Rb-87 is the most widely used quantum computing atom, with large 6.835 GHz hyperfine splitting, convenient 780 nm lasers, and well-understood collisional properties. Rb is the backbone of the Harvard/QuEra and Pasqal neutral-atom ecosystem; Atom Computing instead uses alkaline-earth-like Yb for nuclear-spin qubits. **Rb-85** has a Feshbach resonance for tunable interaction experiments.

Key spectral lines

D1: 794.979 nm | D2: 780.241 nm | $\Gamma/2\pi$ = 6.07 MHz

Qubit transition

Rb-87: $|1,0\rangle \leftrightarrow |2,0\rangle$ "clock" qubit (6835 MHz), or the $|1,-1\rangle \leftrightarrow |2,1\rangle$ pair, first-order field-insensitive near 3.23 G

Data sheets

📄 Rb-87, Steck data sheet (v2.2) 📄 Rb-85, Steck data sheet 📚 Steck alkali data index 🔗 NIST Atomic Spectra Database

Data Sheet

steck.us, Rb-87 data sheet (v2.2)

^{133}Cs, Caesium

Why it's used

Largest hyperfine splitting of the alkalis (9.193 GHz, defines the SI second). Excellent for optical tweezer work: heavy mass → low recoil → tighter confinement. Used in Li-Cs molecular assembly experiments (Hood Lab, Purdue). The 852 nm D2 line sits in a convenient region for diode lasers (DFB, ECDL).

Key spectral lines

D1: 894.593 nm | D2: 852.347 nm | $\Gamma/2\pi$ = 5.23 MHz

Qubit transition

$|3,0\rangle \leftrightarrow |4,0\rangle$ "clock" transition (9193 MHz, field-insensitive at zero B-field)

Data sheets

📄 Cs-133, Steck data sheet 📚 Steck alkali data index 🔗 NIST Atomic Spectra Database

Data Sheet

steck.us, Cs data sheet

^{88}Sr / ^{87}Sr, Strontium

Why it's used

Sr has two laser-cooling stages: the broad 461 nm blue line (Doppler limit 770 μK) and the 689 nm red intercombination line ($\Gamma/2\pi$ = 7.6 kHz, T_D = 0.18 μK). The 698 nm clock transition has a linewidth of ~1 mHz. **Sr-87** (I = 9/2) gives 10 nuclear spin states for SU(N) physics and quantum simulation. Used in world-leading optical lattice clocks (Ye Lab, JILA).

Key spectral lines

Broad: 461 nm ($\Gamma/2\pi$ = 32 MHz, T_D = 770 μK) Narrow: 689 nm ($\Gamma/2\pi$ = 7.6 kHz, T_D = 0.18 μK) Clock: 698 nm (~1 mHz linewidth)

Qubit transition

Sr-87 nuclear spin qubit: $|m_I = -9/2\rangle \leftrightarrow |m_I = -7/2\rangle$ via the 698 nm clock transition

Resources

📄 Boyd PhD thesis (Ye Lab), Sr spectroscopy 📄 Stellmer review, Sr BEC and degenerate gases 🔗 NIST Atomic Spectra, Sr 📦 ARC library (Python)

^{174}Yb / ^{171}Yb, Ytterbium

Why it's used

Yb combines broad (399 nm, $\Gamma/2\pi$ = 28 MHz) and narrow (556 nm, $\Gamma/2\pi$ = 182 kHz) cooling lines with a mHz-linewidth clock transition at 578 nm. **Yb-171** (I = 1/2) gives a very clean two-level nuclear-spin qubit in the electronic ground state. Magic wavelengths at 759 nm make optical lattice clocks insensitive to light shifts.

Key spectral lines

Broad: 399 nm ($\Gamma/2\pi$ = 28 MHz) Narrow: 556 nm ($\Gamma/2\pi$ = 182 kHz) Clock: 578 nm (mHz linewidth)

Qubit transition

Yb-171: $|m_I = +1/2\rangle \leftrightarrow |m_I = -1/2\rangle$ nuclear spin qubit (zero-field insensitive)

Resources

📄 Ludlow et al., Optical atomic clocks (RMP 2015) 📄 Barber et al., Direct excitation of Yb clock 🔗 NIST Atomic Spectra, Yb 📦 ARC library (Python)

^{164}Dy, Dysprosium

Why it's used

Dy-164 has the largest magnetic moment of any element (**10 μ_B**), enabling strong dipolar interactions and anisotropic collisional physics. Used for dipolar BEC, quantum droplets, and supersolid phases. Cooled on a broad 421 nm line; also has intercombination lines at 598 nm and 626 nm.

Key spectral lines

Main: 421 nm (broad) Narrow: 598 nm, 626 nm (intercombination lines)

Resources

📄 Lu et al. (Lev), Strongly dipolar BEC 📄 Chomaz et al., Dipolar physics review (Rep. Prog. Phys. 2023)

^{168}Er, Erbium

Why it's used

Er-168 has a magnetic moment of **7 μ_B** and a rich level structure. Dipolar BEC demonstrated by Ferrier-Barbut/Pfau group (Stuttgart) and Grimm group (Innsbruck). Anisotropic scattering leads to distinctive many-body phases. Being explored for dipolar quantum droplets alongside Dy.

Key spectral lines

Main: 401 nm Narrow: 583 nm (intercombination line)

Resources

📄 Aikawa et al., BEC of erbium 📄 Chomaz et al., Dipolar physics review (Rep. Prog. Phys. 2023)

Research groups · ^{87}Rb, Quantum Computing & Simulation (6 groups)

🔗 LUKIN LAB
Harvard
Rydberg tweezer quantum processor; many-body physics

🔗 GREINER LAB
Harvard
Quantum simulation, Hubbard model, quantum gas microscope

🔗 SAFFMAN LAB
Wisconsin
Rydberg two-qubit gates, trapped-atom qubits

🔗 KAUFMAN LAB
Colorado / JILA
Tweezer arrays, quantum optics with atoms

🔗 BERNIEN LAB
University of Innsbruck / IQOQI (formerly UChicago)
Programmable quantum matter, Rydberg arrays

🔗 WEISS LAB
Penn State
Neutral atom qubits, quantum computing

Calculator output at the page's default settings.

Research groups · ^{133}Cs, Quantum Computing & Molecules (2 groups)

🔗 CHIN LAB
UChicago
Strongly correlated gases, Efimov physics, BEC

🔗 HOOD LAB
Purdue ⋆
Li-Cs molecule assembly, optical tweezers, single-atom control

Research groups · ^{6}Li / ^{7}Li, Fermi Gases & Molecules (4 groups)

🔗 HULET LAB
Rice
Li-6 Fermi gases, BEC-BCS crossover, solitons

🔗 ZWIERLEIN LAB
MIT
Degenerate Fermi gases, fermionic superfluidity

🔗 DEMARCO LAB
UIUC
Fermi-Hubbard model, disordered lattices

🔗 HOOD LAB
Purdue ⋆ Li-Cs molecules in tweezers

Research groups · ^{23}Na & K, Molecules & BEC (3 groups)

🔗 KETTERLE LAB
MIT
First BEC (Nobel 2001); spinor BECs; NaLi molecules

🔗 ZWIERLEIN LAB
MIT
NaLi and NaK ultracold molecules; Fermi gases

🔗 NI LAB
Harvard
Ultracold polar molecules (NaCs, KRb)

Research groups · ^{88}Sr / ^{87}Sr, Optical Clocks & Simulation (5 groups)

🔗 YE LAB
JILA / Colorado
World-leading optical lattice clock; Sr tweezer arrays; quantum simulation

🔗 KILLIAN LAB
Rice
Sr BEC and tweezer arrays, Rydberg excitation

🔗 THOMPSON LAB
Princeton Sr cavity QED, quantum networking

🔗 COVEY LAB
UIUC
Tweezer arrays with alkaline-earth atoms

🔗 REY LAB (THEORY)
JILA / Colorado
AMO theory for Sr, Yb quantum simulation

Research groups · ^{174}Yb / ^{171}Yb, Clocks & Quantum Computing (4 groups)

🔗 YE LAB
JILA / Colorado
Yb optical lattice clock; quantum simulation with Yb

🔗 KAUFMAN LAB
Colorado / JILA Yb tweezer arrays; quantum optics

🔗 SPIELMAN LAB
NIST / Maryland
Optical lattices, synthetic gauge fields

🔗 THOMPSON LAB
Princeton Yb cavity QED

Research groups · ^{164}Dy / ^{168}Er, Dipolar Physics (1 group)

🔗 LEV LAB
Stanford
Dy and Er dipolar gases, quantum magnetism

These are representative US groups, organized by their primary atom. The full directory, with a map, is the AMO Group Finder.

1.3 All 15 isotopes at a glance

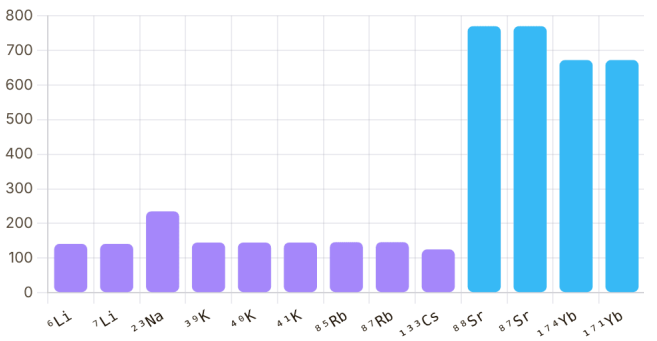
ATOM	Z	A	STATISTICS	FAMILY	COOL. λ (NM)	Γ/2Π (MHZ)	T_D (MK)	2E_R/K_B (MK)	NUCLEAR SPIN I	HF SPLIT (GHZ)	KEY NOTES
^{6}Li	3	6	FERMION	ALKALI	671.0	5.87	141	7.08	1	0.228	Strong Feshbach res.; Li-Cs molecules
^{7}Li	3	7	BOSON	ALKALI	671.0	5.87	141	6.07	3/2	0.804	Early BEC; light mass → large recoil
^{23}Na	11	23	BOSON	ALKALI	589.2	9.80	235	2.40	3/2	1.772	Early BEC (Ketterle 1995, Nobel 2001)
^{39}K	19	39	BOSON	ALKALI	767.0	6.04	145	0.84	3/2	0.462	Feshbach res.; Fermi-Hubbard model
^{40}K	19	40	FERMION	ALKALI	767.0	6.04	145	0.82	4	1.286	Fermionic isotope; degenerate Fermi gas
^{41}K	19	41	BOSON	ALKALI	767.0	6.04	145	0.80	3/2	0.254	Less common; BEC demonstrated

^{85}Rb	37	85	BOSON	ALKALI	780.2	6.07	146	0.37	5/2	3.036	Feshbach res. for tuneable interac-tions
^{87}Rb	37	87	BOSON	ALKALI	780.2	6.07	146	0.36	3/2	6.835	Most popular qubit (6.8 GHz clock)
^{133}Cs	55	133	BOSON	ALKALI	852.3	5.23	125	0.20	7/2	9.193	Defines the SI second; Li-Cs mole-cules
^{88}Sr	38	88	BOSON	ALKALINE EARTH	461.0	32.0	770	0.46	0	—	Narrow 689 nm line → T_D = 0.18 μK; optical clock
^{87}Sr	38	87	FERMION	ALKALINE EARTH	461.0	32.0	770	0.46	9/2	—	10 nu-clear spin states; SU(N) physics
^{174}Yb	70	174	BOSON	ALKALINE EARTH	399.0	28.0	672	0.36	0	—	556 nm narrow line → T_D = 4.4 μK; mHz clock transition
^{171}Yb	70	171	FERMION	ALKALINE EARTH	399.0	28.0	672	0.36	1/2	—	Effective spin-1/2 qubit; clock QC
^{164}Dy	66	164	BOSON	MAGNETIC	421.0	—	—	—	0	—	Largest magnetic moment (10 μ_B); dipolar physics
^{168}Er	68	168	BOSON	MAGNETIC	401.0	—	—	—	0	—	Large magnetic moment (7 μ_B); aniso-tropic interac-tions

T_D = Doppler temperature limit | 2E_r/k_B = absorption-plus-emission recoil scale [single recoil is half this value] | HF = ground-state hyperfine splitting |, = not applicable or varies by isotope. Broad-line T_D shown for all atoms; for alkaline-earth atoms the listed recoil scale is for the narrow intercombination line used for final cooling.

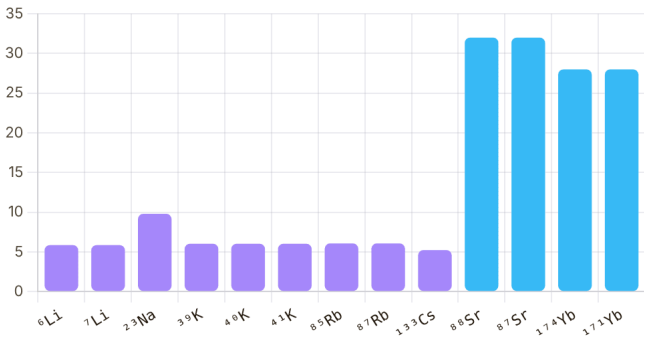
1.4 Visual Comparison

Doppler Temperature T_D (μK), broad cooling line



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Natural Linewidth Γ/2π (MHz), broad cooling line



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Alkali Alkaline Earth

> **Note on alkaline-earth narrow lines:** Sr-88 and Sr-87 have a narrow 689 nm intercombination line with Γ/2π = 7.6 kHz, giving a Doppler limit of **T_D = 0.18 μK**, far below the broad-line value shown above. Yb has a 556 nm intercombination line with Γ/2π = 182 kHz, giving **T_D = 4.4 μK**, far below the broad-line value shown in the table. These narrow lines enable sub-Doppler temperatures without requiring sub-Doppler stages (polarization-gradient cooling etc.).

1.5 Essential Resources & Tools

📄 Data Sheets

Alkali Data Sheets

📄 Li-6 properties, Gehm (NCSU) 📄 Na-23, Steck 📄 K properties, Tiecke — source offline 📄 Rb-85, Steck 📄 Rb-87, Steck 📄 Cs-133, Steck 📚 All Steck sheets →

🔧 Databases & Tools

Spectroscopy Databases

🔗 NIST Atomic Spectra Database 🔗 NIST Physical Reference Data

Python Libraries

📦 ARC, Alkali Rydberg Calculator 📦 QuTiP, Open quantum systems 📦 pylcp, Laser cooling physics

arXiv searches

🔍 quant-ph: neutral atom qubits 🔍 cond-mat.quant-gas (recent)

📖 Textbooks & Reviews

Textbooks

Metcalf & van der Straten, *Laser Cooling and Trapping* (1999) Foot, *Atomic Physics* (2005) Pethick & Smith, *BEC in Dilute Gases* (2008)

Review Articles

📄 Saffman, Rydberg atom QC (RMP 2010) 📄 Kaufman & Ni, Tweezer arrays (Nat. Phys. 2021) 📄 Ludlow et al., Optical clocks (RMP 2015) 📄 Chomaz et al., Dipolar physics (Rep. Prog. Phys. 2023) 📄 Browaeys & Lahaye, Rydberg physics (Nat. Phys. 2020)

Data compiled from Steck data sheets, NIST ASD, and primary literature · All links open in a new tab

CHAPTER 1 SUMMARY

Atomic structure sets the optical, magnetic, and thermal scales of the entire lab. With a target species chosen, we next design the laser architecture needed to drive its cooling, repumping, and trapping transitions.

Problems for Chapter 1

1. USE THE TOOL Use "Compare two" for ^{87}Rb and ^{133}Cs. Tabulate the cooling wavelength, linewidth, Doppler limit and recoil temperature for each.
2. BY HAND Check one Doppler limit with $T_D = \hbar\Gamma/2k_B$ and one recoil temperature with $T_r = \hbar^2k^2/(2m\,k_B)$ (the library's single-photon convention; some texts quote twice this). Explain from the formulas why lithium has the highest recoil temperature of the alkalis.
3. DESIGN Choose a species for (a) a tweezer quantum computer with a nuclear-spin qubit, (b) an optical clock and (c) a dipolar quantum gas. Justify each choice with two numbers from the library.

Use the tool: on the interactive page. By hand: from the formulas in this chapter. Design: a decision a lab would make. Worked solutions are available to instructors on request (see Part V).

CHAPTER 2 · BUILD STEP 02

2

Laser Systems & Wavelength Architecture

💡 Laser System Planner · Design tool

Select an atom species and get the complete laser stack: required wavelengths, source types, power budgets, fiber compatibility, AOM/control notes, and a source → beam control → delivery → chamber diagram.

> *Once the atomic transitions are fixed (Chapter 1), we must assemble a coherent set of laser sources covering every cooling, repumping, imaging, and trapping wavelength.*

IN THIS CHAPTER Prerequisites: Ch. 1

- Select appropriate laser architectures (ECDL, DFB, Ti:Sapphire, tapered amplifier, SHG/SFG) for target wavelengths and linewidths.
- Budget optical power across MOT beams, Zeeman slowers, optical pumping, and tweezer arrays including AOM and fiber-coupling losses.
- Identify where narrow-linewidth cavities or phase-locked seed/amplifier chains are essential.

WEBSITE COMPANION Live interactive page ↗ · Week 2

2.1 Select Atom Species

^{87}Rb

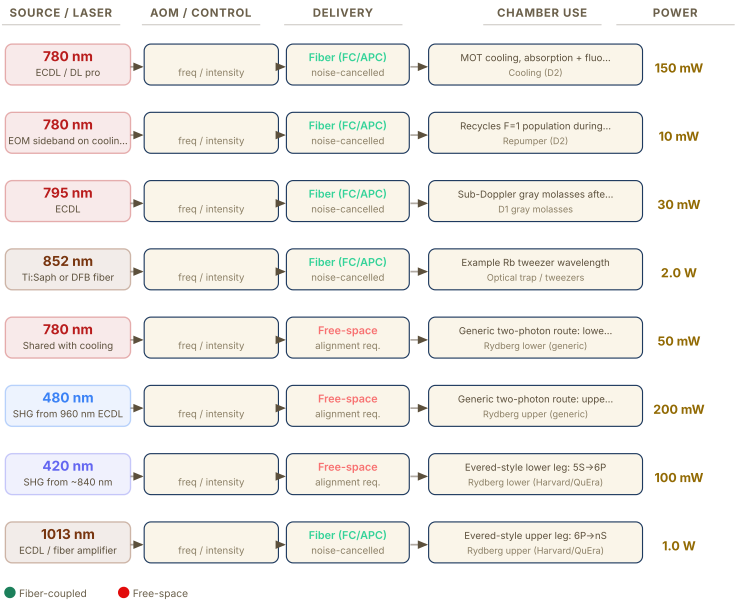
The workhorse of AMO. All cooling lasers are commercially available; Rydberg systems are platform-specific (generic 780+480 nm, Harvard/QuEra-style 420+1013 nm).

2.2 Required Laser Beams

BEAM PURPOSE	WAVELENGTH	SOURCE TYPE	TYPICAL POWER	FIBER?	AOM / CONTROL
Cooling (D2)	780 nm	Commercial ECDL / DL pro	150 mW	Yes ✓	2× AOM (MOT + imaging freq)
Repumper (D2)	780 nm	Commercial EOM sideband on cooling laser	10 mW	Yes ✓	—
D1 gray molasses	795 nm	Commercial ECDL	30 mW	Yes ✓	AOM
Optical trap / tweezers	852 nm	Commercial Ti:Saph or DFB fiber	2.0 W	Yes ✓	AOM (intensity servo)
Rydberg lower (generic)	780 nm	Commercial Shared with cooling	50 mW	Free-space	AOM (pulsed)
Rydberg upper (generic)	480 nm	SHG SHG from 960 nm ECDL	200 mW	Free-space	AOM + fast switch
Rydberg lower (Harvard/QuEra)	420 nm	SHG SHG from ~840 nm	100 mW	Free-space	AOM / fast switch

Rydberg upper (Harvard/QuEra)	1013 nm	Commercial ECDL / fiber amplifier	1.0 W	Yes ✓	AOM / EOM

2.3 System Block Diagram



2.4 Nonlinear Optics Required

SHG: 960 nm → 480 nm

For the common 5P-mediated Rb Rydberg route, 480 nm light is generated by doubling a 960 nm ECDL/fiber source in PPKTP, LBO, or BBO, often in a resonant enhancement cavity. Do not confuse this with the Harvard/QuEra 420+1013 nm gate route.

SHG: ~840 nm → 420 nm

For Harvard/QuEra-style Rb gates through the 6P manifold, the lower Rydberg leg is near 420 nm and is typically generated by frequency doubling near-840 nm light. The companion upper leg is near 1013 nm and is infrared, not UV.

2.5 Practical Warnings

- ⚠️ Do not assume one universal Rb Rydberg laser stack: Harvard/QuEra high-fidelity gates use 420+1013 nm via 6P, while many generic Rb lab designs use 780+480 nm via 5P.
- ⚠️ 480 nm optics: use UV-compatible AR coatings (BBAR 400–700 nm). Standard 780 nm coatings reflect 50% at 480 nm.
- ⚠️ Rydberg blue beam should be free-space — fiber-coupling at >50 mW/μm² risks fiber damage and stimulated scattering.

- ⚠️ D2 repumper via EOM sideband: EOM modulation depth must be calibrated; sidebands also create unwanted frequency components that perturb the MOT if not carefully accounted for.
- ⚠️ ODT wavelength: 852 nm is red-detuned from D2 (780 nm) but far-detuned enough for reasonable trap lifetime. Avoid 795 nm ODT — near D1 resonance causes excess scattering.

Expert tip

Start with the 780 nm system and MOT before buying any Rydberg stack. For a Harvard/QuEra-style gate, budget for 420 nm generation and a high-power/stable 1013 nm beam; for a generic 5P-mediated route, budget for the 480 nm system. Most Rb groups share the repumper via EOM sideband rather than a second laser to reduce optical table complexity.

2.6 References & Further Reading

Wieman & Hollberg (1991), *RSI 62* 1 — using diode lasers in atomic physics; laser line selection and AOM schemes Steck alkali data sheets — D1/D2 wavelengths, linewidths, I_sat for all alkali species Grimm, Weidemüller & Ovchinnikov (2000), *Adv. At. Mol. Opt. Phys. 42* — optical dipole traps: frequency selection and power requirements Foot (2005), *Atomic Physics* — Chapter 9: laser cooling and trapping level diagrams Metcalf & van der Straten (1999), *Laser Cooling and Trapping* — the comprehensive guide to MOT, Zeeman slower, and sub-Doppler laser design Saleh & Teich (2019), *Fundamentals of Photonics* — laser gain media, diode laser physics, AOM modulation

CHAPTER 2 SUMMARY

A well-budgeted laser tree prevents power-starved MOTs and noisy traps later in the build. Before turning those lasers onto atoms, however, the atoms need an ultra-high-vacuum environment free of background collisions.

Problems for Chapter 2

1. USE THE TOOL Generate the laser plan for ^{87}Rb. List every beam the planner proposes with its wavelength and the role it plays in the experiment.
2. BY HAND The repumper drives $F = 1 \rightarrow F' = 2$ while the cooler drives $F = 2 \rightarrow F' = 3$. Using the ground and excited hyperfine splittings, compute the frequency offset between them and the EOM frequency needed to make the repumper as a sideband of the cooling laser.
3. DESIGN Generate the plan for ^{171}Yb and compare it with ^{87}Rb: number of lasers, wavelengths that need frequency doubling, and narrow-line requirements. Which system would you rather build on a fixed budget, and why?

Use the tool: on the interactive page. By hand: from the formulas in this chapter. Design: a decision a lab would make. Worked solutions are available to instructors on request (see Part V).

CHAPTER 3 · BUILD STEP 03

3

Ultra-High Vacuum Engineering

⚗️ Vacuum Systems Guide · Guide

Before you trap a single atom, you must first remove 10^{23} of them. The vacuum system is the silent foundation of every AMO experiment — invisible, yet its failure destroys everything above it in minutes.

> *Laser-cooled atoms trapped at microkelvin temperatures are easily knocked out by room-temperature background gas, making ultra-high vacuum (10^{-10}–10^{-11} mbar) the physical foundation of the apparatus.*

IN THIS CHAPTER Prerequisites: Ch. 1

- Relate background gas pressure and collision cross-sections to trap lifetime in MOTs and optical tweezers.
- Design a staged pumping stack (roughing, turbomolecular, ion, NEG/Ti-sublimation) with appropriate CF hardware and bakeout schedules.
- Diagnose real air leaks versus virtual leaks and outgassing using RGA mass spectra (H_2O, $N_2/O_2 \approx 3.7$, Ar, H_2).

WEBSITE COMPANION Live interactive page ↗ · Notebook: Open in Colab · download vacuum-systems.ipynb · Week 4

10^{-10} **Target pressure (mbar)**	2.5×10⁶ **Molecules/cm³ at UHV**	~700 KM **Mean free path at UHV**	~300 S **Atom lifetime at 10^{-10} mbar**

Calculator output at the page's default settings.

3.1 Why Vacuum?

The physics of the invisible enemy — and 400 years of fighting it

The fundamental problem: background gas collisions

A laser-cooled atom or molecule confined in an optical tweezer is cold — maybe 5 μK — but the residual gas in your chamber is at room temperature (~300 K). When a hot N_2 molecule (moving at ~470 m/s) collides with your precious cold atom, the energy transfer easily exceeds the trap depth, and the atom is ejected in microseconds. The collision rate per trapped atom is:

$$\Gamma_{\mathrm{bg}} = n_{\mathrm{bg}}\,\sigma_{\mathrm{eff}}\,\bar{v} \qquad \Rightarrow \qquad \tau_{\mathrm{trap}} = \frac{1}{\Gamma_{\mathrm{bg}}} = \frac{k_B T}{P\,\sigma_{\mathrm{eff}}\,\bar{v}}$$

where $n_{\mathrm{bg}} = P/k_B T$ is the background number density, $\sigma_{\mathrm{eff}} \approx 3 \times 10^{-14}$ cm² is the effective loss cross-section for a room-temperature molecule hitting a trapped atom, and $\bar{v} \approx 470$ m/s for N_2. This reproduces the measured rule of thumb $\tau \approx (2 \times 10^{-8}\ \mathrm{Torr{\cdot}s})/P$ for Rb (Arpornthip, Sackett & Hughes 2012). At $P = 10^{-8}$ mbar the trap lifetime is only ~3 s. At 10^{-10} mbar it reaches ~300 s — long enough for most experimental sequences.

A very brief history

1643

Torricelli inverts a mercury-filled tube into a dish — the mercury column falls to 76 cm, and the empty space above it is the first reproducible vacuum. "Nature does not abhor a vacuum," he concludes: it weighs about one kilogram per cm².

1654

Otto von Guericke evacuates two bronze hemispheres (30 cm diameter) with his newly invented piston pump. Two teams of eight horses cannot pull them apart — a performance staged for Emperor Ferdinand III. Atmospheric pressure is ~10^5 N/m². The Magdeburg hemispheres remain the most theatrical experiment in physics.

1855

Geissler invents the mercury displacement pump, reaching ~10^{-1} mbar — enough for glowing gas discharge tubes and, eventually, neon signs.

1879

Edison's lightbulb requires 10^{-4} mbar to keep the carbon filament from oxidising. A functioning vacuum system was the literal prerequisite for electric light.

1895–97

Röntgen discovers X-rays and **J. J. Thomson** discovers the electron — both inside vacuum tubes. The two defining discoveries of modern physics happened in evacuated glass vessels.

1960s

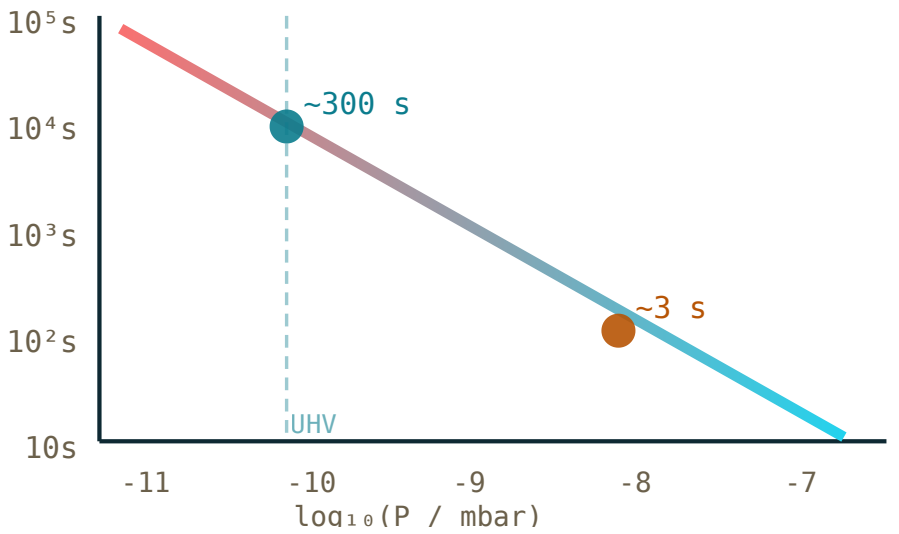
Turbomolecular and ion pumps developed for particle accelerators (CERN, Brookhaven) push the frontier to 10^{-11} mbar. UHV becomes routine in surface science.

1990s–

AMO physics adopts UHV as standard. BEC (Cornell & Wieman, 1995), optical tweezers, and Rydberg arrays all live at 10^{-10}–10^{-11} mbar.

Trap lifetime vs pressure

The difference between 10^{-8} mbar and 10^{-10} mbar is 100× in pressure — and 100× in atom lifetime. Most experiments need *at least* several minutes of trap lifetime to run a meaningful protocol, so UHV is non-negotiable.



💡 **Rule of thumb:** Every decade of pressure improvement multiplies your atom lifetime by 10. Going from 10^{-8} to 10^{-10} mbar turns a ~3-second lifetime into a ~5-minute one.

3.2 The Pressure Landscape

Drag the slider to explore number densities, mean free paths, and what lives at each pressure

Interactive Pressure Explorer

Extreme High Vacuum

PRESSURE	NUMBER DENSITY	MEAN FREE PATH
10^{-10} mbar 10^{-8} Pa 7.5×10^{-11} Torr	**2.5×10^{6} molecules / cm³**	**665 km for N_2 at 293 K**

ATOM LIFETIME

4.8 min background-gas limited

Calculator output at the page's default settings.

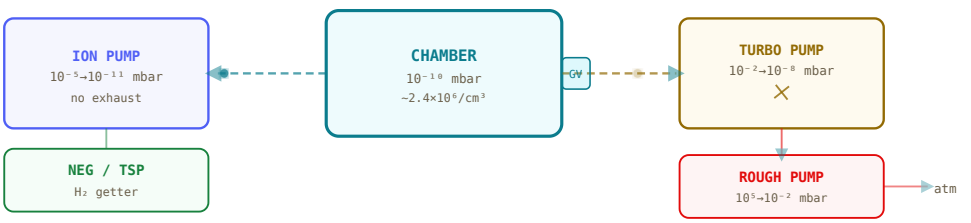
REGIME	PRESSURE RANGE (MBAR)	N (CM^{-3})	MEAN FREE PATH	EXAMPLE APPLICATIONS
Atmosphere	~10^{3}	2.7×10^{19}	~68 nm	Ambient lab air
Low Vacuum	10^{3} – 1	10^{19} – 10^{16}	nm – µm	Rough pumping, CVD
Medium Vacuum	1 – 10^{-3}	10^{16} – 10^{13}	70 µm – 7 cm	Freeze-drying, vacuum ovens
High Vacuum	10^{-3} – 10^{-6}	10^{13} – 10^{10}	7 cm – 70 m	Thin-film deposition, SEMs
Ultra-High Vacuum	10^{-6} – 10^{-9}	10^{10} – 10^{7}	70 m – 70 km	MOTs, surface science, particle accelerators
XHV / AMO	< 10^{-9}	< 10^{7}	> 70 km	Optical tweezers, BEC, Rydberg arrays

$$\lambda = \frac{k_B T}{\sqrt{2}\,\pi\, d^2\, P}$$

Mean free path for N_2 ($d \approx 3.7$ Å, $T = 293$ K). At 10^{-10} mbar: $\lambda \approx 700$ km.

3.3 The Pumping Chain

No single pump spans 15 decades of pressure — you need a staged cascade



Animated gas flow (→) shows the pumping direction. GV = gate valve (isolates the turbo from the chamber during venting or pump failure). Ion pump and NEG/TSP are always-on; turbo + rough are the initial pumpdown path.

🔩

STAGE 1 — ROUGH PUMP (SCROLL OR DIAPHRAGM)

Brings the system from atmosphere to backing pressure

10^{5} → 10^{-2} mbar

The rough pump is first in line. It operates by mechanical compression — gas is drawn in, physically compressed, and pushed out to atmosphere. Two designs dominate AMO labs:

- **Scroll pump** (recommended): Oil-free, two interleaved spiral scrolls orbit to compress gas. No oil mist, no vibration from pistons. Leybold, Agilent, and Edwards make reliable versions used in most tweezer labs.
- **Diaphragm pump**: Even simpler, no sliding contacts — a flexing membrane compresses gas. Lowest ultimate pressure (~5 mbar), best for clean low-throughput applications.
- **Rotary vane** (legacy): Uses oil-lubricated vanes. High throughput but oil backstreaming can contaminate your chamber. Always use a foreline trap if you have one.

Ultimate: **~5×10⁻³ mbar** Speed: **5–30 m³/h** Always backed to **atmosphere**

💡 **Always start with the rough pump.** Never turn on a turbo against atmosphere — the spinning blades will vaporize on contact with air molecules at their operating speed of 50,000–90,000 RPM.

🌀

STAGE 2 — TURBOMOLECULAR PUMP

The workhorse of high and ultra-high vacuum

10^{-2} → 10^{-8} mbar

A turbo pump is essentially a very precisely machined turbine. Blades spin at 30,000–90,000 RPM — fast enough that the blade speed (~300 m/s) is comparable to the mean thermal velocity of residual gas molecules. Molecules that hit the blades are imparted a directed momentum toward the backing port; molecules trying to re-enter from the backing side are physically swept away.

The pumping mechanism is purely kinetic (molecular flow regime). This is why a turbo needs the rough pump to first bring the pressure below ~10 mbar — at higher pressures the mean free path is shorter than the blade spacing and viscous flow prevents the turbo from working.

- **Compression ratio** for N_2: ~10^9; for H_2: ~10^3 (lighter molecules are harder to pump).
- **Tip seals**: Modern turbos use magnetic bearings (vibration-free, no oil). Ideal for optical experiments where vibration couples to trap displacement.
- **Gate valve**: Always place a gate valve between turbo and chamber. If power fails, the gate valve closes and protects your UHV from the backstreaming turbo.

Speed: **50–500 L/s** RPM: **30k–90k** Backed by: **rough pump**

> ⚠️ **Never vent a spinning turbo to atmosphere.** The blades are thin-walled aluminium rotating at near supersonic speeds. A vent valve on the turbo's high-vacuum port controls this — open it slowly through a restricted needle valve after powering down the pump.

⚡

STAGE 3 — ION PUMP

Solid-state, no moving parts, no vibration — the UHV workhorse

10^{-5} → 10^{-11} mbar

Ion pumps have no mechanical parts. They work by ionizing residual gas atoms and then burying the ions inside the pump material. A strong magnetic field (~0.1 T) traps electrons in long helical orbits, massively increasing their path length and ionization cross-section. The resulting ions are accelerated into titanium cathode plates where they are implanted and chemically bound (chemisorption).

- **Pumping mechanism**: Penning discharge. Residual gas is ionized; ions slam into Ti cathodes and are embedded. Reactive gases (O_2, N_2, CO) pump well; noble gases (Ar) pump poorly and can cause "argon instability" — avoid noble gas leaks.
- **Current readout**: The ion pump current is proportional to pressure — it doubles as a gauge. A 20 L/s pump at 10^{-10} mbar draws ~1 nA of current.
- **Magnetic field**: The external magnet stacks make ion pumps heavy (~5–10 kg for a 20 L/s pump). The stray field (~1 mT at 10 cm) can affect magnetic-field sensitive experiments — account for this in your layout.

Speed: **2–200 L/s** Power: **~5 W at UHV** HV supply: **3–7 kV**

✓ **The gold standard:** Once you've baked and reached 10^{-8} mbar with the turbo, close the gate valve, start the ion pump, and leave it running indefinitely. Ion pumps get better with time as they outgas themselves.

🧲

STAGE 4 — NEG PUMP & TITANIUM SUBLIMATION

Chemical getters for H_2 — the dominant UHV residual gas

XHV supplement

Ion pumps do pump H_2, but the effective speed and long-term behavior are strongly pump-design dependent. NEG (Non-Evaporable Getter) pumps and TSP (Titanium Sublimation Pumps) are often added because H_2 is usually the dominant residual gas at UHV and getter surfaces pump it efficiently by sorption into reactive material.

NEG cartridges (e.g., SAES St707: Zr-V-Fe alloy) adsorb H_2, CO, N_2, and O_2 at elevated temperatures or passively near room temperature after activation at ~400°C. They have zero magnetic field, zero vibration, and zero power in passive mode — ideal for sensitivity-limited experiments.

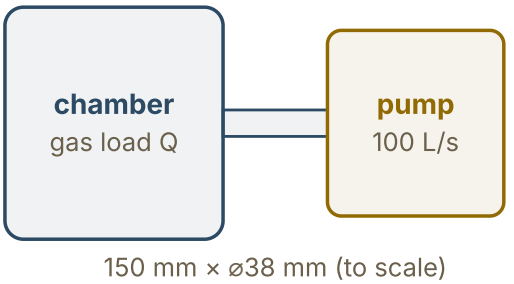
TSP: An electrical current heats a Ti filament until Ti evaporates and coats the chamber walls. The fresh Ti film aggressively chemisorbs nearly all reactive gases. A single filament firing lasts ~30 min; TSPs are typically operated every few days to maintain XHV. They can achieve transient pumping speeds of thousands of L/s.

NEG speed: **5–100 L/s for H_2** TSP filament: **~40 A for 45 s** No moving parts

Interactive

Your pump is only as fast as the pipe in front of it

In UHV the mean free path is far longer than any tube, so a molecule only reaches the pump if its random bounces happen to carry it down the pipe. The pipe has a *conductance* C, and it adds to the pump like resistors in series: $1/S_{\text{eff}} = 1/S + 1/C$. A long, narrow nipple can quietly turn an expensive pump into a small one.



S_eff / S 25 %

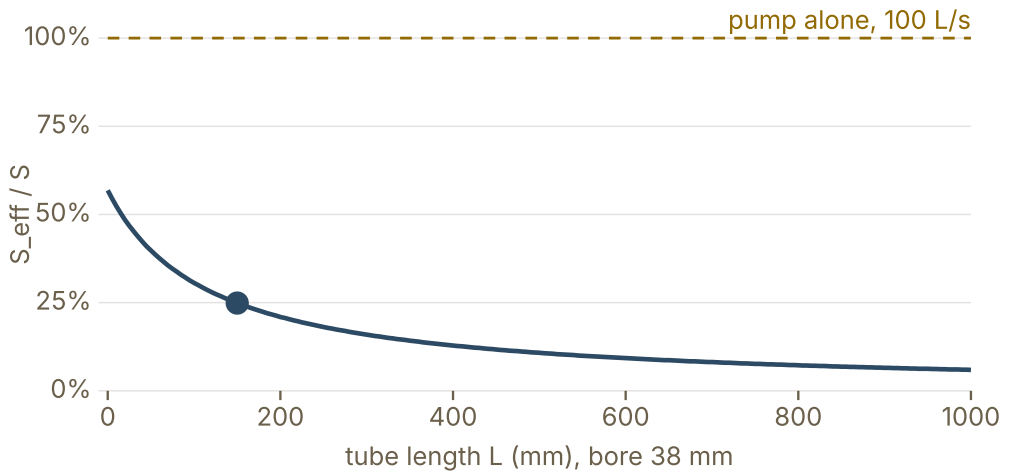
TUBE CONDUCTANCE C
33.1 L/s
aperture 132 + tube 44.3 L/s in series

EFFECTIVE SPEED S_{EFF}
24.9 L/s 25 % of the pump's 100 L/s

CHAMBER PRESSURE VS. PUMP BOLTED ON DIRECTLY
× 4.02
same gas load Q, p = Q/S_{eff}

LENGTH THAT HALVES THE PUMP
16 mm
S_{eff} = S/2 when C = S

Calculator output at the page's default settings.



Molecular flow (mean free path ≫ d), room temperature. Entrance aperture $C_{\text{ap}} = 11.6\,A$ L/s (A in cm²) in series with the long-tube conductance $C_{\text{tube}} = 12.1\,d^3/L$ L/s (d, L in cm), both for air at 20 °C and scaled by $\sqrt{29/M}$ for other gases. Pfeiffer Vacuum, *Vacuum Technology Know-How*, §1.2.8 (Eqs. 1-21, 1-25, 1-31); see also O'Hanlon (below). Adding the aperture to the

tube is the usual approximation for tubes that are not long compared with d; elbows, valves and short tubes need the exact transmission probability (or a Monte Carlo code such as Molflow). Enter S as the pump's speed *for the chosen gas*: ion pumps, for example, pump noble gases much more slowly than N_2.

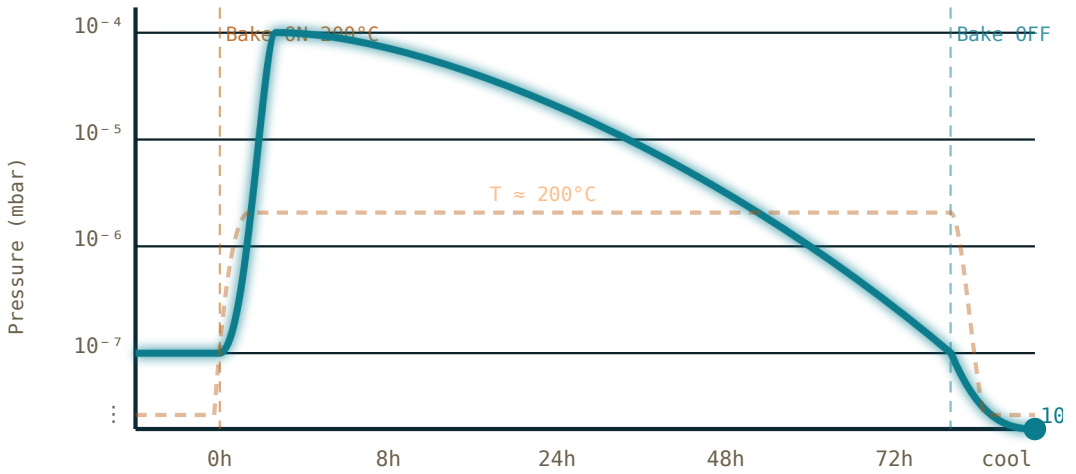
3.4 Baking Out — Cleaning from the Inside

The single most important ritual before first pump-down

Why does every surface hate you?

A freshly machined stainless steel surface has roughly one monolayer of water adsorbed on every exposed face — about 10^{15} H_2O molecules/cm². A typical AMO chamber has perhaps 5,000 cm² of internal surface. That's 5×10^{18} water molecules slowly evaporating into your vacuum over days and weeks. At room temperature, the thermal desorption rate is so slow that it would take *years* to pump all of that water out through the turbo. Baking dramatically accelerates desorption: the outgassing rate scales roughly as e^{-E_a/k_BT} where $E_a \approx 0.5$ eV for physisorbed water. With $E_a \approx 0.5$–1 eV, a modest 150°C bake speeds desorption by a factor of roughly 10^2–10^5.

The bakeout pressure curve (schematic; axis compressed below 10^{-7})



Bakeout protocol

Typical schedule

1. Preparation: Remove all heat-sensitive parts: viton O-rings (max ~200°C), elastomers, electronics boards, optical coatings on unprotected elements. Loosen flange bolts slightly to allow thermal expansion (re-torque during cool-down).

2. Ramp up: Heat to 150–250°C at <5°C/min. Fast ramps crack CF gaskets. Wrap the chamber in aluminium foil + fiberglass heating tape. Thermocouples on flange bodies, not heating tape.

3. Hold: 24–72 hours at temperature. Pressure will spike, then slowly decrease. The turbo (or a dedicated bakeout turbo) should run the whole time. Do not close the gate valve during bake.

4. Cool down: <3°C/min. The biggest pressure drop occurs as you cool below 100°C — remaining water re-adsorbs on any cold surface outside the chamber, leaving your inner surfaces clean.

⚠️ **Viton O-rings** degrade above ~200°C. CF flanges (knife-edge + Cu gasket) are UHV-grade to 450°C and are always preferred.

🚫 **Respect the ion-pump bake rating.** Do not heat non-bakeable magnets, cables, or controllers. Many systems bake the chamber with the ion pump unpowered, sometimes with magnets removed or with bake-rated magnets; check the pump manual before heating above ~100°C. Start the ion pump only after the chamber has cooled and the turbo has pumped to <10^{-7} mbar.

3.5 Materials & Connections

Not everything belongs inside a UHV system — material choice is irreversible

CF vs KF flanges

PROPERTY	CF (CONFLAT)	KF (ISO-KF)
Seal type	Cu knife-edge gasket	Viton O-ring + centering ring
Pressure range	10^{-12} mbar → 10 bar	10^{-8} mbar (practical)
Bakeout	✓ up to 450°C	× max ~120°C (viton)
Reuse	Gasket single-use	Reusable if clean
Use in AMO	Main chamber, sources	Foreline, RGA ports

💡 **CF torque sequence:** Star pattern, 4–5 passes with increasing torque. Final: typically 14–20 N·m for SS bolts on 2¾″ CF. Under-torquing leaks; over-torquing damages the knife edge.

Outgassing rate comparison

MATERIAL	OUTGASSING RATE*	AMO SUITABILITY
304/316L SS (electropolished)	`~10⁻¹² mbar·L/s·cm²`	Excellent
OFHC copper	`~10⁻¹² mbar·L/s·cm²`	Excellent (gaskets)
Fused silica / quartz	`~10⁻¹² mbar·L/s·cm²`	Excellent (windows)

Aluminium (bare)	~10⁻¹⁰ mbar·L/s·cm²	Acceptable; avoid bake
Viton	~10⁻⁸ mbar·L/s·cm²	KF only; not baked
Kapton (wires)	~10⁻⁹ mbar·L/s·cm²	Acceptable in low flux
Regular rubber / PVC	>10⁻⁶ mbar·L/s·cm²	Never use in UHV
Zinc / cadmium	High vapor pressure	Never use

*After 24h bakeout at 200°C. Values highly dependent on surface preparation.

Copper gaskets — handle with care

CF gaskets are OFHC (oxygen-free, high-conductivity) copper — soft enough for the SS knife edges to bite into. Key rules:

- Never reuse a gasket that has been fully torqued. The knife edges work-harden and deform the copper permanently.
- If a flange was slightly loose, inspect the gasket for uneven bite marks — replace if in doubt.
- Nickel-plated copper gaskets exist for higher corrosion resistance; use standard OFHC for baking to 400°C+.
- Store gaskets in sealed bags — oxidised copper is harder and seals less reliably.

What never goes inside a UHV chamber

- **Zinc or cadmium** — vapor pressure too high; will plate out on optics and destroy the vacuum.
- **Regular brass** — contains zinc; dezincifies under vacuum and outgasses.
- **PVC, regular rubber, silicone (RTV)** — extreme outgassing.
- **Fingerprints** — the oils from your skin outgas for weeks. Always use nitrile gloves inside the chamber and clean parts with methanol + acetone before installation.
- **Unprocessed plastics** — PEEK and PTFE are tolerable in small amounts; avoid everything else.

3.6 Reading Your Vacuum

Gauges, RGA, and the art of finding invisible leaks

Gauge types

GAUGE	PRESSURE RANGE	PRINCIPLE	NOTES
Pirani	10^{3} – 10^{-3} mbar	Thermal conductivity of gas (heated wire)	Gas-species dependent; great for rough/medium vacuum readout
Penning (cold cathode)	10^{-2} – 10^{-9} mbar	Ion current in crossed E+B fields	No hot filament = no contamination; reads falsely high in strong B-fields
Bayard-Alpert (hot-ion)	10^{-2} – 10^{-11} mbar	Electron emission → gas ionization → ion current collected	Gold-plated collector reduces X-ray limit; standard UHV gauge
Ion pump current	10^{-5} – 10^{-11} mbar	$I \propto P \times$ pumping speed	Free readout from your existing ion pump supply; calibrate vs gauge
Spinning rotor gauge	10^{-2} – 10^{-7} mbar	Drag on magnetically suspended rotor	Absolute, gas-independent; excellent for calibration

Residual Gas Analyzer (RGA)

Reading the mass spectrum

An RGA is a miniature quadrupole mass spectrometer that ionizes residual gas and sorts ions by mass-to-charge ratio. The peak heights tell you exactly what's in your vacuum. Key signatures:

- **m/z = 2 (H_2):** Often the dominant peak in a clean, baked stainless-steel UHV system. H_2 permeates stainless steel from the bulk, so some residual signal is normal.
- **m/z = 12, 16, 28, 44 (C, O, CO, CO_2):** Hydrocarbon contamination or outgassing from surfaces. Drop significantly after baking.
- **m/z = 18 (H_2O):** High before bakeout, should be much smaller than H_2 after. If water stays high post-bake, suspect a small leak.
- **m/z = 14, 28 (N, N_2):** Elevated N_2 usually means a real air leak. N_2/O_2 ratio ≈ 3.7 (and N_2/Ar ≈ 84) confirms atmospheric air.
- **m/z = 40 (Ar):** Unambiguous real leak indicator — Ar is not outgassed from steel.

Helium Leak Detection

Finding leaks down to 10^{-10} mbar·L/s

The He leak detector is a mass spectrometer tuned to m/z = 4, connected to your system's roughing line or directly to the pump port. He is small (0.26 nm kinetic diameter), inert, and rare in lab air (~5 ppm) — it threads through any crack the atmosphere leaks through, but faster.

Procedure:

- Pump your system to the best vacuum you can achieve.
- Spray He gas gently with a wand (or medical-grade He tubing) around suspect joints — CF flanges, viewports, electrical feedthroughs.
- Wait 5–10 s (He transit time through a small leak).
- A rise in the He signal identifies the leaking joint.
- Start from the top and work down so He (lighter than air) doesn't drift upward into non-leaking regions.

💡 Sensitivity: ≈ 10^{-10} mbar·L/s. A "virtual leak" (trapped volume slowly releasing gas) mimics a real leak but won't respond to He spraying — the key diagnostic distinction.

3.7 Practical Checklist

From assembly to first UHV — a condensed field guide

Assembly phase

1. Clean all parts with acetone then methanol; handle only with nitrile gloves from this point forward.
2. Verify all materials: no zinc, cadmium, PVC, or standard rubber inside the system.
3. Use fresh OFHC Cu gaskets on every CF flange. Torque in star pattern: 3 passes from 3 → 8 → 14 N·m.
4. Install a gate valve between chamber and turbo pump. Verify it closes cleanly.
5. Install an ion gauge on the chamber. Label its filament voltage and emission current settings.
6. Leak check all connections with a He leak detector at rough vacuum (10^{-2} mbar) before proceeding.

Pumpdown phase

7. Start scroll/diaphragm pump. Rough pump to <5×10^{-2} mbar before starting the turbo.
8. Bring turbo to full speed. Open gate valve slowly. Monitor ion gauge.
9. Once below 10^{-6} mbar, begin bakeout ramp: <5°C/min up to 150–200°C.

- **10.** Bake for 24–72h. Run ion gauge intermittently. Monitor RGA: water peak should decrease.
- **11.** Cool at <3°C/min. Below 50°C: close gate valve, start ion pump, turn off turbo.
- **12.** Fire TSP once (if installed). Ion pump current should read <1 nA for a 20 L/s pump at UHV.

✓ **You're at UHV when:** Ion gauge reads <10^{-9} mbar, ion pump current is stable and <10 nA, RGA shows H_2 dominant with no N_2/Ar signature, and your pilot atom lifetime exceeds several minutes.

⚠️ **Common failure modes:** Virtual leak from trapped volumes (blind tapped holes, undrilled screws); fingerprint contamination on optics; bad CF torque — either sequence or final value; or baking too fast and cracking a ceramic feedthrough.

STAGE	START P (MBAR)	TARGET P (MBAR)	TIME	KEY ACTION
Rough pump only	1013	~10^{-2}	10–30 min	Verify no gross leaks
Turbo pump on	10^{-2}	~10^{-6}	1–4 hours	He leak test
Bakeout ramp	10^{-6}	~10^{-4} (spike)	2–5 hours	Remove heat-sensitive parts
Bakeout hold	10^{-4}	~10^{-7}	24–72 hours	Monitor RGA, run turbo
Cool + ion pump	10^{-7}	~10^{-10}	12–24 hours	TSP fire, NEG activate

See also: The vacuum system is just the starting point. Once you're at UHV, the next steps are trapping atoms with MOT design, characterizing your imaging with the Imaging SNR Calculator, measuring the resulting atom temperature via TOF Thermometry, and transferring atoms to tweezers for single-atom thermometry.

3.8 References & Further Reading

Foundational texts, key experimental papers, and online resources for vacuum technology in AMO labs.

O'Hanlon (2003), *A User's Guide to Vacuum Technology* — the standard UHV engineering reference Jousten ed. (2016), *Handbook of Vacuum Technology* — pump types, gauging, materials, bakeout procedures Chambers et al. (1998), *Basic Vacuum Technology* — accessible intro to pump selection and system design Anderson et al. (1995), *Science* 269 198 — first BEC in rubidium: the UHV recipe that made it possible Davis et al. (1995), *PRL* 75 3969 — BEC in sodium; established the UHV/evaporation paradigm Kurt J. Lesker Technical Library — practical guides on CF flanges, feedthroughs, getter pumps Steck alkali data sheets — vapor pressure curves for Rb, Cs, Li, Na, K (essential for dispenser current)

CHAPTER 3 SUMMARY

Achieving 10^{-11} mbar vacuum extends single-atom lifetimes to tens of seconds or minutes. With the chamber evacuated and the lasers chosen, our next step is locking every laser frequency to within a fraction of the atomic linewidth.

Problems for Chapter 3

1. USE THE TOOL With a 100 L/s pump behind a 38 mm bore, 150 mm long tube, find the effective pumping speed at the chamber. Now double the tube length. What limits the speed, the pump or the tube?
2. BY HAND From Section 03, list the pumping chain from atmosphere to ultra-high vacuum with the pressure range of each stage. Explain why an ion pump cannot be started at atmosphere.
3. DESIGN Your tweezer lifetime drops from 20 s to 5 s overnight. Using Section 06, describe how you would use the ion-pump current, a residual-gas analyser and the atoms themselves to tell a real leak from outgassing or a virtual leak.

Use the tool: on the interactive page. By hand: from the formulas in this chapter. Design: a decision a lab would make. Worked solutions are available to instructors on request (see Part V).

CHAPTER 4 · BUILD STEP 04

4

Laser Frequency Stabilization

🔐 Laser Locking Guide · Guide

Free-running diode lasers drift by MHz on minute timescales, comparable to atomic linewidths. Three complementary techniques solve this: SAS for absolute references, beat-note for precise offsets, and PDH for sub-kHz linewidths, plus the vacuum, thermal, and vibration engineering that makes the reference cavity itself stable enough to lock to.

> *Free-running diode lasers can drift by many megahertz over minutes—far wider than a typical 5–6 MHz atomic linewidth—so active frequency locking is required before atoms can be cooled or imaged.*

IN THIS CHAPTER Prerequisites: Ch. 1, Ch. 2

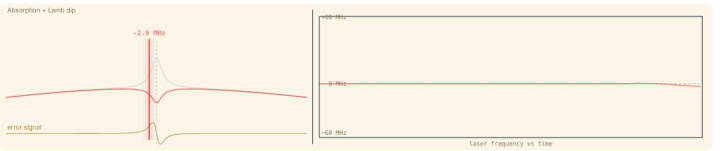
- Generate dispersive error signals using saturated absorption spectroscopy (SAS), Pound-Drever-Hall (PDH) cavity locking, and RF beat-note offset locks.
- Match servo loop bandwidths (current vs. PZT channels) to laser frequency noise spectra.
- Identify common locking failure modes including residual amplitude modulation, etaloning, and acoustic cavity resonances.

WEBSITE COMPANION Live interactive page ↗ · Notebook: Open in Colab · download laser-locking.ipynb · Week 3

Laser frequency stability is foundational to every AMO experiment, a coherence time of 100 μs requires laser linewidth below ~10 kHz, and Rydberg transitions need sub-MHz stability. Saturated absorption spectroscopy (SAS) provides Doppler-free atomic references a few MHz wide (set by the natural linewidth, $\Gamma/2\pi \approx 5$–6 MHz for alkali D lines, plus power broadening); Pound-Drever-Hall (PDH) locking to an optical cavity achieves sub-Hz linewidths by using phase-modulated sidebands to generate a dispersive error signal. This guide explains the physics of both techniques and gives you tools for exploring SAS spectra, PDH error signals, and offset locking, the three pillars of practical laser stabilization. A fourth tab, **Ultrastable Cavities**, goes one level deeper: PDH only works as well as the cavity you lock to, and that cavity's finesse, vacuum, thermal, and vibration engineering deserve their own treatment.

Running example: Cs D_2 at 852 nm ($\Gamma/2\pi = 5.234$ MHz) for SAS and beat-note; 685 nm Cs E_2 quadrupole line ($\Gamma/2\pi \approx 117$ kHz) to motivate PDH locking.

Assumptions: the SAS, beat-note, and PDH formulas are idealized small-signal models. Real locks are limited by detector saturation, residual amplitude modulation, servo bandwidth and phase margin, RF mixer conversion loss, optical feedback, etalon fringes, cavity drift, and acoustic pickup. Use the calculated slopes and SNRs as design targets, then verify with measured error-signal slope and in-loop/out-of-loop spectra.



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

offset: 2.9 MHz

Try: high **P** + low **D** to see ringing · low **P** + high **D** for overdamped crawl · hit **Perturb** while locked

1 / 5

Saturated-Absorption Spectroscopy (SAS) Locking

SAS provides an **absolute optical frequency reference** tied directly to an atomic transition, no external frequency standard needed. The technique exploits the narrow Lamb dip hidden beneath the broad Doppler-broadened absorption profile of a room-temperature vapour cell.

How it works

Step 1, Counter-propagating beams.
A strong pump beam and weak probe beam travel in opposite directions through the vapour cell. An atom moving at velocity v sees the pump Doppler-shifted to $\nu_0(1 - v/c)$ and the probe to $\nu_0(1 + v/c)$.

Step 2, Velocity selectivity.
Only atoms with $v \approx 0$ (zero-velocity class) are simultaneously resonant with both beams at the unshifted line centre ν_0. The pump saturates these atoms, depleting the ground-state population and creating a "Bennett hole" in the velocity distribution.

Step 3, Lamb dip.
The saturated zero-velocity atoms absorb less probe power at ν_0. A narrow Lamb dip appears in probe transmission, with FWHM $\approx \Gamma\sqrt{(1 + I/I_{sat})}$, set by the natural linewidth, not by the 300–500 MHz Doppler envelope.

Step 4, FM error signal.
The laser frequency (or probe EOM) is weakly frequency-modulated. Lock-in demodulation at the modulation frequency yields a dispersive first derivative of the Lamb dip, a zero-crossing exactly at ν_0 to use as the servo error signal.

$$V_{\text{err}} \propto \frac{dS(\nu)}{d\nu} \quad \text{[first-harmonic demodulation]}$$

$$\text{Lamb dip FWHM} \approx \frac{\Gamma}{2\pi}\sqrt{1 + I/I_{\text{sat}}} \quad \text{[power-broadened]}$$

$$\text{Doppler FWHM} = \frac{\nu_0}{c}\sqrt{\frac{8k_{\text{B}}T\ln 2}{m}}$$

$$I_{\text{sat}} = \frac{\pi hc\Gamma}{3\lambda^3} \quad \text{[two-level, natural linewidth]}$$

Crossover resonances

Why they appear

In multi-level atoms (real alkalis), multiple excited states ν_a, ν_b share a common ground state. An atom moving at velocity $v^* = c(\nu_b - \nu_a)/(\nu_a + \nu_b)$ sees the pump resonant with ν_a and the counter-propagating probe resonant with ν_b (or vice versa). This produces an extra Lamb dip at the *midpoint* between the two transitions.

$$\nu_{\text{cross}} = \frac{\nu_a + \nu_b}{2}$$

For pump resonant with ν_a: $v^* = c(\nu_{\text{laser}} - \nu_a)/\nu_a$
Probe seen by same atom at v^*: $\nu_{\text{probe}} = \nu_{\text{laser}} + \nu_a v^*/c = \nu_b$ ✓

Key properties

Crossover features are often **stronger and narrower** than true Lamb dips because *two* velocity classes contribute, one class has pump resonant with v_a and probe with v_b; a second class has the roles reversed. Both pathways deplete the same ground state, doubling the effective saturation.

In Cs D_2, the crossover at (F=4→F'=3 + F=4→F'=5)/2 is the most commonly used lock point: it is sharp, high-contrast, and well-separated from neighbours.

💡 **Practical:** Cs D_2 has 3 allowed transitions (F=4→F'=3,4,5) plus 3 crossovers at (3+4)/2, (4+5)/2, and (3+5)/2. Map them with a wavemeter before choosing your lock point, crossovers are ~2× stronger in many cases.

Modulation Transfer Spectroscopy (MTS), superior alternative

How MTS differs from FM-SAS

In standard FM-SAS, the *laser frequency* is modulated, the derivative of the entire absorption profile is detected, including the broad Doppler background. Residual amplitude modulation (RAM) and etalon fringes superimpose on the useful signal.

In MTS, the *pump beam* is phase-modulated at Ω via an EOM. Through **four-wave mixing** in the resonant medium, the pump modulation transfers to the probe beam *only* at atomic resonances, the Doppler background does not participate. Demodulating the probe at Ω gives a *dispersive signal on a near-zero baseline*.

Four-wave mixing origin

The modulated pump (fields at ω±Ω) drives a coherent population grating at Ω. A counter-propagating probe scatters off this grating, acquiring sidebands at ω±Ω. Beat detection at Ω produces the dispersive error signal without a lock-in amplifier.

Advantages: no Doppler pedestal, immune to laser FM-to-AM conversion (RAM), 5–20× better SNR than FM-SAS in practice.

Limitation: requires a pump EOM and collinear geometry. Not all transitions support efficient four-wave mixing (weak ones don't).

METHOD	BASELINE	TYPICAL SNR	SETUP COMPLEXITY
Direct absorption	Doppler envelope	Low (~1:100)	Simplest, beamsplitter + PD
FM-SAS (current/freq. mod.)	Doppler slope + Lamb dip derivative	Medium (~1:20)	Modulate laser current, demodulate PD
MTS (pump EOM)	**Near-zero (flat)**	**High (~1:3)**	Pump EOM + demodulate probe at Ω

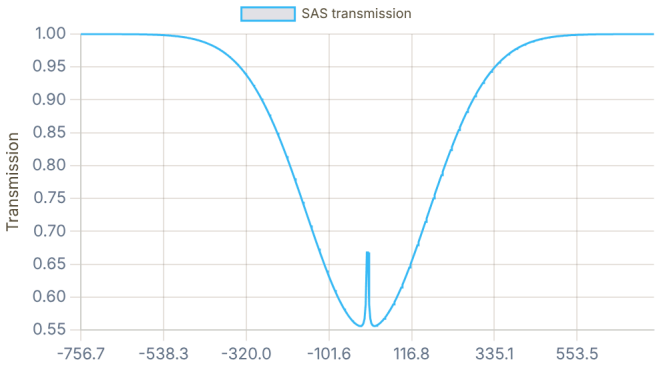
Species-specific considerations

SPECIES	WAVELENGTH	Γ/2Π	I_{sat}	CELL TEMP.	KEY NOTES
Cs	852 nm (D_2)	5.23 MHz	1.1 mW/cm²	25–40 °C	Rich HFS; (3+5)/2 crossover is best lock; 9192 MHz GS splitting (primary standard)
^{87}Rb	780 nm (D_2)	6.07 MHz	1.67 mW/cm²	25–50 °C	Most used BEC species; 6.835 GHz GS split; easy SAS at room temperature
Na	589 nm (D_2)	9.80 MHz	6.3 mW/cm²	100–130 °C	High I_{sat} → needs more pump power; yellow → dye laser or SHG; 1772 MHz GS split
^{6}Li	671 nm (D_1+D_2)	5.87 MHz	2.54 mW/cm²	300–400 °C	Ground hyperfine splitting is 228 MHz; D_1/D_2 fine-structure splitting is ~10 GHz; excited-state HFS is only a few MHz (barely resolved); **heated oven cell essential**; MTS strongly preferred; congested spectrum
^{40}K	767 nm (D_2)	6.04 MHz	1.75 mW/cm²	60–100 °C	Natural abundance 0.012% → use enriched cell or long path length; 1286 MHz GS split

⚠ **Why SAS fails for weak (E2, M1) transitions:** forbidden or weak transitions usually give poor room-temperature SAS signals because the oscillator strength, vapor pressure, optical pumping dynamics, and transit-time/collision broadening work against a clean Lamb dip. Do not estimate their visibility from D-line saturation-intensity intuition alone. In practice, PDH locking to a high-finesse cavity or direct spectroscopy in a cold-atom/lattice system is the cleaner route.

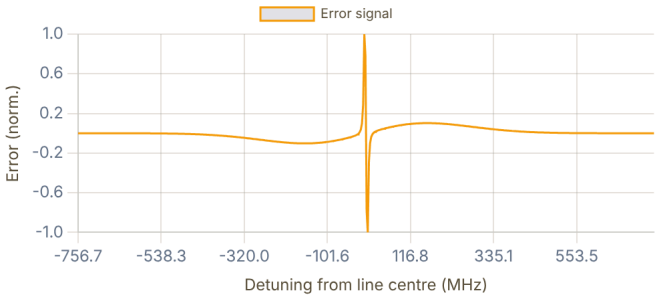
Calculator, Doppler, Lamb dip, saturation & pump power

SAS TRANSMISSION (DOPPLER + LAMB DIP)



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

FM ERROR SIGNAL DS/DN



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

- **Vapour cell temperature:** 25–60 °C for Cs/Rb; 300–400 °C for Li (oven cell). Higher T → stronger absorption but broader Doppler, more collisional broadening, and risk of alkali deposition on windows.
- **Pump power:** needs $I \gtrsim I_{sat}$ to produce a visible Lamb dip. Excess power broadens the dip (FWHM $\propto \sqrt{(1 + I/I_{sat})}$). Optimum: 3–10× I_{sat}.
- **Pump:probe ratio:** probe should be ~10–20% of pump to avoid probe saturation. Route pump through a retro-reflector or PBS to generate counter-propagating beams.
- **Polarisation:** σ^+ pump + σ^- probe (or vice versa) works well for MTS; linear polarisations work for FM-SAS. Tilt beams 5–10 mrad to avoid pump back-scatter into probe path.
- **Modulation frequency:** 1–50 MHz typical. Higher pushes 1/f noise below shot-noise floor. For FM-SAS: Ω should be ≤ Lamb dip FWHM for maximum slope. MTS can use Ω up to ~Γ/2 for an optimum.
- **Crossover resonances:** often strongest features (2× signal contribution); Cs F=4→(3+5)/2 and Rb F=2→(2+3)/2 are common lock points.

- **Li-6 special case:** the ground hyperfine splitting is 228 MHz, while the D_1/D_2 fine-structure splitting is ~10 GHz and the excited-state hyperfine splittings are only a few MHz. Use a heated cell at ~350°C. MTS is often preferred. If D1 and D2 lasers are both needed, lock one laser to the cell and offset/beat-note lock the other at the required fine-structure plus hyperfine offset.
- **SAS fails for E2/M1 transitions** where I_{sat} is ≫ W/cm², use PDH cavity locking.

Beat-Note (Offset) Locking

Many experiments need two lasers separated by a **precise and stable frequency offset** , for example, a cooling laser and repumper pair, or two Raman beams. Beat-note locking stabilises the *difference frequency* between a well-stabilised master laser and a slave laser, without locking each independently to an atomic reference.

Frequency discriminator vs optical phase-locked loop (OPLL)

Frequency discriminator lock

An RF chain converts the instantaneous beat frequency to a voltage. Error signal: $V_{err} \propto (\nu_{beat} - \nu_{ref})$. The servo corrects *frequency* but does not track *phase* — slave and master are not phase-coherent.

Servo bandwidth: a few hundred kHz suffices.
Slave linewidth: inherits master's long-term frequency stability; short-term phase wanders freely.
Applications: cooling + repumper pairs, probe lasers, any non-interferometric use where absolute phase doesn't matter.

Optical phase-locked loop (OPLL)

The servo tracks the *phase* of the beat: $d(\varphi_{beat} - \varphi_{ref})/dt = 0$. The slave becomes a **phase-coherent offset copy** of the master: their relative phase is constant and set entirely by the RF reference.

Servo bandwidth: must exceed the free-running laser linewidth (500 kHz – 5 MHz for ECDLs → bandwidth ≥ 5–20 MHz).
Slave linewidth: transfers master's linewidth at the offset.
Applications: Raman spectroscopy, atom interferometry, STIRAP, EIT, optical lattice clocks, quantum gates.

💡 OPLL requires fast current modulation (not just PZT), careful RF cable length matching to minimise loop delay, and typically a DDS-driven RF reference for agile frequency control.

Physics of the beat note

$$E_1(t) = E_{01}\cos(2\pi\nu_1 t + \varphi_1) \qquad E_2(t) = E_{02}\cos(2\pi\nu_2 t + \varphi_2)$$

$$i_{PD}(t) \propto E_{01}\cdot E_{02}\cdot\cos[2\pi(\nu_2 - \nu_1)t + (\varphi_2 - \varphi_1)]$$

Freq. discriminator: $V_{err} \propto (\nu_2 - \nu_1) - \nu_{ref}$

OPLL: $\frac{d}{dt}[\varphi_{beat} - \varphi_{ref}] = 0 \;\Rightarrow\; \nu_2 - \nu_1 \equiv \nu_{ref}$ (phase-coherent)

Signal-to-noise ratio & detector requirements

Beat SNR formula

The beat photocurrent has a DC component $I_{dc} = R(P_1+P_2)$ and an AC component at $\Delta\nu$. Shot-noise-limited SNR in bandwidth B:

$$SNR_amp = 2R\sqrt{(P_1P_2)} \,/\, \sqrt{(2eI_dcB)}$$

where R is responsivity (A/W), e is the electron charge, and the numerator is the beat-current amplitude. In RF power terms at 50 Ω, using RMS beat current:

$$P_RF\ (\mathrm{dBm}) = 10\ \log_{10}[2R^2P_1P_2Z_0 \,/\, 1\ \mathrm{mW}]$$

Rule of thumb: need ≥ 20 dB SNR (linear factor 100) for a reliable lock. With 0.5 mW on each arm and R = 0.5 A/W, the beat note carries $P_RF = 2R^2P_1P_2Z_0 \approx -22$ dBm into 50 Ω (before any amplifier); the SNR then depends on the noise floor within your detection bandwidth, and is typically ample.

BEAT FREQ.	REQ. PD BW	DETECTOR TYPE	EXAMPLE
1–100 MHz	≥ 300 MHz	Si PIN	Thorlabs FDS010
100 MHz–1 GHz	≥ 1.5 GHz	Si/InGaAs PIN fast	Hamamatsu G4176
1–3 GHz	≥ 5 GHz	InGaAs PIN	Thorlabs DET08CFC
3–10 GHz	≥ 15 GHz	InGaAs balanced	Discovery DSC-R402
>10 GHz	≥ 25 GHz	UTC-PD or resonant	NTT IOD-PMF-22

💡 PD bandwidth ≥ 3–5× beat frequency. Add a 10–20 dB low-noise RF amplifier (e.g. Mini-Circuits ZX60-3018G) before the mixer, it pays off immediately in SNR for sub-mW optical powers.

Phase noise & coherence transfer

Phase noise in a frequency lock

In a frequency discriminator lock, the slave phase drifts freely relative to the master outside the servo bandwidth. The slave linewidth is roughly equal to the free-running slave linewidth reduced by the servo suppression within its bandwidth.

Phase noise PSD $S_\varphi(f)$: characterises phase fluctuations at offset frequency f from the carrier. The servo suppresses $S_\varphi(f)$ for f below the unity-gain frequency, at the cost of a "servo bump" near the bandwidth edge.

Coherence transfer in OPLL

The OPLL forces $\varphi_{slave}(t) = \varphi_{master}(t) + \varphi_{ref}(t)$. If φ_{ref} comes from a low-noise DDS or RF synthesiser, the slave's coherence length equals the master's, the two beams interfere with long coherence even separated by GHz.

Raman transitions require $\Delta\varphi = \varphi_1 - \varphi_2$ stable over the pulse duration (µs–ms scale). With OPLL, $\Delta\varphi$ is set by the RF synthesiser ($\ll 1$ mrad noise), enabling high-contrast Rabi oscillations.

AOM-based frequency offsets

When to use an AOM for offsets

For offsets <30 MHz or when the beat falls in 1/f noise, route one arm through a **double-pass AOM** at $\nu_{AOM} \approx$ 80–110 MHz. This shifts the optical frequency by $2\nu_{AOM} \approx$ 160–220 MHz, placing the beat note well into the shot-noise-limited region.

AOM + beat combination

Route one arm through a double-pass AOM, then beat against master. The beat is $|\Delta\nu_{laser} + 2\nu_{AOM}|$. Lock to that beat. Then:

v_slave = v_master + v_RF_ref − 2v_AOM

Tuning the AOM RF frequency scans the slave without changing the lock point — powerful for spectroscopy scans without mode hops.

Calculator, beat-note frequency, SNR & detector recommendation

Commercial electronics

Vescent D2-125

Laser servo controller for SAS, PDH, and general dispersive locks. Accepts an error or PD input, provides fast current + slow PZT feedback (servo BW up to ~10 MHz), and includes built-in ramp/lock-in circuitry. Widely used in AMO labs as a single-box servo.

Toptica FALC 110

Fast analog laser controller. Bandwidth: DC–10 MHz (current), ~10–30 kHz (PZT / auxiliary). Accepts any external error signal, used when you generate your own discriminator or PDH error. High dynamic range, differential inputs reduce RF pickup. Standard in high-finesse PDH setups and OPLLs.

Vescent D2-135 / D2-105

D2-135: Offset Phase Lock Servo dedicated to beat-note / OPLL locking (10 MHz–9.5 GHz with external fast PD). D2-105: ultra-low-noise Laser Controller (current and temperature driver, not a frequency-locking servo).

Applications & typical offset frequencies

APPLICATION	SPECIES	OFFSET	LOCK TYPE	NOTES
Cooling + repumper	^{6}Li	228 MHz (D_1/D_2)	Freq. discriminator	Lock D_2 repumper to D_1 SAS-locked master
Cooling + repumper	^{87}Rb	6835 MHz	OPLL or chain	Large offset → 10 GHz PD or divide-by-N prescaler chain
Raman pair (2-photon)	^{87}Rb / Cs	6.8 / 9.2 GHz	**OPLL required**	Phase coherence essential for Rabi contrast
EIT / STIRAP	Rb / Cs / Li	100 MHz – 1 GHz	OPLL preferred	Phase noise limits dark-state contrast and transfer efficiency
Probe laser offset	Any	Arbitrary	Freq. discriminator	AOM-accessible range: 40–400 MHz; convenient for scanning
PDH transfer cavity	Any	FSR multiple	Freq. discriminator	Lock two lasers to same ULE cavity via beat against anchor laser

- **Master must already be well-locked** (SAS or PDH). The slave inherits the master's long-term stability, no reference means the whole chain drifts.
- **Fibre combiner vs free-space:** single-mode 50:50 fibre combiner gives automatic mode overlap; free-space PBS requires careful polarisation alignment. Fibre adds ~3 dB loss per arm but produces cleaner spectra.
- **RF amplifier chain:** in the usual low-signal beat-note setup, amplify before the mixer so the LO/RF ports meet the mixer drive specification, often +10 to +20 dBm LO. A 20 dB low-noise amp pays off immediately in SNR for sub-mW optical powers.
- **OPLL bandwidth requirement:** servo bandwidth must exceed the free-running laser linewidth (typically 0.5–5 MHz for ECDLs). Fast current modulation (not PZT alone) is mandatory for OPLL.
- **Very small offsets (<30 MHz):** beat note falls within 1/f noise. Use a double-pass AOM to shift one arm by 160–220 MHz before detecting; lock to the shifted beat note.

- **Offsets >3 GHz:** need fast InGaAs PD (≥ 10 GHz BW) or a microwave prescaler chain to divide the beat to a more manageable frequency. For >10 GHz, consider a frequency comb reference.
- **AOM frequency scanning:** scan the AOM RF frequency to scan the slave laser continuously while keeping the beat lock engaged, ideal for high-resolution spectroscopy without mode hops.
- **RF synthesiser sets the offset:** changing ν_ref tunes the slave without touching the master. A DDS (direct digital synthesiser) allows agile, phase-continuous frequency changes.

📖 This tab includes the PDH derivation, Fabry–Pérot cavity formulas, error-signal intuition, circuit diagram, and practical locking questions. Jump directly to the PDH section: PDH Locking Guide →

Pound–Drever–Hall (PDH) Cavity Locking

When no convenient atomic reference exists, or when the required short-term linewidth is narrower than SAS can provide, the laser is locked to a **high-finesse Fabry–Pérot cavity**. The PDH technique generates an error signal from the reflected field using RF phase modulation, achieving much higher signal slope than simple transmission locking.

$$\nu_{\mathrm{FSR}} = \frac{c}{2nL} \quad [\text{free spectral range}]$$

$$\mathcal{F} = \frac{\pi\sqrt{R}}{1-R} \approx \frac{\pi}{1-R} \quad [R \approx 1]$$

$$\delta\nu_{\mathrm{cav}} = \frac{\nu_{\mathrm{FSR}}}{\mathcal{F}} \quad [\text{cavity linewidth FWHM}]$$

$$\frac{\delta\nu}{\nu} = -\frac{\delta L}{L} = -\alpha_{\mathrm{CTE}}\,\delta T \quad [\text{fractional freq. drift}]$$

PDH error signal, how it works

Step 1, Phase modulation. An EOM imprints weak sidebands at ±Ω on the carrier field:
$E_{\mathrm{in}} \approx E_0 e^{i\omega t} + \frac{\beta}{2} E_0 e^{i(\omega+\Omega)t} - \frac{\beta}{2} E_0 e^{i(\omega-\Omega)t} \quad (\beta \ll 1)$

Step 2, Frequency placement. Choose Ω ≫ δν_cav so the sidebands sit outside the cavity resonance, they reflect unchanged. The carrier acquires a strongly frequency-dependent phase shift near resonance.

Step 3, Demodulation. Detect the reflected field on a fast photodiode. The carrier beats against the sidebands at Ω. Demodulate at Ω → bipolar error signal:

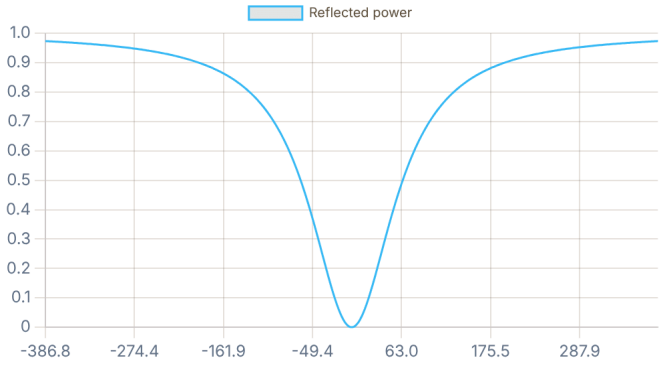
$$V_{\mathrm{err}} \propto \mathrm{Im}[r(\omega)] \quad [\text{imaginary part of reflection coefficient}]$$

Near resonance: $V_{\mathrm{err}} \approx K \cdot \delta$ where $\delta = \omega - \omega_{\mathrm{cavity}}$
Discriminator slope:

$K = 8\sqrt{P_{\mathrm{carrier}} \cdot P_{\mathrm{sideband}}}/\delta\nu_{\mathrm{cav}} = 8\sqrt{P_c P_s}\,\mathcal{F}/\nu_{\mathrm{FSR}}$

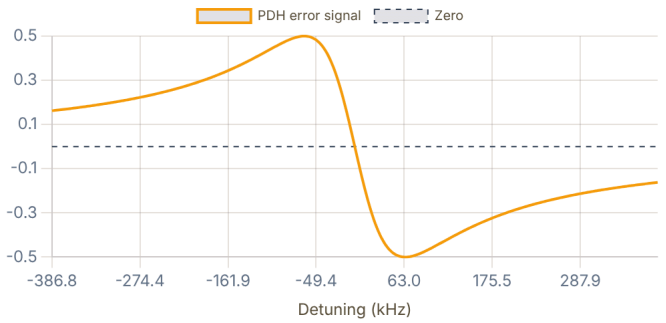
(Black 2001)
($P_{\mathrm{sideband}} \approx \beta^2 P_{\mathrm{carrier}}/4$)

CAVITY REFLECTION (AIRY DIP)



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

PDH ERROR SIGNAL IM[R(Ω)]

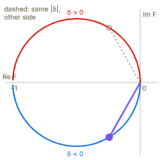


Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Interactive

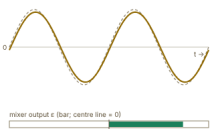
WHY PDH READS THE PHASE, NOT THE POWER

Sweep the laser through the cavity line and watch the reflected carrier field $F(\delta)$ in the complex plane. Its *size* is the same at $+\delta$ and $-\delta$ (that is the symmetric Airy dip above), so power alone can't tell the servo which way to push. Its *phase* runs the other way round the circle on each side. Beating the carrier against the sidebands turns that phase into an RF signal whose sign flips at resonance.



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

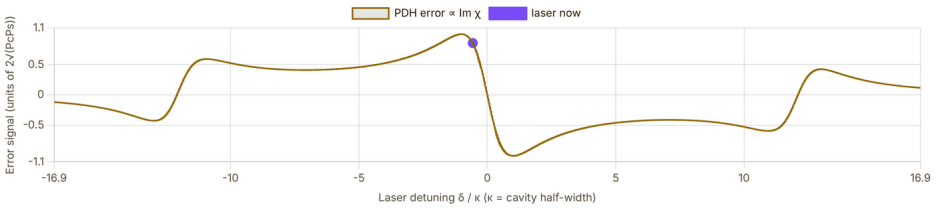
Reflected carrier $F(\delta)$ in the complex plane (κ = cavity half-width)



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Photodiode signal at Ω over two RF periods (solid) and the mixer's local oscillator (dashed)

reflected power |F|² = 0.243 (same at δ = 0.57κ) · phase arg F = -120°
RF beat: -0.071 cos Ωt + 0.854 sin Ωt → mixer output ε ∝ 0.854 (laser below resonance: error positive, push up)



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Model: Black's PDH signal $\epsilon = 2\sqrt{P_c P_s}\,\mathrm{Im}\,\chi$, $\chi = F(\delta)F^*(\delta+\Omega) - F^*(\delta)F(\delta-\Omega)$, for a lossless, impedance-matched cavity near one resonance, $F(\delta) = i(\delta/\kappa)/(1 - i\delta/\kappa)$, the same Lorentzian as the two charts above. The photodiode term at Ω is $\propto \mathrm{Re}[\chi]\cos\Omega t + \mathrm{Im}[\chi]\sin\Omega t$ and the mixer keeps the $\sin\Omega t$ part. When $\Omega \gg \delta\nu_{\mathrm{cav}}$ the sidebands reflect with $F \approx -1$ and $\epsilon \approx -4\sqrt{P_c P_s}\,\mathrm{Im}F(\delta)$, which is linear in δ near resonance while the reflected power $|F|^2 \approx (\delta/\kappa)^2$ has no first-order term at all. The slider's default Ω is a teaching value chosen so all three features fit on one plot; the button switches to the calculator's own ratio. Source: E. D. Black, *Am. J. Phys.* 69, 79 (2001).

Typical optical chain:
diode laser → prism pair (astigmatism corr.) → optical isolator (30–40 dB)
→ EOM (phase mod. at Ω) → mode-matching telescope → ULE cavity

Reflected beam: PBS + QWP → fast PD → RF mixer at Ω → LPF
Error signal: fast path → laser current (high BW, small range)
slow path → PZT (low BW, large range)

- **ULE zero-crossing temperature:** ULE has CTE ≈ 0 near a specific temperature (5–25 °C depending on blank). Operating at the zero-crossing dramatically reduces thermal drift.

- **Vacuum and vibration isolation:** cavity must be in vacuum (P < 10^{-5} mbar) to avoid refractive-index fluctuations. Mount on vibration-isolated platform.
- **Two-stage servo:** fast path (current) compensates high-frequency noise; slow path (PZT) compensates slow drift.
- **PDH does not provide absolute frequency:** the cavity resonance drifts. Beat the locked laser against a frequency comb or SAS reference for absolute knowledge.
- **Finesse measurement:** scan laser across a resonance and fit Airy function, or measure the (intensity) ring-down time τ_c = 𝓕/(2π·FSR).

📖 New: a dedicated deep-dive on **building** the ultrastable reference cavity that a PDH lock stabilizes to, the vacuum, thermal, and vibration engineering that the servo electronics alone don't address. Based on Jamie A. Boyd & Thierry Lahaye, *Am. J. Phys.* 92, 50 (2024) →

Why the Cavity Itself Is the Hard Part

The PDH tab shows how to derive a dispersive error signal from a Fabry–Pérot cavity, but the quality of the reference you lock to is set entirely by the mechanical, thermal, and vacuum design of the cavity spacer, not by the servo electronics. A useful organizing constraint: to hold the laser's fractional frequency stability below a target value, the cavity's fractional refractive-index and length fluctuations must satisfy

$$\frac{\delta\nu}{\nu} \;\le\; \frac{\delta n}{n} + \frac{\delta L}{L} \;\lesssim\; 10^{-9} \quad \text{[target for sub-MHz stability at optical frequency]}$$

Every design choice below, vacuum level, spacer material and temperature, and support geometry, exists to drive one of the two terms on the right toward zero.

Finesse, free spectral range, and mode structure

$$\nu_{\mathrm{FSR}} = \frac{c}{2nL} \quad \text{[free spectral range]}$$

$$\mathcal{F} = \frac{\pi\sqrt{r_1 r_2}}{1 - r_1 r_2} = \frac{\nu_{\mathrm{FSR}}}{\delta\nu_{\mathrm{cav}}} = 2\pi\,\tau\,\nu_{\mathrm{FSR}} \quad [\tau = 1/e \text{ ring-down time}]$$

Build intuition first: $|r| = 98.4\% \Rightarrow \mathcal{F} \approx 100 \Rightarrow \delta\nu_{\mathrm{cav}} \approx 15$ MHz; push the mirrors to $|r| = 99.992\% \Rightarrow \mathcal{F} \approx 20{,}000 \Rightarrow \delta\nu_{\mathrm{cav}} \approx 75$ kHz. A real ultrastable cavity in the paper (10 cm long, 780 nm) reaches $\mathcal{F} = 25{,}200$, FSR ≈ 1.5 GHz, linewidth ≈ 60 kHz, and a measured ring-down time $\tau \approx 2.6$–$2.7\ \mu$s, matching the manufacturer's spec.

Geometry matters as much as reflectivity. Higher-order transverse modes (HOMs) pick up an extra Gouy phase per round trip and resonate at

$$\frac{\nu_{mnp}}{\nu_{\mathrm{FSR}}} = p + \frac{m+n+1}{\pi}\arccos\left[\sqrt{\left(1-\tfrac{L}{R_1}\right)\left(1-\tfrac{L}{R_2}\right)}\right]$$

A **confocal** cavity ($L = R_1 = R_2$) is a bad choice, its HOMs land exactly at FSR/2, degenerate with the neighboring TEM_{00} resonance. The standard practical choice is **planoconvex** (one flat mirror, one curved), with L/R tuned so HOMs sit roughly 100 MHz from the TEM_{00} peak, far enough to avoid mode-overlap-induced lock errors.

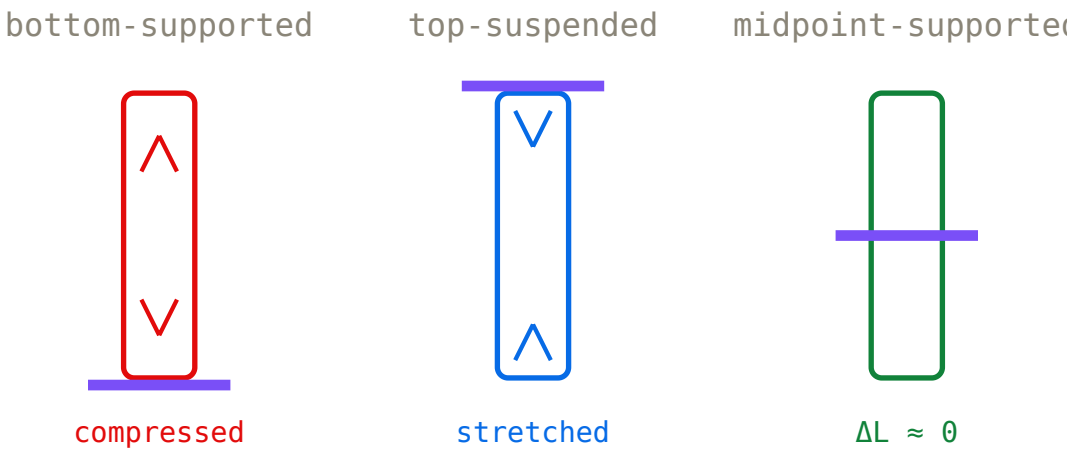
Controlling δn/n: the vacuum requirement

Air has $n - 1 \approx 3 \times 10^{-4}$ at STP, and atmospheric pressure alone drifts by roughly 10% day to day, so an unevacuated cavity sees $\delta n/n \sim 10^{-5}$: four orders of magnitude too large for the 10^{-9} budget above. Pumping to below 10^{-3} mbar already fixes this; in practice, ion pumps (no moving or spinning parts, so they inject no vibration) routinely hold $\sim 10^{-7}$ mbar with no bakeout, pushing $\delta n/n \sim 10^{-14}$, fully negligible.

Controlling δL/L: the ULE zero-crossing

$\delta L/L = \alpha(T)\,\delta T$, so the spacer's thermal-expansion coefficient matters as much as the temperature control itself. ULE (ultra-low-expansion) glass is a SiO_2/TiO_2 mix engineered so that $\alpha(T)$ crosses zero at a specific **zero-crossing temperature** T_{ZC}, chosen a few degrees above room temperature purely so a resistive heater, not a cooler, can hold the setpoint (the paper's own cavity: $T_{ZC} = 34.49\,^{\circ}\mathrm{C}$, held with a Wavelength Electronics LFI-3751 heater controller). At $T_{ZC} \pm 10$ mK: $\delta L/L \sim 10^{-14}$ for a 10 cm spacer. Even 0.5 °C off the zero crossing it's still only $\sim 10^{-11}$, which is exactly why a simple heater-plus-PID loop suffices instead of a full cryostat.

The single most non-obvious idea in cavity design: where you hold it matters more than what you hold it with. Under vertical acceleration a, a cavity spacer of length L supported at its *bottom* is compressed, $\Delta L = a\rho L^2/2E$ (ρ = density, E = Young's modulus), while one suspended from the *top* stretches by the same amount. Support it instead at the horizontal plane through its **midpoint** and the two effects cancel to first order, $\Delta L \approx 0$. This is why real ultrastable cavities are mounted horizontally on specific nodal support points, often Viton (fluoroelastomer) balls or pads for damping and thermal isolation, rather than wherever is mechanically convenient.



The fundamental floor: thermal (Brownian) noise

> Vacuum, thermal control, and vibration isolation all fix *drift*, they don't touch the fundamental floor. The spacer, substrates, and especially the dielectric mirror *coatings* undergo genuine thermodynamic (Brownian) motion that scales as
>
> $\sqrt{T}$
>
> , this is noise, not drift, and no amount of the engineering above removes it. State of the art gets around it by going cryogenic: single-crystal silicon spacers operated at their own zero-crossing temperature (~16 K), which both suppresses Brownian noise directly (lower T) and removes the amorphous-material aging that limits room-temperature ULE cavities over the long term.

Mode matching and intracavity power

$$z_R = \sqrt{L(R_2 - L)}, \qquad w_0 = \sqrt{\frac{z_R \lambda}{\pi}} \quad \text{[planoconvex cavity, flat mirror at } z = 0\text{]}$$

$$\frac{I_{\mathrm{cav}}}{I_{\mathrm{inc}}} = \frac{\mathcal{F}}{\pi} \quad \text{[on-resonance build-up factor]}$$

Worked example from the paper: $L = 100$ mm, $R_2 = 500$ mm $\Rightarrow z_R = 200$ mm $\Rightarrow w_0 = 223\ \mu$m at 780 nm. Their practical recipe: a fiber collimator (~100 µm waist) feeding a single $f = 150$ mm lens, placed 204 mm from the collimator and 421 mm from the cavity's flat mirror, matches this eigenmode.

> **Common mistake:** for $\mathcal{F} = 25{,}200$, the build-up factor is $\mathcal{F}/\pi \approx 8{,}000\times$, so as little as ~100 µW incident can already build up to nearly 1 W circulating. Too much incident power into a high-finesse cavity can heat or damage the mirror coating, always check the build-up factor before raising input power. Diagnose mode matching with a transmission camera (look directly for HOM spatial patterns) and the reflected-signal dip depth (aim for ≥80% of the fully-reflected level).

A concrete PDH lock, start to finish

Example RF chain (Boyd & Lahaye's setup)

20 MHz modulation via a fiber EOM, ~350 µW incident power, function-generator LO (Tektronix AFG3022C), photodiode (Thorlabs PDA10-EC), RF amplifier (Mini-Circuits ZFL-500+), mixer (Mini-Circuits ZAD-1H+), and servo (Vescent D2-125), the same servo family used on this page's Beat-Note tab.

1. Sweep the cavity resonance and adjust the demodulation LO phase to minimize/optimize the reflected sideband beat.

2. Add a 90° phase shift from that optimum, the phase that minimizes the in-phase sideband beat is 90° away from the phase that maximizes the actual PDH error-signal slope.

3. Zero the DC offset of the error signal.

4. Slowly shrink the scan amplitude while re-centering on the target resonance (usually the fundamental TEM_{00} peak).

5. Engage the lock and watch the transmitted-power photodiode, confirm it sits at the peak, not on a nearby HOM wing, and stays there.

6. Tune the PID gains for stability without ringing (see the P/D sliders in the lock animator above for the same intuition).

Measuring finesse independently: cavity ring-down

$$I(t) = I_0\, e^{-t/\tau}, \qquad \mathcal{F} = 2\pi\, \tau\, \nu_{\mathrm{FSR}}$$

Abruptly extinguish or divert the input light (AOM or fast shutter) while locked and resonant, then fit the transmitted intensity decay. Boyd & Lahaye measured $\tau = 2.6\ \mu$s with FSR $= 1.5$ GHz $\Rightarrow \mathcal{F} \approx 25{,}000$, matching the mirror manufacturer's spec, a nice independent, in-lab cross-check rather than trusting a spec sheet number alone.

Typical design requirements at a glance

PARAMETER	REQUIREMENT	TYPICAL VALUE (PAPER'S CAVITY)
Mirror reflectivity	as high as practical	$\mathcal{F}$ ≈ 25,200 measured (equivalent R ≈ 99.988%)
Vacuum pressure	< 10^{-3} mbar minimum	~10^{-7} mbar (ion pump)
Temperature stability	hold near T_{ZC}	±10 mK
Spacer material	near-zero CTE	ULE glass (T_{ZC} = 34.49 °C here)
Spacer length	typical	10 cm
Resulting linewidth	kHz-level or below	~60–75 kHz
Drift	slow, characterizable	~hundreds of kHz/day
Fundamental floor	thermal (Brownian) noise in coatings	overcome only cryogenically (~16 K, Si)

- **Cross-check the spec sheet:** always measure finesse independently by ring-down rather than trusting the mirror coating vendor's number alone.

- **Mount at the nodal points:** Viton balls or pads at the calculated midpoint support locations damp vibration and thermally isolate the spacer at the same time.
- **Budget the HOM spacing, not just the finesse:** a planoconvex geometry with HOMs ~100 MHz from TEM_{00} avoids mode-overlap lock errors that a purely high-finesse, badly-chosen-geometry cavity would suffer.
- **Ion pumps over turbo/scroll pumps** wherever possible: no moving parts means no vibrational or acoustic noise injected into the very cavity you're trying to keep still.
- **This tab complements, not replaces, the PDH tab:** the PDH tab covers the servo and error-signal electronics; this tab covers the mechanical, thermal, and vacuum design of what you're locking to. If you're looking for atom-cavity coupling (cavity QED) rather than laser frequency stabilization, see the Cavity QED Coupling Lab instead, a different physics problem that happens to share the word "cavity."

The Locking Hierarchy

In practice, the three techniques form a hierarchy. SAS provides the absolute anchor; beat-note locks propagate stability to other lasers at controllable offsets; PDH locks provide narrow-linewidth operation wherever no atomic reference is available.

TECHNIQUE	ABSOLUTE?	TYPICAL LINEWIDTH	TUNABLE OFFSET?	BEST FOR	LIMITATION
SAS lock	✓ Atomic line	100 kHz – 1 MHz	× Fixed to transition	Primary absolute reference (D lines)	Needs strong transition; vapour-cell lines only
Beat-note lock	Via master	Same as master	✓ RF synthesiser	Cooling/repump pairs, Raman beams	Needs pre-stabilised master; fast PD + RF chain
PDH cavity lock	× Cavity drifts	1 Hz – 10 kHz	Via AOM after lock	Narrow-linewidth spectroscopy, weak transitions	Thermal cavity drift; expensive; needs vacuum

Typical laboratory hierarchy

Layer 0 (Absolute): D-line laser → SAS lock → ν known to ~1 MHz Layer 1 (Derived): Laser B → beat-note lock to Layer 0 → $\nu_0 \pm \nu_RF$ Laser C → beat-note lock to Layer 0 → $\nu_0 \pm \nu_RF'$ Layer 2 (Narrow): Spectroscopy laser → PDH lock to ULE cavity → ~kHz linewidth (Long-term drift corrected by beating against a Layer 0 laser)

Key equations at a glance

SAS: $V_{\text{err}} \propto dS(\nu)/d\nu$, Lamb dip FWHM

$\approx (\Gamma/2\pi)\sqrt{1 + I/I_{\text{sat}}}$

Beat-note: $i_{\text{PD}} \propto \cos[2\pi(\nu_2 - \nu_1)t]$, $V_{\text{err}} \propto (\nu_2 - \nu_1) - \nu_{\text{ref}}$

PDH: $V_{\text{err}} \propto \text{Im}[r(\omega)] \approx K \cdot \delta$,

$$\delta\nu_{\text{cav}} = \frac{c}{2L\mathcal{F}}, \qquad \frac{\delta\nu}{\nu} = -\frac{\delta L}{L} = -\alpha_{\text{CTE}}\,\delta T$$

Error signal slopes compared

TECHNIQUE	SLOPE K (TYPICAL)	NOTES
SAS (direct lock-in)	~0.1–1 mV/MHz	Limited by Doppler background contrast
SAS (modulation transfer)	~1–10 mV/MHz	Better baseline; uses four-wave mixing
Beat-note (freq. discriminator)	~1–10 mV/MHz	Scales with RF power and mixer gain

PDH (high finesse)	~10–1000 mV/MHz	Scales as √(P_c·P_s)/δν_cav = √(P_c·P_s)·ℱ/FSR

Our lab: Cs + Li system

LAB LASER HIERARCHY (CS/LI TWEEZER EXPERIMENT)

- **852 nm Cs D_2:** SAS-locked in Cs vapour cell → primary absolute reference.
- **685 nm Cs E_2 ($6S \rightarrow 5D_{5/2}$):** PDH-locked to ULE cavity (L = 77.5 mm, $\mathcal{F} \approx 1.5\times10^4$). FSR ≈ 1.93 GHz, δν_cav ≈ 130 kHz, laser linewidth ≈ 1 kHz. Thermal drift ≈ 2.5 kHz per 10 mK.
- **671 nm Li D_1:** SAS-locked in heated Li vapour cell.
- **671 nm Li D_2:** Beat-note locked to Li D_1 (Vescent D2-135), offset set by RF synthesiser.

4.1 Sources

Black (2001), an introduction to Pound-Drever-Hall locking Drever et al. (1983), *Appl. Phys. B* 31 97 — laser phase and frequency stabilization using an optical resonator Steck alkali data sheets for Cs, Rb, Li, Na, K line constants NIST Atomic Spectra Database Wieman & Hollberg (1991), *RSI* 62 1 — using diode lasers in atomic physics; the definitive practical guide Hänsch & Couillaud (1980), *Opt. Comm.* 35 441 — polarization spectroscopy for cavity locking without RF modulation Ludlow et al. (2015), *Rev. Mod. Phys.* 87 637 — optical atomic clocks; covers ultra-narrow laser stabilization at the frontier Salomon, Hils & Hall (1988), *JOSA B* 5 1576 — laser stabilization at the millihertz level Boyd & Lahaye (2024), *Am. J. Phys.* 92 50 — a basic introduction to ultrastable optical cavities for laser stabilization; source for this page's Ultrastable Cavities tab (peer-reviewed) Boyd & Lahaye, arXiv:2308.08015 — free preprint of the above

CHAPTER 4 SUMMARY

Locked lasers provide the frequency stability required for reproducible detuning and narrow-line cooling. Next, we verify that the polarization of each beam matches the atomic selection rules.

Problems for Chapter 4

1. USE THE TOOL In the saturated-absorption tab, raise the vapour-cell temperature from 300 K to 330 K. Describe what happens to the Doppler background and the sub-Doppler features, and explain both.
2. BY HAND Two beams of 1 mW each fall on a photodiode with responsivity 0.5 A/W. Compute the amplitude of the beat-note photocurrent, $2R\sqrt{(P_1 P_2)}$, and compare with the beat-note calculator.
3. DESIGN A cavity is 100 mm long with finesse 10,000 at 780 nm. Compute its free spectral range and linewidth, choose a PDH modulation frequency, and justify the choice. Check your numbers with the PDH tab.

Use the tool: on the interactive page. By hand: from the formulas in this chapter. Design: a decision a lab would make. Worked solutions are available to instructors on request (see Part V).

CHAPTER 5 · BUILD STEP 05

5

Polarimetry & Optical State Tomography

🔭 Polarimetry Simulator · Simulator

Interactive learning tool for optical polarimetry. Visualize any polarization state, simulate the rotating-QWP measurement (Schaefer et al. 2007), extract Stokes parameters via Fourier analysis, and understand why S_3 (circular) requires a QWP, not just an HWP.

> *Atomic selection rules (σ^+, σ^-, π) depend critically on polarization purity; a few percent of unintended ellipticity can spoil optical pumping, MOT balance, or magic-polarization tweezers.*

IN THIS CHAPTER Prerequisites: Ch. 2, Ch. 4

- Represent polarization states on the Poincaré sphere using Stokes vectors and propagate them through waveplates and polarizers with Mueller matrices.
- Align quarter- and half-waveplates systematically to convert linear polarization to pure σ± light after mirrors and vacuum viewports.
- Connect optical Stokes polarimetry to single-qubit quantum state tomography and maximum-likelihood density-matrix reconstruction.

WEBSITE COMPANION Live interactive page ↗ · Week 2

Every beam in an AMO experiment has a polarization that must be carefully controlled. σ^+/σ^- selection rules determine optical pumping efficiency into stretched states; π beams drive $\Delta m=0$ transitions for spectroscopy; the Rydberg excitation polarization selects which magnetic sublevel of the Rydberg manifold is excited. The Stokes parameters (S_0, S_1, S_2, S_3) and the Poincaré sphere provide a complete coordinate system for any polarization state, and understanding how waveplates ($\lambda/4$, $\lambda/2$) and polarizing beam splitters transform a Stokes vector is an everyday skill for building and aligning AMO optical setups.

Convention and calibration note: RCP/LCP and σ^+/σ^- labels depend on whether you define handedness looking into the beam or along the propagation direction. This page uses an optics/Stokes convention for visualization. Before mapping to atomic $\Delta m = \pm 1$ selection rules, fix the beam propagation direction, quantization axis, and lab convention. Real polarimetry also needs calibration of QWP retardance, PBS extinction, detector offsets, and angle zero.

1 / 6

Light is an electromagnetic wave. **Polarization** describes how the electric field vector **E** oscillates as the wave propagates. Click below to see the E-field for each type.

Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Linear 0° (Horizontal): E-field oscillates along x-axis. Stokes: S=[1,1,0,0].

Types of Polarization

Linear Polarization

E-field oscillates in a fixed plane at angle α. Created by a linear polarizer or PBS.

$$\mathbf{S} = [1,\ \cos 2\alpha,\ \sin 2\alpha,\ 0]$$

Circular Polarization

E-field rotates in a circle. **RCP**: clockwise looking into beam. Created by linear + QWP at 45°.

$$\mathbf{S} = [1,\ 0,\ 0,\ \pm 1]$$

Elliptical Polarization

Most general case. Described by orientation ψ and ellipticity χ. Linear/circular are special cases.

$$\mathbf{S} = [1,\ \cos 2\chi \cos 2\psi,\ \cos 2\chi \sin 2\psi,\ \sin 2\chi]$$

Key Optical Components

PBS, Polarizing Beam Splitter

Splits light into orthogonal linear polarizations. Extinction ratio: **~1:1,000** (cube) or **~1:100,000** (GT PBS).

Transmit Horizontal Reflect Vertical

HWP, Half-Wave Plate

Phase retardance **π**. Rotates linear polarization by 2θ (fast axis at θ). Flips circular handedness.

Used in Scenario 1 & 2 (rotating HWP + PBS). **Cannot detect S_3!**

QWP, Quarter-Wave Plate

Phase retardance **π/2**. Converts linear ↔ circular. QWP at 45° converts RCP/LCP to linear.

Central element in the **rotating QWP method**, detects all four Stokes parameters.

Lab Scenarios (from your notes)

Scenario 1, Linear polarization input:

SOURCE

Unknown linear

HWP

Half-wave plate

θ: 0 → 180°

PBS

Fixed horizontal

DETECTOR

Sinusoidal I(θ) curve. Formula: I(θ) = ½[1 + cos(4θ − 2α)]. Maximum at θ = α/2 reveals polarization angle.

Scenario 2, Circular polarization input:

SOURCE

Circular pol.

QWP

at 45° (fixed)

HWP

rotating

θ: 0 → 180°

PBS

DETECTOR

Without QWP: circular gives a **flat line**. QWP at 45° converts RCP/LCP → linear, then HWP+PBS sees a sinusoidal curve.

The Stokes formulation represents polarization using **four intensity measurements**, no complex math needed. All values are directly measurable with a power meter.

The Four Parameters

S_0, TOTAL INTENSITY

Total beam power. For normalized states, $S_0 = 1$.

$$S_0 = I_H + I_V$$

S_1 (Q), HORIZONTAL VS VERTICAL

+1 = fully horizontal, −1 = fully vertical, 0 = equal.

$$S_1 = I(0^\circ) - I(90^\circ)$$

S_2 (U), DIAGONAL PREFERENCE

+1 = fully +45°, −1 = fully −45°, 0 = symmetric.

$$S_2 = I(45^\circ) - I(135^\circ)$$

S_3 (V), CIRCULAR HANDEDNESS

+1 = fully RCP, −1 = fully LCP. Cannot be measured with HWP+PBS alone!

$$S_3 = I_{\mathrm{RCP}} - I_{\mathrm{LCP}}$$

Calculator output at the page's default settings.

Poincaré Sphere

Every fully polarized state maps to a point on the unit sphere.

+S_1 equator

Linear 0°

+S_2 equator

Linear +45°

$+S_3$ north

Right Circular

$-S_1$ equator

Linear 90°

$-S_2$ equator

Linear −45°

$-S_3$ south

Left Circular

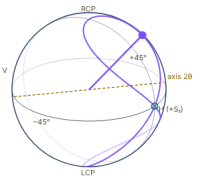
Common States Table

STATE	S_0	S_1	S_2	S_3	LOCATION
Horizontal (0°)	1	+1	0	0	$+S_1$ equator
Vertical (90°)	1	−1	0	0	$-S_1$ equator
Linear +45°	1	0	+1	0	$+S_2$ equator
Linear −45°	1	0	−1	0	$-S_2$ equator
Right Circular (RCP)	1	0	0	+1	North pole
Left Circular (LCP)	1	0	0	−1	South pole
Linear at angle α	1	cos 2α	sin 2α	0	Equator at 2α
Unpolarized	1	0	0	0	Center (inside)

Interactive

Waveplates move points on the sphere

The geometric picture behind every method on this page: a waveplate with retardance δ and fast axis at angle θ **rotates the Stokes vector by δ about an axis in the equator at longitude 2θ**. A half-wave plate (δ = π) reflects linear states across that axis; a quarter-wave plate (δ = π/2) turns linear light at 45° to its axis into circular light at a pole. Turn the plate and the output traces a curve: a double loop around the equator for a HWP, a figure-eight through both poles for a QWP. That curve is exactly the I(θ) signal a rotating-waveplate polarimeter reads out.



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Waveplate

Input

Fast axis θ: **22°**

output S_1 = 0.506, S_2 = 0.500, S_3 = 0.703
orientation ψ = 22.3°, ellipticity χ = 22.3° (elliptical)
○ input · ● output · rotation axis

Mueller matrix of a linear retarder (this page's convention, RCP = $+S_3$); the path is drawn for θ from 0° to 180°. Watch the sweep or drag the slider (the pause button stops the sweep); drag the sphere to turn it.

Fully Polarized Light

$$S_0^2 = S_1^2 + S_2^2 + S_3^2$$

$$\text{DOP} = \frac{\sqrt{S_1^2 + S_2^2 + S_3^2}}{S_0}$$

DOP = 1.0 for fully polarized; DOP = 0.0 for unpolarized light.

Ellipse Parameters

$$\psi = \tfrac{1}{2}\,\text{atan2}(S_2,\ S_1) \quad [\text{orientation},\ 0{-}180^\circ]$$
$$\chi = \tfrac{1}{2}\arcsin(S_3/S_0) \quad [\text{ellipticity},\ \pm 45^\circ]$$

χ = 0 → Linear | χ = ±45° → Circular

Three approaches with increasing sophistication. The rotating QWP method (Schaefer 2007) recovers **all four Stokes parameters** from a single continuous rotation.

Method A: Classical (6 Separate Measurements)

Directly measure I with a linear polarizer at 0°, 45°, 90°, 135°, and with a QWP for circular.

$$S_0 = I(0^\circ) + I(90^\circ)$$
$$S_1 = I(0^\circ) - I(90^\circ)$$
$$S_2 = I(45^\circ) - I(135^\circ)$$
$$S_3 = I_{\text{RCP}} - I_{\text{LCP}}$$

S_3 requires a QWP at 45° before the detector.

> **Limitation:** Six separate alignments. Drift between shots introduces error. Tedious for S_3.

Method B: Rotating HWP + PBS

SOURCE

HWP

θ: 0 → 360°

PBS

Fixed horizontal

DETECTOR

$$I(\theta) = \tfrac{1}{2}\left[S_0 + S_1\cos(4\theta) + S_2\sin(4\theta)\right]$$

No S_3 term — circular component is invisible. For circular input [S_1=S_2=0]: I(θ) = S_0/2 = flat line.

> **Key limitation:** HWP flips S_3 but PBS transmits only linear, so S_3 cancels out. To detect circular, you need a fixed QWP at 45° first (Scenario 2), or use Method C below.

Method C: Rotating QWP + LP (Schaefer 2007) ★

SOURCE

QWP

Fast axis at θ

θ: 0 → 360°

LP

Linear polarizer

Fixed (horiz.)

DETECTOR

$$I(\theta) = \tfrac{1}{2}\left[S_0 + S_1\cos^2(2\theta) + S_2\sin(2\theta)\cos(2\theta) - S_3\sin(2\theta)\right]$$

Using trig identities (Schaefer 2007 expansion):

$$I(\theta) = a_0 + b_2\sin(2\theta) + a_4\cos(4\theta) + b_4\sin(4\theta)$$

where:

$$\begin{aligned} a_0 &= S_0/2 + S_1/4 && \text{(DC offset)} \\ b_2 &= -S_3/2 && \text{(2nd harmonic encodes } S_3\text{!)} \\ a_4 &= S_1/4 && \text{(4th harmonic cosine)} \\ b_4 &= S_2/4 && \text{(4th harmonic sine)} \end{aligned}$$

Recovery:

$$S_1 = 4a_4 \qquad S_2 = 4b_4 \qquad S_3 = -2b_2 \qquad S_0 = 2a_0 - 2a_4$$

Why this wins: S_3 appears in the 2nd harmonic, unique to the QWP method. A single continuous rotation gives all four parameters via Fourier analysis. Noise averages out across the full sweep.

Comparison

METHOD	S_1	S_2	S_3	NOTES
Classical (6 shots)	✓	✓	✓	6 separate alignments, drift-prone
Rotating HWP + PBS	✓	✓	×	Cannot detect S_3 (circular)
Rotating QWP + LP ★	✓	✓	✓	Single rotation, Fourier fit, all 4 params

Set any input polarization state, choose the measurement method, and see the exact I(θ) curve live. Fourier analysis extracts all Stokes parameters automatically.

E-Field (xy-plane view)

I(Θ), INTENSITY VS. WAVEPLATE ANGLE
Raw measurement signal. The Plotly toolbar (top-right) lets you download as PNG or SVG.

EXTRACTED STOKES PARAMETERS
Recovered via Fourier analysis. Compare with input.

POINCARÉ SPHERE
Current state (cyan dot).

FOURIER DECOMPOSITION OF I(Θ)
The I(θ) curve broken into its harmonic components. Each encodes a Stokes parameter.

Calculator output at the page's default settings.

Mueller matrices operate on Stokes vectors [S_0,S_1,S_2,S_3]. Jones matrices operate on the complex E-field vector [Ex, Ey·e^iδ].

Mueller Matrices

Horizontal Linear Polarizer

$$M_{\text{LP}} = \frac{1}{2}\begin{pmatrix} 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}$$

HWP (fast axis at θ)

$$M_{\text{HWP}}(\theta) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos 4\theta & \sin 4\theta & 0 \\ 0 & \sin 4\theta & -\cos 4\theta & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix}$$

Note: $S_3 \rightarrow -S_3$ (flips circular handedness, but PBS still loses it).

QWP (fast axis at θ)

$$M_{\text{QWP}}(\theta) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos^2 2\theta & \sin 2\theta \cos 2\theta & -\sin 2\theta \\ 0 & \sin 2\theta \cos 2\theta & \sin^2 2\theta & \cos 2\theta \\ 0 & \sin 2\theta & -\cos 2\theta & 0 \end{pmatrix}$$

This matrix generates the $S_3 \rightarrow \sin(2\theta)$ term in I(θ).

Jones Vectors & Matrices

$$\begin{aligned}\text{H linear: } &[1,\ 0]^\top \\ \text{V linear: } &[0,\ 1]^\top \\ +45^\circ: &\tfrac{1}{\sqrt{2}}[1,\ 1]^\top \\ -45^\circ: &\tfrac{1}{\sqrt{2}}[1,\ -1]^\top \\ \text{RCP: } &\tfrac{1}{\sqrt{2}}[1,\ +i]^\top \\ \text{LCP: } &\tfrac{1}{\sqrt{2}}[1,\ -i]^\top\end{aligned}$$

Handedness conventions differ between textbooks; on this page RCP is the state with $S_3 = +1$, consistent with the Jones → Stokes relations below.

Jones → Stokes

For $[E_x,\ E_y e^{i\delta}]$ normalized to $S_0 = 1$:

$$\begin{aligned}S_0 &= |E_x|^2 + |E_y|^2 \\ S_1 &= |E_x|^2 - |E_y|^2 \\ S_2 &= 2|E_x||E_y|\cos\delta \\ S_3 &= 2|E_x||E_y|\sin\delta\end{aligned}$$

Inverse (S_0=1):

$$\begin{aligned}|E_x| &= \sqrt{(1+S_1)/2} \\ |E_y| &= \sqrt{(1-S_1)/2} \\ \delta &= \operatorname{atan2}(S_3,\ S_2)\end{aligned}$$

Key Formulas Cheatsheet

FORMULA	WHAT IT GIVES
$I_{QWP}(\theta) = \tfrac{1}{2}[S_0+S_1\cos^2 2\theta+S_2\sin 2\theta\cos 2\theta-S_3\sin 2\theta]$	Rotating QWP + LP intensity
$I_{HWP}(\theta) = \tfrac{1}{2}[S_0+S_1\cos 4\theta+S_2\sin 4\theta]$	Rotating HWP + PBS (no S_3 term!)
$I(\theta) = a_0 + b_2\sin 2\theta + a_4\cos 4\theta + b_4\sin 4\theta$	QWP Fourier expansion
$S_1=4a_4,\ S_2=4b_4,\ S_3=-2b_2,\ S_0=2a_0-2a_4$	Stokes from Fourier coefficients
$\mathrm{DOP} = \sqrt{S_1^2+S_2^2+S_3^2}/S_0$	Degree of polarization
$\psi = \tfrac{1}{2}\cdot\operatorname{atan2}(S_2,S_1)$	Polarization orientation (0–180°)
$\chi = \tfrac{1}{2}\cdot\arcsin(S_3/S_0\cdot\mathrm{DOP})$	Ellipticity angle (±45°)
$I_{linear}(\theta)\vert_{HWP} = \tfrac{1}{2}[1+\cos(4\theta-2\alpha)]$	HWP scan for linear pol. at angle α

Step-by-Step: Rotating QWP Method

Step 1, Setup

Place rotating QWP (motorized or manual) followed by a fixed horizontal linear polarizer. Connect output to a power meter or photodiode.

Step 2, Measure I(θ)

Rotate QWP 0°→360°, recording intensity at each angle. Need ≥50 points for good SNR. More points = better Fourier resolution.

STEP 3, FOURIER EXTRACT

Fit $I(\theta) = a_0 + b_2\sin 2\theta + a_4\cos 4\theta + b_4\sin 4\theta$. Apply: $S_1=4a_4$, $S_2=4b_4$, $S_3=-2b_2$, $S_0=2a_0-2a_4$.

📖 New: everything on this page for measuring a classical polarization state is, unchanged, the recipe for measuring a single-photon **polarization qubit**. Based on James, Kwiat, Munro & White, *Phys. Rev.* A 64, 052312 (2001).

A Stokes vector is a Bloch vector

Write a single photon's polarization state as a qubit, $|\psi\rangle = \alpha|H\rangle + \beta|V\rangle$. Its density matrix $\hat{\rho} = |\psi\rangle\langle\psi|$ (or a mixed state for partially polarized light) is a 2×2 Hermitian, unit-trace matrix, exactly the object the Poincaré sphere in the Stokes Parameters tab already visualizes. Expand it in the Pauli basis $\hat{\sigma}_0 = \hat{I}, \hat{\sigma}_1, \hat{\sigma}_2, \hat{\sigma}_3$ and the Stokes parameters fall out as the expansion coefficients:

$$\hat{\rho} = \frac{1}{2}\sum_{i=0}^{3}\frac{S_i}{S_0}\hat{\sigma}_i$$

The normalized Stokes vector $(S_1/S_0, S_2/S_0, S_3/S_0)$ is the Bloch vector. Degree of polarization is Bloch-vector length; a fully mixed (unpolarized) beam sits at the Poincaré sphere's center, exactly like the maximally mixed qubit state $\hat{\rho} = \hat{I}/2$.

The minimal measurement set for one qubit

The continuous rotating-QWP method above is built for a classical beam with a power meter, plenty of photons to average over every angle. A single photon gives one click per shot, so photonic qubit tomography instead uses four fixed projective measurements:

$$\hat{\mu}_0 = |H\rangle\langle H| + |V\rangle\langle V|, \quad \hat{\mu}_1 = |H\rangle\langle H|, \quad \hat{\mu}_2 = |D\rangle\langle D|, \quad \hat{\mu}_3 = |R\rangle\langle R|$$

μ_0 (both PBS outputs summed, so n_0 is the total count N) fixes the normalization; μ_1, μ_2, μ_3 are projections onto H, +45° diagonal, and right-circular, the same three axes this page already calls the H/V, D/A, and R/L bases. Count n_i photons in each setting, then:

$$S_0 = n_0, \quad S_1 = 2n_1 - n_0, \quad S_2 = 2n_2 - n_0, \quad S_3 = 2n_3 - n_0$$

(James et al. 2001 write $S_0 = 2n_0$, $S_i = 2(n_i - n_0)$; there μ_0 is a 50% neutral filter, so their n_0 is half of the total used here.)

Feed these into $\hat{\rho} = \frac{1}{2}\sum_i (S_i/S_0)\hat{\sigma}_i$ above and you have a full single-qubit tomographic reconstruction, four measurement settings, no rotation stage required.

Why the raw estimate can be unphysical

With finite photon counts, statistical noise on $n_0..n_3$ can push the linear-inversion $\hat{\rho}$ above to a matrix with a *negative* eigenvalue, not a valid quantum state. James et al.'s fix, now the field standard, is **maximum-likelihood estimation**: parametrize $\hat{\rho} = \hat{T}^\dagger\hat{T}/\mathrm{Tr}(\hat{T}^\dagger\hat{T})$ for a lower-triangular $\hat{T}$ (this form is positive semidefinite by construction for *any* $\hat{T}$) and numerically maximize the likelihood of the observed counts over $\hat{T}$'s free parameters. The result is always a physical density matrix.

The same recipe, on an atomic qubit

If you've built the rotating-QWP polarimeter above, you already understand atomic-qubit state tomography. Swap μ_1 (H projection) for a bare population measurement in $\{|0\rangle, |1\rangle\}$; swap μ_2, μ_3 (D and R projections) for a population measurement preceded by a resonant $\pi/2$ pulse at phase 0 and phase $\pi/2$, rotating the X- or Y-axis of the atomic Bloch sphere onto the measurement axis. Same four measurement settings, same $\hat{\rho} = \frac{1}{2}\sum_i (S_i/S_0)\hat{\sigma}_i$ reconstruction, same maximum-likelihood fix for shot-noise-induced unphysical estimates.

Worked check

The **+45°** preset in the Interactive Simulator tab has Stokes vector $(S_1, S_2, S_3)/S_0 = (0, 1, 0)$. Its density matrix is $\hat{\rho} = \frac{1}{2}\begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix}$, a pure state, $\mathrm{Tr}(\hat{\rho}^2) = 1$, eigenvalues $\{1, 0\}$. Try the **Elliptical** preset, $(S_1, S_2, S_3)/S_0 = (0.5, 0.5, 0.707)$: since $S_1^2 + S_2^2 + S_3^2 \approx S_0^2$, it is still (to rounding) a pure state sitting on the Poincaré sphere's surface, not a partially mixed one inside it.

5.1 Sources

Schaefer et al. (2007), measuring Stokes parameters with a rotating quarter-wave plate PEER-REVIEWED Collett (2005), Field Guide to Polarization (SPIE) TEXTBOOK Goldstein, Polarized Light (CRC Press), Stokes and Mueller calculus TEXTBOOK James, Kwiat, Munro & White (2001), measurement of qubits — the Stokes/Poincaré formalism on this page is single-qubit tomography; source for the Qubit Tomography tab PEER-REVIEWED James, Kwiat, Munro & White, arXiv:quant-ph/0103121 — free preprint of the above

CHAPTER 5 SUMMARY

Clean polarization control ensures that every beam drives only its intended Zeeman sublevels. With vacuum, locked lasers, and verified polarizations in place, we are ready to trap our first cloud of atoms.

Problems for Chapter 5

1. USE THE TOOL Send horizontal light through a quarter-wave plate with its fast axis at 45°. Predict the Stokes vector, then check it in the simulator.
2. BY HAND Show with Jones matrices that a half-wave plate at angle θ rotates linear polarisation by 2θ. Confirm with the simulator at θ = 22.5°.
3. DESIGN Optical pumping into a stretched state needs pure σ^+ light along the quantisation axis. Describe how you would measure and optimise circular purity with a polarising beam splitter, a quarter-wave plate and a power meter, and what value of S_3 you would aim for.

Use the tool: on the interactive page. By hand: from the formulas in this chapter. Design: a decision a lab would make. Worked solutions are available to instructors on request (see Part V).

CHAPTER 6 · BUILD STEP 06

6

Magneto-Optical Traps

🧲 MOT Designer · Calculator

Full-parameter magneto-optical trap calculator — from laser detuning and field gradient to damping coefficient, spring constant, capture velocity, and scattering rate. Includes a live detuning optimization curve, species comparison, and the full cooling pathway.

> *Combining red-detuned σ^+/σ^- laser beams (Chapters 4–5) with a quadrupole magnetic field gradient creates the magneto-optical trap (MOT)—the workhorse source of cold atoms for almost every AMO experiment.*

IN THIS CHAPTER Prerequisites: Ch. 1, Ch. 4, Ch. 5

- Compute the scattering force, velocity damping coefficient, restoring spring constant and capture velocity of a MOT from the two-level scattering-rate model, and know where that model stops being accurate.
- Choose beam waist, detuning, saturation parameter, and magnetic gradient to maximize capture velocity and atom number.
- Understand density limits from radiation trapping and light-assisted collisions, and plan the transition to sub-Doppler cooling.

WEBSITE COMPANION Live interactive page ↗ · Notebook: Open in Colab · download mot-designer.ipynb · Week 5

A magneto-optical trap (MOT) is the first stage of every cold-atom experiment. Three pairs of counter-propagating red-detuned beams with a magnetic quadrupole field simultaneously cool and confine atoms via radiation pressure. Understanding α (damping), κ (spring constant), and v_c (capture velocity) is essential for optimizing loading rates and temperature before transfer into tweezers, a lattice, or an evaporative cooling stage.

Model validity: this calculator uses a 1D, low-velocity, two-level scattering-force model. Real alkali MOTs are multilevel systems with Clebsch-Gordan factors, optical pumping, polarization gradients, beam imbalance, reabsorption, and sub-Doppler physics. Treat α, κ, and capture velocity as scaling estimates, not calibrated predictions for a loaded cloud.

6.1 MOT Calculator

DAMPING COEFF. A **0.0678 10⁻²⁰ kg/s** F = −αv	SPRING CONSTANT K **0.111 10⁻¹⁷ N/m** F = −κr	MOT FREQUENCY **441.6 Hz** ω/2π = √(κ/m)/(2π)
CAPTURE VELOCITY **7.10 m/s** v_c ≈ |Δ|/k	DOPPLER LIMIT T_D **146 μK** ħΓ/2k_B	MAX ACCELERATION **5.61e+4 m/s²** ħkΓs/2m(1+s)
SCATTERING RATE R_SC **1.73 10⁶ photons/s** Γ/2 · s/(1+s+(2Δ/Γ)²)	DAMPING TIME T **0.2 ms** τ = m/α	CLOUD RADIUS Σ **58** μm (rms σ; 1/e² radius = 2σ) · √(k_BT_MOT/κ)

Calculator output at the page's default settings.

MOT force in 1D (two-beam, low-velocity limit)

Two counter-propagating beams in a quadrupole B-field create both velocity-dependent damping and position-dependent restoring forces. The effective detuning for each beam shifts as $\Delta_{\text{eff}} = \Delta \pm kv \pm \mu_{\text{eff}} B(r)/\hbar$.

$$F_{\text{MOT}} \approx -\alpha v - \kappa r$$

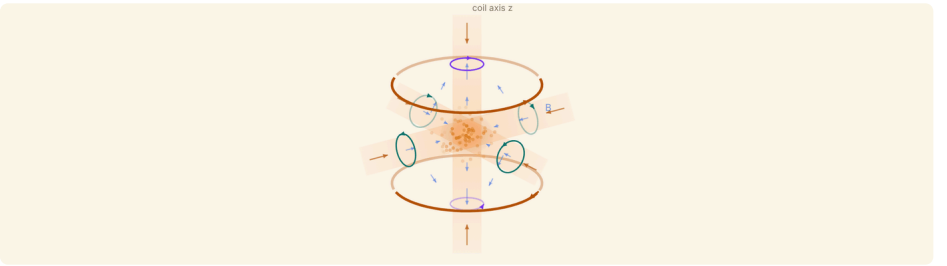
$$\alpha = \frac{8\hbar k^2 s |\Delta|/\Gamma}{(1+s+(2\Delta/\Gamma)^2)^2} \qquad \kappa = \frac{\mu_{\text{eff}}}{\hbar k} \cdot \frac{dB}{dz} \cdot \alpha$$

Optimal damping:

$|\Delta|_{\text{opt}} \approx \Gamma/(2\sqrt{3}) \approx 0.29\,\Gamma$

(low-s limit)

Live MOT Visualization · 3D



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

σ_z ≈ 58 μm · σ_r ≈ 81 μm · drawn at display limit

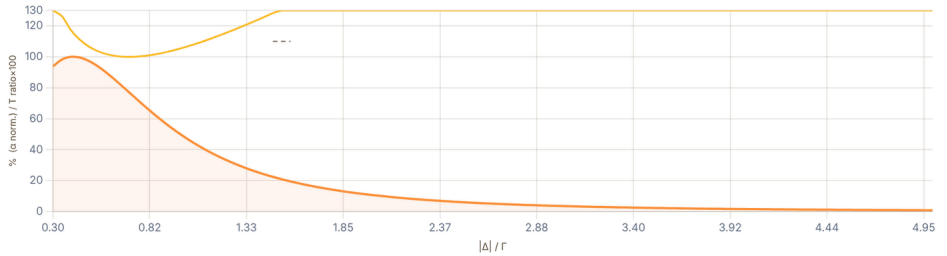
Coils are anti-Helmholtz (opposite currents), giving B = G(−x/2, −y/2, z): the gradient along the coil axis is twice the radial one, so the cloud is squeezed along z ($\sigma_z = \sigma_r/\sqrt{2}$). Rings show each beam's circular polarization about its own direction: the axial pair has the opposite handedness to the four radial beams. Geometry is schematic; the cloud size follows the calculator.

> **Lab tip:** when loading atoms, start with a large detuning (~2–3Γ) and high gradient for maximum capture volume, then compress (increase detuning to cut radiation pressure from rescattered light + increase gradient) for higher density before the sub-Doppler stage. This "cMOT" step is often 10–50 ms.

6.2 Detuning Optimization

Damping α and steady-state temperature both depend on |Δ|/Γ. The orange curve shows where your current detuning sits on the tradeoff — high α wants small detuning, low temperature wants |Δ|~Γ/2. Practical MOTs live at |Δ| = 1.5–2.5Γ, a compromise between the two.

A (% OF MAX) AND T_MOT / T_MIN VS |Δ|/Γ — LIVE OPERATING POINT SHOWN (T_MIN = √(1+S)·T_DOPPLER)



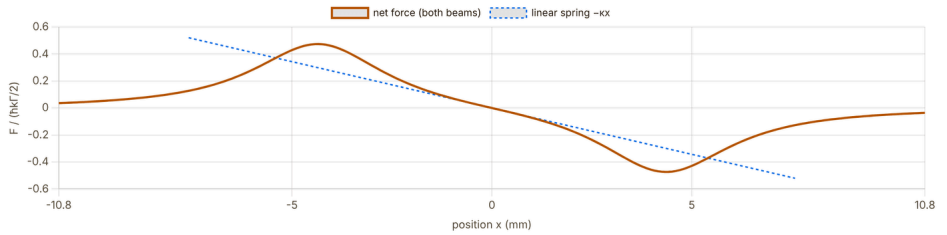
Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Damping α (% of αmax) T_MOT / T_min (×100) Your operating point

Interactive

WHY A MOT TRAPS: THE RESTORING FORCE ACROSS THE TRAP (ATOM AT REST)

The field gradient Zeeman-shifts the transition by ±μ′B′x/ħ, so for an atom displaced to +x the beam pushing it back toward the centre is closer to resonance than the beam pushing it out. Near the centre the net force is linear, $F \approx -\kappa x$ (dotted: the calculator's spring constant). It is strongest where the Zeeman shift cancels the detuning, $x_{res} = \hbar|\Delta|/(\mu' B')$, and fades beyond, so a larger detuning or a gentler gradient makes a wider trap with a weaker spring.



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

$x_{res} = \hbar|\Delta|/(\mu' B')$ = 4.33 mm · beam radius w = 10 mm · peak force 0.47 ħkΓ/2 · κ = 1.11e-18 N/m

Optimal α detuning

$|\Delta| \approx 0.29\,\Gamma$

Maximum damping at low s. Shifts to $\Gamma\sqrt{((1+s)/12)} \approx 0.29\Gamma\sqrt{(1+s)}$ with saturation. Rarely used in practice — capture velocity is too small.

Minimum temperature

$|\Delta| = \Gamma/2$

Doppler limit $T_D = \hbar\Gamma/2k_B$ reached here at $s \to 0$. At finite s, minimum shifts to $|\Delta| = (\Gamma/2)\sqrt{(1+s)}$.

Practical range

1.5 – 2.5 Γ

Balances adequate damping, reasonable capture velocity, and tolerable cloud temperature. Most alkali MOTs operate here.

Lab tip: if your MOT fluorescence is high but the cloud stays hot, try moving toward larger $|\Delta|$. If loading is slow, try a larger detuning and larger beams: the capture velocity grows with $|\Delta|$ ($v_c \approx |\Delta|/k$) and the capture region with x_{res} and w — but the damping and the spring get weaker.

6.3 Species at a Glance

Key MOT parameters for all five species at your current detuning and saturation settings. Highlighted row is your selected atom.

ATOM	Λ (NM)	Γ/2Π (MHZ)	T_DOPPLER (MK)	E_REC/K_B (NK)	V_C AT |Δ| (M/S)	A/A_RB
^{87}Rb	780.2	6.065	146	181	7.10	ref
^{133}Cs	852.3	5.234	126	99	6.69	0.84×
^{23}Na	589.2	9.795	235	1200	8.66	1.75×
^{7}Li	671.0	5.8724	141	3031	5.91	1.35×
^{39}K	766.7	6.035	145	418	6.94	1.04×

Li note: Li's nearly degenerate D1/D2 lines cause mixing in the MOT — gray molasses on D1 is the standard sub-Doppler step. **K note:** K39's small ground-state hyperfine splitting (~461 MHz) requires careful repumper alignment to avoid optical pumping out of the cycling transition.

6.4 From MOT to Experiment — The Cooling Cascade

The MOT is stage one. Here's how temperature, atom number, and density evolve from the atomic source all the way to a tweezer array or quantum gas.

🔥

Atomic Source

300 K → ~100 mK

Oven / 2D MOT + Zeeman slower. Slows atoms from ~300 m/s to capture velocity.

→

🧲

3D MOT

100 – 300 μK

10^7–10^9 atoms. mm-scale cloud. Density ~10^{10}/cm³. Limited by re-absorption.

→

🗜️

cMOT

30 – 100 μK

Ramp: ↑ gradient, ↓ repumper, ↑ |Δ|. Higher density for better transfer. ~10–50 ms.

→

🌫️

Gray Molasses / PGC

2 – 20 μK

Sub-Doppler. Dark states (gray molasses) or Sisyphus (PGC). B-field zeroed to <50 mG.

→

🎯

Tweezer / Lattice

5 – 30 μK in trap

Single-atom loading (~50%). Site-resolved imaging. Sideband cool to $\bar{n} < 0.1$.

→

⚛️

Experiment

< 1 μK

Rydberg gates, quantum simulation, BEC, or quantum gas microscopy.

Sub-Doppler: why it works

Polarization gradients create spatially varying light shifts. Atoms climbing a potential hill scatter a photon and are reborn at the bottom — net energy removed per cycle. For lin⊥lin beams (Sisyphus; $\sigma^+\sigma^-$ beams cool by a related motion-induced orientation effect) this gives $T \propto U_0/k_B$ where $U_0 \ll \hbar\Gamma$, reaching 1–20 μK. Gray molasses (D1 line) exploits dark states to further suppress heating.

Tweezer loading: the 50% barrier

Gaussian tweezers load atoms from the MOT stochastically. Parity projection — two atoms collide and both leave via light-assisted collisions — caps loading at ~50% per site. Blue-detuned light pulses or resolved sideband cooling can push fidelity above 90% by deterministic preparation after loading.

6.5 Physics Reference

Position- and velocity-dependent forces from two beams

For a moving atom at position r in the quadrupole field, the effective detuning seen by each beam is $\Delta_{\text{eff}} = \Delta \pm kv \pm \mu_{\text{eff}}B(r)/\hbar$. Expanding to first order in kv and $\mu_{\text{eff}}B/\hbar|\Delta|$ gives the linear damping + spring-constant picture.

$$R_{\text{sc}} = \frac{\Gamma}{2} \cdot \frac{s}{1 + s + (2\Delta_{\text{eff}}/\Gamma)^2}$$

Net force ($|kv| \ll |\Delta|$): $F = F_+ - F_- \approx -\alpha v - \kappa r$

$$\alpha = \frac{8\hbar k^2 s|\Delta|/\Gamma}{(1 + s + (2\Delta/\Gamma)^2)^2} \quad [\text{1D, 2 beams}]$$

$$\kappa = \frac{\mu_{\text{eff}}}{\hbar k} \cdot \frac{dB}{dz} \cdot \alpha$$

Optimal α:

$|\Delta| = \Gamma/(2\sqrt{3}) \approx 0.29\,\Gamma$

(low-s limit)

Sign convention: this page uses Δ = ω_L − ω_0, so red detuning is Δ < 0. Some AMO texts use Δ = ω_0 − ω_L, flipping the sign of every term in the damping algebra. Always confirm the convention before comparing formulas.

v_c sets the flux captured from the background vapor

Only atoms slower than v_c ≈ |Δ|/k can be captured before they traverse the beam. The loading rate L scales strongly with v_c and beam area A: L ∝ n_bg · v_c⁴ · A (Maxwell-Boltzmann tail). Larger beams and larger detuning dramatically increase loading — at the cost of a hotter cloud.

Capture velocity: $v_c \approx |\Delta|/k$ (velocity at which Doppler shift = detuning)

$$L \approx n_{\text{bg}}\,\bar{v}\left(\frac{v_c}{\bar{v}}\right)^4 \cdot \pi w^2 \quad \bar{v} = \sqrt{8k_B T_{\text{oven}}/\pi m}$$

Stopping distance: $d = v_c^2/(2a_{\max})$, must satisfy $d \lesssim w$ (beam radius)
Maximum deceleration: $a_{\max} = \hbar k \Gamma s/(2m(1+s))$

For Rb87 at Δ = −1.5Γ: v_c ≈ 7 m/s and a_max ≈ 10^5 m/s². Stopping distance ~0.25 mm, comfortably within a 10 mm beam. Background pressure of ~10^{-8} Torr gives loading rates of 10^7–10^8 atoms/s for a standard vapor-cell MOT.

Polarization gradients and dark states break the Doppler limit

The Doppler limit T_D = ħΓ/2k_B assumes a two-level atom. Multi-level alkalis in polarization-gradient fields experience sub-Doppler cooling (Sisyphus in lin⊥lin, motion-induced orientation in $\sigma^+\sigma^-$) or gray-molasses cooling (D1 line, dark states), reaching temperatures 10–100× below T_D.

Sisyphus (PGC): $k_B T_{\mathrm{PGC}} \sim U_0 \propto \frac{\hbar\Omega_R^2}{|\Delta|}$; requires $|\Delta| \gg \Gamma$, $B = 0$
Gray molasses (D1): exploits Λ dark states; $T \sim 5$–$20\,\mu$K for Rb, Cs, Li
Sideband cooling (tweezer): $T \sim \hbar\omega_{\mathrm{trap}}/k_B$; $\bar{n} < 0.1$ achievable
Comparison (Rb87): $T_D = 146\,\mu$K → PGC ∼ 10–30 μK → SBC ∼ 1 μK

Key requirement for all sub-Doppler cooling: residual B-field <50–100 mG at the cloud. Eddy currents from switching MOT coils typically require 1–10 ms of wait time before the molasses phase begins.

Peak density is set by reabsorption, not atom number

Scattered photons from atoms near the cloud center are reabsorbed by outer atoms, creating an outward radiation pressure that limits peak density. This gives a density-independent equilibrium: adding more atoms increases cloud size, not density.

Peak density limit: $n_{\max} \approx \frac{\kappa}{\sigma_\infty \hbar k \Gamma/2}$ (multiple-scattering limit)
Typical Rb87 MOT: $n_{\mathrm{peak}} \sim 10^{10}$–$10^{11}$ cm^{-3}
cMOT fix: reducing repumper intensity pumps atoms to dark $F = 1$ ground state → less reabsorption → 5–10× density gain
Dark-SPOT MOT: shadow mask blocks repumper at center; density ∼ 10^{12} cm^{-3}

6.6 References

📄 Metcalf & van der Straten — Laser Cooling and Trapping (textbook) TEXTBOOK 📄 Chu (1998) — Nobel Lecture: The manipulation of neutral particles PEER-REVIEWED 📄 Raab et al. (1987) — Trapping of neutral sodium atoms with radiation pressure PEER-REVIEWED 📄 Anderson et al. (1995) — Observation of BEC in a dilute atomic vapor PEER-REVIEWED 📄 Steck — Quantum and Atom Optics lecture notes (free, comprehensive) TEXTBOOK

CHAPTER 6 SUMMARY

A well-tuned MOT loads 10^6–10^9 atoms at tens to hundreds of microkelvin, serving as the cold reservoir for optical dipole traps and tweezers. To orchestrate MOT loading, ramps, and imaging with microsecond precision, we next build the experiment control system.

Problems for Chapter 6

1. USE THE TOOL For ^{87}Rb with the default beams, step the detuning from 1 Γ to 3 Γ. Record the capture velocity and the Doppler temperature at each step and describe the trade-off.
2. BY HAND Check one temperature with $T = (\hbar\Gamma/4k_B)(1 + s + 4\delta^2/\Gamma^2)/(2|\delta|/\Gamma)$ and confirm that it is minimised at $|\delta| = \Gamma/2$ when $s \ll 1$.
3. DESIGN Use Section 03 to compare a lithium and a cesium MOT. Which parameters must change, and why does a lithium MOT usually need a separate cooling stage afterwards?

Use the tool: on the interactive page. By hand: from the formulas in this chapter. Design: a decision a lab would make. Worked solutions are available to instructors on request (see Part V).

CHAPTER 7 · BUILD STEP 07

7

Experimental Control & Timing Systems

🎛️ Experiment Control Systems · Guide

Every shot of an AMO experiment is a list of a few hundred events on dozens of channels, timed to nanoseconds and repeated thousands of times. This page is what a new student needs to run that machine: the signals (TTL, DAC, ADC, DDS, AWG), how timing really works, how pulses are engineered, which systems labs use, how the whole experiment is driven from Python, and how a group shares one code base with git.

> *Running a cold-atom sequence requires coordinating dozens of shutters, AOMs, coils, and cameras on microsecond timescales—a job handled by deterministic hardware timing systems.*

IN THIS CHAPTER Prerequisites: Ch. 4, Ch. 6

- Distinguish the roles of TTL digital outputs, DAC analog ramps, DDS/AWG RF synthesizers, and ADC/camera triggers in an AMO shot.
- Account for physical hardware latencies (mechanical shutters, coil eddy currents, AOM rise times) when compiling timing sequences.
- Structure reproducible Python experiment scripts (ARTIQ/Sinara, labscript) with version control and parameter logging.

WEBSITE COMPANION Live interactive page ↗ · Notebook: Open in Colab · download control-systems.ipynb · Week 7

7.1 What the control system has to do

An experiment is a sequence of switches, ramps and triggers. The control system turns a description of that sequence into voltages at the right instant, every shot, and records what happened.

Hardware-timed

Time comes from a clock, not the PC

A Python script on a desktop cannot promise when a line of code runs: the operating system can pause it for milliseconds. So the PC only *describes* the sequence; a dedicated timing device (an FPGA, a pulse generator or a DAQ card with its own sample clock) plays it out. Everything below follows from this split.

Repeatable

Thousands of identical shots

Most measurements average many shots or scan one parameter while everything else stays fixed. A useful system lets you change one number (a detuning, a hold time) between shots and guarantees that nothing else moved.

Recorded

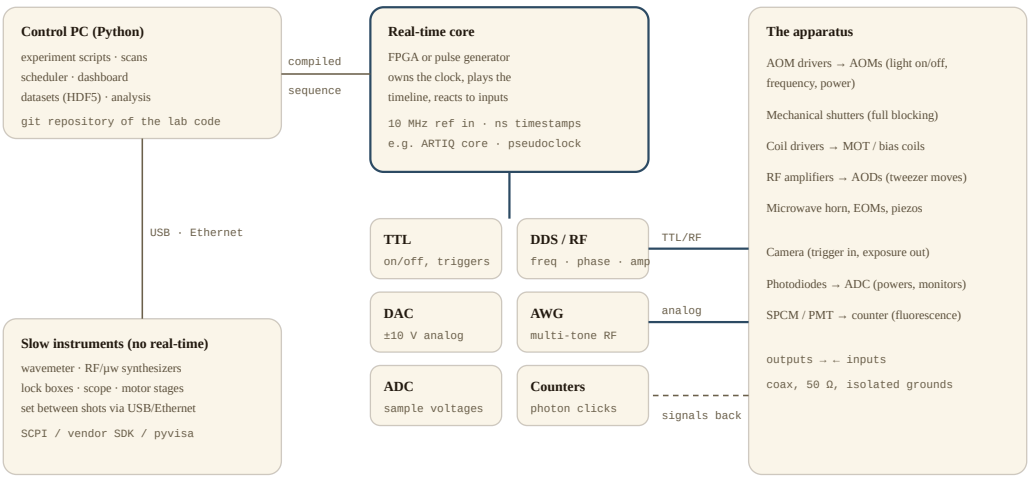
Every shot leaves a record

Months later someone will ask what the imaging detuning was on a given run. The parameters, the code version and the data should be stored together, shot by shot, automatically.

Four numbers describe any timing system. *Resolution*: the smallest step in which an edge can be placed. *Jitter*: how much an edge wanders from shot to shot. *Latency*: how long between an input (a photon count) and the earliest output that can react to it. *Update rate*: how often an analog value can change. A system can be excellent at one and poor at another, so ask for all four.

7.2 Anatomy: from a Python script to photons

The same layout appears in almost every lab, whatever the brand: a PC that writes and schedules sequences, a real-time core that owns the clock, output and input cards, and the lab hardware those cards drive.



Real-time links in the page accent colour; dashed = signals coming back. Slow instruments never sit in the real-time path: they are set between shots, so their latency does not matter.

What lives where

On the PC: the experiment description, parameter scans, the queue of experiments, plots, the data files, and drivers for slow instruments.

On the real-time core: the timeline itself, plus any decision that must be made within microseconds (count photons, then decide whether to fire the next pulse).

Why the split matters

Anything that crosses back to the PC during a shot is slow and unpredictable (milliseconds). ARTIQ's manual makes the point concretely: after a call back to the host, the next real-time event must be pushed into the future again or it raises an underflow. Keep timing-critical logic on the core.

7.3 The signals you will wire

Five kinds of signal cover nearly everything. The examples are real hardware specifications, from the Sinara (ARTIQ) datasheets and the NI PCIe-6363 manual, so you can see the order of magnitude a modern system offers.

SIGNAL	WHAT IT DRIVES	TYPICAL NUMBERS (EXAMPLE HARDWARE)	WHAT BITES PEOPLE
TTL (digital)	AOM RF switches, shutters, camera and scope triggers, coil switches, gating counters	Sinara 8-channel BNC/SMA cards: minimum pulse 3 ns (typ. 5 ns), ~350 ps peak jitter, galvanically isolated in banks of four, 50 Ω or high-Z per channel. ARTIQ places edges on a 1 ns grid.	Logic levels (5 V vs 3.3 V), driving a 50 Ω load, unterminated cables ringing, a TTL line that is a ground loop.
DDS / RF	AOM and EOM frequency, phase and amplitude; Raman and microwave tones (with a synthesizer)	Urukul (AD9910): 4 channels, 1 GS/s, output <1 to >400 MHz, 0.25 Hz frequency steps, 14-bit amplitude, 0–31.5 dB attenuator, RF switch with 70 dB isolation. Mirny: 53 MHz–4 GHz synthesizer.	Phase continuity when you change frequency; images and spurs; forgetting that a frequency change takes time to program.
DAC (analog out)	Coil-driver setpoints, laser and tweezer power (via AOM amplitude), piezo offsets, servo setpoints	Zotino: 32 × 16-bit, ±10 V, 1 MS/s shared, 75 kHz bandwidth. Fastino: 32 × 16-bit, up to 2.55 MS/s per channel with current gateware. NI PCIe-6363: 4 × 16-bit, 2.86 MS/s (one channel) to 1.25 MS/s (all four).	16 bits over 20 V is 305 µV per step; the output filter (here 75 kHz) rounds fast steps; the driver downstream is usually slower than the DAC.
ADC (analog in)	Photodiode powers, error signals, monitor voltages, intensity servo inputs	Sampler: 8 × 16-bit, 1.5 MS/s simultaneous, programmable gain from ±10 V to ±10 mV. NI PCIe-6363: 32 inputs, 16-bit, 2 MS/s single-channel (1 MS/s aggregate).	Aliasing (filter before you sample), ground offsets, input range set too wide so the signal uses a few bits.
AWG (fast arbitrary RF)	Multi-tone AOD drive for tweezer arrays and rearrangement, shaped Raman or Rydberg pulses	Phaser: 4 × 16-bit DAC channels at 1.25 GS/s with IQ upconversion (0.3–4.8 GHz LO). Standalone AWG cards are common for tweezer rearrangement.	Memory and streaming bandwidth, phase relationships between tones, intermodulation in the RF amplifier.

Inputs that make decisions

A counter on a TTL input turns single-photon detector clicks into a number inside the shot. The Sinara TTL datasheet quotes a gateware edge counter of up to 125 million events per second. That is how "is there an atom in this tweezer?" becomes a decision in microseconds instead of a camera frame later.

Cameras

Cameras are triggered by TTL and usually report back an "exposure active" output. Wire that output to a scope or an input at least once: the real exposure window is rarely exactly where the trigger edge suggests. Frame grabbers (Sinara's Grabber reads CameraLink) can extract regions of interest in real time.

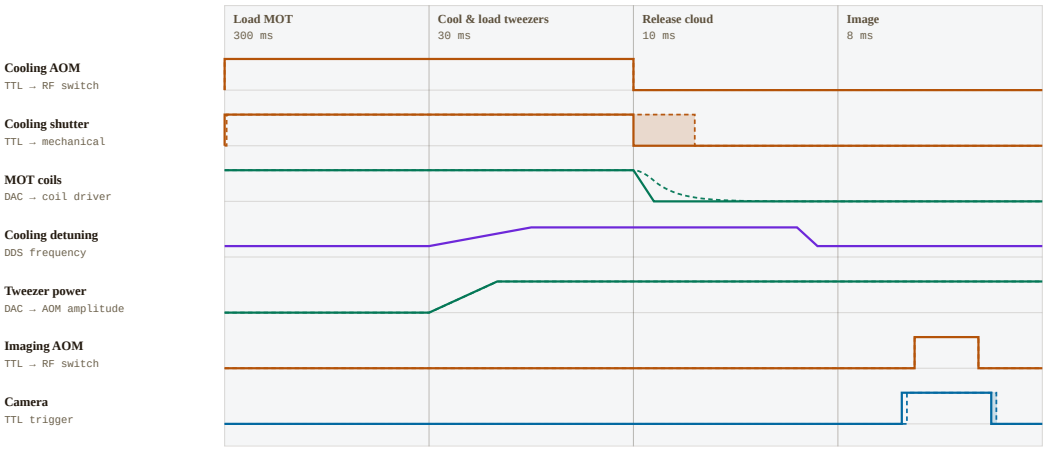
Signal integrity in one paragraph. Use coax and terminate fast lines in 50 Ω at the receiving end (or deliberately use high-Z and short cables). A signal travels roughly 5 ns per metre of coax (velocity factor ≈ 0.66), so a 3 m cable is a 15 ns delay. Keep the control system's ground and the high-current coil supplies apart; isolated TTL cards exist for this, though their datasheet notes that isolation adds jitter, so a non-isolated card is better for the most timing-critical lines. Look at every new signal on a scope before trusting it.

7.4 Timing: resolution, jitter, latency and hardware delays

The control system can place an edge to the nanosecond. The hardware it drives cannot respond that fast, and each device lags by a different amount. Good sequences are written for what the atoms see, not for what the TTL line does.

One shot, channel by channel

A tweezer-loading shot. Solid lines are what the control system outputs; dashed lines are what the hardware actually does. Click a channel name for what it is and what goes wrong. Each phase has its own time scale.



programmed output physical response Shaded: the lag between them

> Pick **Real delays** and look at the shutter and the coils: the light keeps leaking and the field keeps decaying after the command. **Delays compensated** moves each command earlier by that device's measured lag.

Delay values are illustrative (shutter: a few ms of delay plus sub-ms jitter; coils: an L/R decay of about a millisecond; AOM: tens to hundreds of ns, invisible at this scale). Measure your own: the numbers depend on the shutter, the coil inductance and the driver.

How ARTIQ keeps time

Code on the core moves a timeline cursor, `now_mu`, forward with `delay()`; outputs are timestamped with the cursor and queued in FIFOs; the hardware fires each event when the wall clock reaches its timestamp. The machine unit is normally 1 ns; the coarse RTIO clock is commonly 125 MHz, so events are processed in 8 ns cycles with 1 ns fine placement.

The gap between cursor and wall clock is the *slack*. If code falls behind and submits an event whose time has passed, ARTIQ raises `RTIOUnderflow`. Two events on the same channel at the same coarse time collide and are dropped.

How a pseudoclock keeps time

The labscript suite compiles the whole shot in Python, then plays it from a *pseudoclock*: a pulse generator (a SpinCore PulseBlaster, or the Raspberry-Pi-Pico-based PrawnBlaster) that ticks only when something must change. DAQ cards and DDS boards update on those ticks. Each shot is stored as an HDF5 file holding the globals and the compiled instructions.

The resolution is set by the pulse generator's clock (labscript encodes 20 ns for a PulseBlaster-USB, 4 ns for the ESR-PRO-500), and a few clock cycles are the minimum interval between instructions.

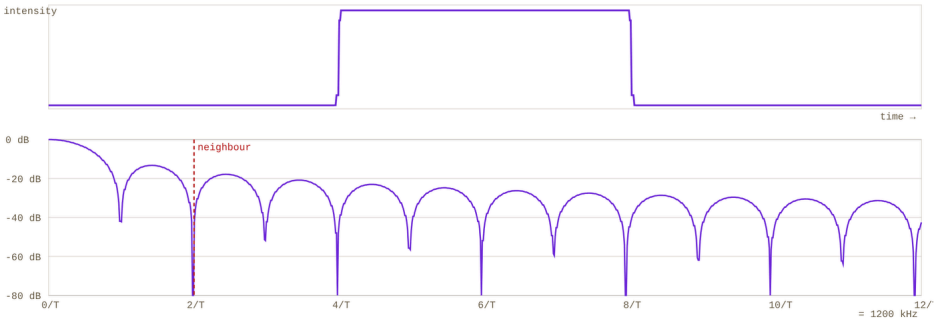
Speed limits that are not the control system. An AOM switches in roughly the time sound takes to cross the beam (next section's calculator); a mechanical shutter needs milliseconds and its delay drifts as it wears; a coil's current decays with its L/R time and eddy currents in nearby metal add a slower tail; a camera has its own trigger latency. Measure each with a photodiode or a pick-up coil, write the delays into one place in the code, and compensate there.

7.5 Engineering pulses and ramps

A light pulse is not just "on for T". Its shape sets its spectrum, and the spectrum decides which transitions it drives. Square pulses are easy to program and spectrally dirty; shaped pulses cost a little length and buy a lot of cleanliness.

Pulse shape and its spectrum

Top: the intensity envelope in time. Bottom: its power spectrum in dB against detuning, in units of 1/T and in kHz for the chosen length. The marker shows a neighbouring transition (for example a motional sideband at the trap frequency) and how much the pulse drives it.



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

First side lobe **-13.3 dB** −20 dB half-width **90.6 kHz** (0.91/T) Drive at neighbour (200 kHz) **-42.2 dB**

Computed in the page with a numerical Fourier transform of the envelope (the AOM rise is a Gaussian smoothing of the edges). All shapes have the same total length T, so smooth shapes have a wider central lobe but far lower side lobes: about −13 dB for a square pulse, about −58 dB for Blackman.

Where shaping matters

- Raman and sideband pulses next to a carrier.
- Rydberg pulses, where leakage drives unwanted states (see the gate error budget).
- Fast switching of trap light, where a hard edge heats the atom.

How shapes are made

- DDS amplitude profiles or on-chip RAM (the AD9910 in Urukul has a waveform RAM for this).
- A DAC driving a voltage-controlled attenuator or the AOM driver's amplitude input.
- An AWG that synthesises the shaped RF directly.

Ramps

- Linear ramps for detunings and fields.
- Exponential or piecewise ramps for evaporation.
- Smooth (minimum-jerk or $\sin^2$) profiles for moving tweezers, so atoms are not heated by sudden acceleration.
- Enough points: a DAC at 1 MS/s gives 1000 steps per millisecond.

7.6 How fine, how fast: a signal budget

Before choosing hardware or writing a sequence, check that each channel can resolve what the physics needs. Change any input; every result updates.

Formulas: DAC step = range / 2^{bits}; DDS step = f_{clk} / 2^N; AOM optical rise time $\tau \approx 0.65\ d/V$ and modulation bandwidth $\approx 0.35/\tau$ (Isomet application note AN2600; acoustic velocities from the same note); coax delay ≈ 5 ns/m for a velocity factor of 0.66.

7.7 Which system? ARTIQ, labscript, NI and the rest

You will usually inherit a system rather than choose one. Knowing the families helps you read a new lab's code and judge what it can and cannot do.

ARTIQ + Sinara

An FPGA that runs your Python

Started in 2013 by NIST's Ion Storage Group with M-Labs; open source (LGPLv3+ software, CERN-OHL hardware). Experiments are Python classes; methods marked `@kernel` are compiled and run on the FPGA core (Kasli: Artix-7, 12 slots for Sinara cards), the rest runs on the PC. Distributed RTIO (DRTIO) links several crates over fibre. A 2023 review of open quantum hardware counted almost a thousand experiments using ARTIQ/Sinara.

Best at: nanosecond timing, real-time decisions (count, then act), large channel counts. **Costs you:** learning the timeline model and its errors.

labscript suite

Compile the whole shot, then play it

From Monash (Starkey et al., *Rev. Sci. Instrum.* 84, 085111, 2013). You write the shot in Python; *runmanager* fills in parameters and builds one HDF5 file per shot; BLACS runs the hardware; *lyse* analyses shots as they arrive; *runviewer* plots the compiled instructions. It drives off-the-shelf hardware: PulseBlaster and PrawnBlaster pseudoclocks, NI-DAQmx cards, NovaTech DDS, many cameras.

Best at: mixing commercial hardware, a clean record of every shot. **Costs you:** no decisions inside a shot.

NI DAQ + LabVIEW or Python

Buffered, hardware-timed cards

Multifunction cards such as the PCIe-6363 (4 analog outputs, 32 inputs, 48 digital lines with 32 hardware-timed up to 10 MHz, a 100 MHz timebase) are a common first system, programmed in LabVIEW or from Python with the official `nidaqmx` package. Several cards share a clock and a start trigger to act as one.

Best at: availability and drivers. **Costs you:** slower digital timing, and sequencing logic you must build yourself.

RFSoC and commercial controllers

When the RF is the experiment

Boards built on Xilinx RFSoC (the open QICK project reached over 40 labs within two years) and commercial controllers from Quantum Machines, Zurich Instruments, Keysight and others generate GHz waveforms directly with fast feedback. They are common in superconducting-qubit labs and increasingly used for Rydberg and Raman control. Pulser (from Pasqal) describes neutral-atom pulse sequences at a higher level.

Best at: wideband shaped RF. **Costs you:** price or closed software.

7.8 Running the experiment from Python

Whatever the hardware, the pattern is the same: name the channels once, describe one shot in terms of physical parameters, scan those parameters from outside, and save everything. The sketches below show that pattern in four common stacks; device names come from your lab's configuration.

```
from artiq.experiment import *

class LoadAndImage(EnvExperiment):
    def build(self):
        self.setattr_device("core")
        self.setattr_device("ttl_cooling")     # names are defined in device_db.py
        self.setattr_device("ttl_camera")
        self.setattr_device("urukul0_ch0")     # imaging AOM (DDS + RF switch)
        self.setattr_device("zotino0")         # DAC: MOT coil-driver setpoint
        self.setattr_argument("t_load", NumberValue(200*ms, unit="ms"))
        self.setattr_argument("detuning", NumberValue(-10*MHz, unit="MHz"))

    @kernel
    def run(self):
        self.core.reset()
        self.urukul0_ch0.cpld.init()
        self.urukul0_ch0.init()
        self.zotino0.init()
        delay(1*ms)                            # build up slack before the first event

        self.zotino0.set_dac([2.5], [0])       # coils on
        self.ttl_cooling.on()
        delay(self.t_load)                     # load the MOT
        self.ttl_cooling.off()
        self.zotino0.set_dac([0.0], [0])       # coils off
        delay(2*ms)                            # wait for the field to decay (measured!)

        self.urukul0_ch0.set(80*MHz + self.detuning, amplitude=0.8)
        with parallel:                         # camera and light open together
            self.ttl_camera.pulse(30*us)
            self.urukul0_ch0.sw.pulse(20*us)
```

Sketches, not drop-in code: check the API version of your ARTIQ, labscript or nidaqmx release. The ARTIQ calls follow the ARTIQ manual (experiment classes, `@kernel`, `setattr_argument`, `gate_rising` / `count`); the labscript calls follow its device API; the NI example uses `nidaqmx`'s hardware-timed sample clock and start trigger.

Scans

Expose physical parameters as arguments or globals and scan them from the dashboard or runmanager. Randomise or interleave the scan order so a slow drift (a laser, the room temperature) does not masquerade as a signal.

Data

Both ARTIQ and labscript write HDF5. Store the parameters, the code version, the raw data and the analysis results of each shot together. ARTIQ's master keeps an HDF5 record of every experiment it runs.

Monitoring

Log slow signals (laser powers, wavemeter readings, lab temperature, vacuum pressure) continuously, with timestamps, into a database the analysis can query. When a dataset looks odd, the first question is what else changed.

7.9 Working as a team: git and GitHub

An apparatus outlives every student who works on it. Its code is written by many people, often at 2 a.m., and has to keep working. Version control is not optional.

The minimum

- One repository for the experiment: device configuration, sequences, drivers, analysis. Data stays out of it.
- Commit small changes with messages that say why ("compensate 3.1 ms shutter delay, measured 2026-11-02").
- Never leave uncommitted edits on the lab computer overnight; the next person cannot tell what changed.
- Write the commit hash into every data file, so any result can be traced to the exact code that produced it.

How groups work

- A protected main branch that always runs; new features on branches; changes merged by pull request after someone else reads them.
- Issues for broken hardware and planned work, so knowledge does not live in one person's head.
- Pin the software environment (ARTIQ uses Nix; conda or pip lock files elsewhere) so an update cannot silently change the machine.
- A few automated tests that compile every sequence, run on each push.

```
# a typical day
git switch main && git pull               # start from what everyone agreed on
git switch -c imaging-detuning-scan       # your change gets its own branch
# ... edit, run a test shot ...
git add sequences/imaging.py
git commit -m "Scan imaging detuning; add 3.1 ms shutter lead (measured)"
git push -u origin imaging-detuning-scan  # open a pull request; a labmate reviews it

# in the experiment, record the code version with the data
rev = subprocess.check_output(["git", "rev-parse", "HEAD"]).decode().strip()
```

Repository layout that works: `device_db.py` or `connection_table.py` (hardware), `sequences/` (shots), `lib/` (shared building blocks such as "image atoms"), `drivers/`, `analysis/`, `tests/`, and a `README` that tells a new student how to run one shot. For the wider habits, see Wilson et al., "Good enough practices in scientific computing" (*PLoS Comput. Biol.* 2017).

7.10 Checklist before you trust a new sequence

- **Look at it.** Put the key TTLs and analog lines on a four-channel scope, triggered by the start of the shot, and compare with what you meant.
- **Measure the delays** of every shutter, AOM and coil with a photodiode or pick-up coil, and compensate them in one place in the code.
- **Check the slack or the compile output.** Underflows, collisions and "instruction too short" warnings are bugs, not noise.
- **Make it safe when it stops.** If a sequence crashes, lasers, coils and RF should end in a safe state; coils must not stay at full current.
- **Record everything:** parameters, code hash, slow-instrument settings and environmental logs, per shot.
- **Change one thing at a time** and keep a short lab-notebook line for each change; your future self is the main reader.

> Next in the build: the tweezer array this control system will drive. The laser locks and day-to-day lab handbook cover the servos and routines that run alongside it.

7.11 References and further reading

ARTIQ manual (M-Labs). Experiment structure, kernels and RPCs, the RTIO timeline (cursor, slack, underflow, collisions), management system.

Sinara hardware and datasheets (M-Labs). Kasli, TTL cards, Urukul, Mirny, Zotino, Fastino, Sampler, Phaser, Grabber; the numbers in Section 03.

G. Kasprowicz et al., "Urukul – Open-source frequency synthesizer module for quantum physics," *Int. J. Electron. Telecommun.* 68, 123 (2022).

P. T. Starkey et al., "A scripted control system for autonomous hardware-timed experiments," *Rev. Sci. Instrum.* 84, 085111 (2013), and the labscript suite documentation.

N. Shammah et al., "Open hardware solutions in quantum technology," arXiv:2309.17233 (2023), a survey of ARTIQ/Sinara, QICK, QubiC and related tools PREPRINT

National Instruments, PCIe-6363 specifications, and the nidaqmx Python package.

Isomet, "Basic AO modulator equations" (AN2600): rise time $\tau \approx 0.65\, d/V$, bandwidth $\approx 0.35/\tau$, acoustic velocities.

G. Wilson et al., "Good enough practices in scientific computing," *PLoS Comput. Biol.* 13, e1005510 (2017).

CHAPTER 7 SUMMARY

Deterministic timing turns static optical benches into programmable experimental cycles. With control hardware ready, we can now focus tightly confined optical tweezers to trap individual atoms.

Problems for Chapter 7

1. USE THE TOOL In the shot diagram, switch from “Real delays” to “Delays compensated”. Which channels moved, by how much, and why does the mechanical shutter need the largest correction?

2. USE THE TOOL With T = 10 µs and a neighbouring transition 200 kHz away, compare square and Blackman pulses. Find the shortest Blackman pulse that keeps the drive at the neighbour below −40 dB.
3. BY HAND A 16-bit DAC with a ±10 V range drives a coil with 1 G/V. What is the smallest field step? Is it fine enough for a 1 mG field-noise budget? Check with the signal-budget calculator.

Use the tool: on the interactive page. By hand: from the formulas in this chapter. Design: a decision a lab would make. Worked solutions are available to instructors on request (see Part V).

CHAPTER 8 · BUILD STEP 08

8

Optical Tweezer Arrays

🔦 Tweezer Array Design Lab · Design lab

Design the trap before you build the waveform: estimate waist, depth, trap frequencies, Lamb-Dicke parameter, off-resonant scattering, site spacing, total power, and array crosstalk from one coherent set of knobs.

> *Focusing a far-off-resonant laser through a high-NA objective (NA $\gtrsim 0.5$) creates micron-scale dipole traps small enough that light-assisted collisions leave at most one atom per site.*

IN THIS CHAPTER Prerequisites: Ch. 1, Ch. 6, Ch. 7

- Calculate dipole trap depth, radial/axial trap frequencies, Lamb-Dicke parameter, and off-resonant photon scattering rate from wavelength, waist, and power.
- Compare static SLM hologram arrays with dynamic AOD tweezers for defect-free rearrangement and shuttling.
- Select magic wavelengths for clock or cooling transitions to cancel differential AC Stark shifts across the array.

WEBSITE COMPANION Live interactive page ↗ · Notebook: Open in Colab · download tweezer-designer.ipynb · Week 8

Tweezer-design workflow

8.1 Design one trap, then scale it into an array

The page now follows the actual design order: choose atom and trap wavelength, check single-site physics, then ask whether the array spacing and total power are sane.

01
Choose atom + wavelength
Start with species, resonance, trap wavelength, NA, and per-site power.

02
Check one tweezer
Read depth, waist, radial/axial frequencies, scattering, and Lamb-Dicke η.

03
Scale to an array
Set Nx, Ny, and spacing to estimate total optical power and crosstalk.

Calculator output at the page's default settings.

04 **Reality check**

Use the reference table and notes before trusting two-level estimates near magic wavelengths.

8.2 Single-Tweezer Physics

A tightly focused, far-red-detuned laser beam creates a conservative dipole potential. In the far-off-resonance limit ($|\Delta| \gg \Gamma$), the trap depth and photon scattering rate from the trap light follow directly from the two-level interaction:

$$U_0 = -\frac{3\pi c^2}{2\omega_0^3}\left(\frac{\Gamma}{\omega_0-\omega}+\frac{\Gamma}{\omega_0+\omega}\right) I_0 \;\approx\; \frac{3\pi c^2}{2\omega_0^3}\,\frac{\Gamma}{\Delta}\,I_0$$

$$\Gamma_{\mathrm{sc}} = \frac{3\pi c^2}{2\hbar\omega_0^3}\left(\frac{\omega}{\omega_0}\right)^3\left(\frac{\Gamma}{\omega_0-\omega}+\frac{\Gamma}{\omega_0+\omega}\right)^2 I_0 \;\approx\; \frac{3\pi c^2}{2\hbar\omega_0^3}\,\frac{\Gamma^2}{\Delta^2}\,I_0$$

(Grimm, Weidemüller & Ovchinnikov 2000, Eqs. 10–12; the right-hand forms drop the counter-rotating term, which the calculator keeps because it adds ~20–30% to the depth of Li or Na at 1064 nm). Here $I_0 = 2P/(\pi w_0^2)$ is the peak intensity, $\Delta = \omega - \omega_0$ (negative, so $U_0 < 0$, for red detuning), and w_0 is the **1/e² intensity radius** of the Gaussian that best matches a diffraction-limited focus of a uniformly filled pupil: $w_0 = (\sqrt{2}/\pi)\,\lambda/\mathrm{NA} \approx 0.45\lambda/\mathrm{NA}$. This Gaussian has the same peak intensity and the same curvature at the centre as the Airy spot, so it gives the right depth and radial frequency. (The Airy FWHM, $0.51\lambda/\mathrm{NA}$, is a *diameter*; using it as w_0 underestimates the depth by ~25%.) If the trap beam underfills the objective, the focus is larger: use $w_0 = \lambda/(\pi\theta)$ with θ the $1/e^2$ half-angle of the converging beam. For a filled pupil the paraxial Airy focus is also longer along z than this Gaussian (effective $z_R \approx 1.1\lambda/\mathrm{NA}^2$ instead of $\pi w_0^2/\lambda \approx 0.64\lambda/\mathrm{NA}^2$), so read ω_z below as an upper estimate. The key trade-off: $|U_0| \propto 1/|\Delta|$ while $\Gamma_{\mathrm{sc}} \propto 1/\Delta^2$ — deeper detuning buys a lower scattering rate at the cost of needing more power. The harmonic frequencies near the potential minimum and the Lamb-Dicke parameter are:

$$\omega_r = \sqrt{\frac{4|U_0|}{m w_0^2}} \qquad \omega_z = \sqrt{\frac{2|U_0|}{m z_R^2}} \qquad \eta = k_{\mathrm{probe}}\sqrt{\frac{\hbar}{2m\omega_r}}$$

where $z_R = \pi w_0^2/\lambda$ is the Rayleigh range. Ground-state sideband cooling needs the Lamb-Dicke condition $\eta \ll 1$ and resolved sidebands: ω_r larger than the linewidth of the cooling transition (the effective Raman or optical-pumping linewidth, set by the cooling lasers, not by this calculator). Trap-light scattering Γ_{sc} is a separate, slow heating background.

1064 nm

Cs, Li, Na common tweezer

850 nm

Common Rb tweezer

~1 mK

Typical trap depth

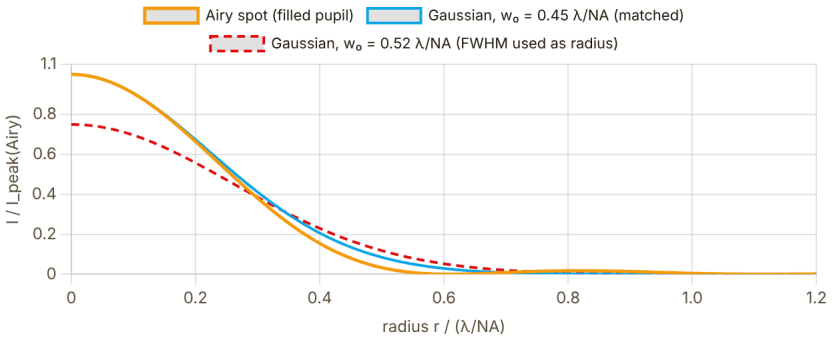
50–200 kHz

Radial frequency $\omega_r/2\pi$

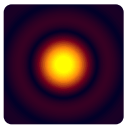
Interactive

Diameter or radius? The focal spot the formulas assume

A uniformly filled lens focuses to an Airy pattern, $I(r) \propto [2J_1(x)/x]^2$ with $x = 2\pi \cdot NA \cdot r/\lambda$. Its full width at half maximum is 0.51 λ/NA, a *diameter*. The trap formulas need the $1/e^2$ *radius* w_0 of a Gaussian, and the Gaussian with the same peak intensity and the same curvature at the centre has $w_0 = (\sqrt{2}/\pi)\,\lambda/NA \approx 0.45\,\lambda/NA$ (solid). Mistaking the FWHM for w_0 (dashed, 0.52 λ/NA) spreads the same power over a wider spot: the peak, and so the trap depth and ω_r, come out 25% low. The Airy rings carry the remaining ~16% of the power, which the Gaussian model ignores.

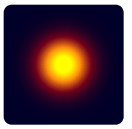


Snapshot of an interactive figure at its default settings. On the website it responds to the controls.



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Airy spot (√I scale)



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

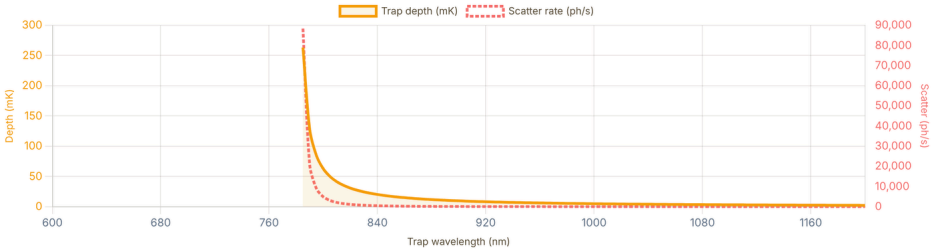
Matched Gaussian

Radius in units of λ/NA, so the picture holds for every wavelength and objective. Same total power for all three curves; intensity normalised to the Airy peak.

8.3 Trap Parameter Calculator

Two-level approximation using the dominant optical transition (D2 for alkali, $^1P_1$ for Yb/Sr). Accurate to ~10–20% for tweezers detuned more than 50 nm from resonance. For Yb/Sr near magic wavelengths, consult full polarizability calculations.

Trap depth & scatter rate vs trap wavelength — at current NA and power. Only red-detuned side shown (λ > resonance).



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

8.4 Array Geometry

Tweezer arrays are produced by a spatial light modulator (SLM) imposing a holographic phase on the trapping beam, or by acousto-optic deflectors (AODs) for dynamic reconfiguration. Cross-talk requires site spacing $d \gtrsim 3w_0$.

Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

5 µm

> To scale · w_0 = 0.70 µm · z_R = 1.79 µm · d = 5.0 µm · beams shown to ±3 z_R ($w = \sqrt{10}\, w_0$ = 2.20 µm)

Sites: **25**

Total power: **750 mW**

d/w_0: **7.2**

✓ Negligible crosstalk

8.5 Species & Wavelength Reference

Representative operating points (order of magnitude; a few mW per trap and $w_0 \approx 0.8$–1.5 µm). Depth and ω_r scale with your power and waist, so use the calculator above for real numbers. η is $k\sqrt{(\hbar/2m\omega_r)}$ on the resonance line listed, at the ω_r in the table. Magic wavelengths eliminate differential AC Stark shifts on clock transitions — essential for AEL qubit coherence.

SPECIES	RESONANCE	TYPICAL Λ_TRAP	DEPTH (MK)	Ω_r/2Π (KHZ)	H	NOTES
Rb-87	780 nm D2	850 nm	~0.8	~80	~0.22	Most common; Ti:Sapph or diode

Cs-133	852 nm D2	938 nm / 1064 nm	~1.0	~50	~0.20	Heavy mass helps; Lukin / Thompson groups
Li-6	671 nm D2	1064 nm	~0.5	~150	~0.70	Light mass → high ω_r; fermion
Li-7	671 nm D2	1064 nm	~0.5	~140	~0.67	Boson partner to Li-6
Na-23	589 nm D2	1064 nm	~0.7	~100	~0.50	Needs fiber amplifier; Greiner / Zwierlein
K-39	767 nm D2	1064 nm	~0.6	~85	~0.32	Feshbach resonances; Hulet / Salomon
Yb-171	399 nm $^1P_1$	532 nm / 759 nm★	~1.2	~45	~0.40	★759 nm magic for clock; AEL qubit
Sr-88/87	461 nm $^1P_1$	813 nm★	~0.9	~55	~0.44	★813 nm magic for $^1S_0 \rightarrow {}^3P_0$; Ye / Kaufman

Magic wavelength traps for AEL atoms (Yb, Sr)

For alkaline-earth-like atoms used as optical clock qubits (^{171}Yb, ^{87}Sr), the trap wavelength must be a *magic wavelength* where the ground state $^1S_0$ and the clock state $^3P_0$ experience the same AC Stark shift, eliminating differential light shifts that otherwise dephase the qubit.

Yb-171: magic wavelengths near 759.35 nm (578 nm clock line) and ~1036 nm (556 nm intercombination line). Sr-87/88: magic near 813.4 nm (698 nm clock line) and 914 nm (689 nm intercombination line). Outside these wavelengths a fractional intensity noise $\delta I/I$ translates directly to qubit frequency noise $\delta\nu = (\alpha_1 - \alpha_0)\cdot\delta I/(2h\varepsilon_0 c)$ — untenable for quantum computing coherence times.

Alkali atoms (Rb, Cs, Li) have no accessible clock transition — use whatever wavelength gives the best depth/scatter trade-off. For Rb, 850 nm is a common sweet spot: far enough from 780 nm D2 for low scatter, close enough to maintain reasonable trap depth per mW.

SLM vs AOD array generation

SLM (spatial light modulator): holographic phase pattern imprinted on a liquid-crystal SLM, then Fourier transformed by an objective. Can generate arbitrary 2D or 3D geometries — triangular, honeycomb, Kagome — but updating the pattern takes ~20 ms (one-sided liquid crystal). Efficiency (fraction of power in target sites) typically 40–70%; remaining light must be blocked or redirected.

AOD (acousto-optic deflector): two crossed AODs can deflect to any (x,y) position in ~µs by changing the RF frequency. Enables rapid rearrangement ("sorting") of atoms to fill defects — key for scaling. In time-averaged operation the effective power per site scales roughly as 1/N, while multi-tone static operation trades power uniformity against RF bandwidth and diffraction efficiency. Combining a static pattern with a dynamic steering beam is common when fast rearrangement is needed.

8.6 References & Further Reading

Key papers and reviews for optical tweezer physics, array assembly, and single-atom experiments.

Foundational theory & review articles

- **Grimm, Weidemüller & Ovchinnikov (2000)** — "Optical dipole traps for neutral atoms." *Adv. Atom. Mol. Opt. Phys.* **42**, 95–170. The canonical derivation of the dipole force, trap depth, and scattering rate starting from the two-level model. Every formula in this calculator follows directly from §2. DOI →
- **Kaufman & Ni (2021)** — "Quantum science with optical tweezer arrays of ultracold atoms and molecules." *Nature Physics* **17**, 1324–1333. The modern comprehensive review: tweezer physics, array assembly, Rydberg gates, AEL qubits, molecules. Start here if you want the full current picture. DOI →
- **Saffman, Walker & Mølmer (2010)** — "Quantum information with Rydberg atoms." *Rev. Mod. Phys.* **82**, 2313. Reviews tweezer arrays as the hardware platform for Rydberg two-qubit gates; covers blockade radius, C_6 coefficients, and fidelity requirements. DOI →
- **Metcalf & van der Straten, *Laser Cooling and Trapping* (1999)** — Springer. Chapter 11 covers the dipole force and optical lattice/tweezer potentials from first principles. Good reference for the Rayleigh range geometry and adiabaticity conditions.
- **Foot, *Atomic Physics* (2005)** — Oxford University Press. Chapter 9 "Laser cooling and trapping" includes a clear derivation of dipole trap depth vs detuning and the photon scattering rate trade-off.

Key experimental papers — single atoms & arrays

- **Schlosser et al. (2001)** — "Sub-Poissonian loading of single atoms in a microscopic dipole trap." *Nature* **411**, 1024–1027. First demonstration of single-atom tweezer trapping via collisional blockade. Established the sub-Poissonian loading mechanism and measured trap frequencies from parametric excitation. DOI →
- **Barredo et al. (2016)** — "An atom-by-atom assembler of defect-free arbitrary two-dimensional atomic arrays." *Science* **354**, 1021–1023. SLM + moving tweezer assembly of defect-free 2D arrays. The method used in most neutral-atom QC experiments today. DOI →
- **Endres et al. (2016)** — "Atom-by-atom assembly of defect-free one-dimensional cold atom arrays." *Science* **354**, 1024–1027. Lukin, Greiner and Vuletić groups (Harvard/MIT): complementary approach that sorts atoms by moving the tweezers themselves. DOI →
- **Kaufman et al. (2014)** — "Two-particle quantum interference in tunnel-coupled optical tweezers." *Science* **345**, 306–309. Hong-Ou-Mandel interference of two Rb atoms in adjacent tweezers; demonstrates ground-state preparation and Lamb-Dicke regime operation. DOI →
- **Norcia et al. (2019)** — "Seconds-scale coherence on an optical clock transition in a tweezer array." *Science* **366**, 93–97. Kaufman group (JILA): strontium atoms in a magic-wavelength tweezer array, with coherence on the optical clock transition lasting seconds. DOI →
- **Evered et al. (2023)** — "High-fidelity parallel entangling gates on a neutral-atom quantum computer." *Nature* **622**, 268–272. Lukin group (Harvard): 99.5% two-qubit gate fidelity in an Rb tweezer array; key Rb benchmark for neutral-atom gates. DOI →

Online resources & atomic data

- **Steck atomic data pages** (steck.us/alkalidata) — precise spectroscopic constants, hyperfine structure, D1/D2 line parameters, oscillator strengths, and saturation intensities for Rb, Cs, Na, Li, K. The authoritative online reference for any AMO calculation.
- **NIST Atomic Spectra Database** (physics.nist.gov/asd) — transition wavelengths, A-coefficients, and energy levels for all elements including Yb and Sr.
- **Lukin group resources** (lukin.physics.harvard.edu) — open-access papers on tweezer arrays, Rydberg gates, and atom assembly protocols.

CHAPTER 8 SUMMARY

Optical tweezers give single-atom spatial control across hundreds to thousands of sites. Because a high-NA objective is unforgiving of optical imperfections, we next characterize and cancel wavefront aberrations.

Problems for Chapter 8

1. USE THE TOOL For ^{87}Rb at the default wavelength and power, record the waist, trap depth, radial and axial frequencies and the Lamb–Dicke parameter. Then lower the NA by 30% and record them again.
2. BY HAND Using $U \propto P/w_0^2$, $\omega_r \propto \sqrt{U}/w_0$ and $\omega_z \propto \sqrt{U}/z_R$, predict the changes you saw in problem 1 from the change in waist alone.
3. DESIGN Design a 10 × 10 array: choose the spacing, check the cross-talk criterion, and compute the total power you need. Is your design limited by the laser or by the optics?

Use the tool: on the interactive page. By hand: from the formulas in this chapter. Design: a decision a lab would make. Worked solutions are available to instructors on request (see Part V).

CHAPTER 9 · BUILD STEP 09

9

Wavefront Aberrations & Zernike Modes

🌊 Zernike Wavefront Lab · Simulator

Learn the optical language behind imperfect objectives, shaped tweezers, SLM holograms, and PSFs. Start from one Zernike mode, combine aberrations, watch the Strehl ratio fall, then use Fourier optics to connect pupil phase to the light pattern in the focal plane.

> At NA > 0.5, *even a fraction of a wavelength of coma, astigmatism, or spherical aberration from vacuum viewports degrades the Strehl ratio and destroys uniform tweezer depths across the field of view.*

IN THIS CHAPTER Prerequisites: Ch. 8

- Decompose pupil wavefront errors into orthogonal Zernike polynomials and recognize their signatures in the point-spread function (PSF).
- Quantify how RMS wavefront error lowers the Strehl ratio and widens the effective tweezer waist.
- Use spatial light modulator (SLM) phase masks and Gerchberg-Saxton algorithms to equalize trap depths and flatten aberrations in situ.

WEBSITE COMPANION Live interactive page ↗ · Week 8

9.1 What problem is this page solving?

A high-NA objective, SLM, vacuum viewport, or imperfect relay does not just make light dimmer; it changes the phase across the pupil. Zernike modes are the clean vocabulary for naming those phase errors. Once you know the mode, you can predict the PSF damage, estimate Strehl, and decide whether to shim, realign, pre-compensate on an SLM, or stop worrying.

pupil phase → focal field $W(\rho,\theta)=\Sigma c_n^m Z_n^m$ Strehl$\approx\exp[-(2\pi\sigma)^2]$

1. **Name the aberration**

Start with one Zernike mode. Learn what defocus, astigmatism, coma, and spherical aberration look like.

2. **Quantify optical damage**

Combine modes, compute RMS wavefront error, and watch the Strehl ratio and PSF respond.

3. **Connect to AMO hardware**

Use SLM and Fourier-optics panels to see how phase masks become tweezers, OAM beams, or distorted spots.

4. **Correct or synthesize**

Use Gerchberg-Saxton for phase retrieval, then use the mode atlas as a reference sheet.

5 chapters on this page. Open one from a card, or from the chapter bar below, which stays pinned while you read.

Single atoms **PSF quality is survival and fidelity.**

Aberrations broaden the focus, lower collection efficiency, and make single-atom images harder to classify.

Tweezers **Wavefront error becomes trap error.**

A distorted phase front changes waist, depth, spacing, and uniformity across large tweezer arrays.

SLMs **Holograms are intentional aberrations.**

The same math used to diagnose bad optics also creates LG beams, OAM modes, and programmable spot arrays.

1 / 5

Start with one aberration at a time.

Move through radial order n and azimuthal order m. The goal is pattern recognition: when a real image looks comet-tailed, stretched, or ringy, you should know which knob to suspect.

intuition first

Mathematical convention used here (OSA/ANSI)

Orthonormal over the unit disk ρ ≤ 1 in polar coordinates (ρ, θ):

$$Z_n^m(\rho,\theta) = \begin{cases} \sqrt{2(n+1)}\, R_n^{|m|}(\rho)\cos(m\theta) & m > 0 \\ \sqrt{2(n+1)}\, R_n^{|m|}(\rho)\sin(|m|\theta) & m < 0 \\ \sqrt{n+1}\, R_n^0(\rho) & m = 0 \end{cases}$$

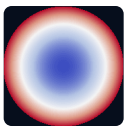
$$R_n^{|m|}(\rho) = \sum_{s=0}^{(n-|m|)/2} \frac{(-1)^s\,(n-s)!\;\rho^{n-2s}}{s!\,\left(\frac{n+|m|}{2}-s\right)!\,\left(\frac{n-|m|}{2}-s\right)!}$$

Orthonormality: ∫|Z_n^m|² ρ dρ dθ / π = 1. Any wavefront W(ρ,θ) = Σ c_n^m Z_n^m; RMS WFE = √(Σ c²).

Standard mode names

Z_0^+0 Piston Z_1^-1 Tilt Y Z_1^+1 Tilt X Z_2^-2 Astig 45° Z_2^+0 Defocus Z_2^+2 Astig 0° Z_3^-3 Trefoil Y Z_3^-1 Coma Y Z_3^+1 Coma X Z_3^+3 Trefoil X Z_4^-4 Tetrafoil Y Z_4^-2 Sec.Astig Y Z_4^+0 Prim.Spherical Z_4^+2 Sec.Astig X Z_4^+4 Tetrafoil X Z_5^-5 Pentafoil Y Z_5^-3 Sec.Trefoil Y Z_5^-1 Sec.Coma Y Z_5^+1 Sec.Coma X Z_5^+3 Sec.Trefoil X Z_5^+5 Pentafoil X Z_6^+0 Sec.Spherical

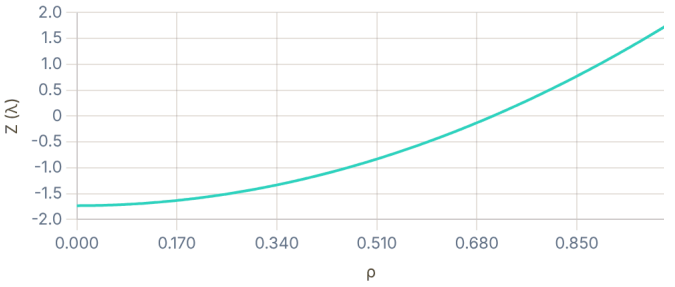
Z_N^M (OSA, UNITS Λ)



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

RdBu: blue=min, white=0, red=max

RADIAL CROSS-SECTION (Θ = 0 FOR M ≥ 0; Θ = Π/2|M| FOR M < 0)



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Now combine modes like a real optical system.

Real objectives rarely have a single pure aberration. Add terms, read RMS wavefront error, and use the PSF/Strehl output as the practical verdict.

diagnose

Wavefront from Zernike Expansion

Any wavefront over a circular aperture: $\mathbf{W(\rho,\theta) = \Sigma\ c_n\hat{}m \cdot Z_n\hat{}m(\rho,\theta)}$, coefficients in waves (λ). RMS WFE = $\sqrt{(\Sigma\ c^2)}$, summed over every term except piston (a constant phase changes nothing). Strehl ratio (Maréchal approximation):

$S \approx \exp[-(2\pi \cdot \sigma_W)^2]$

Strehl > 0.8 → diffraction-limited by the Maréchal criterion. The PSF is computed via 2D FFT of the pupil function: $P(\rho,\theta) = A(\rho) \cdot \exp[i \cdot 2\pi \cdot W(\rho,\theta)]$.

WAVEFRONT W(P,Θ) (Λ)



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

RdBu colormap

PSF = |FT{P}|²

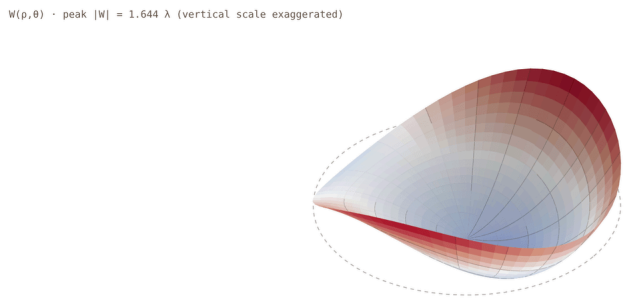


Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Inferno · field ±16 λ/D

IN 3D: THE SAME WAVEFRONT AND PSF AS SURFACES

Defocus is a bowl (ρ^2), astigmatism a saddle ($\rho^2 \cos 2\theta$), coma an S-shaped fold along one axis ($(3\rho^3 - 2\rho)\cos\theta$). Switch to the PSF to see what that shape does at the focus.



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Height ∝ W: the largest |W| on the grid, 1.644 λ, is drawn at half the aperture radius, so the vertical scale is exaggerated (schematic); colours are the map's RdBu scale. RMS 0.615 λ, peak-to-valley 2.715 λ.

Turn pupil phase into focal-plane light.

This is the bridge from aberration language to AMO hardware: SLM phase masks, vortex beams, diffraction orders, and tweezer-array Fourier optics.

hardware bridge

SLM Phase → Far-Field Intensity

A spatial light modulator (SLM) imprints a programmable phase φ(x,y) on an input Gaussian beam. In the focal plane of a lens, the field is the Fourier transform of the modulated input:

$E_{ff}(u,v) = \mathrm{FT}\{ E_{in}(x,y) \cdot \exp[i\varphi(x,y)] \} \rightarrow I_{ff} = |E_{ff}|^2$

Helical phase $\varphi = l\cdot\theta$ converts a Gaussian to a Laguerre–Gaussian LG_0^l donut (for $l \neq 0$), carrying orbital angular momentum $l\cdot\hbar$ per photon.

INPUT INTENSITY $|E_{IN}|^2$



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Inferno colormap

SLM PHASE Φ (MOD 2Π)



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

HSV cyclic: 0 → 2π = red → red

FAR-FIELD INTENSITY $|FT\{E\}|^2$



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Inferno · field ±16 λ/D

Laguerre–Gaussian LG_0^l beams

Carry OAM l·ħ per photon. At the beam waist (p=0): $|LG_0^l| \propto (r\sqrt{2}/w)^{|l|} \cdot \exp(-r^2/w^2)$ Intensity ring at $r = w\sqrt{(|l|/2)}$. Larger ring for larger |l|.

SLM recipe for LG_0^l:
1. Input: Gaussian beam
2. Phase: $\varphi = l \cdot \theta$ (helical/vortex phase)
3. Far field $\approx LG_0^l$ (donut for $l \neq 0$)

Blazed grating $l \cdot \theta + k \cdot x$: shifts donut off-axis, separates OAM orders spatially

L	NAME	FAR-FIELD SHAPE
0	Gaussian (LG_0^0)	Central Airy peak
±1	First-order vortex	Single donut ring
±2	Second-order vortex	Larger donut ring
±3	Third-order vortex	Even larger ring
0 + defocus	Defocused Gaussian	Broadened central peak
0 + astigmatism	Astigmatic Gaussian	Elongated / elliptical
0 + grating	Blazed diffraction	Off-axis Gaussian spot

Ask the inverse question: what phase makes this light?

Gerchberg-Saxton is the workhorse idea behind many SLM tweezer-array holograms: alternate between planes, keep the desired amplitudes, update the phase.

inverse design

Gerchberg-Saxton (GS) Algorithm — iterative phase retrieval

Given a desired far-field intensity T(u,v) (e.g. a reconfigurable tweezer array), GS finds the SLM phase pattern φ(x,y) that produces it. The algorithm alternates between the SLM plane and the target plane, enforcing known amplitude constraints in each:



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

TARGET INTENSITY

Desired tweezer pattern

SLM PHASE Φ (MOD 2Π)

Computed phase hologram

RECONSTRUCTED FAR-FIELD

$|\mathcal{F}\{A\, e^{i\varphi}\}|^2$ at each iteration (A = Gaussian input beam)

CONVERGENCE: NORMALISED RMSE VS ITERATION

Algorithm details

Step 1 (SLM plane): Combine input Gaussian amplitude A(x,y) with current phase estimate $\varphi_k(x,y)$.
Step 2 (FFT → target plane): Propagate the field to the focal plane using a 2D discrete Fourier transform.
Step 3 (target constraint): Replace the amplitude with $\sqrt{T(u,v)}$ while keeping the retrieved phase arg(FT{E}).
Step 4 (IFFT → SLM plane): Propagate back; extract new phase φ_{k+1}.
Typically converges in 20–50 iterations. Efficiency (fraction of power in target spots) reaches 70–90% for typical arrays; adding a blazed grating shifts the 1st order off-axis to separate it from undiffracted light.

GS convergence criterion: RMSE =

$$\sqrt{\frac{1}{P}\sum_i \left(|E_i|^2 - T_i\right)^2}$$

/ max(T) (normalised to the peak of the scaled target)

Use this as the cheat sheet.

Once you understand the workflow, the atlas is a fast visual reference for identifying modes by eye and remembering the OSA/ANSI indexing structure.

reference

Zernike Pyramid, all modes at a glance

Every Zernike polynomial Z_n^m arranged in the standard pyramid (OSA/ANSI order). Columns index the azimuthal frequency m, rows the radial order n. Cosine modes (m > 0) on the right, sine modes (m < 0) on the left. **Blue** = negative wavefront, **white** = zero, **red** = positive (RdBu colormap).

9.2 Sources

ANSI/OSA convention family for Zernike indexing and optical-aberration notation Noll (1976), *J. Opt. Soc. Am.* 66 207 — Zernike polynomials and atmospheric turbulence Gerchberg & Saxton (1972), *Optik* 35 237 — phase retrieval / hologram algorithm (no DOI) Allen et al. (1992), PRA 45 8185 — orbital angular momentum of light and its interaction with matter; Laguerre-Gaussian modes Born & Wolf (1999), *Principles of Optics* — Chapter 9 covers Zernike polynomials and diffraction theory in full depth Tyson (2010), *Principles of Adaptive Optics* — wavefront sensing, Zernike decomposition, DM control in AO systems Padgett & Allen (2000), *Contemp. Phys.* 41 275 — light beams carrying orbital angular momentum and SLM holography

CHAPTER 9 SUMMARY

Correcting low-order Zernike aberrations restores diffraction-limited tweezers and crisp single-atom fluorescence imaging. As array size and shot counts grow, manual tuning becomes unsustainable—leading directly to laboratory automation.

Problems for Chapter 9

1. USE THE TOOL Build a wavefront with 0.5 λ of coma and read the Strehl ratio. Compare it with the Maréchal estimate $S \approx \exp[-(2\pi\sigma)^2]$, where σ is the rms wavefront error in waves.
2. USE THE TOOL Add an OAM order of l = 1 to a blazed grating and compute the far field. Explain the shape you see.

3. DESIGN Run the Gerchberg–Saxton algorithm for a target pattern with 10, 30 and 100 iterations. How does the uniformity change, and how would you check uniformity on the atoms rather than on a camera?

Use the tool: on the interactive page. By hand: from the formulas in this chapter. Design: a decision a lab would make. Worked solutions are available to instructors on request (see Part V).

CHAPTER 10 · BUILD STEP 10

10

Laboratory Automation & Closed-Loop Optimization

🤖 Automating the Experiment · Guide

A well-run AMO experiment looks after itself for most of the night. Servos hold powers, fields and frequencies. Scripts relock lasers and recalibrate whatever drifts. Optimisers tune one stage of the sequence while you sleep. This page is the practical guide: feedback and feedforward, beam-pointing stabilisation, relock and calibration loops, a workflow for optimising one part of the experiment, where machine learning has already been used, and what AI agents can and cannot be trusted with.

> *A modern neutral-atom experiment has dozens of drifting knobs—laser powers, beam pointing, coil offsets, and ramp timings—that must stay locked over hours of data taking.*

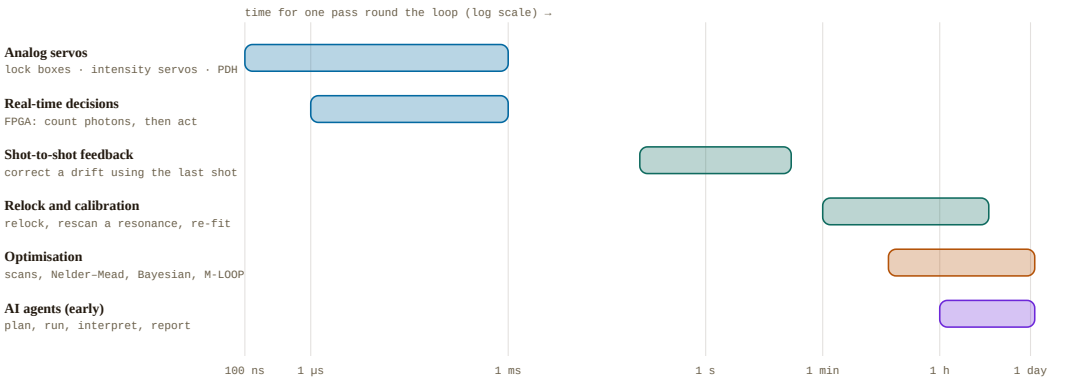
IN THIS CHAPTER Prerequisites: Ch. 4, Ch. 7, Ch. 9

- Contrast feedback servos (PID intensity and beam-pointing stabilization) with feedforward compensation for repeatable disturbances.
- Apply derivative-free optimizers (Nelder-Mead, Gaussian-process Bayesian optimization, M-LOOP) to tune MOT loading, cooling, and imaging stages.
- Design automated relock scripts, telemetry dashboards, and safe guardrails for unattended operation.

WEBSITE COMPANION Live interactive page ↗ · Notebook: Open in Colab · download automation.ipynb · Week 13

10.1 What to automate, and at which timescale

Automation is a stack of loops. Each one senses something, decides, and acts, and each closes at its own speed. The fast loops are electronics; the slower ones are scripts on the control PC; the slowest are optimisers and, since 2024, AI agents. A loop belongs at the slowest level that still keeps up with the disturbance it fights.



Blue: electronics. Teal, amber and violet: software on the control PC. Bars show orders of magnitude for one pass round each loop, not limits. A PDH lock corrects in microseconds; a resonance recalibration takes a few minutes of shots; an optimisation run can take a night.

Automate first

What you do by hand every day

Relocking a laser, re-peaking a fibre, checking a resonance, re-measuring a π-pulse. These are well understood, easy to verify, and they eat hours of student time. Automating them also makes the data more consistent, because the check is done the same way every time.

Automate carefully

What can break hardware

Coil currents, high laser powers, RF amplifiers, vacuum valves and anything that heats. Limits for these belong in hardware (interlocks, current limits, thermal switches), with software limits on top. A loop that rails at full scale must fail safe, not try harder.

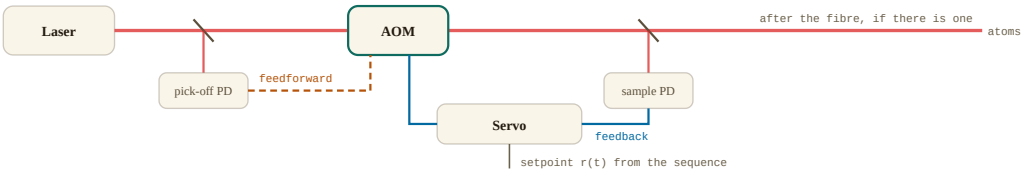
Not yet

What you do not understand

An optimiser maximises the number you give it, including its flaws. If the "atom number" can be inflated by a background count, a saturated camera or a mis-set threshold, the optimiser will find that first. Understand the measurement before you let a machine push on it.

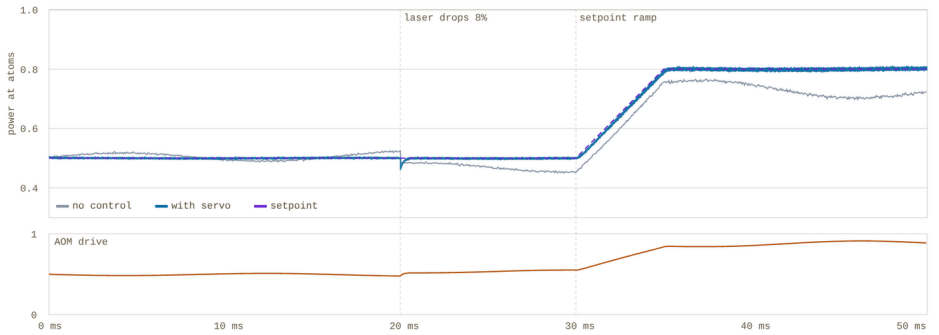
10.2 Feedback and feedforward: holding a laser's power

The first servo most students build. A photodiode samples the beam after the AOM, a servo compares it with a setpoint and adjusts the AOM's RF amplitude. Feedback corrects whatever it sees, but only after it sees it and only up to its bandwidth. Feedforward acts before an error appears, but only on what you can predict or measure, and only as well as it is calibrated. Real systems use both.



Power-servo simulator

50 ms of a beam whose laser power has mains ripple, slow drift, a sudden 8% drop at 20 ms and fast noise. The setpoint holds, then ramps up at 30 ms. Grey: what reaches the atoms with no control. Coloured: with your servo.



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

rms power error **0.65%** 0.5–5 ms after the drop **0.27%** worst error on the ramp **3.1%** stability limit (I) **≈ 167 kHz**

> **Feedback only.** The integrator removes the drift and recovers from the drop in roughly $1/(2\pi f_c) \approx 156$ µs. It always reacts after the fact, so the ramp lags slightly. Setpoint feedforward removes that lag.

Simulated in the page at 0.25 µs steps. Laser power: 3% mains ripple, slow drift of about 2%, an 8% drop at 20 ms, 0.3% fast noise. Photodiodes: 0.1% noise. The AOM is a 0.5 µs first-order response after the loop delay, and its transmission is clamped between 0 and 1. The pick-off feedforward assumes a 5% calibration error. The numbers are illustrative; the behaviour is general.

Rules of thumb

- **Sample the light where the atoms are.** After the fibre, if there is one; fibre coupling and polarisation drift change the power after the AOM.
- **The loop is only as quiet as its photodiode.** Sensor noise is written onto the beam. Check with a second, out-of-loop photodiode: the in-loop one always looks good.
- **Delay sets the bandwidth.** A pure integrator with a total delay τ goes unstable near $f \approx 1/(4\tau)$, where the delay has used up its 90° of phase margin. Cables, the acoustic travel time in the AOM and digital processing all count.
- **Leave headroom.** Run the AOM below full diffraction so the servo can push up when the laser sags. At the limit the servo saturates and the power simply follows the laser.

Practical details

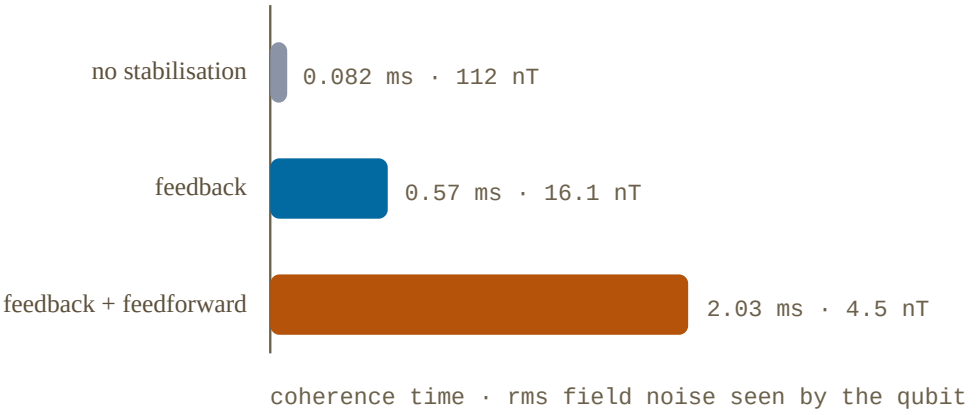
- **Pulsed beams.** A servo cannot regulate light that is off. Use a sample-and-hold (servo while on, hold the drive while off) or a digital loop that measures each pulse and corrects the next shot.
- **Which actuator.** Feeding back to a diode laser's current also moves its frequency, so the AOM is the usual actuator. An EOM between polarisers ("noise eater") is faster still.
- **Setpoint feedforward** is a calibration table: measured power against AOM drive. Then a ramp is followed immediately and the servo only cleans up the residual.
- **Digital or analog.** FPGA servos (ARTIQ's SU-Servo combines ADC, servo and DDS amplitude in one loop) make the setpoint part of the sequence; analog PI boxes are simpler and have less delay.

10.3 Fields, frequencies and the mains

The same two ideas, feedback for what you cannot predict and feedforward for what you can, hold magnetic fields and laser frequencies steady. The mains is the textbook case for feedforward: its noise repeats every cycle.

Feedforward on the mains: a trapped-ion example

Merkel et al. stabilised a 14.6 mT field for a $^{43}Ca^{+}$ clock qubit with a fluxgate sensor in a feedback loop, then added feedforward that cancels the 50 Hz mains harmonics. Qubit coherence time, from their Ramsey measurements:



Values from Table II of Merkel et al., arXiv:1808.03310. Their abstract quotes 4.3 nT rms (0.29 ppm) for the stabilised field against more than 100 nT without stabilisation; the coherence time improved 25-fold.

What else is done this way

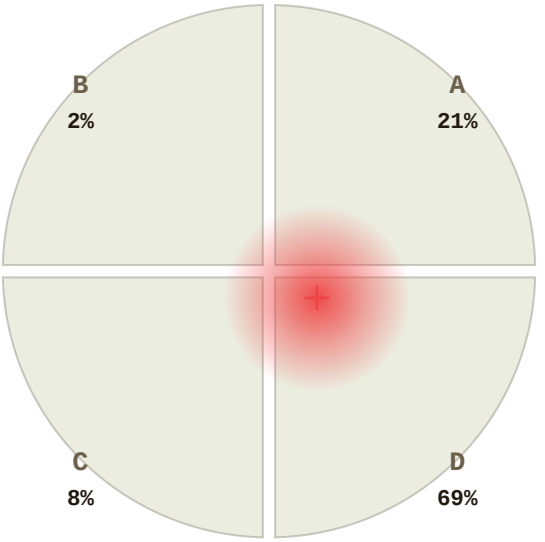
- **Line triggering.** Start every shot at the same phase of the mains. The residual 50/60 Hz field is then the same in every shot, so it shifts a resonance by a fixed amount you can calibrate, instead of broadening it.
- **Laser frequency, fast.** A PDH lock cannot suppress noise above its bandwidth. Chao et al. fed the PDH error signal *forward* to the light after a delay fibre and report phase-noise suppression around a few MHz about four orders of magnitude better than PDH feedback alone (*Optica* 2024). An adaptive version reports ≥40 dB around 2 MHz with up to 50 MHz bandwidth (preprint).
- **Laser frequency, slow.** A software loop reads a wavemeter or a reference cavity and nudges a piezo or setpoint every second or so. It keeps a laser in range for the fast lock; see the laser locking guide.
- **Fields from known sources.** If your own sequence switches a nearby coil, its stray field arrives at a known time. Feed the expected change forward to the compensation coils rather than waiting for the magnetometer to see it.

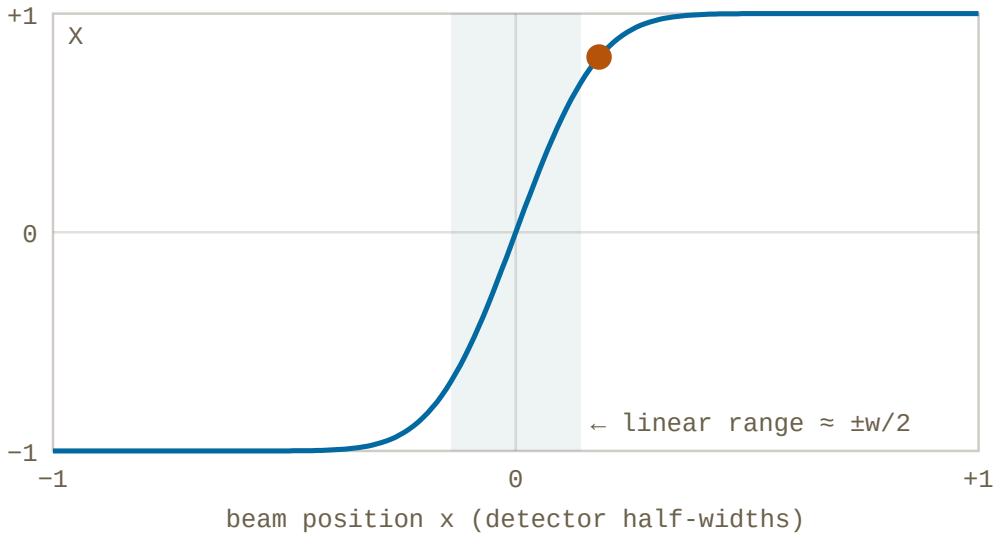
10.4 Keeping beams pointed: quadrant detectors and auto-alignment

Beams wander as the lab warms and cools, as mounts relax and as air moves. A position-sensitive detector turns that wander into an error signal and a piezo-driven mirror turns it back. The same hardware, with a smarter loop, can re-couple a fibre on its own.

A quadrant photodiode as an error signal

Drag the spot. The four quadrant signals give a horizontal error X = [(A+D) − (B+C)] / Σ and a vertical error Y = [(A+B) − (C+D)] / Σ, while Σ itself is the total power. Turn on drift to see the beam wander, then lock it.





X **0.801** Y **-0.525** Σ **1.00** (power) Offset **0.74 w** Loop **open**

For a Gaussian spot of $1/e^2$ radius w on an ideal split detector, X = erf($\sqrt{2}$·x/w): linear only for displacements well below w, then saturating at ±1. A small spot is sensitive but has a short range; a large spot the reverse. Detector size is in units of the half-width of the detector; gaps between quadrants are ignored.

Four degrees of freedom

A beam has a position and an angle in each transverse direction. One detector sees position; a second placed at the focus of a lens sees angle. Two detectors and two piezo mirrors close all four loops. One detector and one mirror only fixes the spot at the detector, and may tilt the beam elsewhere.

Commercial stabilisers

Closed-loop beam stabilisers (MRC Systems is a well-known example) combine piezo-driven steering mirrors with four-quadrant diodes and a controller. They are common in front of long beam paths, high-power fibres, SLMs and AODs, where a small angle becomes a large shift of a tweezer array.

Without a detector

When you only care about the throughput (a fibre), use the throughput itself: a photodiode after the fibre and motorised mirror mounts, driven by a search or an optimiser. Slower, but it optimises exactly what you need.

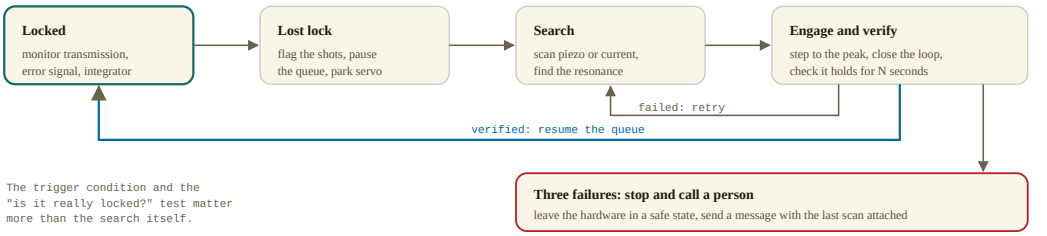
Automatic fibre coupling: three published approaches

APPROACH	HARDWARE	METHOD	REPORTED RESULT
Model-based (Geng et al., 2024)	Two 2D position-sensitive detectors, two piezo mirror mounts, a microcontroller	A linearised geometric model of the beam path, run as a state machine	Recovers the maximum coupling in about 10 s, even from a beam misaligned to zero coupling PREPRINT
Optimiser (Mathew et al., 2020)	Stepper motors on the knobs of two kinematic mounts, a photodiode after the fibre, a Raspberry Pi	Gaussian-process optimiser (M-LOOP)	About 20 minutes for 200 runs; coupling as good as, or slightly better than, an experienced experimenter
Reinforcement learning (Richtmann et al., 2024)	Motorised mirror mounts with backlash and lost steps	Model-free RL agent trained on the real setup	90% coupling, comparable to a human expert; training took nearly four days at about one second per step PREPRINT

The lesson generalises. If you can write a model of the system, a model-based loop is fastest. An optimiser needs no model but costs hundreds of measurements. Learning a policy costs the most, and pays off only when the task repeats under conditions a fixed model cannot capture.

10.5 Auto-relock and calibration loops

The most valuable automation in most labs is unglamorous: noticing that a lock has dropped, relocking it, and re-measuring the parameters that drift. Done well, it turns a machine that needs a student watching it into one that runs overnight.



Detecting a lost lock

- Cavity transmission or a beat note drops below a threshold.
- The error signal's rms grows, or the integrator output approaches its rail (it will soon run out of range even if still locked).
- A wavemeter reading leaves its window.
- The atoms tell you: fluorescence or loading rate collapses. Treat this as a symptom to diagnose, not as a lock signal on its own.

Make it honest

- **Flag, don't average.** Every shot records the lock state of every laser; shots taken while anything was unlocked are marked bad, not silently averaged in.
- **Log every relock.** A laser that relocks ten times a night is telling you something about its temperature or its current.
- **Give up gracefully.** A relock script that keeps trying forever can drive a laser to a mode hop or a piezo to its limit.

Calibrations worth running on a schedule

WHAT DRIFTS	CALIBRATION EXPERIMENT	WHAT THE SCRIPT UPDATES
A transition frequency (field drift, light shifts)	Short spectroscopy scan across the line, fit the centre	The detuning parameter used by every later experiment
Rabi frequency (beam power, pointing, polarisation)	Rabi flop or a pulse-length scan, fit the period	The π-pulse time or amplitude
Bias magnetic field	Microwave or Raman spectroscopy on a field-sensitive transition	Compensation-coil setpoints
Tweezer-to-tweezer uniformity	Per-site light shift, loading or trap-frequency measurement	Per-tone amplitudes on the AWG (see the tweezer designer)
Imaging threshold	Histogram of counts per site, fit the two peaks	The atom/no-atom threshold for each site

Calibrations depend on each other. A π-pulse calibration is meaningless if the resonance has moved. Google's superconducting-qubit group described organising calibrations as a directed acyclic graph: each node knows which others it depends on, checks them first, and only recalibrates what is out of date (Kelly et al., arXiv:1803.03226). The idea transfers directly to atoms. The cadence of each calibration should come from how fast that quantity actually drifts, which your monitoring logs (Section 10) will tell you.

10.6 Workflow: optimising one stage of the experiment

Say the tweezer-loading probability has dropped and you want to re-optimise the stage that loads atoms from the MOT into the tweezers. The steps below work for any stage, with or without machine learning.

1. **Define one number to maximise**, measured at the end of the stage you are tuning (loading probability, atom number, temperature, survival). If you care about two things, combine them deliberately, for example atom number with a penalty above a temperature limit.*Why: the optimiser will maximise exactly this number, flaws included (Section 01).*
2. **Measure the noise of that number.** Repeat the current settings 10–20 times. For a loading probability p from N shots the statistical error is $\sqrt{p(1-p)/N}$: about 5% at p = 0.5 with 100 shots per point.*Why: if two settings differ by less than this, no optimiser can tell them apart. It sets how many shots each point needs.*

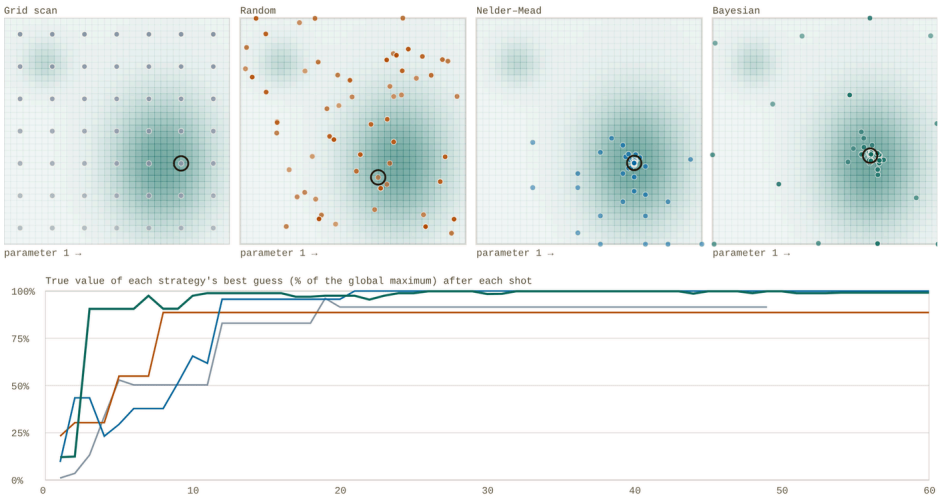
3. **Choose the parameters and safe bounds.** Detunings, intensities, field gradients, durations. Start with the few you believe matter; bound each so that every combination is safe for the hardware.*Why: cost grows with the number of parameters, and an optimiser will try the corners of the box.*
4. **Freeze and monitor everything else.** Lock states, laser powers, fields and the lab temperature should be logged alongside every shot.*Why: a drift during the run looks to the optimiser like a change in the landscape.*
5. **Look before you optimise.** One-dimensional scans of each parameter around the current point show which ones matter and how smooth the landscape is.*Why: if a 1D scan already finds the problem (a detuning off by a linewidth), you are done.*
6. **Run the optimiser with reference points.** Every so often re-measure the original settings.*Why: if the reference drifts, the experiment changed under you, and the run must be corrected or restarted.*
7. **Validate the winner.** Re-measure the best settings many times, interleaved with the old ones. The best single measurement in a noisy run is biased high (the "winner's curse").*Why: the improvement you report should survive a fair comparison.*
8. **Understand it.** Look at what the optimiser changed, and use its model of the landscape (if it has one) to see which parameters are sensitive.*Why: a surprising optimum often reveals a misaligned beam or a wrong calibration, which is worth more than the gain itself.*
9. **Write it down and schedule it.** Commit the new values with a note (see control systems: git). If the optimum drifts, turn the stage into a periodic calibration (Section 05).*Why: next month's student needs to know what changed and why.*

10.7 Choosing an optimiser

Every shot costs seconds and every measurement is noisy, so the question is not which optimiser is cleverest but which finds a good setting in the fewest shots, without being fooled by noise. The race below runs four common strategies on the same toy landscape.

Optimiser race on a noisy two-parameter landscape

Each panel shows the true landscape (stronger colour is better) and the points each strategy measured. The chart below shows, after each shot, how good the strategy's current best guess really is. Raise the noise to see which methods get fooled.



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Grid scan Random Nelder–Mead Bayesian (Gaussian process)

Grid scan **92%** Random **89%** Nelder–Mead **100%** Bayesian **99%** after **60** shots

A toy landscape, not a real MOT: a broad main peak, a smaller second peak and a gentle slope, with Gaussian measurement noise in units of the peak height. Grid, random and Nelder–Mead report the best point they measured; the Bayesian optimiser reports the measured point its model rates highest, which is what protects it from lucky noisy shots. The grid uses the largest square number of shots that fits the budget. Results vary with the noise seed, so press "New noise" a few times before drawing conclusions.

Scans

Grid and 1D scans

Easy to understand and to plot, and the right first look. The cost is k^d points for k values of d parameters, so a grid is hopeless beyond two or three parameters.

Simplex

Nelder–Mead

Needs no model and few points per step. It was the workhorse of early online optimisation, but it can mistake noise for a slope and stall on a ridge, and it only finds the nearest peak.

Model-based

Bayesian optimisation

Fits a Gaussian process (a smooth model with error bars) to every shot so far, then measures where the model says improvement is likely. Efficient for expensive, noisy shots and tens of parameters; Vendeiro et al. tuned up to 55.

Population

Evolutionary methods

Differential evolution and CMA-ES keep a population of settings and breed better ones. Robust to noise and local peaks, but they spend more shots; good when shots are fast.

Neural networks

Neural-network learners

Replace the Gaussian process with a network when there are many parameters. Tranter et al. tuned a 63-parameter MOT compression sequence this way.

Policies

Reinforcement learning

For a different task: not one best set of numbers, but a policy that reacts to what it observes during the shot. Costly to train; usually trained in simulation first.

Tools. M-LOOP (Machine-Learner Online Optimization Package, from the group behind Wigley et al.) is built for exactly this job: it offers Gaussian-process, neural-net, differential-evolution, Nelder–Mead and random learners and talks to an experiment through files or a Python interface. General libraries such as scikit-optimize or BoTorch work too if you write the interface yourself. Either way the optimiser only proposes settings; your control system runs them (see control systems).

10.8 Machine learning in AMO labs: what has been done

Machine learning entered cold-atom labs as an optimiser for sequences too long to tune by hand, then spread to readout, control and alignment. These are representative papers, with what each one actually reports.

PAPER	WHAT WAS AUTOMATED	METHOD	REPORTED RESULT
Wigley et al., *Sci. Rep.* 6, 25890 (2016)	Evaporation ramp for a Bose–Einstein condensate	Gaussian-process online optimiser (the origin of M-LOOP)	A ramp producing high-quality BECs in 10 times fewer iterations than the previous online optimiser (Nelder–Mead)
Tranter et al., *Nat. Commun.* 9, 4360 (2018)	MOT compression sequence for maximum optical depth	Stochastic artificial neural network; 63 parameters (21 time bins × 3 controls)	Optical depth 970 ± 20 against 535 ± 8 for the best human-designed sequence
Barker et al., *Mach. Learn.: Sci. Technol.* 1, 015007 (2020)	Laser cooling and evaporation together, for a BEC	Differential evolution, Gaussian process and neural network compared	Four times the BEC atom number of the manually optimised settings; a BEC from fully randomised starting parameters
Vendeiro et al., *Phys. Rev. Research* 4, 043216 (2022)	Fast BEC production via Raman cooling in an optical trap	Bayesian optimisation of up to 55 control parameters	BECs of 2.8×10^3 ^{87}Rb atoms from a room-temperature gas in 575 ms
Reinschmidt et al., *Nat. Commun.* 15, 8532 (2024)	Live control of a ^{87}Rb MOT (detuning and field gradient) from fluorescence images	Deep reinforcement learning, trained in simulation and transferred to the experiment	Adaptive loading control; compensated a sudden detuning step and met atom-number targets on demand
Phuttitarn et al., *Phys. Rev. Applied* 22, 024011 (2024)	State readout of a Cs tweezer array	Convolutional neural networks on the camera images	Up to 56% (single-site) and 32% (whole-array) lower readout error; crosstalk correlation reduced by up to 78.5%
Lee & Chae, arXiv:2401.04893	Choosing tweezer moves for defect-free rearrangement	Reinforcement learning (PPO), simulations of 50–300 tweezers	A simulation study, compared against A* path search; no experiment PREPRINT

Reading these papers. Notice what they compare against. "Better than our manual settings" depends on how hard the humans tried; "fewer shots than Nelder–Mead" is a cleaner benchmark. The lasting value is often in what the optimiser found: Tranter et al.'s sequences and Vendeiro et al.'s fast Raman-cooling route are now physics, not just numbers.

10.9 AI agents and language models in the lab

Since late 2024 several groups have connected large language models to real experiments, letting them plan, run and interpret measurements. These are early results, mostly preprints, and the architecture of each system is as instructive as the result.

Superconducting qubits

k-agents (Cao et al.)

LLM agents turn a lab procedure written in plain language into a state machine of experiments, run it, and use separate inspection agents to judge each result before moving on. They ran calibrations autonomously for hours on a 16-qubit processor, reaching 99.87% single-qubit gate fidelity after recalibration.

arXiv:2412.07978; journal version in *Patterns* (2025).

Cold atoms

QCopilot (Sha et al.)

A team of LLM agents (decision maker, experimenter, analyst, fault diagnoser, recorder) ran Bayesian optimisation on a Rb cold-atom apparatus and diagnosed faults. The authors report 10^8 sub-μK atoms without human intervention within a few hours, about 100 times faster than manual tuning.

arXiv:2508.05421 PREPRINT

NV centres

Agentic quantum sensing (Isogawa et al.)

An LLM agent planned and interpreted NV-centre measurements, deciding for itself to run a CPMG follow-up after a weak Ramsey feature. It had no direct hardware access: it wrote requests to a queue that deterministic control software checked and executed.

arXiv:2607.25145 PREPRINT

Where LLMs help today

- Writing and explaining instrument drivers, analysis scripts and sequence code, which you then test like a new student's code.
- Reading the night's logs and summarising what failed and when.
- Turning a procedure you already do by hand into a script or a checklist.
- Searching manuals and datasheets for the setting you need.

Guardrails

- **Hardware interlocks stay independent of all software**, agent or not.
- **The agent proposes, the control system disposes.** Requests go through a queue that checks every parameter against fixed, safe ranges (the design in the NV-centre paper above).
- **A person approves anything irreversible or hazardous**: high laser power, coil currents, valves, bake-outs.
- **Every action is logged** with the reasoning the agent gave, so a bad night can be reconstructed.

10.10 The infrastructure that makes automation safe

Automation without monitoring is a machine that fails quietly. These four pieces are what let a group trust an experiment that runs while nobody is watching.

Log every slow variable

Laser powers, lock states, wavemeter readings, temperatures and humidity, vacuum pressures, coil currents: sampled continuously into a time-series database with dashboards (a common choice is InfluxDB with Grafana). When the atom number dropped at 3 a.m., the logs show that the lab temperature rose at 2:40.

Alert a person

A message to the group chat or a phone when a lock fails repeatedly, a pressure rises, cooling water stops, or a calibration lands far from its usual value. Alerts should be rare and specific, or people learn to ignore them.

One source of truth for parameters

Calibrated values (resonances, π-times, thresholds) live in one database or file with history. Every shot records which values it used, and every automatic update is logged with the data that justified it.

Test it without atoms

Run new automation against a simulator or with the lasers blocked before letting it loose. Check what it does when a sensor is unplugged, a reading is NaN or a device does not answer: those are the failures that happen at night.

10.11 Checklist before you let it run overnight

- **Every loop has limits** on its output and a defined safe state when it gives up.
- **Every shot records** the lock states, the parameters used and the code version.
- **The figure of merit has been checked** by hand: it cannot be raised by a background, a saturated camera or a mis-set threshold.
- **Reference measurements are interleaved** so drift can be told apart from improvement.
- **Someone gets a message** if it stops, and knows what to do.
- **In the morning, validate** before you believe: re-measure the winner against the old settings.

> Next in the build: coupling atoms to a cavity. The control systems page covers the hardware these loops run on, and the lab handbook the hands-on alignment that comes before any of this.

10.12 References and further reading

Feedback, feedforward and stabilisation

B. Merkel et al., "Magnetic field stabilization system for atomic physics experiments," arXiv:1808.03310. Fluxgate feedback plus mains feedforward; the numbers in Section 03.

Y.-X. Chao et al., "Pound-Drever-Hall feedforward: laser phase noise suppression beyond feedback," *Optica* 11, 945 (2024), arXiv:2309.09759.

Y.-X. Chao et al., "Robust high-frequency laser phase noise suppression by adaptive Pound-Drever-Hall feedforward," arXiv:2407.19642 PREPRINT

MRC Systems beam position stabiliser, Laser Focus World product note (2010): piezo steering mirrors with four-quadrant diodes.

Auto-alignment

Y. Geng et al., "A compact and open-source microcontroller-based rapid auto-alignment system," arXiv:2405.02306 PREPRINT

R. S. Mathew et al., "The Raspberry Pi auto-aligner: machine learning for automated alignment of laser beams," arXiv:2010.15246.

L. Richtmann et al., "Model-free reinforcement learning with noisy actions for automated experimental control in optics," arXiv:2405.15421 PREPRINT

Calibration and optimisation

J. Kelly et al., "Physical qubit calibration on a directed acyclic graph," arXiv:1803.03226 PREPRINT

P. B. Wigley et al., "Fast machine-learning online optimization of ultra-cold-atom experiments," *Sci. Rep.* 6, 25890 (2016); M-LOOP documentation.

A. D. Tranter et al., "Multiparameter optimisation of a magneto-optical trap using deep learning," *Nat. Commun.* 9, 4360 (2018).

A. J. Barker et al., "Applying machine learning optimization methods to the production of a quantum gas," *Mach. Learn.: Sci. Technol.* 1, 015007 (2020).

Z. Vendeiro et al., "Machine-learning-accelerated Bose-Einstein condensation," *Phys. Rev. Research* 4, 043216 (2022).

M. Reinschmidt et al., "Reinforcement learning in cold atom experiments," *Nat. Commun.* 15, 8532 (2024).

L. Phuttitarn et al., "Enhanced measurement of neutral atom qubits with machine learning," *Phys. Rev. Applied* 22, 024011 (2024).

Y. Lee and E. Chae, "Machine learning-enhanced optical tweezers for defect-free rearrangement," arXiv:2401.04893 PREPRINT

AI agents

S. Cao et al., "Agents for self-driving laboratories applied to quantum computing," arXiv:2412.07978; journal version "Automating quantum computing laboratory experiments with an agent-based AI framework," *Patterns* 6 (2025).

R. Sha et al., "LLM-based multi-agent copilot for quantum sensor," arXiv:2508.05421 PREPRINT

T. Isogawa et al., "Agentic AI for scientific reasoning in autonomous quantum sensing experiments," arXiv:2607.25145 PREPRINT

CHAPTER 10 SUMMARY

Automation stabilizes daily performance and frees researchers to focus on physics rather than knob-turning. To add a strong photonic interface to our trapped atoms, we next introduce optical cavities.

Problems for Chapter 10

1. USE THE TOOL In the power-servo simulator with an integral (I) servo, find the bandwidth at which the loop starts to oscillate for loop delays of 0.5, 1 and 2 μs.
2. BY HAND Compare your answers in problem 1 with $f \approx 1/4\tau$, where τ is the total delay including the AOM response. Explain the phase-margin argument behind it.
3. DESIGN Follow the nine-step workflow in Section 06 to write an optimisation plan for one stage of an experiment you know: the figure of merit, its noise, the parameters and bounds, and the optimiser you would choose.

Use the tool: on the interactive page. By hand: from the formulas in this chapter. Design: a decision a lab would make. Worked solutions are available to instructors on request (see Part V).

CHAPTER 11 · BUILD STEP 11

11

Cavity Quantum Electrodynamics

💎 Cavity QED Coupling Lab · Calculator

Turn mirror spacing, finesse, mode waist, and atomic linewidth into the language of cavity QED: g_0, κ, γ, cooperativity, Purcell enhancement, β-factor, and the strong-coupling boundary. The goal is not just a number — it is intuition for what the cavity is doing to one atom and one photon.

> *Placing a trapped atom inside a high-finesse optical resonator channels spontaneous emission into a single spatial mode, unlocking fast nondestructive readout and high-efficiency photonic interconnects.*

IN THIS CHAPTER

Prerequisites: Ch. 1, Ch. 8

- Compute the single-atom coupling rate g_0, cavity decay rate κ, atomic decay rate γ, cooperativity $C = g_0^2/(\kappa\gamma)$, and Purcell factor.
- Distinguish the strong-coupling, bad-cavity (Purcell), and weak-coupling regimes and their signatures in cavity transmission spectra.
- Evaluate single-photon extraction efficiency $\beta = C/(1+C)$ and outcoupling trade-offs for quantum networking.

WEBSITE COMPANION Live interactive page ↗ · Notebook: Open in Colab · download cavity-qed.ipynb · Week 14

Cavity QED workflow

11.1 Turn geometry into a physical regime.

A cavity is not just a shiny box around an atom. Use the page in this order: compute the three rates, compare them, then decide whether the system is a spectrometer, a Purcell interface, or a strong-coupling device.

01
Enter geometry
Length, waist, finesse, linewidth, and transition wavelength set g_0, κ, and γ.

02
Read the rates
Compare coherent exchange to cavity leakage and atomic free-space decay.

03
Tune finesse
See how raising finesse lowers κ until the strong-coupling threshold is crossed.

04
Classify regime
Strong coupling, Purcell enhancement, or weak coupling each imply different experiments.

05
Compare references
Check your numbers against real cavity QED systems from the literature.

06
Check intuition
Common mistakes to catch before trusting a design.

Calculator output at the page's default settings.

g_0

How strongly one photon talks to one atom

Large dipole moment, small mode volume, and placing the atom at an antinode all increase the coherent exchange rate.

$g_0 \propto \sqrt{(\Gamma / V)}$

κ

How fast the cavity forgets the photon

High finesse and longer photon lifetime lower κ. This makes coherent oscillation visible, but also narrows the bandwidth.

$\kappa/2\pi = \mathrm{FSR} / 2\mathcal{F}$

γ

How fast the atom leaks to free space

γ is the unavoidable competitor. Cooperativity asks: does the cavity channel beat all the other optical modes?

$C = g_0^2 / \kappa\gamma$

11.2 Cavity & Atom Parameters

> Start from a geometry preset if you want intuition fast. These are not sacred designs — they are plausible regimes: macroscopic reference cavity, short strong-coupling cavity, fiber Fabry-Pérot, and a deliberately bad cavity used for Purcell thinking.

11.3 Results

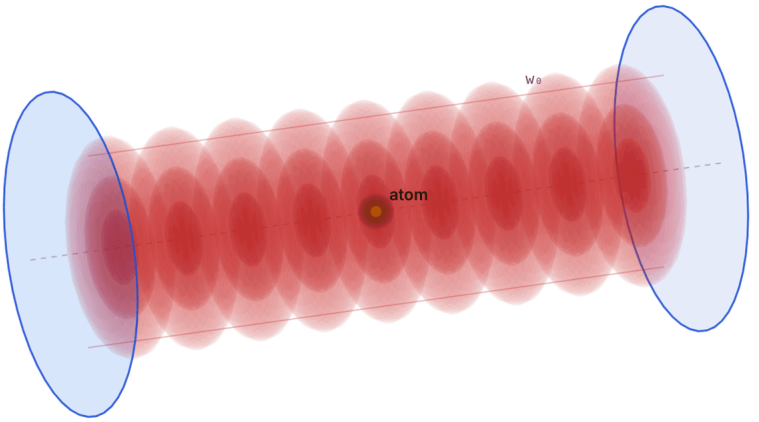
⚡ Strong coupling — $g_0/2\pi$ = 9.26 MHz > $\kappa/2\pi$ = 3.00 MHz. Vacuum Rabi splitting $2g_0/2\pi$ = 18.5 MHz is observable. C = 9.4.

9.26 **$g_0/2\pi$ MHz**	3.00 **$\kappa/2\pi$ MHz**	3.03 **$\gamma/2\pi$ MHz**
9.42 **Cooperativity *C* dimensionless**	9.42 **Purcell F_p** = C	4.13E+6 **Mode vol. V ×(λ/2)³**
299.792 **FSR GHz**	5.996 **Cav. linewidth MHz (FWHM)**	5.37E-2 **Sat. photons n_0 photons**
0.212 **Critical atoms N_0 atoms**	90.4% **β-factor exact 44.9% (§04)**	16194 **$\mathcal{F}$ for κ=g_0 threshold**

Calculator output at the page's default settings.

$$g_0 = \xi\sqrt{\frac{3\pi\Gamma c^3}{2\omega^2 V}} \qquad V = \frac{\pi w_0^2 L}{4} \qquad C = \frac{g_0^2}{\kappa\gamma} \qquad \kappa = \frac{\mathrm{FSR}}{2\mathcal{F}} \qquad \gamma = \frac{\Gamma}{2}$$

Convention used here: rates are displayed as angular rates divided by 2π. The calculator uses a TEM_{00} standing-wave mode volume and treats ξ as an amplitude-level overlap factor. For real experiments, multiply by Clebsch-Gordan coefficients, polarization projections, antinode position, and transverse mode overlap.



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Where the atom sits matters

The g_0 above assumes the atom is on the cavity axis at an antinode. Move it and the coupling falls as $g = g_0 \cdot \cos(kx) \cdot e^{-r^2/w_0^2}$, so the cooperativity falls as the square.

Along the axis **antinode**

Off axis, r / w_0 **0.00**

9.256$g/2\pi$ at the atom (MHz)

1.000g / g_0

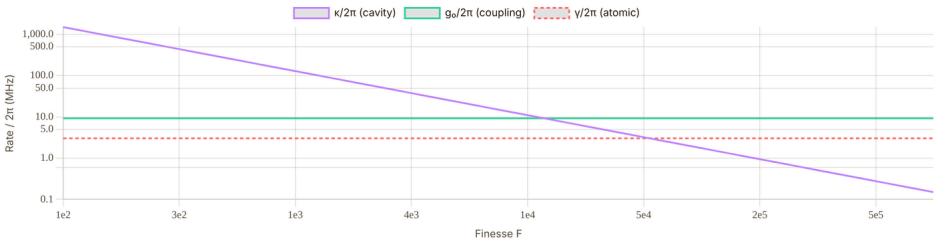
9.424cooperativity C at the atom

90.4%$\beta = C/(1+C)$

Schematic along the axis: 9 antinodes drawn; your cavity has $2nL/\lambda \approx 1{,}282$. Radial scale is w_0. Near-waist formula ($z_R \gg L$).

11.4 Coupling Rates vs Finesse

> Increasing finesse reduces κ (cavity leaks more slowly) while g_0 and γ remain constant. Strong coupling is reached when the **purple κ line drops below g_0**. The vertical dashed line marks the minimum finesse for strong coupling with the current cavity geometry.



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

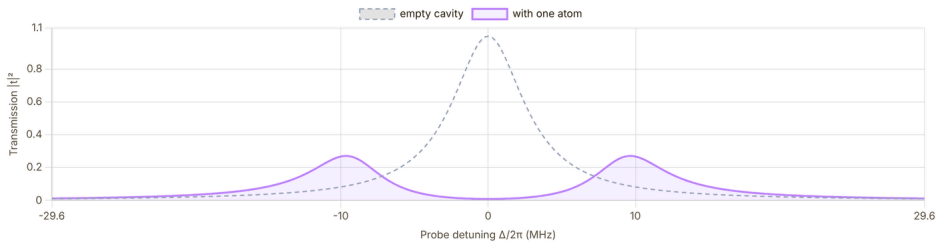
Interactive

What you would see: the cavity transmission spectrum

Scan a weak probe laser across the cavity resonance (atom and cavity on resonance) and record the transmitted light. The empty cavity gives one Lorentzian of half-width κ. Add one atom and the response becomes

$$t(\Delta) = \frac{\kappa}{\kappa - i\Delta + \dfrac{g^2}{\gamma - i\Delta}}, \qquad |t(0)|^2 = \frac{1}{(1+C)^2}, \quad C = \frac{g^2}{\kappa\gamma}$$

(weak drive; κ and γ are the half-widths used above). With **C > 1** the atom punches a hole in the middle of the line; with **g > κ, γ** the line splits into two dressed atom–cavity states at $\Delta \approx \pm g$: the *vacuum Rabi splitting*, 2g, set by a single atom and zero photons. Move the atom off the antinode (smaller g/g_0) to watch the splitting close.



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Design knobs: what should you actually change?

Make w_0 smaller

Most powerful lever for g_0 because `g₀ ∝ 1/w₀` and `C ∝ 1/w₀²`. The price is mirror curvature, clipping loss, alignment sensitivity, and atom-surface constraints.

Make L shorter

Boosts g_0 through smaller mode volume and raises FSR. But short cavities are mechanically harder and can limit optical access.

Increase finesse

Reduces κ without changing g_0. This helps strong coupling, but very high finesse narrows bandwidth and makes lock acquisition unforgiving.

Choose the transition

Broad lines give larger g_0 but larger γ too. Narrow lines can have beautiful spectra but often require extremely small mode volume to matter.

A tiny history of making empty space non-empty

1946

Purcell realizes that spontaneous emission is not immutable; the environment changes the rate.

1963

Jaynes–Cummings gives the clean one-atom, one-mode model that still powers the intuition.

1992

Rempe–Thompson–Kimble observe optical normal-mode splitting from atom-cavity coupling.

2005

Birnbaum et al. demonstrate photon blockade: one atom can make one photon block the next.

today

Fiber cavities, nanophotonics, and tweezer arrays aim to turn cavity QED into scalable photon interfaces.

11.5 Watch One Photon Leave

Interactive

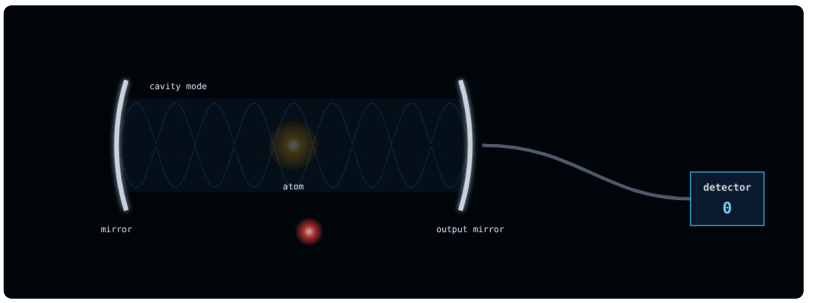
Excite the atom with a π-pulse and let it decay. Its one quantum of energy can leave two ways: out through the cavity mirror (where a fiber can collect it) or sideways into free space. Each run below is one quantum trajectory; run enough of them and the fraction that reach the detector is your **photon collection efficiency**. In the single-excitation subspace, with κ and γ the half-widths used above,

$$\dot{c}_e = -\gamma\, c_e - i g\, c_1, \qquad \dot{c}_1 = -\kappa\, c_1 - i g\, c_e,$$

photons leave through the mirrors at rate $2\kappa|c_1|^2$ and into free space at $2\gamma|c_e|^2$. Integrating these exactly gives the probability that the photon comes out of the cavity:

$$P_{\text{cav}} = \frac{C}{1+C}\cdot\frac{\kappa}{\kappa+\gamma}, \qquad C = \frac{g^2}{\kappa\gamma}.$$

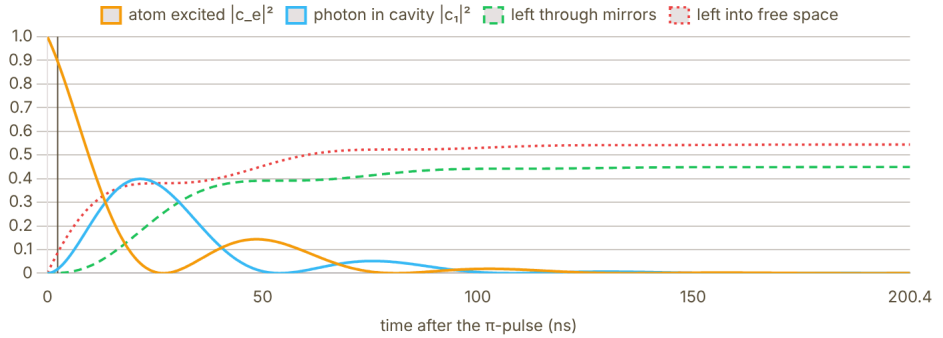
The β-factor C/(1+C) in the results above is its bad-cavity limit (κ ≫ γ). In strong coupling with κ comparable to γ the photon sloshes back and forth and can still leak out sideways, so β overestimates what you collect.



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

1 s on screen = 19 ns

Where the excitation goes (ensemble)



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

INTO THE CAVITY MODE, P_{CAV}	COLLECTED AT THE DETECTOR	BAD-CAVITY B = C/(1+C)
44.9% exact, C/(1+C)·κ/(κ+γ)	**40.4% P_cav × 90% output mirror**	**90.4% overestimates by 45.5 points here**

COOPERATIVITY C	MEAN EMISSION TIME
9.42 strong coupling	**26 ns free-space lifetime 26 ns**

Calculator output at the page's default settings.

11.6 Coupling Regimes

When $g_0 > \kappa$ and $g_0 > \gamma$, the atom–cavity system enters **strong coupling**. The cavity-QED Hamiltonian $H = \hbar g_0(a^\dagger\sigma^- + a\sigma^\dagger)$ drives coherent energy exchange between atom and photon faster than either can decay. The normal modes of the coupled system (the *dressed states* or *polaritons*) are split by $2g_0$ — directly observable as a doublet in the cavity transmission spectrum (vacuum Rabi splitting).

Achieving strong coupling requires simultaneously large g_0 (tight mode, short cavity) and small κ (high finesse mirrors). The product gives the condition $\mathcal{F} > \mathrm{FSR}/(2g_0)$ when both rates are expressed in Hz, i.e., a minimum finesse scales with the free spectral range and inversely with coupling.

$$\text{Strong coupling:}\quad g_0 > \kappa,\ \gamma \quad\Leftrightarrow\quad \mathcal{F} > \frac{\pi c}{2Lg_0} \quad \text{for angular } g_0$$

In the **bad-cavity limit** ($\kappa \gg g_0, \kappa \gg \gamma$) the photon leaks out before a coherent Rabi oscillation completes. However, when $C > 1$ the cavity still dramatically enhances the spontaneous emission rate into the cavity mode. With $C = g_0^2/(\kappa\gamma)$ (κ, γ half-widths), the cavity adds a decay channel $2g_0^2/\kappa = C\Gamma$, so the total emission rate becomes $\Gamma_{tot} = \Gamma(1 + C)$ and the fraction of photons emitted into the cavity mode is $\beta = C/(1 + C)$. This is the **Purcell effect**, with Purcell factor $F_P = C$. (Papers that use Kimble's $C_1 = g_0^2/(2\kappa\gamma)$ write the same results as $2C_1$ and $2C_1/(1 + 2C_1)$.) Applications: efficient single-photon sources, deterministic atom–photon entanglement.

When $C < 1$, the cavity offers negligible enhancement of atom–photon interactions. The atom emits primarily into free-space modes. This regime is typical for initial cavity characterization experiments and for cavities used primarily for filtering or as reference resonators (e.g., PDH locking). Increasing finesse, reducing mode volume (shorter L or smaller w_0), or choosing a higher-Γ transition all increase C.

11.7 Reference Systems

GROUP	ATOM / Λ	L (MM)	FINESSE	w_0 (MM)	$G_0/2\Pi$ (MHZ)	$K/2\Pi$ (MHZ)	$\Gamma/2\Pi$ (MHZ)	C	REGIME
Kimble, Caltech (2005)	Cs, 852 nm	0.042	460 000	23	34	4.1	2.6	108	Strong
Hood Lab, Purdue (ongoing)	Cs-133, 852 nm	—	—	—	—	—	2.6	—	—

$(g_0, \kappa, \gamma)/2\pi = (34, 4.1, 2.6)$ MHz for the Caltech Cs cavity (Kimble group); $\gamma = \Gamma/2$ from Steck data sheets. $C = g_0^2/(\kappa\gamma)$ is computed from those rates (papers using $C_1 = g_0^2/2\kappa\gamma$ quote half). The calculator above reproduces $g_0 \approx 34$ MHz and $\kappa \approx 4$ MHz from the listed L, $\mathcal{F}$ and w_0.

11.8 Common Mistakes

Three ways cavity numbers quietly go wrong

Mixing Γ and γ

Many papers quote $\Gamma/2\pi$ as the full natural linewidth. This calculator uses $\gamma = \Gamma/2$ as the dipole decay rate in `C = g₀²/(κγ)`.

Forgetting 2π

If g, κ, and γ are all consistently angular rates or all consistently divided by 2π, ratios are fine. Trouble starts when MHz and rad/s are mixed in one equation.

Treating ξ as a fudge factor

ξ should mean something physical: standing-wave position, polarization, Clebsch-Gordan coefficient, and transverse overlap. If ξ = 0.2, ask which imperfection caused it.

🏔️

Hood Lab context: Our Li–Cs tweezer experiments don't use a cavity, but cavity QED is a natural next step — coupling a single Cs atom in a tweezer to a fiber Fabry–Pérot cavity would allow single-photon-level readout without destruction. The key challenge is that the Cs D2 line ($\gamma/2\pi = 2.6$ MHz) needs a small waist and high finesse: at $L = 0.5$ mm and $w_0 = 25$ µm, reaching $C = 1$ already takes $\mathcal{F} \approx 4\,400$, and $\kappa = g_0$ takes $\mathcal{F} \approx 1.6\times10^4$ (strict strong coupling also needs $g_0 > \gamma$, which holds here: 9.4 vs 2.6 MHz).

📚 Key References

📄 Birnbaum et al. (2005). Photon blockade in an optical cavity with one trapped atom. *Nature* 436, 87–90. doi:10.1038/nature03804

📄 Thompson et al. (1992). Observation of normal-mode splitting for an atom in an optical cavity. *Phys. Rev. Lett.* 68, 1132. doi:10.1103/PhysRevLett.68.1132

📄 Kimble, H.J. (1998). Strong interactions of single atoms and photons in cavity QED. *Phys. Scr.* T76, 127. doi:10.1238/Physica.Topical.076a00127

📄 Haroche & Raimond (2006). *Exploring the Quantum: Atoms, Cavities, and Photons.* Oxford University Press. (Nobel Prize 2012)

📄 Reiserer & Rempe (2015). Cavity-based quantum networks with single atoms and optical photons. *Rev. Mod. Phys.* 87, 1379. doi:10.1103/RevModPhys.87.1379

📄 Purcell (1946). Spontaneous emission probabilities at radio frequencies. *Phys. Rev.* 69, 681. doi:10.1103/PhysRev.69.674.2 historical

📄 Jaynes & Cummings (1963). Comparison of quantum and semiclassical radiation theories. *Proc. IEEE* 51, 89. doi:10.1109/PROC.1963.1664 theory

CHAPTER 11 SUMMARY

High-cooperativity cavities bridge localized atomic memories and flying photons. Rounding out Part I, we consolidate the practical alignment, RF, fiber, and troubleshooting habits that keep an AMO lab running smoothly.

Problems for Chapter 11

1. USE THE TOOL Enter a 10 mm cavity with finesse 10,000 and a 20 μm waist for ^{87}Rb at 780 nm. Record g, κ, γ, the cooperativity C and the regime.
2. BY HAND Check g_0 with the page's formula $g_0 = \sqrt{(3\pi\Gamma c^3/2\omega^2 V)}$, $V = \pi w_0^2 L/4$, and verify $C = g_0^2/\kappa\gamma$ from your values, taking care with the half-width conventions for κ and γ.
3. DESIGN In the photon race, compare the measured fraction of photons through the output mirror with $C/(1 + C)\cdot\kappa/(\kappa + \gamma)$. What finesse and waist would you need for 90% collection?

Use the tool: on the interactive page. By hand: from the formulas in this chapter. Design: a decision a lab would make. Worked solutions are available to instructors on request (see Part V).

CHAPTER 12 · BUILD STEP 12

12

AMO Laboratory Operations Handbook

🔬 AMO Lab Operations Handbook · Guide

A field guide for the operations that make an AMO experiment work: align light, prepare atoms, drive RF and microwaves, lock lasers, debug failures, and turn lab intuition into numbers.

> *Beyond individual subsystems, day-to-day experimental progress depends on disciplined optical alignment, fiber coupling, AOM driving, magnetic-field control, and systematic debugging.*

IN THIS CHAPTER Prerequisites: Ch. 2, Ch. 4, Ch. 5

- Execute repeatable beam-walking, single-mode fiber coupling, and AOM double-pass alignment workflows.
- Manage RF drive chains, magnetic field coils, eddy currents, and bucking coils cleanly.
- Apply structured fault-isolation checklists when MOTs disappear, locks drop, or fluorescence SNR degrades.

WEBSITE COMPANION Live interactive page ↗ · Week 4

Use it as an operations handbook

12.1 Start from the lab problem, then open the matching drawer.

AMO technique pages are easiest to use when organized by what you are trying to do at the table, not by component names. The six chapters below follow those actions, in this order.

01 **Align light**

Fiber coupling, AOM geometry, polarimetry, beam delivery, and the diagnostics that reveal bad modes.

02 **Prepare atoms**

Optical pumping, stretched states, repump logic, and the selection rules that decide where population goes.

03 **Drive transitions**

RF antennas, microwaves, impedance thinking, safety, and what power actually reaches the atoms.

04 **Optimize & simulate**

Bayesian optimization and QuTiP models for navigating a high-dimensional, expensive-to-measure parameter space.

05 **Lock & route lasers**

Build the four ECDL wavelengths a Li-Cs tweezer experiment needs, from Littrow lasers to fiber delivery.

06 **Debug & fix**

Tacit lab knowledge from years of debugging — organized by AOM/EOM, fiber, locking, pumping, tweezers, imaging, RF, and field control.

6 chapters on this page. Open one from a card, or from the chapter bar below, which stays pinned while you read.

1 / 6

1 · Optical Fiber Coupling

Mode-matching visualiser — beam waist animated into fiber



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Almost every laser beam in an AMO experiment is delivered to the optical table via optical fiber. This is not merely a convenience: fibers **spatially filter** the beam, **mechanically decouple** the laser table from the experiment table, and allow laser sources to be swapped without realigning downstream optics. The practical consequence is that *fiber coupling efficiency determines beam quality at the atom.*

Three fiber types appear in a typical lab:

- **Single-mode (SM)**, supports only the fundamental HE_{11} mode; any higher-order input content is rejected. Essential wherever spatial coherence matters.
- **Polarisation-maintaining (PM)**, SM fiber with stress rods that introduce birefringence; preserves the input polarisation when aligned to the principal axis. Every beam requiring a defined polarisation at the atom must travel on PM fiber.
- **Multimode (MM)**, large core, easy to couple; used for wavemeter pick-offs, diagnostics, and detection paths where polarisation is unimportant.

A lens of focal length **f** converts an input Gaussian beam of $1/e^2$ radius **w_in** into a focused waist:

$$w_0 \approx \frac{\lambda f}{\pi w_{\text{in}}}$$

The coupling condition is $\mathbf{w_0 = w_f}$, where **w_f = MFD/2** is the fiber mode radius. Solving for the required focal length:

$$f = \frac{\pi w_{\text{in}} w_f}{\lambda}$$

Trade-off: shorter f → tighter focus, higher peak efficiency, stricter alignment tolerances. Longer f → more robust over temperature drifts.

Target: 50–80% coupling into SM/PM fiber; >80% into MM. PM fiber extinction ratio target: >20 dB ($P_\parallel/P_\perp > 100$), meaning <1% power in the wrong polarisation axis.

Mode-matching calculator

Formula: f = π w_in w_f / λ · Verify with VNA or throughput measurement

1. Use a PBS to prepare clean linear polarisation upstream.
2. Rotate a HWP after the PBS to align the polarisation to the fiber's keyed (slow/fast) axis.
3. Optimise coupling for maximum throughput.
4. Measure extinction ratio (ER) at the output through an analyser PBS.
5. Iterate HWP angle to maximise ER.

Target ER > 20 dB. Residual polarisation impurity (leakage into the orthogonal axis) is the dominant limit on optical-pumping fidelity and on the σ^+–σ^- purity of gray-molasses imaging beams.

THORLABS, SM & PM FIBERS

SM-780HP · PM780-HP

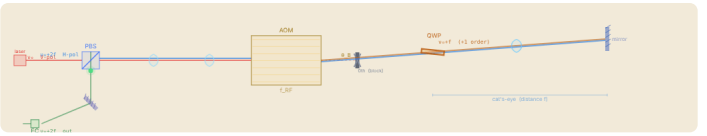
THORLABS, COUPLING LENSES

C220TMD · C230TMD · C240TMD

Calculator output at the page's default settings.

2 · Double-Pass Acousto-Optic Modulators (AOMs)

Cat's-eye double-pass AOM (Fig. 3). Laser → PBS → lenses → AOM: 0th order blocked, +1st order (ν_0+f) deflected at Bragg angle θ_B → QWP → cat's-eye lens ↔ mirror → retroreflects → AOM again (+f) → lens → PBS. After the double QWP pass the polarisation is rotated by 90°, so the PBS sends it out of its other port to the FC. Output: ν_0+2f, direction unchanged, frequency-agile.



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

> An AOM diffracts light off a moving acoustic grating, shifting the optical frequency by f_RF while deflecting the beam by the Bragg angle. The problem with a single pass: the Bragg angle depends on f_RF, so *tuning the frequency also steers the beam.* This coupling is unacceptable for precision experiments.
>
> The solution is **double-passing**: retroreflect the first-order beam back through the same crystal. The frequency shift doubles to **2 f_RF** and the angular deviations cancel exactly, the output beam direction is *independent* of f_RF. Every tunable beam in the experiment uses a double-pass AOM.

The most compact double-pass geometry:

1. AOM crystal → first-order beam deflected by Bragg angle
2. Lens (f = one focal length from crystal) collimates and focuses
3. QWP + mirror at one focal length from lens, retroreflects beam
4. QWP traversed twice → **polarisation rotated 90°**
5. Second pass through AOM → frequency shifted by another f_RF
6. PBS separates the returning 2f_RF beam (now orthogonal polarisation) from the input; an iris blocks the zero order

The cat's-eye geometry makes retroreflection **insensitive to mirror tilt**, which is why it is preferred over simpler flat-mirror configurations.

Typical performance

~90%

Single-pass eff.

~80%

Double-pass eff.

40–50 dB

Extinction (AOM off)

PROPERTY	LONGITUDINAL	SHEAR-WAVE
Acoustic mode	Atoms ∥ propagation	Atoms ⊥ propagation
Sound velocity	~4–6 km/s (TeO_2: 4.2)	~0.6–3 km/s (TeO_2 slow shear: 0.62)
Deflection per MHz (θ = λf/v)	Smaller	Larger (~7× in TeO_2; AOD use)
Resolvable spots (Δf·D/v)	Few	Many (slow v → long access time)
Diffraction eff.	Lower at peak	Higher at peak

In the lab: Cooling/imaging beams → longitudinal AOM (fixed frequency shift, fast switching). Tweezer array generation → slow-shear TeO_2 AOD, because the slow acoustic wave gives a large deflection range and many resolvable spots to position hundreds of tweezers (at the cost of slower switching).

3 · Polarimetry, Stokes Parameters & the Poincaré Sphere

Polarisation purity is one of the most critical and least forgiving parameters in AMO experiments. A **1% admixture** of the wrong circular component in an optical-pumping beam can reduce state-preparation fidelity from >99% to <95%. A few percent of linear polarisation in a σ^+–σ^- gray-molasses beam opens decoherence pathways. Measuring polarisation quantitatively, not just qualitatively, is essential.

Stokes parameters

The complete polarisation state is described by four intensity measurements:

PARAMETER	MEANING
S_0 = I_total	Total power
S_1 = I_H − I_V	Horizontal vs. vertical linear
S_2 = I_{+45} − I_{-45}	Diagonal linear
S_3 = I_RCP − I_LCP	Right vs. left circular

$$P = \frac{\sqrt{S_1^2 + S_2^2 + S_3^2}}{S_0}$$

P = degree of polarisation (1 = fully polarised)

Rotating-QWP measurement method

Place a QWP on a rotation stage before a PBS analyser. As QWP angle θ is swept, transmitted intensity follows:

$$I(\theta) = \tfrac{1}{2}(A + B\sin 2\theta + C\cos 4\theta + D\sin 4\theta)$$

$$S_0 = A - C \;\mid\; S_1 = 2C \;\mid\; S_2 = 2D \;\mid\; S_3 = B$$

Use 50–100 evenly spaced angles for reliable results.

QWP retardance calibration

$$\delta = 2\arccos\left(\sqrt{\frac{I_{\min}}{I_{\max}}}\right)$$

Measure between two PBS ports with linear input. Any deviation from 90° is corrected in the fitting procedure. **Calibrate every QWP** in optical-pumping and imaging paths.

Poincaré Sphere: RCP (σ^+) sits at the north pole ($S_3 = +1$), LCP (σ^-) at the south pole ($S_3 = -1$). All linear states lie on the equator. Elliptical states are on the surface. A pure state has $|S| = S_0$; partially polarised light has $|S| < S_0$.

4 · Interferometers: Fabry–Pérot, Michelson & Mach–Zehnder

These three come up in almost every AMO lab conversation, often in the same sentence: a **Fabry–Pérot** cavity locks the laser, a **Michelson** wavemeter tells you what wavelength it's actually at, and a **Mach–Zehnder** sequence is how the atoms themselves become the interferometer. The three differ mainly in *what happens to the light between the splitting and the recombining*.

Fabry–Pérot (cavity)

Two parallel, partially-reflective mirrors facing each other. Light doesn't split into separate arms at all — it bounces back and forth between the same two mirrors, interfering with itself on every round trip. Resonances occur when the round-trip path is an integer number of wavelengths.

$$\mathrm{FSR} = \frac{c}{2L}, \qquad \mathcal{F} = \frac{\pi\sqrt{R}}{1-R}, \qquad \delta\nu_{\mathrm{FWHM}} = \frac{\mathrm{FSR}}{\mathcal{F}}$$

L = mirror spacing, R = mirror reflectivity, $\mathcal{F}$ = finesse. Higher finesse → narrower, sharper resonances → a more sensitive frequency discriminator.

Where it shows up: PDH reference cavities for laser locking, spectrum analyzers for checking single-mode operation, mode-cleaning/filter cavities, cavity QED (this page's Cavity QED Coupling Lab is entirely built on this geometry).

Michelson

A beamsplitter sends light down two arms, each ending in a mirror that **retroreflects** the light straight back the way it came. Both beams recombine at the same beamsplitter they started from. The signature feature: each arm is traversed twice, there and back.

$$I \;\propto\; 1 + \cos\!\left(\frac{4\pi\,\Delta L}{\lambda}\right)$$

ΔL = arm-length difference. The factor of 4π (not 2π) comes from light making the round trip in each arm.

Where it shows up: wavemeters (fringe-counting against a stabilized reference laser to read out an unknown wavelength), coherence-length measurements, beam-pointing and optical-path diagnostics.

Mach–Zehnder

A first beamsplitter sends light down two **separate, one-way** arms — nothing is retroreflected. A second beamsplitter recombines them, producing two distinct output ports whose relative intensity encodes the phase difference accumulated between the arms. Because each arm is a genuine single-pass path, you can put something *inside* it — a vapor cell, a phase modulator, or a freely-falling atom.

$$P_{\mathrm{out}} = \tfrac{1}{2}\big[1 + \cos(\Delta\phi)\big], \qquad \Delta\phi = \phi_1 - \phi_2$$

Where it shows up: this is the geometry behind **atom interferometry** — a π/2–π–π/2 sequence of Raman or Bragg pulses splits, redirects, and recombines an atomic wavepacket in momentum space exactly like a beamsplitter–mirror–beamsplitter sequence does for light. That's how cold-atom gravimeters (measuring g), gyroscopes (rotation via the Sagnac effect),

and equivalence-principle tests work. It's also the standard geometry for fiber/EOM phase-to-amplitude readout and for differential probe-vs-reference measurements.

INTERFEROMETER	WHAT HAPPENS BETWEEN BS AND RECOMBINATION	TYPICAL AMO ROLE
Fabry–Pérot	Light resonates between the same two mirrors, no separate arms	Laser locking, spectrum analysis, cavity QED
Michelson	Two arms, each retroreflected back to the starting beamsplitter	Wavemeters, coherence length, alignment
Mach–Zehnder	Two separate one-way arms, recombined at a second beamsplitter	Atom interferometry, EOM readout, differential probing

Quick day-to-day distinguisher: if the two arms retroreflect back through the same beamsplitter, it's a Michelson. If light in each arm only ever travels one-way before hitting a *second* beamsplitter, it's a Mach–Zehnder. When a cold-atom physicist says "the interferometer" with no qualifier, they almost always mean the Mach–Zehnder atom-optics sequence.

5 · Laser Frequency Stabilisation

🔐

Laser frequency stabilisation (SAS, beat-note, and PDH cavity locking) is covered in full on the dedicated **Laser Locking** page, including interactive error-signal simulators, a lock-hierarchy diagram, and cavity design calculators.

Laser Locking page →

Optical pumping drives an atom into a specific Zeeman sublevel via repeated photon absorption and spontaneous emission cycles. In a tweezer experiment, its purposes are twofold:

1. **State preparation**, put the atom in a single, well-defined $|F, m_F\rangle$ state before any coherent manipulation.
2. **Detection preparation**, define the initial condition for state-selective fluorescence imaging.

Imperfect optical pumping is a direct source of systematic error in lifetime measurements, qubit state detection, and gate fidelity.

Cesium: pumping to $|F=4, m_F=+4\rangle$

Scheme: σ^+-polarised light driving **$F=4 \rightarrow F'=4$** has no allowed absorption for an atom already in $m_F=+4$ (which would require $\Delta m_F=+1$, but no $m'_F=+5$ exists in the excited state). That state is *dark*; all other m_F sublevels are continuously depopulated until the atom accumulates in $m_F=+4$.

Repumper: A beam resonant with $F=3 \rightarrow F'=4$ returns any population that decayed to $F=3$.

Why the magnetic-field direction matters: The quantisation axis is defined by the **local B-field**, not the laser k-vector. If the bias field is misaligned with the beam, the σ^+ in the lab frame decomposes into σ^+, π, and σ^- in the atom's frame. Apply a bias field of ~6 G **along the optical axis** of the pumping beam.

Diagnostic: the depumping-ratio test

1. Drive $F=4 \rightarrow F'=4$ *without* the repumper, atoms in $F=4$ scatter photons and heat out of the trap.
2. Under aligned B-field: atoms in the dark state $|4,+4\rangle$ survive for ~1 ms (off-resonant scattering only).
3. Under deliberately misaligned B-field (~45°): σ^+ acquires σ^- component → atoms depumped in ~10 µs.
4. **Target depumping ratio > 100** (aligned/misaligned survival times).

Adjustment knobs: laser frequency (exact line centre), QWP orientation (polarisation purity), bias field direction. Together these achieve >99% pumping fidelity.

Quick diagnosis checklist:
✅ ER of pumping beam fiber > 20 dB
✅ Bias field coil along pump beam axis
✅ Pump laser on $F=4 \rightarrow F'=4$ (not $F=3 \rightarrow F'=4$)
✅ Repumper on $F=3 \rightarrow F'=4$
✅ Depumping ratio > 100
✅ Pump pulse duration > 5 × ($1/\Gamma_scatter$)

BIAS COIL DRIVERS

Wavelength Electronics · B&K Precision

KEY OPTICS

Thorlabs PBS cubes · Thorlabs QWPs · Thorlabs HWPs

Calculator output at the page's default settings.

Lithium: D2-line optical pumping

Challenge: The ^{6}Li D2 excited-state hyperfine splitting is only ~5 MHz, comparable to Γ. Different F′ components overlap, making it impossible to address a single excited hyperfine level cleanly. In practice, a combination of $F=3/2 \rightarrow F'=3/2$ and $F=1/2 \rightarrow F'=3/2$ light is used.

Procedure

1. Coarse beam alignment to MOT on the diagonal camera.
2. Fine alignment on a single trapped atom.
3. Verify resonance by scanning laser frequency over the atom-loss signal.
4. Characterise fidelity with the depumping-ratio test (same as Cs).

Advantage: Ground-state hyperfine splitting of 228 MHz ≫ D2 linewidth → state-selective probing is clean and free of cross-talk.

Practical note: In a dual-species (Li+Cs) setup, the Li pumping beams co-propagate with the Cs beams; beam geometry and PM-fiber infrastructure are shared, so only minor frequency and alignment adjustments are needed once Cs pumping is optimised.

REFERENCES

Gehm (2003), Properties of ^{6}Li (NCSU)
Steck, Cs D-line data

Coherent control of atomic hyperfine states requires delivering oscillating magnetic fields at precise frequencies, 228 MHz for ^{6}Li ground-state splitting, ~76 MHz in the Paschen-Back (high-field) regime, 9.2 GHz for Cs. The engineering challenge is delivering enough B-field amplitude at the atom while fitting inside an existing vacuum apparatus. Safety-critical high-current electronics (Feshbach coils) demand dedicated hardware interlocks that are independent of the computer control system.

Loop antenna for Li hyperfine transitions

A **single-turn loop** radiates primarily through its magnetic dipole field. The radiation resistance of a small loop of area A at frequency f is:

$$R_{\text{rad}} \approx 20\pi^2 \left(\frac{4\pi A}{\lambda^2}\right)^2 \ \Omega$$

For any loop that fits near a vacuum cell, R_rad ≪ 50 Ω. Efficient power delivery from a 50 Ω source therefore requires an **impedance-matching network**.

Three approaches tested (with VNA)

1. **Capacitive loading**, series or parallel capacitor shifts resonance. *Result: parallel ~47 pF proved most effective for 76 MHz, compact geometry.*
2. **Transmission-line stub matching**, moves impedance on Smith chart. Works in principle, but spurious resonances can complicate things.
3. **Discrete LC networks**, more design freedom but more components.

Achieved: reflection minimum ~10 dB, sufficient B-field at the atom with ~100 W amplifier.

Lessons learned

- Lead length (connector to loop) contributes parasitic inductance at 100 MHz, non-negligible.
- Simulate with **SimSmith** before building every iteration.
- The resonant frequency scales inversely with circumference at fixed inductance.

Loop antenna quick estimate

Approximate formulas (single turn of 1 mm diameter wire, $L = \mu_0 R[\ln(8R/a) - 2]$); stray capacitance and leads shift the match, so always verify with a VNA.

TOOLS & INSTRUMENTS

Siglent SVA1015X · SimSmith · Mini-Circuits

Feshbach coil safety interlocks

Feshbach coils produce fields of order **1000 G** by carrying large DC currents. Sudden current interruption generates inductive voltage spikes that can damage power supplies and coils.

Hardware interlock logic (essential rules)

- Monitor: coil temperature, current level, supply voltage
- On threshold exceeded: **ramp down smoothly**, do not switch off abruptly
- All interlock logic implemented in **relay hardware**, independent of computer control
- Interlock circuit must be untriggerable by software bugs

> ⚠️ **Never bypass a hardware interlock for convenience.** Inductive spikes from sudden coil switching have destroyed power supplies and cracked coil epoxy in labs worldwide. The hardware interlock is not optional. A software interlock fails if the computer crashes. A relay cannot be bypassed by software.

HIGH-CURRENT SUPPLIES

Kepco BOP bipolar · iSeg · AMETEK Programmable Power

> Optimising an ultracold-atom experiment means navigating a high-dimensional continuous parameter space where each measurement takes seconds to minutes. A MOT has ≥5 knobs; a tweezer experiment adds trap depth, cooling polarisation, B-field direction. Naïve grid search over 5 variables × 10 points = 10^5 trials, months of run time.

How it works

Gaussian-Process Regression (GPR): a non-parametric Bayesian model that maintains a probabilistic map of the response surface (e.g. atom survival vs. laser detuning + power). At each step it returns a predicted value **and an uncertainty**; unexplored regions have high uncertainty.

Acquisition function: selects the next measurement point by trading off exploitation (sample near current optimum) and exploration (reduce uncertainty in poorly sampled regions).

ACQUISITION FUNCTION	WHEN TO USE
Expected Improvement (EI)	Default; works well near optimum
Upper Confidence Bound (UCB)	When signal is absent, need broad search
Probability of Improvement (PI)	Conservative; avoids risk

Workflow

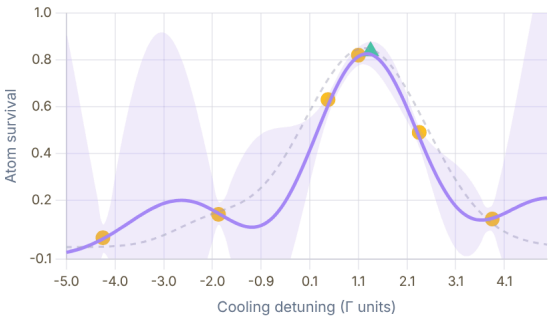
1. Collect 5–20 initial points from Latin hypercube or random design
2. Train GPR model; inspect posterior mean + uncertainty
3. Select next point by maximising EI/UCB
4. Run experiment, append data, retrain, repeat

> **Convergence:** 30–100 evaluations for 4–8 dimensional spaces (vs. thousands for grid search).

SOFTWARE

scikit-learn GPR · BayesianOptimization (Python) · Meta Ax platform · Frazier (2018)

GP posterior demo, atom survival vs. cooling detuning



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Grey dotted = true function · Purple band = GP ±2σ · Yellow = observations · Green dashed = next query point

The master-equation simulations in the Li and Cs chapters are implemented in **QAtomTweezer**, a purpose-built Python library on top of QuTiP. The key idea is to treat a single atom with *full hyperfine structure* coupled to a **1D harmonic oscillator** representing one motional axis of the tweezer.

Physical model

$$\hat{H} = \hat{H}_{\text{internal}} \otimes \hat{H}_{\text{HO}}$$

- **H_internal:** all hyperfine and Zeeman sublevels (typically 12–24 states for D1 of an alkali)
- **H_HO:** harmonic oscillator, truncated at N_HO = 10–20 Fock states
- Total dimension: d_int × N_HO (e.g. 12 × 12 = 144 for Li D1)

Key feature: the full matrix exponential is used for the recoil operator R̂ = e^(iη(â+â†)) (not the Lamb-Dicke expansion), so the code is valid *beyond* the strict Lamb-Dicke regime.

validation cross-checks

1. Fock distribution P(n), fit to Boltzmann to extract T_eff
2. Excited-state fraction p_e, compare to measured photon rate
3. Temperature minimum, verify location in 2D parameter scans

Code structure

```
# Entry points:
QAtomTweezer.py
QAtomTweezer_SingleLevel.py

# Main callable:
SteadyStateTweezer(
    x,      # [δ1, δ2, Ω1, Ω2, φ1, φ2]
    wh,     # trap freq in units of Γ
    Nh,     # HO truncation
    atom,   # AtomSettings object
    eta,    # Lamb-Dicke parameter
    pol,    # polarisation config
)
# Returns: ⟨n⟩, P(n), p_e
```

Performance

1–5 s

per steady-state eval (144×144)

10–30 min

2D scan 50×50 (joblib)

SOFTWARE STACK

QuTiP documentation · joblib (parallelism) · Johansson et al. (2013)

AMO experiments for Li–Cs tweezer work require four distinct laser wavelengths, each with its own technical demands. All are produced by **external-cavity diode lasers (ECDLs)** in Littrow configuration, except the 1064 nm tweezer which uses a commercial Nd:YAG/fiber laser.

A bare semiconductor diode lases on multiple longitudinal modes. Two tuning mechanisms in a bare diode:

- Injection current: $\partial\nu/\partial I$ ~ 1–3 GHz/mA (fast but noisy)
- Temperature: $\partial\nu/\partial T$ ~ −20 to −40 GHz/K (slow, hysteretic)

Neither alone provides the narrow linewidth (<100 kHz) or mode-hop-free tuning range (>1 GHz) required.

The Littrow ECDL solution

- Holographic grating at Littrow angle feeds first-order diffraction back into the diode
- Selects one longitudinal mode of the extended cavity
- Grating angle (PZT-tuned) + injection current → mode-hop-free tuning over 1–2 GHz
- AR coating on front facet suppresses internal Fabry-Pérot resonances

~100 kHz

Free-running linewidth

1–2 GHz

Mode-hop-free scan

~30 MHz/V

PZT tuning

Common failure modes

- **Mode hops:** usually from temperature drift; cure with better temperature control
- **Reduced output:** check AR coating; diodes degrade with age and excessive current
- **Multiple modes:** grating feedback misaligned; realign while monitoring on wavemeter
- **Linewidth broadening:** current noise from noisy driver; use low-noise current source

Alignment tips

1. Set temperature for approximate target wavelength
2. Coarsely align grating to first-order feedback with IR card
3. Monitor wavelength on wavemeter; find single-mode region
4. Maximise mode-hop-free range by co-scanning PZT + current (feed-forward)

ECDL VENDORS

Toptica DL Pro · Sacher LION · MOGLabs
CEL · RIO Photonics

Transitions

D1 (670.992 nm) and D2 (670.977 nm)

Role in experiment

MOT (D2), Zeeman slower (D2), Λ-GM cooling and tweezer imaging (D1)

Linewidth

`Γ/2π = 5.87 MHz` (D1 and D2 nearly identical)

Key challenge

D1 and D2 lines separated by only ~10 GHz, both needed simultaneously. D2 laser offset-locked to D1. Lightest alkali: single-recoil scale ≈3.5 μK, so each scatter cycle heats strongly unless cooling is simultaneous.

Locking strategy

D1: SAS locked to vapour cell. D2: beat-note locked to D1 (Vescent D2-135).

Laser source

Toptica or home-built ECDL at 671 nm; TA amplifier often needed for MOT power

DIODE SOURCES

Thorlabs L671P200 · Eagleyard Photonics · Laser Components

Transitions

D2 (852.347 nm); D1 at 894.6 nm also used in some labs

Role

MOT, fluorescence imaging, repumper

Linewidth

`Γ/2π = 5.23 MHz`

Key challenge

Large hyperfine splitting (9.193 GHz) means cooler and repumper must be offset by ~9.2 GHz; use an AOM chain or separate laser with beat lock.

Locking strategy

SAS locked directly to Cs D2 line in vapour cell.

DIODE SOURCES

Thorlabs L852P150 · Eagleyard Photonics · Laser Components

Transitions

$6S_{1/2} \rightarrow 5D_{5/2}$ (electric quadrupole, $\Gamma/2\pi \approx 117.6$ kHz)

Role

Narrow-line sideband cooling of Cs in tweezer; excited-state lifetime measurement

Linewidth

`Γ/2π = 117.6 kHz` (resolved sidebands at typical trap frequency)

Key challenge

Forbidden E2 transition → I_sat ≈ 2.3 W/cm² (much higher than D-lines). No SAS possible. Requires PDH lock to ULE cavity for <1 kHz linewidth. Astigmatism correction needed (prism pair before cavity).

Locking strategy

PDH locked to ULE cavity (L=77.5 mm, F≈15 000, linewidth ~100 kHz). Laser linewidth ~1 kHz.

DIODE SOURCE

Thorlabs L685P010

Transitions

Far-detuned (no resonant absorption); acts as conservative dipole trap potential

Role

Creates the optical tweezer potential; all atoms trapped in the 1064 nm focus

What matters here

Intensity noise, not frequency noise. Intensity noise at trap frequencies (kHz) causes parametric heating.

RIN target

`< −130 dBc/Hz` at trap sidebands. Needs intensity stabilisation (AOM servo on pick-off PD).

Locking strategy

Free-running (stable Nd:YAG or fiber laser); intensity servo via AOM feedback.

LASER SOURCES

Thorlabs Nd:YAG

Tacit knowledge that isn't in any textbook — collected from years of debugging in cold-atom and tweezer labs. Each section covers a specific experimental area.

⚠️

AOM double-pass: retroreflect off the AOM crystal, not the fiber coupler. The retro mirror must send the beam back through the AOM at exactly the same transverse position. If your retro mirror is after a fiber launch, you lose the double-pass frequency shift — the AOM shifts the beam twice only when both passes hit the active aperture.

Diagnostic: block the retro and check for a single-pass diffracted spot. Restore retro and verify the RF power needed for 80% efficiency drops by roughly half compared to single-pass.

⚠️

AOM efficiency drops if the beam is not focused inside the crystal. The acoustic wave has a finite Bragg acceptance angle. A collimated 3 mm beam has a much smaller angular spread than one focused to 100 µm — tight focus gives better efficiency for high-frequency AOMs (>200 MHz) because the Bragg condition is satisfied over a wider range of angles.

Rule of thumb: focus to a spot size where the Rayleigh range ≈ acoustic transit length (~5–10 mm for most crystal heights).

⚠️

EOM V_π is not constant across drive frequencies. At 9.2 GHz (Cs hyperfine), a resonant EOM has dramatically lower V_π than at 20 MHz (PDH). Driving an EOM off-resonance or above its bandwidth produces almost no modulation — always check the EOM spec sheet for the resonant frequency and 3 dB bandwidth.

⚠️

PM fiber after EOM scrambles your polarization state. An EOM requires linearly polarized input aligned to its crystal axis. If you launch into PM fiber before the EOM, check that the PM fast axis aligns to the EOM crystal axis, not just to some arbitrary horizontal. A quarter-wave rotated PM fiber coupling completely kills the EOM modulation depth.

⚠️

Coupling efficiency >80% but PER is only 17 dB? Your beam is elliptical or astigmatic. Polarization extinction ratio (PER) in a PM fiber requires the input beam to be perfectly linearly polarized and Gaussian. An astigmatic beam (common after an AOM) couples light into both PM axes — producing high efficiency but poor PER. Add a cylindrical lens pair to correct astigmatism before coupling into PM fiber.

⚠️

Power instability through PM fiber: check the patch cable stress, not the laser. A bent PM fiber or a clamp pressing on the cable applies stress birefringence, rotating the polarization axis and beating with the PM alignment. Slow power drifts often trace to a fiber cable that rests against a moving part of the table (shutter actuator, beam dump, cooling water line) rather than the laser itself.

⚠️

Mode-matching to a fiber: optimize tip/tilt first, then z-position of the lens. Counter-intuitive but true: the lens transverse position sets the input angle, and the lens z-position sets the mode waist at the fiber face. Iterate tip/tilt → z-lens → tip/tilt. Trying to optimize all three simultaneously gives slow convergence.

⚠️

Don't clean fiber connectors with compressed air — use IPA and a reel cleaner. Compressed air blows particles onto the ferrule face. Use a single-use reel cleaner (Cletop-S) for APC/PC connectors. Inspect with a fiber scope before mating — a single 1-µm particle on the 9-µm core is enough to degrade coupling efficiency from 85% to 50%.

⚠️

SAS crossover resonances are not real transitions — don't lock to them by accident. A crossover peak appears halfway between two real transitions that share a common ground state. It's the tallest feature on the SAS spectrum (coherent population transfer), so it's tempting to lock there. But the F'=1/F'=3 crossover, for example, sits 211.8 MHz below the ^{87}Rb F=2→F'=3 cycling transition, not 0 MHz.

Verification: compare your SAS spectrum peak positions against published Rb/Cs hyperfine intervals. For Rb87 D2 from F=2: F'=2 and F'=3 lie 156.9 and 423.6 MHz above F'=1 (excited splittings 156.9 and 266.7 MHz); crossovers fall midway, at 78.5, 211.8 and 290.3 MHz above F'=1.

⚠️

PDH error signal polarity depends on which sideband is higher in frequency. If you flip the LO phase by 180° or swap the mixer IF/RF ports, the error signal flips sign — and your PI controller drives the laser away from resonance instead of toward it. Always verify locking direction: give the laser a small push (tap the table) and confirm the lock corrects, not amplifies, the perturbation.

⚠️

Servo bump tells you where your loop bandwidth is — not where you want it. The characteristic bump in the in-loop noise PSD at the servo bandwidth is not an artifact to filter away — it's diagnostic. Flattening it by lowering gain reduces bandwidth and worsens noise. Instead, improve the phase margin by adding a lead compensator or reducing cable length in the feedback path.

⚠️

High-finesse cavity drift: check temperature, not the laser. A 10 cm Zerodur spacer with CTE ~2×10^{-8}/K expands by ~2 nm/K, and $\delta\nu/\nu = \alpha\cdot\delta T$: a 1 mK drift shifts the resonance by ~8 kHz at 780 nm, a 0.1 K drift by ~0.8 MHz. If your PDH-locked laser frequency drifts slowly over minutes, the cavity is thermally drifting — improve vacuum can temperature stabilization before touching laser current.

⚠️

Quantization axis must be defined before the pump beam arrives. If the bias magnetic field is not on (or is lower than the stray-field level ~100 mG), the atom has no preferred axis. The σ^+ photon pumps into different m_F levels depending on local stray field direction. Always verify the bias field is on and stable before running the OP sequence.

⚠️

Re-pumping power too high scrambles your state preparation. The repump beam ($F=1 \rightarrow F'=2$ for Rb87) should be weak enough to recycle dark atoms without broadening the resonance for the main pump. Too much repump power heats the sample by preventing atoms from dwelling in the $F=1$ dark state, and can repump atoms out of the target m_F state.

⚠️

OP efficiency check: use a short probe pulse, not absorption imaging. To verify m_F state purity, do a quick Stern-Gerlach separation with a field gradient before imaging. Absorption imaging along the same OP beam axis can misrepresent state purity because the imaging beam itself pumps atoms during exposure. A short (10 μs) probe pulse after Stern-Gerlach separation resolves individual m_F populations cleanly.

⚠️

Intensity noise heats your atoms. Frequency noise does not (for far-detuned traps). A 1064 nm tweezer far from any transition is an intensity-noise-limited trap. The trap depth $U_0 \propto I$, so fractional intensity noise $\delta I/I$ directly converts to fractional trap-depth noise, driving parametric heating at $2\omega_{trap}$. Frequency noise at 1064 nm has no first-order effect on the trap depth (it only shifts absolute energy by the tiny light-shift gradient).

Action: use a low-noise laser source at 1064 nm and servo the intensity with an AOM + photodiode + PI loop. Target RIN < −140 dBc/Hz at $2\omega_{trap}$ (~2×100 kHz = 200 kHz).

⚠️

SLM-generated tweezer arrays: check diffraction efficiency, not just spot count. An SLM can project 100 spots, but if your hologram algorithm doesn't compensate for the sinc envelope of the SLM pixel pitch, outer spots in the array are 30–40% dimmer than central ones. Run a weighted Gerchberg-Saxton (WGS) algorithm that measures spot intensities in a feedback loop and reweights the target amplitudes accordingly.

⚠️

Evaporation ramp that works in a crossed ODT fails in a single-beam tweezer. A crossed ODT has equal radial and axial trapping in two dimensions. A single-beam tweezer has radial frequency ω_r >> axial ω_z — atoms escape axially first as you ramp down power. Design the evaporation ramp by monitoring axial size with in-situ imaging, not just total atom number.

⚠️

Camera background subtraction: take dark frames at the same temperature as your science frames. EMCCD and sCMOS dark current is strongly temperature-dependent (~2× per 6–7°C). If your camera warms up between your background run and your data run, the dark floor changes, creating a systematic offset in atom number. Let the camera thermalize for at least 30 minutes and take background frames immediately before (or interleaved with) science shots.

⚠️

Single-atom fluorescence: saturation parameter determines photon budget, not just pulse duration. At high saturation $s \gg 1$, the scattering rate plateaus at $\Gamma/2$. Going from $s=10$ to $s=100$ only gives you ~9% more photons ($R \propto s/(1+s)$) and ~9% more recoil heating, while the extra power adds stray light and light shifts. The optimal single-atom imaging uses $s \approx 1$–5 with a short pulse, not maximal power.

Rule of thumb: for Rb87 D2 with NA=0.5 objective, $s=2$ on resonance scatters ~1.3×10^4 photons in 1 ms (~2.5×10^4 recoil kicks, so cooling during imaging matters), of which ~400 are detected — sufficient for single-atom detection with EMCCD EM gain ~300.

⚠️

Charge smear on EMCCD: don't readout during the fluorescence pulse. EMCCD vertical shift (frame transfer) takes ~1 ms. If your shutter timing overlaps with the shift, every row smears fluorescence into adjacent rows, broadening PSFs and reducing per-pixel SNR. Use a mechanical shutter or an AOM to gate the light, not just the camera trigger.

⚠️

Calibrate your RF π-pulse time at the same bias field you use for science. The Zeeman shift moves the transition by ~0.7 MHz/G for Rb87 $|F{=}1,m_F{=}{-1}\rangle \rightarrow |F{=}2,m_F{=}0\rangle$. If you calibrate the π-pulse at 1 G and run science at 3 G, you're off by 1.4 MHz — and instead of a π-pulse you're doing a Ramsey dephasing experiment. Always recalibrate Ω_R and the resonance frequency when you change the bias field.

⚠️

Ground loop in the RF circuit kills coherence time. An RF ground loop between your signal generator, amplifier, and antenna creates a circulating current that modulates the local magnetic field at 50/60 Hz and its harmonics. Use a single-point ground, float the amplifier chassis on rubber feet, and check T_2^* with a Ramsey sequence: if you see T_2^* oscillate at 50 Hz, you have a ground loop.

⚠️

Avoid antenna resonances: design for broadband, not peak efficiency. A resonant antenna at 6.835 GHz (Rb87 clock transition) is efficient but narrow-band. Any drift in the antenna resonance (thermal expansion of coax, SMA connector loosening) shifts the coupling and changes your effective Rabi frequency. A broadband impedance-matched antenna ($s11 < -15$ dB over ±100 MHz) is more robust for day-to-day reproducibility.

⚠️

Eddy currents from fast field switching limit how quickly you can apply a Feshbach field. A large metallic vacuum chamber acts as a shorted transformer: when you ramp the Feshbach coil current rapidly (>1 A/ms), induced eddy currents oppose the change and the actual field at the atoms lags behind the coil current by the chamber's L/R time constant (typically 1–10 ms). Always measure the actual field delay with RF spectroscopy before assuming the field follows your current waveform.

📖

New: design the coil to not couple to the chamber, rather than fighting the eddy currents after the fact. Kell et al. (2021) get a Feshbach field switch of 36 G in 3 μs, with a chamber present or removed making *no measurable difference* to the switching speed, by making the mutual inductance to the chamber negligible in the first place: two concentric, oppositely-wound "bucking" coils confine the field close to the coil itself instead of spreading out to the chamber walls, and a 3D-printed plastic coil mount adds zero eddy-current path of its own. A power MOSFET with an RC snubber does the actual current switching (30 ns switch time on a 6 μH, 0.1 Ω coil). If your chamber's L/R delay is already a problem, a compact bucking-coil retrofit is usually cheaper than an active compensation loop.

Kell, Link, Breyer, Hoffmann, Köhl & Gao, "A compact and fast magnetic coil for the interaction manipulation of quantum gases with Feshbach resonances," *Rev. Sci. Instrum.* 92, 093202 (2021). arXiv:2103.05273 →

⚠️

Stray-field zeroing: do it with atoms at the operating temperature, not room temperature. Mu-metal shielding has temperature-dependent permeability. The field zero you find at room temperature can shift by 20–50 mG when the chamber bakes out or the surrounding equipment heats up during a long experimental run. Zero the field with atoms by minimizing the Ramsey fringe decay rate at the exact operating conditions.

⚠️

MOT coil polarity: anti-Helmholtz for trapping, Helmholtz for bias field — never mix up. Anti-Helmholtz coils produce a gradient (zero at center, $\nabla B \neq 0$) — correct for MOT. Helmholtz produces a uniform field. If you accidentally run anti-Helmholtz during optical pumping, the atoms experience a spatially varying quantization axis and your OP efficiency drops to ~50% at best. Double-check the relay/H-bridge state before every sequence that switches between gradient and bias.

🔬

Hood Lab note: Most of the pitfalls above were discovered the hard way — by running a beautiful experiment for three days and then realizing the state preparation was wrong. The most common root cause is *not verifying assumptions*: always ask "how do I know the field is on?", "how do I know the fiber is PM-aligned?", "how do I know the AOM is double-passing correctly?" Build a short diagnostic sequence for each subsystem and run it at the start of every session.

12.2 References & Further Reading

Primary Source & Textbooks

Metcalf & van der Straten, *Laser Cooling and Trapping* (1999)

Foot, *Atomic Physics* (OUP, 2005)

Saleh & Teich, *Fundamentals of Photonics* (Wiley)

Pozar, *Microwave Engineering* (Wiley, 4th ed.)

Key Papers

📄 Black (2001), PDH locking tutorial (Am. J. Phys.) 📄 Johansson et al. (2013), QuTiP 2 📄 Kaufman & Ni (2021), Tweezer arrays review (Nat. Phys.) 📄 Frazier (2018), Bayesian optimisation tutorial 📄 Kell et al. (2021), fast Feshbach coil design (Rev. Sci. Instrum.)

Vendor App Notes & Tools

🔗 Thorlabs fiber coupling guide 🔗 AA Opto AOM selection guide 🔗 Vescent laser locking app note 🔗 SimSmith, Smith chart tool 🔗 QuTiP documentation 🔗 scikit-learn GP regression

CHAPTER 12 SUMMARY

With the full physical apparatus built, stabilized, and operational, Part I is complete. In Part II, we put the lab to work: cooling, imaging, and measuring clouds and single atoms.

Problems for Chapter 12

1. USE THE TOOL Use the fibre-coupling calculator for 780 nm light, a 1 mm input beam and a 5 µm mode-field diameter. What focal length does it suggest?
2. BY HAND Derive the same focal length from $f = \pi w_{in} w_f / \lambda$ with $w_f = MFD/2$, and estimate how much coupling you lose if you only have a lens 30% longer.
3. DESIGN Open the debugging drawer and pick one pitfall you have met (or could meet) in a real lab. Write the five-step check you would follow, in the order you would follow it.

Use the tool: on the interactive page. By hand: from the formulas in this chapter. Design: a decision a lab would make. Worked solutions are available to instructors on request (see Part V).

PART II

Trap, Image and Cool

Running a first experiment: cool atoms with light, image the cloud and measure its temperature, see one atom, measure its temperature, cool it to the motional ground state, run a full cycle, and keep the everyday numbers at hand.

13	Laser Cooling Forces & Sub-Doppler Mechanisms	Laser Cooling Simulator
14	Absorption & Dispersive Cloud Imaging	Absorption Imaging Lab
15	Time-of-Flight Thermometry	MOT Temperature
16	Single-Atom Fluorescence Imaging	Single-atom Imaging
17	Release-and-Recapture Thermometry	Single-atom Temperature
18	Ground-State Cooling of Trapped Atoms	Single-atom Cooling
19	End-to-End Tweezer Experimental Cycle	Virtual Tweezer Lab
20	Everyday AMO Reference Calculations	Quick Lab Console

CHAPTER 13 · EXPERIMENT STEP 01

13

Laser Cooling Forces & Sub-Doppler Mechanisms

❄️ Laser Cooling Simulator · Simulator

From room temperature to quantum ground state. Explore Doppler cooling, sub-Doppler gray molasses, and resolved-sideband cooling: the three stages that bring atoms to the quantum limit.

> *Part II shifts from building hardware to running experiments. We start by dissecting the optical forces that cool atomic clouds from room temperature down to microkelvin temperatures.*

IN THIS CHAPTER Prerequisites: Ch. 1, Ch. 6

- Compute velocity-dependent radiation pressure forces, momentum diffusion, and the Doppler temperature limit $T_D = \hbar\Gamma/(2k_B)$.
- Understand how multi-level Zeeman structure enables Sisyphus (polarization-gradient) cooling well below the Doppler limit.
- Compare scattering rates, capture velocities, and equilibrium temperatures across alkali species.

WEBSITE COMPANION Live interactive page ↗ · Notebook: Open in Colab · download cooling-simulator.ipynb · Week 5

Cooling landscape

13.1 Separate free-space cooling from trapped-atom cooling.

This page starts with velocity-space Doppler cooling, then shows why sub-Doppler and sideband methods are needed as experiments move toward optical tweezers and Lamb-Dicke physics.

01
Doppler damping
Use detuning and saturation to see a thermal velocity distribution narrow. · Open now ↓

02 **Sub-Doppler mechanisms**
Connect gray molasses and Sisyphus cooling to dark states, light shifts, and recoil scales.

03
Sideband cooling
Move from velocity damping to motional-state cooling in a harmonic trap.

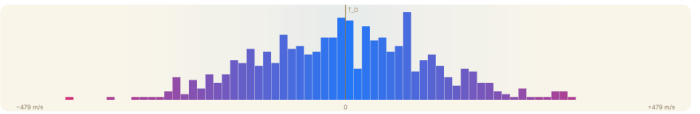
04
Choose a method
Compare linewidth, recoil, trap frequency, and required laser complexity.

Calculator output at the page's default settings.

6 chapters on this page. Open one from a card, or from the chapter bar below, which stays pinned while you read.

Live Doppler Cooling Simulator

Drag the sliders — watch the velocity distribution narrow as atoms cool

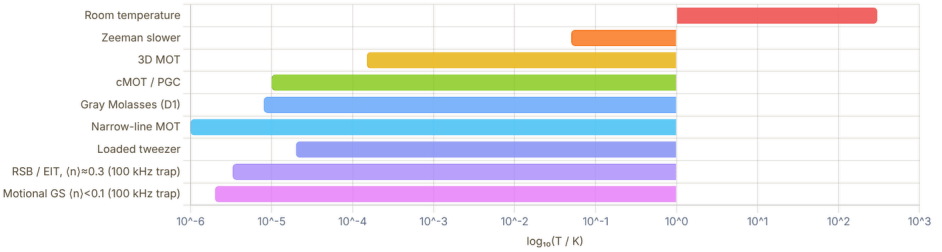


Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

T ≈ 183 K | = Doppler limit ← hot | velocity | hot →

Temperature Ladder

From room temperature to motional ground state, 9 orders of magnitude



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Why Cold Atoms for Quantum Computing?

Gate fidelity for a two-qubit Rydberg entangling gate scales as

$$\mathcal{F} \approx 1 - \alpha\langle n\rangle - \beta/\Omega_R$$

where ⟨n⟩ is the mean vibrational quantum number and Ω_R is the Rabi frequency. Each cooling stage suppresses thermal motion, reducing state-preparation errors from >10% to <0.5%. Resolved-sideband cooling to ⟨n⟩ < 0.1 is now routine in tweezer arrays.

1 / 6

Radiation Pressure Force

A two-level atom moving with velocity v along a laser beam of detuning δ = ω_L − ω_0 experiences a Doppler-shifted resonance at ω_0 − kv. The scattering rate

$$\Gamma_{sc} = \frac{\Gamma}{2} \cdot \frac{s_0}{1 + s_0 + (2\delta'/\Gamma)^2}$$

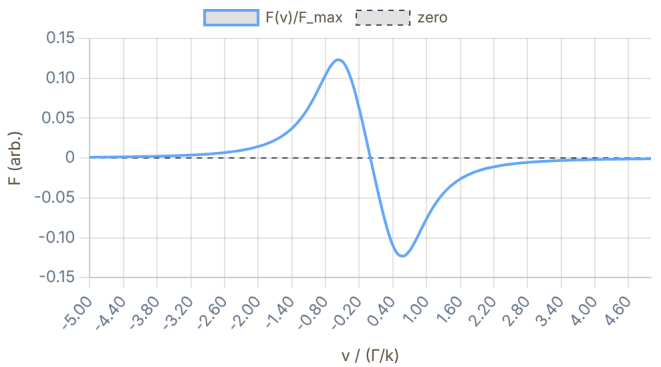
where δ' = δ + kv is the effective detuning and $s_0 = I/I_{sat}$. With counter-propagating beams the net force becomes velocity-dependent:

$$F(v) \approx -\alpha v, \qquad \alpha = 8\hbar k^2 \frac{s_0\,|\delta|/\Gamma}{(1 + s_0 + (2\delta/\Gamma)^2)^2}$$

This viscous force (α > 0 for red detuning δ < 0) cools atoms. The minimum Doppler temperature is set by the balance between velocity damping and momentum diffusion:

$$T_D = \frac{\hbar\Gamma}{2k_B}$$

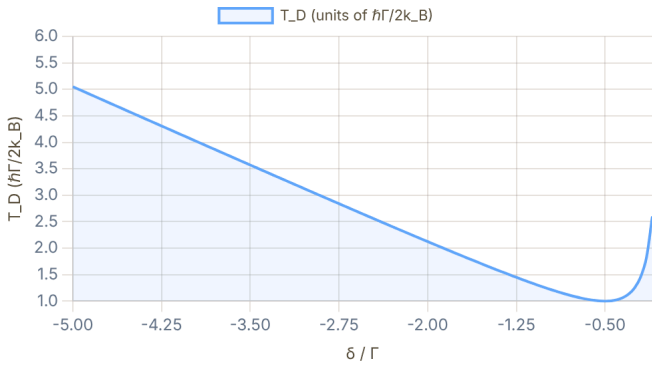
FORCE CURVE F(V), INTERACTIVE



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

α (damping) = 0.318 F_max·k/Γ

T_DOPPLER VS DETUNING



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Minimum $T_D = \hbar\Gamma/(2k_B)$ at $\delta = -\Gamma/2$

Doppler Limit in Practice

For Cs D2 line ($\Gamma/2\pi = 5.234$ MHz): $T_D = 125$ μK. For ^{6}Li D2 ($\Gamma/2\pi = 5.87$ MHz): $T_D = 141$ μK. Sub-Doppler effects can break this limit, reaching the single-photon recoil scale $E_r/k_B = \hbar^2k^2/(2mk_B)$.

A 3D molasses (6 beams ±x, ±y, ±z) provides isotropic cooling. Adding a magnetic field gradient creates a magneto-optical trap (MOT) that both cools and traps atoms.

Doppler Cooling Parameters by Species

^{6}LI (671 NM)

MOT → cMOT → GM

$\Gamma/2\pi = 5.87$ MHz. $T_D = 141$ μK. Unresolved excited hyperfine structure, GM essential for sub-Doppler cooling. Single recoil $E_r/k_B = 3.5$ μK; one scatter cycle is roughly 7.1 μK without cooling.

^{87}RB (780 NM)

MOT → PGC → evaporation

$\Gamma/2\pi = 6.065$ MHz. $T_D = 146$ μK. Well-resolved F levels. Sub-Doppler polarisation-gradient cooling (PGC) to ~4 μK.

^{133}CS (852 NM)

MOT → cMOT → GM → RSB

$\Gamma/2\pi = 5.234$ MHz. $T_D = 125$ μK. Strong Sisyphus cooling. Tweezer arrays with RSB cooling to $\langle n \rangle < 0.1$.

^{23}NA (589 NM)

Zeeman slower → MOT → PGC

$\Gamma/2\pi = 9.795$ MHz. $T_D = 235$ μK. Yellow D2 line. First atom laser (MIT 1997). High evaporation efficiency.

^{88}Sr (689 nm)

Blue MOT → Red MOT → magic tweezer

Narrow-line intercombination: $\Gamma/2\pi$ = 7.5 kHz. T_D = 180 nK. Direct sub-μK MOT. Magic wavelength at 813 nm for optical lattice clock.

^{174}Yb (556 nm)

Blue MOT → Green MOT → lattice

Narrow $^1S_0 \rightarrow {}^3P_1$: $\Gamma/2\pi$ = 182 kHz. T_D = 4.4 μK. Nuclear spin I=0. Ideal for SU(N) fermions (^{171}Yb) and optical clocks.

Sisyphus Cooling (Polarisation-Gradient)

In a standing wave with spatially varying polarisation (lin⊥lin), the light-shifted sublevel energies vary with position on the scale of $\lambda/4$. An atom in |m_F = +F⟩ moving in the +x direction climbs a potential hill, losing kinetic energy. At the top it optically pumps to |m_F = −F⟩ (lowest potential at that location), and the cycle repeats, Sisyphus loses kinetic energy each half-cycle.

$$k_B T_{\text{Sisyphus}} \sim U_0 \propto \frac{I}{|\delta|} \qquad (|\delta| \gg \Gamma)$$

The temperature is set by the depth U_0 of the light-shift potential the atom climbs, so it falls linearly with intensity and with $1/|\delta|$ (C. Cohen-Tannoudji, Nobel lecture, Rev. Mod. Phys. 70, 707 (1998)), until it reaches a floor of order tens of recoil energies. The single-photon recoil energy E_r = $\hbar^2k^2/(2m)$ sets the microscopic scale for random kicks. For Cs D2, E_r/k_B ≈ 99 nK. In practice, polarization-gradient cooling reaches μK-scale temperatures, often tens of recoil energies rather than a universal 2–4 recoil limit.

Λ-Enhanced Gray Molasses (GM) for ^{6}Li

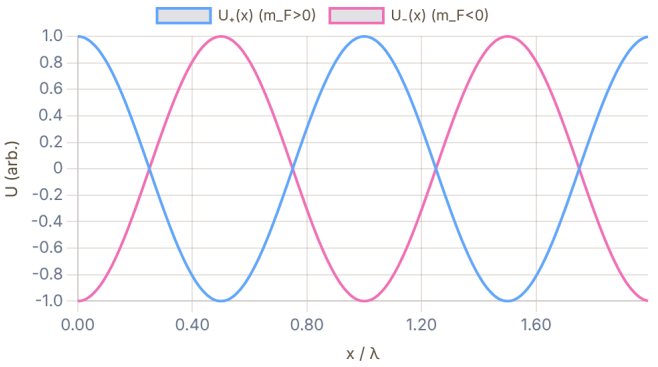
^{6}Li has an unresolved excited-state hyperfine structure (Δ_HF < Γ), making Sisyphus PGC inefficient. Instead, a **Λ-system gray molasses** (D1 line at 671 nm) exploits a dark state between |F=1/2⟩ and |F=3/2⟩:

$$|\text{dark}\rangle = \cos\theta\,|F=1/2\rangle - \sin\theta\,|F=3/2\rangle$$

Dark-state atoms are velocity-selected: slow atoms stay dark (low scatter), fast atoms couple to the bright state and get cooled. GM achieves T ≈ 40–60 μK for ^{6}Li, well below T_D = 141 μK.

Our experiment: GM on the D1 line at $\delta/\Gamma \approx +2$, reaching T ≈ 50 μK before loading into the optical tweezer array.

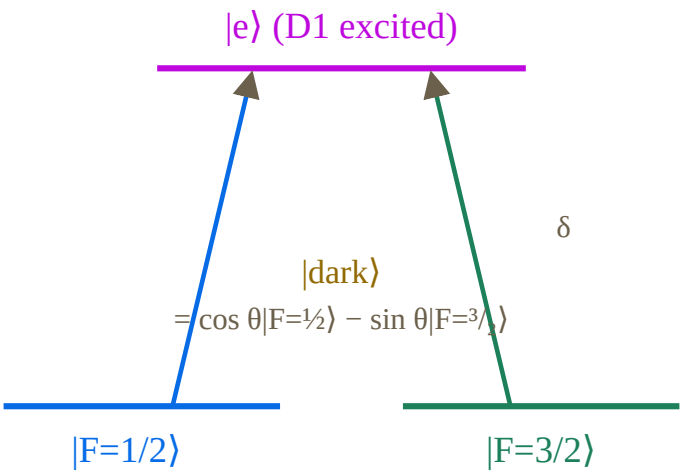
SISYPHUS POTENTIAL (LIN⊥LIN)



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Sublevel energies U±(x) and optical pumping at crests

Λ SYSTEM, GRAY MOLASSES LEVEL DIAGRAM



D1 Λ-system: dark state protected from scattering

⁶LI GRAY MOLASSES: EXPERIMENTAL RESULT

By applying a D1 gray molasses (Ω/2π ≈ 3 Γ, δ/2π ≈ +2 Γ) for 3 ms after the MOT switch-off, we achieved T ≈ 50 μK from T_MOT ≈ 200 μK, a 4× improvement in temperature that directly reduces the motional-state occupation after loading into the tweezer: ⟨n⟩_initial ∝ T/ω_trap.

Harmonic Trap & Sideband Structure

An atom in a harmonic potential with trap frequency ω_trap/2π has quantised motional states |n⟩. The motional sidebands appear at:

$$\omega_{\text{carrier}} \pm n \cdot \omega_{\text{trap}}/2\pi$$

In the Lamb-Dicke regime (η² = E_r/ħω ≪ 1), transition rates on the red sideband (RSB, ω − ω_trap) go as η²n, while the blue sideband (BSB) goes as η²(n+1).

$$\eta = k\sqrt{\frac{\hbar}{2m\omega_{\text{trap}}}} = \sqrt{\frac{E_r}{\hbar\omega_{\text{trap}}}}$$

Resolved Sideband Cooling Cycle

Drive the RSB: |g,n⟩ → |e,n−1⟩. Spontaneous emission returns mostly to |g,n−1⟩ (Lamb-Dicke suppression). Each cycle removes one vibrational quantum:

$$\langle \dot{n} \rangle = -\Gamma_{\text{cool}} \cdot n + \Gamma_{\text{heat}}$$

Steady-state: ⟨n⟩_ss = Γ_heat/Γ_cool ≈ ¼(Γ/ω)²(α + ¼) for resolved sidebands (ω ≫ Γ), independent of η because both rates scale as η² (α ≈ 1/3–1 is the emission-direction factor). In practice ⟨n⟩_ss < 0.1 is achievable.

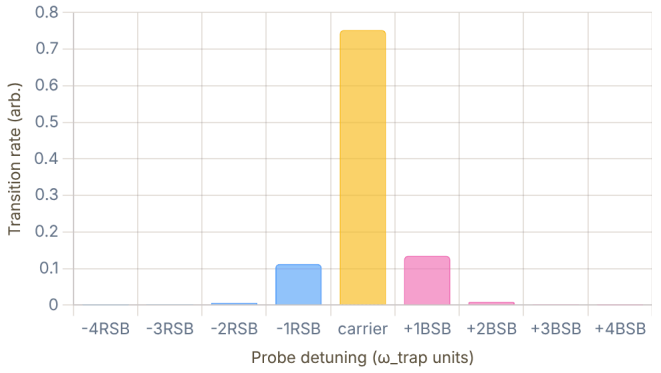
EIT-Assisted Cooling

Electromagnetically Induced Transparency (EIT) cooling uses a probe + coupling beam to engineer a narrow absorption feature at the RSB frequency. The dark state suppresses carrier scattering while the RSB absorption is enhanced, achieving faster cooling rates than conventional RSB cooling.

$$\langle n \rangle_{\text{EIT}} \approx \frac{1}{4}\left(\frac{\gamma_{\text{dark}}}{\omega_{\text{trap}}}\right)^2$$

EIT is especially useful for motional frequencies ω_trap < Γ (unresolved sideband regime), enabling ground-state cooling in shallower traps.

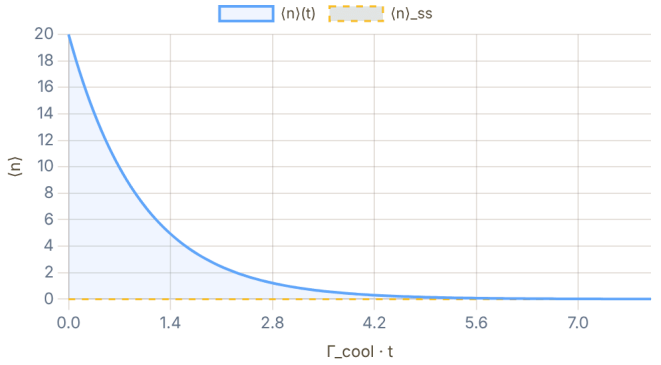
SIDEBAND SPECTRUM, INTERACTIVE



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

RSB/BSB asymmetry ≈ ⟨n⟩/(⟨n⟩+1), thermometry probe

⟨N⟩(T) COOLING EVOLUTION



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

$$\langle n\rangle(t) = \langle n\rangle_{\mathrm{ss}} + (n_0 - \langle n\rangle_{\mathrm{ss}})\, e^{-\Gamma_{\mathrm{cool}} t}$$

All Cooling Methods at a Glance

METHOD	STAGE	T_MIN	LIMIT	TYPICAL ATOMS	NOTES
Zeeman Slower	Pre-cooling	~10 mK	Doppler	Most alkalis, Sr, Yb	Slows hot atomic beam before MOT
3D MOT	Stage 1	~100–200 μK	Doppler	All	Combines cooling + 3D spatial trapping
Compressed MOT (cMOT)	Stage 1b	~50 μK	Doppler / PGC	Cs, Rb, Li	Ramp B-field + detuning before sub-Doppler
Polarisation-Gradient (PGC)	Stage 2	2–10 μK	Photon recoil	Rb, Cs, K, Na	lin⊥lin or σ+σ− standing waves, no B-field
D1 Gray Molasses (GM)	Stage 2	~40–60 μK	Photon recoil	^{6}Li, K (unresolved HF)	Λ-system dark state, works on D1 line
Narrow-line MOT (red)	Stage 2	~1 μK	Photon recoil	Sr (689 nm), Yb (556 nm)	Single-photon recoil kicks visible; direct sub-μK

Resolved Sideband (RSB)	Stage 3	⟨n⟩ < 0.1	Lamb-Dicke, η	Cs, Rb, Ca^+, Mg^+ in tweezers/lattices	Requires resolved ω_trap > Γ; closes on RSB
EIT Cooling	Stage 3	⟨n⟩ < 0.1	Dark-state linewidth	Ca^+, Mg^+, Sr, neutral atoms	Dark resonance engineered at RSB; faster rate

TEMPERATURE SCALE COMPARISON (LOG_{10} T / K)



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Doppler Limit

$T_D = \hbar\Gamma/(2k_B)$. Set by the balance between laser cooling force and random recoil kicks from spontaneous emission. Scales with linewidth.

Recoil Limit

Single recoil: $E_r/k_B = \hbar^2k^2/(2mk_B)$. One absorption plus spontaneous emission gives a two-recoil heating scale $\approx 2E_r/k_B$. For Cs D2: $E_r/k_B \approx 99$ nK and $2E_r/k_B \approx 198$ nK.

Lamb-Dicke Limit

In a trap with resolved sidebands ($\omega \gg \Gamma$): $\langle n\rangle_{ss} \to \frac{1}{4}(\Gamma/\omega)^2(\alpha + \frac{1}{4})$, so $\langle n\rangle \to 0$ as $\Gamma/\omega \to 0$. The Lamb-Dicke parameter η sets the cooling *rate* ($\propto \eta^2$) and must be small for the expansion to hold, but it cancels from the floor.

pylcp, Laser Cooling Physics

`pylcp` (Python Laser Cooling Physics) is an open-source package for simulating laser-atom interactions. It supports rate-equation, optical Bloch equation, and Hund's case (a) Hamiltonian approaches.

```
import pylcp import numpy as np # Define laser beams (3D molasses) laser_beams = pylcp.laserBeams([ {'kvec': np.array([1,0,0]), 's': 0.5, 'delta': -0.5}, {'kvec': np.array([-1,0,0]), 's': 0.5, 'delta': -0.5}, # ... ±y, ±z beams ], beam_type=pylcp.infinitePlaneWaveBeam) # Define atom (e.g., Rb F=2 → F'=3) atom = pylcp.hamiltonians.singleF( F=2, Fp=3, gamma=1, dipole_matrix_element=1 ) # Integrate force equation eqn = pylcp.rateeq(laser_beams, magField, atom) eqn.set_initial_pop(np.array([1/5]*5)) sol = eqn.equilibrium_populations(r=np.zeros(3), v=np.array([0.1,0,0]))
```

arc, Alkali Rydberg Calculator

`arc` provides atomic structure data, dipole matrix elements, and Rydberg state properties for alkali atoms. Useful for computing polarizabilities, C_6 coefficients, and Rydberg blockade radii.

```
from arc import Caesium atom = Caesium() # Get natural linewidth of D2 line gamma = atom.getTransitionRate(6,0,0.5, 6,1,1.5) print(f"Γ/2π = {gamma/2/np.pi/1e6:.3f} MHz") # Rydberg C6 coefficient for 70S1/2 C6 = atom.getC6term(70,0,0.5, 70,0,0.5) R_b = (abs(C6) / (hbar * 2*np.pi*1e6)) ** (1/6) print(f"Blockade radius = {R_b*1e6:.1f} μm")
```

QuTiP, Quantum Toolbox

For master equation simulation of the density matrix ρ under sideband cooling (Lindblad form):

$$\dot{\rho} = -i[H_{\text{eff}}, \rho] + \sum_k \left(L_k \rho L_k^\dagger - \tfrac{1}{2}\{L_k^\dagger L_k, \rho\} \right)$$

import qutip as qt N = 30 # Fock space truncation a = qt.destroy(N) H = eta * Omega * (a + a.dag()) # RSB drive c_ops = [np.sqrt(gamma) * a] # decay result = qt.mesolve(H, rho0, tlist, c_ops, e_ops=[a.dag()*a])

ADDITIONAL LIBRARIES

AtomicUnits.jl (Julia), unit conversions for atomic physics. **QuantumOptics.jl**, fast master-equation solvers in Julia. **MOLSCAT**, molecular scattering for Feshbach resonances. **COMSOL**, FEM for magnetic trap geometry and electrode design.

- [1] T. W. Hänsch and A. L. Schawlow, "Cooling of gases by laser radiation," *Opt. Commun.* 13, 68 (1975).
- [2] D. J. Wineland and H. Dehmelt, "Proposed $10^{14}\ \delta\nu < \nu$ laser fluorescence spectroscopy on Tl^+ mono-ion oscillator," *Bull. Am. Phys. Soc.* 20, 637 (1975).
- [3] S. Chu et al., "Three-dimensional viscous confinement and cooling of atoms by resonance radiation pressure," *Phys. Rev. Lett.* 55, 48 (1985). Nobel Prize 1997.
- [4] P. D. Lett et al., "Observation of atoms laser cooled below the Doppler limit," *Phys. Rev. Lett.* 61, 169 (1988).
- [5] J. Dalibard and C. Cohen-Tannoudji, "Laser cooling below the Doppler limit by polarization gradients: simple theoretical models," *J. Opt. Soc. Am. B* 6, 2023 (1989).
- [6] A. Aspect et al., "Laser cooling below single-photon recoil by velocity-selective coherent population trapping," *Phys. Rev. Lett.* 61, 826 (1988).
- [7] F. Sievers et al., "Simultaneous sub-Doppler laser cooling of fermionic ^{6}Li and ^{40}K on the D1 line," *Phys. Rev. A* 91, 023426 (2015).
- [8] I. Diedrich et al., "Laser cooling to the zero-point energy of motion," *Phys. Rev. Lett.* 62, 403 (1989). First sideband cooling.
- [9] A. M. Kaufman et al., "Cooling a single atom in an optical tweezer to its quantum ground state," *Phys. Rev. X* 2, 041014 (2012).
- [10] J. D. Thompson et al., "Coherence and Raman sideband cooling of a single atom in an optical tweezer," *Phys. Rev. Lett.* 110, 133001 (2013).
- [11] G. Morigi, J. Eschner, and C. H. Keitel, "Ground state laser cooling using electromagnetically induced transparency," *Phys. Rev. Lett.* 85, 4458 (2000).
- [12] B. J. Lester et al., "Rapid production of uniformly filled atomic arrays," *Phys. Rev. Lett.* 115, 073003 (2015).

CHAPTER 13 SUMMARY

Laser cooling compresses phase-space density by many orders of magnitude. To verify atom number and cloud geometry experimentally, our first diagnostic is absorption imaging.

Problems for Chapter 13

1. USE THE TOOL In the Doppler tab at $s_0 = 0.1$, vary δ/Γ and find where the damping (the slope of F(v) at v = 0) is largest. Show that it peaks near $|\delta| = (\Gamma/2)\sqrt{(1 + s_0)/3}$, which tends to $\Gamma/2\sqrt{3}$ at low intensity, not at the $\delta = -\Gamma/2$ that minimises the Doppler temperature, and explain why the two differ.
2. USE THE TOOL In the sideband tab with $\eta = 0.15$ and $\Gamma/\omega = 0.05$, cool from $\langle n \rangle = 20$. Then raise Γ/ω to 0.5 and repeat. Explain why the final $\langle n \rangle$ gets worse.

3. DESIGN Lithium's excited-state hyperfine structure is unresolved, so ordinary sub-Doppler cooling works poorly. Use the comparison tab to choose a cooling route for ^{6}Li and for ^{133}Cs, and explain the difference.

Use the tool: on the interactive page. By hand: from the formulas in this chapter. Design: a decision a lab would make. Worked solutions are available to instructors on request (see Part V).

CHAPTER 14 · EXPERIMENT STEP 02

14

Absorption & Dispersive Cloud Imaging

Absorption Imaging Lab · Imaging lab

Turn shadow images into optical depth, column density, atom number, and a realistic noise floor. Built for cold-atom ensembles where saturation, detuning, camera noise, and fringes decide whether the number you report is trustworthy.

> *Once a cold cloud is trapped (Chapters 6 and 13), shining a weak resonant probe beam through it casts a shadow whose optical depth directly encodes the column density and atom number.*

IN THIS CHAPTER Prerequisites: Ch. 1, Ch. 6, Ch. 13

- Extract optical depth (OD), column density, and total atom number using the Beer-Lambert law with high-intensity saturation corrections.
- Minimize fringe artifacts and shot noise through proper reference-frame subtraction and probe fluence selection.
- Identify when dense clouds require off-resonant phase-contrast or dispersive imaging instead of resonant absorption.

WEBSITE COMPANION Live interactive page ↗ · Notebook: Open in Colab · download absorption-imaging.ipynb · Week 6

Absorption-imaging workflow

14.1 Read the page like an experimental sequence

Start with the three images, then use the calculator to decide whether shot noise, read noise, saturation, or fringes dominate your atom-number estimate.

01
Build OD
Use atom, probe, reference, and dark images to compute optical depth.

02 **Convert to atoms**

Choose species, detuning, pixel size, and magnification to get column density and atom number.

03
Budget noise
Separate photon shot noise from camera read noise and identify the minimum detectable OD.

04 **Check species data**

Resonant cross section and saturation intensity for the cycling imaging transition, per species.

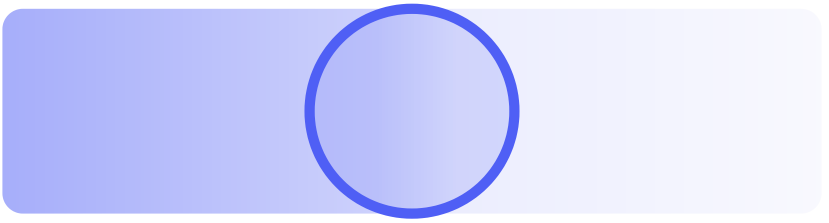
05 **Check failure modes**

Saturation, optical pumping, finite linewidth, and fringes usually dominate before algebra does.

14.2 How Absorption Imaging Works

A resonant (or near-resonant) probe beam passes through the atom cloud. Atoms scatter photons out of the beam, leaving a shadow on the camera. Three images are taken: probe with atoms, probe without atoms, and a dark frame (no light). The optical depth (OD) is:

① Probe + atoms



I_shadow

I_shadow = I_probe · e^(−OD)

② Probe only



I_probe

Reference (no atoms)

③ Dark frame

I_dark

Background subtraction

$$\mathrm{OD} = -\ln\left(\frac{I_{\mathrm{shadow}} - I_{\mathrm{dark}}}{I_{\mathrm{probe}} - I_{\mathrm{dark}}}\right) \qquad n_{2D} = \frac{\mathrm{OD}}{\sigma(\Delta)} \qquad N = \frac{\mathrm{OD} \cdot A_{\mathrm{pixel,obj}}}{\sigma} \cdot N_{\mathrm{pixels}}$$

where $\sigma(\Delta) = \sigma_0/(1 + 4\Delta^2/\Gamma^2)$ is the detuning-dependent cross section and $\sigma_0 = 3\lambda^2/(2\pi)$ for a closed two-level transition. The object-plane pixel size is $p_{\mathrm{obj}} = p_{\mathrm{cam}}/M$ (camera pixel size divided by magnification). At high probe intensities ($I \gtrsim 0.1 \cdot I_{\mathrm{sat}}$) saturation correction is needed:

$$\mathrm{OD}_{\mathrm{true}} = \mathrm{OD}_{\mathrm{meas}} + \frac{I_{\mathrm{probe}} - I_{\mathrm{shadow}}}{I_{\mathrm{sat}}}$$

3λ²/2π

Resonant cross section σ_0

OD ~ 1–3

Sweet spot for SNR

√(2/N_ph)

Shot noise on OD

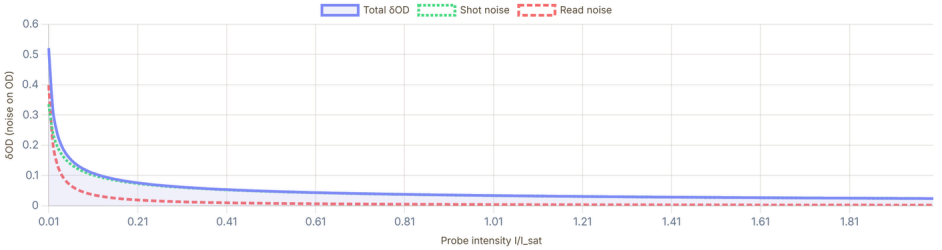
0.01–0.1

Typical fringe noise (OD)

14.3 Imaging Calculator

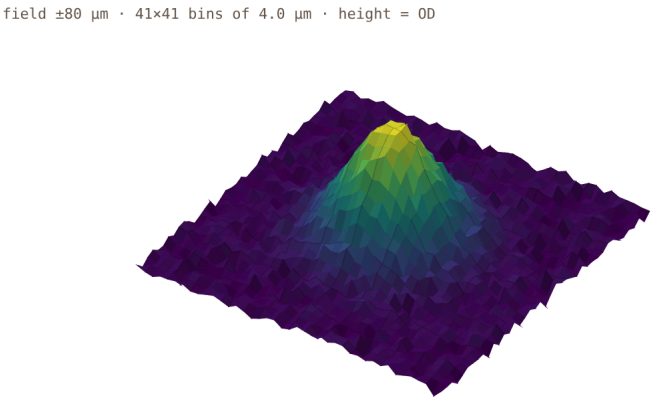
Set your imaging parameters to get the effective cross section, probe photon budget, OD noise, and minimum detectable atom number. Uses cycling transition cross sections.

OD noise (δOD) vs probe intensity — shot noise falls as $1/\sqrt{s}$, read noise as $1/s$, so noise alone always favours more light; saturation (badge above) sets the practical upper limit.



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Interactive **What the camera gives you: a simulated OD image** — a Gaussian cloud of N atoms (rms size σ) imaged with the settings above, plus each pixel's δOD noise from the calculator. Height is the measured optical depth. A faint cloud can look like noise pixel by pixel yet still give a good atom number, because summing over the cloud averages the noise down ($\sqrt{N_{pix}}$).



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Model: $OD(x,y) = \sigma \cdot n_{2D}(x,y)$ with $\sigma = \sigma_0/(1 + 4\Delta^2/\Gamma^2)$ and $n_{2D} = N\, e^{-r^2/2\sigma^2}/(2\pi\sigma^2)$, sampled on object-plane pixels (binned when the cloud spans more than 41 of them, which lowers δOD per bin by the bin width in pixels). Saturation is assumed corrected; real images also lose accuracy above OD ≈ 3, where almost no probe light gets through.

14.4 Species Reference

Resonant cross section and saturation intensity for the cycling imaging transition. OD column shows the expected optical depth for 1000 atoms uniformly distributed in a 10 µm radius disk.

SPECIES	Λ (NM)	Σ_0 ($\times 10^{-13}$ M²)	I_SAT (MW/CM²)	OD / 1000 ATOMS	IMAGING TRANSITION
Rb-87	780.24	2.907	1.669	~0.93	D2 F=2→F'=3 (cycling)
Cs-133	852.35	3.469	1.102	~1.10	D2 F=4→F'=5 (cycling)
Li-6	670.98	2.150	2.54	~0.68	D2 (approx cycling)
Li-7	670.96	2.149	2.54	~0.68	D2 (approx cycling)
Na-23	589.16	1.657	6.26	~0.53	D2 F=2→F'=3 (cycling)

K-39	766.70	2.807	1.75	~0.89	D2 F=2→F'=3 (cycling)
K-40	766.70	2.807	1.75	~0.89	D2 F=9/2→F'=11/2

14.5 Practical Considerations

The standard OD formula $-\ln(I_{\text{out}}/I_{\text{in}})$ underestimates atom number when $I/I_{\text{sat}} \gtrsim 0.1$, because saturation reduces the absorption cross section. The corrected formula adds back the "consumed photons" term:

$$\text{OD}_{\text{true}} = -\ln\left(\frac{I_{\text{out}}}{I_{\text{in}}}\right) + \frac{I_{\text{in}} - I_{\text{out}}}{I_{\text{sat}}}$$

This is exact within the two-level model. In practice, use $I/I_{\text{sat}} \sim 0.1$–$0.5$ to balance photon count (SNR) against saturation distortion. At $I = I_{\text{sat}}$ uncorrected OD can be badly biased unless the saturation term is included; at $I = 0.1\, I_{\text{sat}}$ the correction is usually small but still worth tracking. If you can't apply the correction (e.g., unknown I_{sat} for mixed states), use $I \ll I_{\text{sat}}$.

For high-field imaging (Paschen-Back regime, B > 100 G), the cycling transition breaks down and the effective cross section must be computed from the actual magnetic sublevel populations. Imaging at the stretched states $|F, m_F = \pm F\rangle$ restores a near-cycling cross section.

At large probe photon counts (N_ph > 1000/pixel), shot noise $\delta\text{OD} = \sqrt{2/N_{\text{ph}}}$ becomes negligible — but OD noise floors around 0.01–0.1 remain due to *imaging fringes*: interference patterns from dust, optical surfaces, or air currents that shift slightly between the probe and reference frames.

PCA / eigenface removal: take a library of reference images (no atoms), decompose into principal components, and subtract the best fit from each probe image. This suppresses common fringe modes, reducing effective noise to ~0.002–0.01 OD in well-optimized setups. Implementation: `ref_lib @ pinv(ref_lib) @ probe` gives the optimal projection.

Practical tips: take probe and reference within the same camera frame (jitter-limited by vibration, not drift), enclose beam path, use antireflection-coated optics, keep the probe pulse < 50 µs to freeze air fluctuations, and take the reference frame immediately (within 1 ms) after the probe frame.

For **single atoms in tweezers**, absorption imaging is typically not used — even spread over a diffraction-limited ~1 µm² spot, one atom's resonant OD is at most $\sigma_0/A \approx 0.3$, and it absorbs only a handful of photons before it is heated out — far too few against probe shot noise. Instead, *fluorescence imaging* (collecting photons scattered by the atom during a detection pulse) is the standard, giving ~50–200 detected photons per atom and a detection fidelity > 99%.

Absorption imaging is most powerful for **ensemble experiments**: MOTs (10^6–10^9 atoms), BECs (10^3–10^6 atoms), and Fermi gases, where the column density OD is 0.1–5 across a ~10–100 µm cloud. For atom numbers below ~500 in a MOT-scale trap, fluorescence wins even for ensembles.

Everything above destroys the sample: resonant photons that make the shadow also scatter the atoms out of the trap, and the calculator's saturation correction only pushes the OD ceiling higher, it doesn't remove it. **Phase-contrast (dispersive) imaging** sidesteps both limits by probing far off resonance and detecting the atoms' effect on the light's *phase* rather than its amplitude, so almost no photons scatter and the cloud survives the shot.

$$\phi(x, y) \approx -\frac{\Gamma}{4\Delta}\,\text{OD}(x, y) \qquad [|\Delta| \gg \Gamma,\ \text{OD}]$$

A phase plate (or, in modern setups, a spatial light modulator) at the imaging system's Fourier plane converts this phase shift into an intensity signal a camera can record, the same trick used in Zernike phase-contrast microscopy. Because ϕ scales as $1/\Delta$ while photon scattering (and saturation) fall off much faster with detuning, there is a wide detuning window where the phase signal is still easily measurable but the sample is essentially unperturbed.

Andrews et al. (1996) introduced this to cold-atom physics to solve a problem absorption imaging cannot: watching the *same* Bose-Einstein condensate evolve over time. A single absorption shot destroys the cloud, so any dynamics study needed a fresh cloud per data point; phase-contrast imaging let them take repeated, non-destructive snapshots of one condensate.

Frometa Fernandez et al. (2025) push the technique to the opposite regime this calculator's OD sweet-spot warns you about: very dense clouds. Using a circular SLM phase mask (phase shift $\pi/3$, far detunings of −298 and −1010 MHz), they validated phase-contrast imaging up to resonant optical depth OD ≈ 64, densities well above λ^{-3}, where absorption imaging's two-level model (and this page's saturation correction) has already broken down from collective and many-body effects.

Non-destructive

Repeated shots, one cloud

$\varphi \propto 1/\Delta$

Phase signal vs detuning

OD ≈ 64

Validated dense-cloud limit (2025)

SLM / phase plate

Added hardware vs absorption

14.6 References & Further Reading

Key papers on absorption imaging technique, saturation correction, fringe removal, and ensemble thermometry.

- **Reinaudi et al. (2007)** — "Strong saturation absorption imaging of dense clouds of ultracold atoms." *Opt. Lett.* **32**, 3143–3145. Derives the saturation-corrected OD formula used in this calculator. Essential reading before imaging any cloud with OD > 1 or I > 0.1 I_sat. DOI →
- **Ockeloen et al. (2010)** — "Detection of small atom numbers through image processing." *Phys. Rev. A* **82**, 061606(R). Introduces PCA/eigenface fringe removal for absorption images. Demonstrates detection of ~10 atoms in a BEC with standard CCD cameras. Includes MATLAB code outline. DOI →
- **Hueck et al. (2017)** — "Calibrating high intensity absorption imaging of ultracold atoms." *Opt. Express* **25**, 8670. Rigorous treatment of saturation correction beyond the two-level model, including sublevel populations and polarization effects. Provides calibration procedure for I_sat in the actual experimental setup. DOI →
- **Schmidutz et al. (2014)** — "Quantum Joule-Thomson Effect in a Saturated Homogeneous Bose Gas." *Phys. Rev. Lett.* **112**, 040403. Appendix contains a concise, pedagogical summary of the three-image absorption technique and noise sources — useful as a quick reference alongside this calculator. DOI →
- **Ketterle, Durfee & Stamper-Kurn (1999)** — "Making, probing and understanding Bose-Einstein condensates." In: *Bose-Einstein Condensation in Atomic Gases*, Proc. Int. School of Physics "Enrico Fermi." arXiv:cond-mat/9904034. Sections 3–5 give a complete treatment of absorption imaging, cross sections, and cloud reconstruction for BEC experiments. The standard reference for BEC groups. arXiv →
- **Andrews et al. (1996)** — "Direct, Nondestructive Observation of a Bose Condensate." *Science* **273**, 84–87. Introduces phase-contrast imaging to cold-atom physics: the far-detuned, non-destructive alternative to everything above. Source for the new "Beyond Absorption" note in Practical Considerations. DOI →
- **Frometa Fernandez et al. (2025)** — "Phase-contrast imaging of a dense atomic cloud." *Rev. Sci. Instrum.* **96**, 083203. Validates phase-contrast imaging up to resonant OD ≈ 64 using an SLM phase mask, well past where absorption imaging's two-level model breaks down. Free preprint below. arXiv →
- **Steck atomic data pages** (steck.us/alkalidata) — comprehensive tables of σ_0, I_sat, hyperfine A/B coefficients, transition dipole matrix elements, and Clebsch-Gordan factors for the cycling transition. Values used in this calculator are taken directly from these tables.
- **Steck, "Rubidium-87 D Line Data" (2019)** — steck.us/alkalidata/rubidium87numbers.pdf — Table 7 gives cycling-transition $\sigma_0 = 2.907\times10^{-13}$ m² and I_sat = 1.669 mW/cm² for the F=2→F'=3 D2 line used in this calculator.

- **NIST Atomic Spectra Database** (physics.nist.gov/asd) — Einstein A-coefficients and oscillator strengths for transitions not covered by Steck (e.g., Sr, Yb, Dy, Er).
- **Loftus et al. (2004)** — "Narrow line cooling and momentum-space crystals." *Phys. Rev. A* **70**, 063413. Contains saturation intensity and cross section data for the Sr $^1S_0 \rightarrow {}^3P_1$ intercombination line, useful if imaging Sr on the narrow line. DOI →
- **Metcalf & van der Straten, *Laser Cooling and Trapping* (1999)** — Springer. Chapter 2 derives the resonant cross section $\sigma_0 = 3\lambda^2/2\pi$ from the oscillator model and discusses the role of degeneracy, polarization, and hyperfine pumping in real atoms.
- **Foot, *Atomic Physics* (2005)** — Oxford University Press. Chapter 8 "Doppler cooling" and Chapter 9 contain derivations of the scattering force, relevant for understanding the probe-atom interaction during the imaging pulse.
- **Pethick & Smith, *Bose-Einstein Condensation in Dilute Gases* (2002, 2008)** — Cambridge University Press. Appendix covers absorption imaging as a diagnostic tool for BEC density profiles.
- **Inguscio & Fallani, *Atomic Physics: Precise Measurements and Ultracold Matter* (2013)** — Oxford University Press. Chapter 4 covers time-of-flight expansion and absorption imaging together, useful for connecting this calculator to the TOF thermometry tool.

CHAPTER 14 SUMMARY

Absorption imaging gives calibrated atom numbers and spatial profiles in a single camera frame. Letting the cloud expand ballistically before taking that image immediately turns it into a thermometer.

Problems for Chapter 14

1. USE THE TOOL For ^{87}Rb, run the imaging calculator at $I/I_{sat} = 0.1$ and 1.0 with the other settings fixed. How do the effective cross-section, the OD noise and the minimum detectable atom number change, and why?
2. BY HAND With the resonant cross-section $\sigma_0 = 3\lambda^2/2\pi$, compute the peak optical depth of a Gaussian cloud of 10^6 atoms with an rms radius of 100 µm.
3. BY HAND Each scattered photon gives the atom a recoil velocity $v_r = \hbar k/m$. How many photons does it take to Doppler-shift a ^{87}Rb atom by $\Gamma/2$, and what does that imply for the probe pulse length?

Use the tool: on the interactive page. By hand: from the formulas in this chapter. Design: a decision a lab would make. Worked solutions are available to instructors on request (see Part V).

CHAPTER 15 · EXPERIMENT STEP 03

15

Time-of-Flight Thermometry

🌡️ MOT Temperature · Calculator

Ballistic expansion calculator for ultracold atom clouds, extract temperature from σ² vs t² fits, compute de Broglie wavelength, phase-space density, and compare to fundamental cooling limits.

> *Switching off the trap and taking absorption images (Chapter 14) after variable expansion delays maps the cloud's velocity distribution directly onto space.*

IN THIS CHAPTER Prerequisites: Ch. 13, Ch. 14

- Fit cloud width $\sigma^2(t) = \sigma_0^2 + (k_B T / m)\, t^2$ to extract cloud temperature independent of initial size.
- Compute peak phase-space density (PSD) and compare measured temperatures against Doppler and recoil scales.
- Avoid systematic errors from pixel magnification calibration, gravitational sag, and finite camera depth of field.

WEBSITE COMPANION Live interactive page ↗ · Notebook: Open in Colab · download tof-calculator.ipynb · Week 6

Time-of-flight (TOF) thermometry is the most direct way to measure ultracold atom cloud temperatures. After suddenly switching off the trap, the cloud expands ballistically: its width grows as $\sigma^2(t) = \sigma_0^2 + (k_B T/m)t^2$. A linear fit to σ^2 vs t^2 gives temperature without any model of the trap. This calculator fits your TOF data, returns temperature with uncertainty, and computes thermodynamic quantities including phase-space density, letting you track your evaporation ramp and identify when you cross the BEC threshold (PSD ≥ 2.612).

15.1 Cloud Expansion σ(t)

TOF RANGE

Ballistic expansion law

After release from a trap, the atom cloud expands freely. For a thermal gas with 1D width σ_0 at t=0:

$$\sigma^2(t) = \sigma_0^2 + \frac{k_B T}{m} t^2$$

$$\Rightarrow T = \frac{m\,(\sigma^2(t) - \sigma_0^2)}{k_B\, t^2}$$

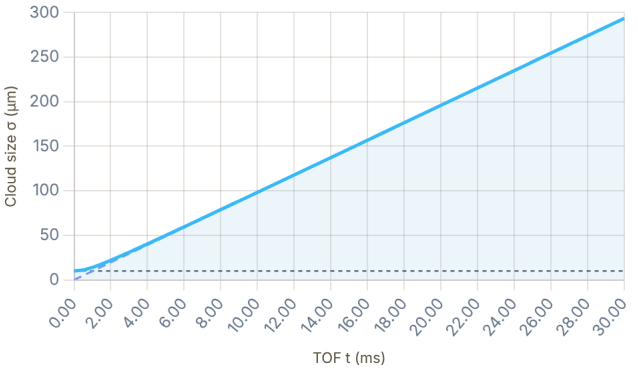
For harmonic trap (ω):

$\sigma_0 = \sqrt{k_B T/(m\omega^2)}$

$$\Rightarrow \sigma^2(t) = \frac{k_B T}{m}\left(\frac{1}{\omega^2} + t^2\right)$$

The slope of σ^2 vs t^2 gives $k_B T/m$ directly, plot $\sigma^2(t^2)$ and fit a line. Intercept = σ_0^2, slope = $k_B T/m$.

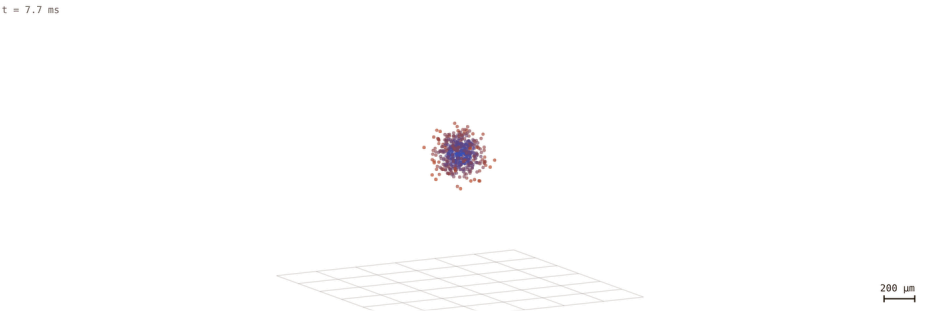
Σ(T), CLOUD RADIUS VS TIME OF FLIGHT



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

σ(t), full expansion Free expansion only (σ_0=0) σ_0 (in-situ limit)

Live TOF expansion · watch the cloud grow (and fall) after release



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

t = 7.7 ms | σ = 75.6 μm

> Colour = speed (blue slow → orange fast): the fast atoms end up at the edge, which is why σ^2 grows as $(k_B T/m)\cdot t^2$. 500 atoms sampled from a Gaussian cloud of size σ_0 and a Maxwell–Boltzmann velocity distribution at T; follows the sliders above.

Interactive

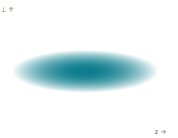
Who remembers the trap? Thermal cloud vs. condensate

Release both from the same cigar-shaped trap. A thermal gas expands with the same velocity spread in every direction, so after a few trap periods it is **round** whatever the trap looked like (and its temperature, not its shape, is what TOF measures). A Bose–Einstein condensate expands fastest along the direction it was squeezed hardest, so its shape **inverts**: the long axis becomes the short one. That inversion was one of the signatures of the first condensates in 1995.



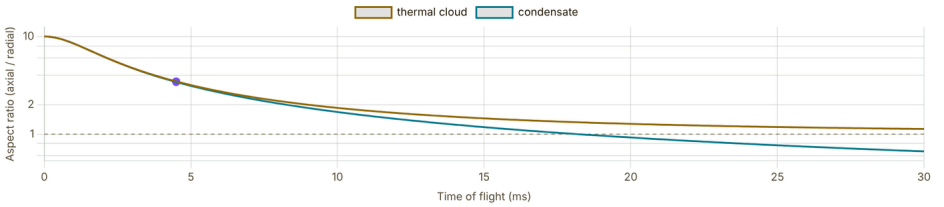
Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Thermal cloud (Gaussian, ballistic)



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Condensate (Thomas–Fermi, hydrodynamic)



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

t = 4.5 ms ($\omega_\perp t$ = 2.8): axial/radial size, thermal 3.47, condensate 3.42.

$\omega_\perp$ is the trap-frequency slider above; $\omega_z = \omega_\perp/(\text{aspect})$; z is drawn horizontally. Each image is scaled to fit its panel (shape only: the condensate's absolute size needs the atom number and scattering length). Thermal: $\sigma_j^2(t) = \frac{k_B T}{m}\left(\omega_j^{-2} + t^2\right)$, so the aspect ratio does not depend on T or m. Condensate: Thomas–Fermi radii $R_j(t) = R_j(0)\,\lambda_j(t)$ with $\ddot{\lambda}_j = \omega_j^2/(\lambda_j \lambda_x \lambda_y \lambda_z)$, $\lambda_j(0) = 1$, integrated numerically, and $R_z(0)/R_\perp(0) = \omega_\perp/\omega_z$ (Castin & Dum, PRL 77, 5315 (1996)). Gravity, interactions in the thermal cloud and the expansion of a partly condensed (bimodal) cloud are left out.

15.2 Fit Temperature from Data

Enter measured cloud widths σ at known TOF times t. The tool performs a linear least-squares fit of σ^2 vs t^2 to extract T and σ_0. You need at least 2 points.

MEASURED DATA

#	T (MS)	Σ (MM)
1		
2		
3		
4		
5		

FIT RESULTS

FITTED TEMPERATURE
1.983 μK

FITTED Σ_0
17.8 μm (in-situ size)

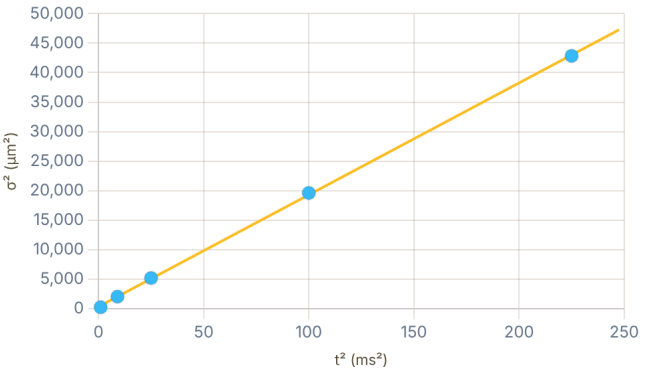
SLOPE (K_BT/M)
189.6956 μm²/ms²

R^2 (GOODNESS OF FIT)
0.9998
closer to 1 = better fit

Calculator output at the page's default settings.

Fitted T = 1.983 μK. Excellent R^2, good linear fit.

Σ^2 VS T^2, LINEAR FIT



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

15.3 Cooling Limit Reference

DOPPLER TEMPERATURE
145.5 μK
ħΓ/(2k_B)

SINGLE RECOIL SCALE
181.0 nK
$\hbar^2k^2/(2mk_B)$

THERMAL SIZE IN TRAP
15.57 μm
at ω/2π=100 Hz

Calculator output at the page's default settings.

T_Doppler = ħΓ / (2k_B)

In a 1D optical molasses, the equilibrium between laser cooling and heating from photon recoil occurs at the Doppler temperature. It depends only on the natural linewidth Γ and is independent of atomic mass. Typical values: Rb 145 μK, Cs 125 μK, Na 235 μK.

$$T_{\mathrm{DOPPLER}} = \frac{\hbar\Gamma}{2K_{\mathrm{B}}}$$

87

$\Gamma/2\pi = 6.065 \cdot T_D = 145.6\,\mu$

E_rec/k_B = ħ²k² / (2m k_B)

Each absorbed or emitted photon gives a momentum kick of order ħk. The single-photon recoil energy is $E_{\mathrm{rec}} = \hbar^2k^2/2m$; an absorption plus spontaneous-emission scatter cycle deposits an order-$2E_{\mathrm{rec}}$ heating scale when no cooling mechanism removes that energy. Sub-recoil cooling (VSCPT, Raman cooling, EIT) can go below this limit.

$$\frac{E_{\mathrm{rec}}}{k_{\mathrm{B}}} = \frac{\hbar^2k^2}{2mk_{\mathrm{B}}}$$

Example (^{87}Rb, $\lambda = 780$ nm): $E_{\mathrm{rec}}/k_{\mathrm{B}} \approx 181$ nK; one scatter cycle is roughly 362 nK without cooling.

PSD = n × λ_dB³

Phase-space density quantifies the occupation of quantum states. When PSD ≈ 2.612 (the Riemann ζ(3/2) value), a Bose gas in a uniform box reaches Bose-Einstein condensation. In a harmonic trap the critical condition is (k_BT_c/ħω) = (N/ζ(3))^(1/3).

$$\lambda_{\mathrm{dB}} = \frac{h}{\sqrt{2\pi mk_{\mathrm{B}}T}}$$

$$\mathrm{PSD} = n_{\mathrm{peak}}\lambda_{\mathrm{dB}}^3 = \frac{N}{(2\pi)^{3/2}\sigma^3} \cdot \left(\frac{h}{\sqrt{2\pi mk_{\mathrm{B}}T}}\right)^3$$

BEC threshold (uniform 3D box): PSD $= 2.612$
BEC critical T (harmonic trap):

$$T_c = \frac{\hbar\omega}{k_{\mathrm{B}}}\left(\frac{N}{\zeta(3)}\right)^{1/3} \qquad \zeta(3) \approx 1.202$$

15.4 References

📄 Phillips (1998), Nobel Lecture: Laser cooling and trapping of neutral atoms (Rev. Mod. Phys.) 📄 Cornell & Wieman (2002), Nobel Lecture: Bose-Einstein condensation in a dilute gas (Rev. Mod. Phys.) 📄 Bloch, Dalibard & Zwerger (2008), Many-body physics with ultracold gases (Rev. Mod. Phys.) 📊 Steck, Alkali Data Sheets (mass, Γ, λ, I_sat for all species) 📄 Anderson et al. (1995), Observation of Bose-Einstein Condensation (first BEC, Science)

CHAPTER 15 SUMMARY

Time-of-flight expansion is the gold-standard thermometer for cold atomic ensembles. Next, we zoom in from clouds of millions of atoms to detecting a single trapped atom in an optical tweezer.

Problems for Chapter 15

1. USE THE TOOL Set $T = 10\ \mu K$ and $\sigma_0 = 100\ \mu m$, add five synthetic data points between 2 and 15 ms, and fit the temperature. How close is the fit to the input?
2. BY HAND Using $\sigma^2(t) = \sigma_0^2 + k_B T t^2/m$, find the time at which the cloud size has doubled for your parameters.
3. DESIGN A thermal cloud released from an elongated trap becomes round, while a condensate inverts its aspect ratio. Use the aspect-ratio control to explain why, and say what this means for choosing the time of flight.

Use the tool: on the interactive page. By hand: from the formulas in this chapter. Design: a decision a lab would make. Worked solutions are available to instructors on request (see Part V).

CHAPTER 16 · EXPERIMENT STEP 04

16

Single-Atom Fluorescence Imaging

📷 Single-atom Imaging · Calculator

Fluorescence imaging budget for single atoms in optical tweezers, photon scattering rate, collection efficiency, camera noise model, and detection fidelity for EMCCD and sCMOS systems.

> *Unlike dense clouds, a single atom in a tweezer casts negligible shadow; instead, we collect its scattered fluorescence photons through a high-NA objective onto an EMCCD or sCMOS camera.*

IN THIS CHAPTER

Prerequisites: Ch. 8, Ch. 13

- Budget detected photon counts from scattering rate, objective NA, optical transmission, and camera quantum efficiency.
- Evaluate single-atom detection fidelity from the separation of empty-trap and loaded-trap photon histograms.
- Balance exposure time and imaging intensity to maximize readout fidelity without heating the atom out of the tweezer.

WEBSITE COMPANION Live interactive page ↗ · Notebook: Open in Colab · download imaging-calculator.ipynb · Week 9

Imaging budget workflow

16.1 Follow the photons from atom to threshold.

Single-atom imaging is not one number. Work through the chain in order: how many photons the atom scatters, how many your optics collect, how camera noise broadens the distributions, and how much the bright/dark histograms overlap.

01
Set the imaging transition
Choose species, detuning, saturation, and exposure to compute the scattering rate.

02 **Collect photons**
Set NA, transmission, QE, and camera model to turn scattered photons into detected signal.

03
Read the SNR meter
See the combined signal-to-noise ratio update live from your current inputs.

04
Check exposure time
See how SNR grows with exposure and where the SNR=10 and SNR=50 targets fall.

05
Understand the histogram
Learn what the bright/dark peaks and their overlap mean for detection fidelity.

06
Use the histogram
Choose a threshold and inspect the false-positive/false-negative tails directly.

07
Check the formulas
See the underlying scattering-rate, SNR, and fidelity equations this tool is built on.

Calculator output at the page's default settings.

Common mistake, SNR is not fidelity. A high photon-count SNR is only a proxy. Detection fidelity is the overlap error of the dark-count and bright-count histograms after choosing a threshold, and it can be limited by atom loss, blinking, camera tails, or imperfect state preparation even when the raw SNR looks excellent.

More NA helps twice

Sentence: Larger NA increases collected photons without increasing atom heating.

ηcoll = (1 - cosθ)/2

Failure mode: assuming higher camera gain improves photon collection.

Thresholds are physics

Sentence: The optimal threshold trades false positives against false negatives.

F = 1 - (P0→1 + P1→0)/2

Failure mode: quoting SNR when state loss dominates.

Survival is separate

Sentence: A measurement can identify an atom correctly while still ejecting it.

Freadout ≠ Psurvival

Failure mode: multiplying every error into one unlabeled number.

16.2 Imaging Beam

SCATTERING RATE	SATURATION PARAMETER	BRANCHING / LINEWIDTH
9.53 M photons / s	**1.00 I / I_{sat}**	**6.065 Γ/2π (MHz)**
9.5 k/ms	1.67 mW/cm² actual	τ = 26.2 ns

R_MAX AT SATURATION
19.05 M
photons / s (s→∞, Δ=0) · = Γ/2

Calculator output at the page's default settings.

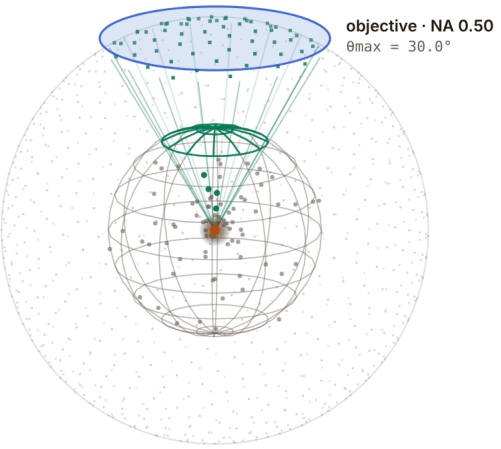
Two-level scattering rate

On resonance (Δ=0) a two-level atom saturates at Γ/2. Off resonance the Lorentzian denominator reduces the rate:

$$R_{sc} = \frac{\Gamma}{2} \cdot \frac{s}{1 + s + (2\Delta/\Gamma)^2}$$

$s = I/I_{sat}$, Γ = natural linewidth, Δ = laser detuning

16.3 Collection & Camera



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

WHERE THE PHOTONS GO

The atom scatters in every direction; the objective only catches the bright cap, $\theta < \theta_{max}$ = arcsin NA. Follows the NA slider above.

Photons caught: 4 of 79 (5.1%)

η here is the geometric fraction only (n = 1, no losses); the calculator then multiplies by transmission T and QE. The photon stream is a Monte Carlo sample of the chosen pattern, so the running tally settles on η.

16.4 Results

Signal-to-Noise Ratio 29.9

COLLECTION EFFICIENCY H	PHOTONS SCATTERED (TOTAL)	PHOTONS DETECTED (SIGNAL)
3.75%	**47.7 k in exposure time**	**1.8 k**
$\eta = \eta_{geo} \times QE \times T \cdot \eta_geo = 6.70\%$		$N_{sig} = R_sc \times t \times \eta$

Calculator output at the page's default settings.

100.000%

Detection fidelity

Excellent (>99.9%)

0.56 ms

Min exposure for SNR=10

target SNR = 10

13.99 ms

Min exposure for SNR=50

target SNR = 50

NOISE BUDGET

SOURCE	Σ^2 (PHOTONS^2)	FRACTION	CONTRIBUTION
Atom shot noise	3582.3	100.0%	
Background noise	1.0	0.0%	
Dark count noise	0.0	0.0%	
Read noise^2	0.0	0.0%	
Total σ^2	3583.3	100%	

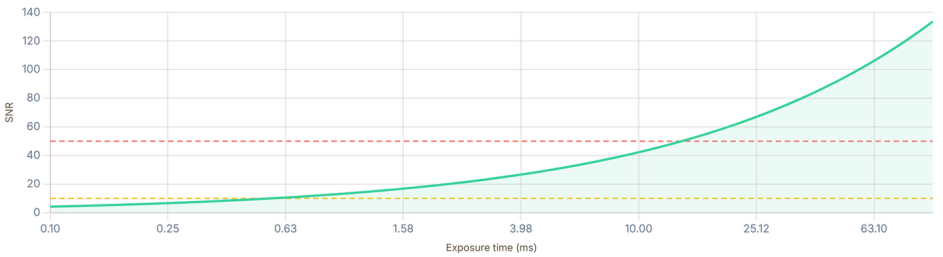
⚠️ At high scattering rates, atom heating from recoil kicks will limit imaging time. Consider gray molasses or EIT cooling during imaging, or reduce s.

🔬

Hood Lab imaging setup: We image ^{87}Rb and ^{133}Cs single atoms in optical tweezers using a NA = 0.6 objective (Mitutoyo G Plan Apo) and an EMCCD camera (Andor iXon). Typical imaging: 780 nm D2 line, detuning $\Delta = -2\Gamma$, saturation $s \approx 2$, exposure 2 ms. We collect ~600 photons from Rb at 99% survival probability with gray molasses applied during imaging. The critical figure is not peak SNR but the *histogram separation* between the 0-atom and 1-atom peaks — we target >8σ separation to achieve detection fidelity > 99.9%.

16.5 SNR vs Exposure Time

SNR AS A FUNCTION OF IMAGING DURATION



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

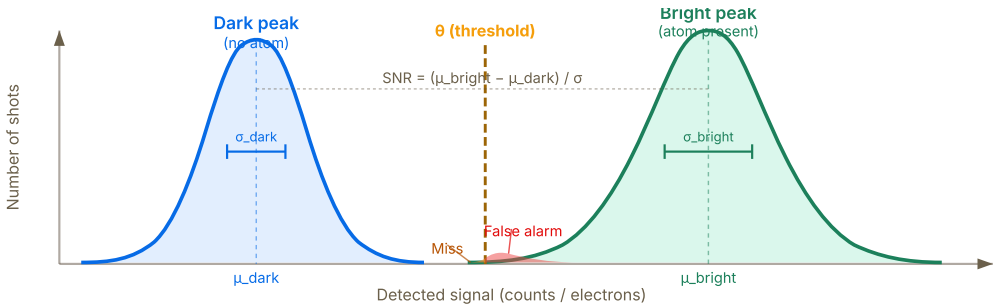
Current settings SNR = 10 target SNR = 50 target

16.6 Imaging Histogram & Detection Fidelity

What is the imaging histogram?

Run your experiment ~500–2000 times. Each shot, sum the camera electrons over the region of interest (typically a 3–7 px ROI around one tweezer site). Plot the distribution of those counts. You get a **bimodal histogram**: a narrow "dark" peak when the site is empty (only background scatter + camera noise), and a broader "bright" peak when an atom is present (atom fluorescence + background). The *overlap* between the two peaks is the fundamental limit on how well you can distinguish atom present from atom absent, this is your imaging fidelity.

ANATOMY OF A SINGLE-ATOM IMAGING HISTOGRAM



θ = threshold. Shots below θ → "no atom"; shots above θ → "atom". Orange = missed atoms (false negative). Red = false alarms (false positive). The overlap area determines error rates.

Peak models

For $N_{\text{det}} \gg 1$ detected photons, both peaks are well-approximated by Gaussians (Central Limit Theorem on Poisson photon statistics + Gaussian read noise):

$$P_{\text{dark}}(n) = \mathcal{N}(\mu_{\text{dark}}, \sigma^2_{\text{dark}})$$
$$\mu_{\text{dark}} = N_{\text{bg}} + N_{\text{dark}}$$
$$\sigma^2_{\text{dark}} = F^2 \mu_{\text{dark}} + (\sigma_{\text{read}}/G)^2$$
$$P_{\text{bright}}(n) = \mathcal{N}(\mu_{\text{bright}}, \sigma^2_{\text{bright}})$$
$$\mu_{\text{bright}} = N_{\text{sig}} + \mu_{\text{dark}}$$
$$\sigma^2_{\text{bright}} = F^2 \mu_{\text{bright}} + (\sigma_{\text{read}}/G)^2$$

$F = \sqrt{2}$

(EMCCD excess noise), $F = 1$ (sCMOS/ideal)

Why σ_bright > σ_dark

In the shot-noise-dominated model, noise grows as $\sqrt{N}$. More photons in the bright peak → larger absolute fluctuations → the bright peak is wider than the dark peak. This asymmetry means the optimal threshold is not exactly at the midpoint — it shifts slightly toward the brighter side to equalize miss and false-alarm probabilities.

ERROR PROBABILITIES & FIDELITY

Given threshold θ, define two error types:

$P_{\text{miss}}(\theta) = \Phi\left(\frac{\theta - \mu_{\text{bright}}}{\sigma_{\text{bright}}}\right)$ [bright peak below θ]

$P_{\text{false}}(\theta) = 1 - \Phi\left(\frac{\theta - \mu_{\text{dark}}}{\sigma_{\text{dark}}}\right)$ [dark peak above θ]

$$\mathcal{F}_{\text{detect}} = 1 - \frac{P_{\text{miss}} + P_{\text{false}}}{2}$$

Optimal threshold: minimise $P_{\text{miss}} + P_{\text{false}}$.
Equal σ: $\theta_{\text{opt}} = (\mu_{\text{bright}} + \mu_{\text{dark}})/2$.
Unequal σ: solve numerically.

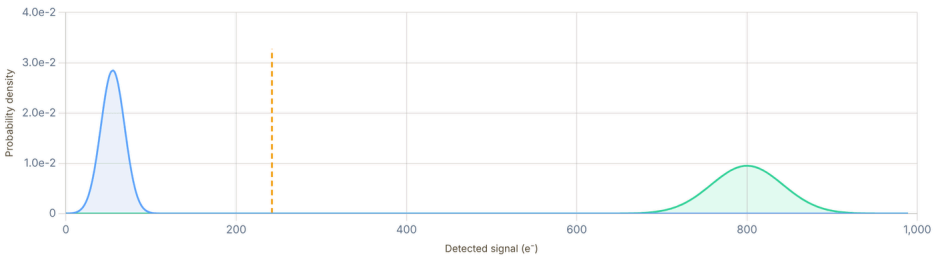
For equal peak widths and $\text{SNR} = (\mu_{\text{bright}} - \mu_{\text{dark}})/\sigma$:
$\mathcal{F} \approx \Phi\left(\frac{\text{SNR}}{2}\right)$. With unequal widths, compute P_{miss} and P_{false} separately.

SNR → Fidelity rule of thumb:
For equal-width Gaussian peaks and SNR defined as peak separation divided by one-peak σ: SNR = 5 → F ≈ 99.4% and SNR = 10 → F > 99.9999%. Real cameras often have unequal bright/dark widths, so histogram fitting is the safer final metric.

INTERACTIVE HISTOGRAM SIMULATOR

Preset:

^{87}Rb 780 nm D_2 fluorescence · EMCCD · 1 ms exposure · NA=0.5 · ~800 photons collected · σ_bright ≈ √(2·800 + 12²) ≈ 42 e⁻ (Poisson shot noise √800 ≈ 28, with its variance doubled by the EMCCD excess noise (factor √2 at high gain), plus read noise in quadrature) · σ_dark ≈ read noise 12 e⁻ + dark counts · With these peaks the page's SNR = (μ_b − μ_d)/√(σ_b² + σ_d²) ≈ 16.8, so the two-Gaussian overlap error at the optimal threshold is ≈ 1×10^{-40} (the fidelity card), far beyond anything a real detector reaches. At this separation real experiments are limited by what the two-Gaussian model leaves out: atom loss or heating during the exposure, and non-Gaussian tails (EMCCD multiplication noise, background fluctuations) · Illustrative parameters



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Dark peak (no atom) Bright peak (atom) Miss zone (orange) False alarm zone (red) Threshold θ

Alkali (^{87}Rb, ^{133}Cs)

Imaging transition: D2 cycling transition (780 nm for Rb, 852 nm for Cs). Closed cycling transition with repumper keeps the atom in the imaging cycle. R_sc → Γ/2 ≈ 1.9×10⁷ photons/s for Rb87 at strong saturation (Γ = 2π × 6.07 MHz); Γ/4 at s = 1 on resonance.

Photon counts: With NA=0.5 (η_geo ≈ 6.7%; total η = η_geo × QE × T ≈ 5%), 5 ms exposure at R_sc ≈ 3×10⁶/s (detuned, molasses-style imaging) → ~750 detected photons. σ_bright ≈ 40e⁻ (EMCCD with excess noise). Well-separated from dark peak (~50 e⁻).

Main challenge: Photon recoil heating (~181 nK per one-way recoil for Rb87; roughly twice that per scatter cycle without cooling). At R_sc = 3×10⁶/s, the atom heats by ~1 K/s (2E_rec·R_sc/k_B), i.e. ~1 mK per ms. Limit imaging to <10 ms or use concurrent Sisyphus/gray molasses cooling (recapture imaging, EIT imaging).

Background: Stray light from near-resonant beams contributes to dark peak. Typical dark peak: μ_dark = 50–200 e⁻ depending on background suppression.

Alkaline-earth-like (^{171}Yb, ^{88}Sr)

Imaging transition: Broad $^1P_1$ transition (399 nm for Yb, 461 nm for Sr, Γ/2π = 29/32 MHz). R_sc can reach Γ/2 ≈ 10⁸ photons/s, about 5× the alkali D2 maximum. Alternative: narrow $^3P_1$ (556 nm Yb, 689 nm Sr) for gentler imaging.

J=0 ground state advantage: Qubit states (nuclear spin) are largely spectator degrees of freedom for electronic cycling transitions such as $^1S_0$ → $^1P_1$. Atoms in the metastable $^3P_0$ state are invisible to 399/461 nm, enabling *shelved* (non-destructive) readout of one qubit basis while leaving the other undisturbed.

Magic wavelength: 759 nm tweezer is magic for $^1S_0/^3P_0$, no differential light shift on the qubit. The imaging beam (399 nm) shifts $^1S_0$ but not $^3P_0$, so imaging the un-shelved population gives state discrimination.

Histogram: Narrower dark peak (lower background at UV wavelengths far from tweezer), tight bright peak. Excellent separation even at 1ms exposure. Atom Computing Yb171: reports >99.9% imaging fidelity.

TWO-GAUSSIAN FIT PROCEDURE

Given your measured histogram H(n), fit the sum of two Gaussians:

$$H(n) = A_0\,\mathcal{N}(n;\,\mu_{\text{dark}}, \sigma^2_{\text{dark}}) + A_1\,\mathcal{N}(n;\,\mu_{\text{bright}}, \sigma^2_{\text{bright}})$$

6 free parameters: $A_0, \mu_{\text{dark}}, \sigma_{\text{dark}}, A_1, \mu_{\text{bright}}, \sigma_{\text{bright}}$. Fit with scipy.optimize.curve_fit or a Gaussian mixture model (sklearn.mixture.GaussianMixture).

Quality check: is $A_0/A_1 \approx$ expected loading fraction? Is

$\sigma_{\text{bright}}/\sigma_{\text{dark}} \approx \sqrt{\mu_{\text{bright}}/\mu_{\text{dark}}}$

? (this tests whether shot noise dominates as expected)

WHEN PEAKS ARE NOT CLEAN GAUSSIANS

Real histograms show deviations from two clean Gaussians in several situations:

- **Atom loss during imaging:** Some fraction of bright shots become mid-sequence dark → bright peak develops a low-side tail. Fix: reduce imaging power or duration, or use concurrent cooling to prevent loss.
- **Multiple occupancy:** If two atoms load into one site, you get a third peak at ~2×μ_bright. Set loading conditions (MOT density, tweezer depth) to suppress. Good diagnostic: multi-peak structure at integer multiples of μ_bright.
- **Background drift:** If background changes shot-to-shot (laser power fluctuations), both peaks broaden. σ_measured² = σ_intrinsic² + σ_drift². Use real-time background subtraction.

- **Non-Poissonian photon statistics:** Sub-Poissonian emission (from e.g. shelving protocols) or super-Poissonian (from intensity noise) → peak widths deviate from

  $\sqrt{\mu}$

  scaling. Check: plot σ² vs μ, should be linear with slope 1 (Poisson) or slope F² (EMCCD).
- **Charge-transfer smear (EMCCD):** Very fast pixel readout + high scattering rates → charge smear across pixels. Solved by choosing appropriate readout speed or using sCMOS.

Bootstrap fidelity uncertainty

After fitting, estimate statistical uncertainty via bootstrapping: resample your histogram N_shots times with replacement, refit, and take the std of the fidelity distribution. Typical uncertainty: δF ≈ 0.01–0.1% for 1000 shots, 0.001–0.01% for 10,000 shots. Note: systematic error from threshold choice is usually larger than statistical error for well-separated peaks.

Definitions used in neutral-atom QC papers

Different papers define "fidelity" slightly differently. For qubit experiments, the total SPAM (state preparation and measurement) error bundles several contributions:

Total SPAM error: $\varepsilon_{\text{SPAM}} = \varepsilon_{\text{prep}} + \varepsilon_{\text{meas}}$

$\varepsilon_{\text{meas}} = (P_{\text{miss}} + P_{\text{false}})/2$, from histogram overlap [what this section computes]

$\varepsilon_{\text{prep}}$, optical pumping error, atom number fluctuations, tweezer loading fraction

$\mathcal{F}_{\text{SPAM}} = 1 - \varepsilon_{\text{SPAM}} \approx 1 - \varepsilon_{\text{meas}} - \varepsilon_{\text{prep}}$

In Rb87 Harvard experiments: $\varepsilon_{\text{meas}} \lesssim 0.1\%$, $\varepsilon_{\text{prep}} \lesssim 0.2\%$ → $\mathcal{F}_{\text{SPAM}} > 99.5\%$

⚠️ **Common confusion:** Evered et al. 2023 (99.5% two-qubit gate) reports gate fidelity after SPAM correction, meaning SPAM errors are subtracted out. The *raw* circuit fidelity is lower. Always check whether a paper's quoted fidelity is SPAM-corrected or not. For practical fault-tolerant thresholds, you need the raw (uncorrected) fidelity.

🔬

Detection fidelity in practice: In tweezer experiments the dominant infidelity source shifts from *photon shot noise* (fixed by increasing NA or exposure) to *atom loss during imaging* (fixed by reducing recoil heating via gray molasses or EIT cooling). For Cs we switch to a 685 nm gray molasses imaging scheme: the cooling rate prevents axial escape while the camera collects photons, enabling ≥ 99.5% detection fidelity with only 1 ms exposure. The limiting factor at that level is *preparation fidelity* (state initialization), not imaging photon budget.

Key Papers — Single-Atom Imaging & Detection

[1] Schlosser et al., *Sub-Poissonian loading of single atoms in a microscopic dipole trap*, Nature 411, 1024 (2001) — first single-atom tweezer imaging. DOI

[2] Bakr et al., *A quantum gas microscope for detecting single atoms in a Hubbard-regime optical lattice*, Nature 462, 74 (2009) — quantum gas microscope, site-resolved imaging. DOI

[3] Fuhrmanek et al., *Free-space lossless state detection of a single trapped atom*, PRL 106, 133003 (2011) — 99.4% detection fidelity benchmark. DOI

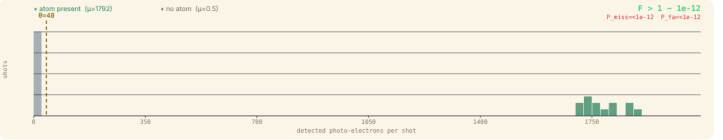
[4] Evered et al., *High-fidelity parallel entangling gates on a neutral-atom quantum computer*, Nature 622, 268 (2023) — detection fidelity analysis from bimodal histogram (Supplementary Methods). DOI

[5] Tuchendler et al., *Energy distribution and cooling of a single atom in an optical tweezer*, PRA 78, 033425 (2008) — atom survival and heating during imaging; release-recapture thermometry. DOI

[6] Manetsch et al., A *tweezer array with 6,100 highly coherent atomic qubits*, Nature 647, 60 (2025) — large-scale tweezer array imaging and loading statistics at >99.9% fidelity. DOI

16.7 Fluorescence Histogram — Threshold Preview

Each frame adds simulated imaging shots for your settings above: Poisson photo-electrons plus the camera's noise (sCMOS: read noise σ_read; EMCCD: gain-register multiplication, the source of the $\sqrt{2}$ excess noise, plus σ_read/G). **Green = atom present** (N_sig + N_bg + N_dark) · **Gray = empty tweezer** (N_bg + N_dark) · dotted curves = the expected histograms · **dashed = optimal threshold θ** (atom if counts ≥ θ). Here detection fidelity = 1 − P(miss) − P(false alarm), the stricter convention; the calculator's fidelity above is 1 − (P_miss + P_false)/2, the average for equal priors.



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

F_det > 1 − 1e-12

13 / 300 shots

P_miss=<1e-12 P_false_alarm=<1e-12

16.8 Physics & Formulas

Geometric collection fraction η_geo

A high-NA objective collects photons into a cone defined by the half-angle θ_max = arcsin(NA/n), where n is the refractive index of the medium (n=1 in air/vacuum).

$$\eta_{\text{geo}} = \frac{1 - \sqrt{1 - \text{NA}^2}}{2} \quad [\text{exact, for } n = 1]$$

$$\approx \frac{\text{NA}^2}{4} \quad [\text{paraxial, NA} \ll 1]$$

Example: $\text{NA} = 0.6 \Rightarrow \eta_{\text{geo}} = 10\%$; $\text{NA} = 0.8 \Rightarrow \eta_{\text{geo}} = 20\%$

Note that isotropic emission into 4π steradians means even NA=0.95 captures only ~34% of photons geometrically. High-NA objectives (0.5–0.85) are typical for tweezer experiments.

Noise variance per pixel per frame

Each camera type has a distinct noise model. The total variance σ²_total determines the denominator of SNR = N_sig / σ_total.

EMCCD (excess noise factor

$F = \sqrt{2}$

):

$$\sigma^2 = F^2(N_{\text{sig}} + N_{\text{bg}} + N_{\text{dark}}) + (\sigma_{\text{read}}/G)^2 = 2(N_{\text{sig}} + N_{\text{bg}} + N_{\text{dark}}) + (\sigma_{\text{read}}/G)^2$$

sCMOS (no multiplication noise):

$$\sigma^2 = N_{\text{sig}} + N_{\text{bg}} + N_{\text{dark}} + \sigma_{\text{read}}^2$$

Ideal: $\sigma^2 = N_{\text{sig}} + N_{\text{bg}} + N_{\text{dark}}$
where G = EM gain, σ_{read} = read noise (e⁻ RMS)

For EMCCD the read noise is divided by EM gain G, making it negligible at high gain (G ≳ 50). The penalty is the $\sqrt{2}$ excess noise factor that increases shot noise variance by 2×. For sCMOS cameras (1–2 e⁻ read noise), read noise is the dominant term at low photon counts.

Binary atom detection

We distinguish "atom present" (Poisson mean μ_{sig}) from "no atom" (Poisson mean μ_{bg}) by choosing a threshold θ. The optimal threshold minimises total error probability.

For Gaussian approximation ($\mu \gg 1$):

$$P_{\text{miss}} = \Phi\left(\frac{-(\mu_{\text{sig}} - \theta)}{\sqrt{\sigma_{\text{sig}}^2}}\right) \quad [\text{false negative}]$$

$$P_{\text{false}} = 1 - \Phi\left(\frac{\theta - \mu_{\text{bg}}}{\sqrt{\sigma_{\text{bg}}^2}}\right) \quad [\text{false positive}]$$

Optimal $\theta \approx (\mu_{\text{sig}} + \mu_{\text{bg}})/2$
Detection fidelity: $\mathcal{F} = 1 - (P_{\text{miss}} + P_{\text{false}})/2$.
If the bright and dark peaks have equal variance and SNR $= (\mu_{\text{sig}} - \mu_{\text{bg}})/\sigma$, then $\mathcal{F} \approx \Phi(\text{SNR}/2)$. The calculator below uses the bright and dark variances separately.

For SNR ≥ 10, fidelity exceeds 99.9%. The key insight is that the error is a Gaussian tail, ½erfc(SNR/2√2), so doubling SNR reduces it dramatically. Background suppression (e.g. dark-field or EIT imaging) improves fidelity by reducing μ_bg without reducing μ_sig.

Photon recoil vs trap depth

Every scattered photon imparts a recoil kick $\hbar k$. At scattering rate R_sc the atom heats at rate dE/dt ≈ 2E_rec × R_sc — one recoil from absorption, one from emission (in a σ^+/σ^- Sisyphus beam, this is partially cancelled; in a σ beam near resonance it accumulates).

$$E_{\text{rec}} = \frac{\hbar^2 k^2}{2m}$$

^{87}Rb D2: $E_{\text{rec}}/k_{\text{B}} \approx 181$ nK, so one absorption plus one emission deposits $\sim 2E_{\text{rec}}/k_B \approx 362$ nK
Heating rate: $\dot{T} \approx 2E_{\text{rec}}R_{\text{sc}}/k_{\text{B}}$ [absorption + spontaneous emission, no cooling]
Max imaging time (U_0 = trap depth): $t_{\text{max}} \sim U_0/(2E_{\text{rec}}R_{\text{sc}})$
Example: $U_0/k_{\text{B}} = 1$ mK, $R_{\text{sc}} = 50$ kHz $\Rightarrow t_{\text{max}} \approx 55$ ms

In practice, tweezer imaging uses repump + imaging beams in a gray molasses or Λ-enhanced dark SPOT configuration to scatter >1000 photons without losing the atom. This pushes fidelity above 99.5% in optimized tweezer experiments (Bergamini 2004; Blodgett et al. 2023; Manetsch et al. 2025).

16.9 References & Further Reading

📄 Bergamini et al. (2004), Holographic generation of microtrap arrays for single atoms by use of a programmable phase modulator, J. Opt. Soc. Am. B 21, 1889 PEER-REVIEWED 📄 Blodgett, Peana, Phatak et al. (2023), Imaging a ^{6}Li atom in an optical tweezer 2000 times with Λ-enhanced gray molasses (PRL 131, 083001) PEER-REVIEWED 📄 Hutzler et al. (2017), Eliminating light shifts for single atom trapping (New J. Phys.) PEER-REVIEWED 📊 Steck, Alkali Data Sheets (D1/D2 line data for all species) TEXTBOOK 📄 Kaufman & Ni (2021), Quantum science with optical tweezer arrays of ultracold atoms and molecules (Nat. Phys.) PEER-REVIEWED

CHAPTER 16 SUMMARY

High-fidelity, low-loss single-atom imaging makes individual tweezers readable and reusable. With single-atom detection working, we can measure an individual atom's temperature via release and recapture.

Problems for Chapter 16

1. USE THE TOOL With the defaults (NA 0.5, 5 ms exposure), record the scattering rate, the collected photons and the detection fidelity.
2. BY HAND Check the collected fraction with $(1 - \cos\theta)/2$, where $\sin\theta = \mathrm{NA}$, and multiply through by the transmission and quantum efficiency.
3. DESIGN Switch to the "Marginal (SNR ≈ 1.6)" preset. Change one thing at a time (exposure, NA, camera type) to reach 99.5% fidelity, and say which change you would make first in a real lab and why.

Use the tool: on the interactive page. By hand: from the formulas in this chapter. Design: a decision a lab would make. Worked solutions are available to instructors on request (see Part V).

CHAPTER 17 · EXPERIMENT STEP 05

17

Release-and-Recapture Thermometry

🎯 Single-atom Temperature · Simulator

Monte Carlo thermometry for optical tweezers. Release atoms, let them fly ballistically under gravity, then recapture, the survival curve encodes the temperature.

> *For a single atom in a tweezer, standard time-of-flight imaging lacks SNR; instead, we briefly extinguish the trap for time Δt, turn it back on, and measure the survival probability.*

IN THIS CHAPTER Prerequisites: Ch. 8, Ch. 16

- Model single-atom ballistic trajectories in a Gaussian tweezer potential via Monte Carlo sampling to generate recapture curves.
- Fit experimental survival versus release time Δt to infer the atom's temperature T relative to trap depth U_0.
- Distinguish what release-and-recapture measures (thermal energy for a known potential) from direct trap-frequency spectroscopy.

WEBSITE COMPANION Live interactive page ↗ · Notebook: Open in Colab · download release-recapture.ipynb · Week 9

Thermometry workflow

17.1 Read the survival curve before worrying about the derivation.

Release-recapture is action-first thermometry: switch off the trap, vary release time, measure survival, and compare against simulated curves to estimate temperature. The physics boxes explain why the curve changes shape.

01
What the curve means
Hot atoms expand faster, so survival drops earlier as release time increases.

02 **Run the simulator**

Set species, trap geometry, depth, waist, and temperature; compare the generated survival curve.

03 **Check species data**

D1/D2 transition data and core polarizability used for the dynamic polarizability calculation.

04 **Check failure modes**

Bad waist, anharmonicity, pointing drift, and nonthermal distributions can bias the result.

What the Survival Curve Means

How thermal atoms escape and the survival probability encodes temperature

Dynamic Polarizability

The trap depth of an optical tweezer is set by the AC Stark shift. For an alkali atom with D1 and D2 resonances, the scalar dynamic polarizability is (in atomic units):

$$\alpha(\omega) = \sum_j \frac{\Delta E_j \, |d_j|^2}{3(\Delta E_j^2 - \omega^2)} + \alpha_{\text{core}}$$

where ΔE_j = E_excited − E_ground in Hartree, d_j is the reduced matrix element in units of $e{\cdot}a_0$, and ω is in atomic units ($E_h/\hbar$). Converting to SI and computing the peak intensity:

$$U_0 = \frac{\alpha_{\text{SI}} \, I_{\text{peak}}}{2\,c\,\varepsilon_0}, \qquad I_{\text{peak}} = \frac{2SP}{\pi w_0^2}$$

where S is the Strehl ratio (1 = perfect focus). The trap frequency is then:

$$\omega_r = \sqrt{\frac{4U_0}{m w_0^2}}, \qquad \omega_z = \sqrt{\frac{2U_0}{m z_R^2}}$$

Thermal Initial Conditions

In the harmonic approximation of the Gaussian well, the thermal position distribution has widths:

$$\sigma_r = \sqrt{\frac{k_B T}{m\omega_r^2}} = \sqrt{\frac{k_B T \, w_0^2}{4U_0}}$$

$$\sigma_z = \sqrt{\frac{k_B T}{m\omega_z^2}} = \sqrt{\frac{k_B T \, z_R^2}{2U_0}}$$

Velocities are sampled from a Maxwell-Boltzmann distribution:

$$v_{x,y,z} \sim \mathcal{N}\left(0,\ \sqrt{k_B T / m}\right)$$

Ballistic Free Flight

During the release time Δt, atoms propagate freely under gravity (along $\hat{z}$):

$$x(\Delta t) = x_0 + v_x\,\Delta t$$

$$y(\Delta t) = y_0 + v_y\,\Delta t$$

$$z(\Delta t) = z_0 + v_z\,\Delta t - \tfrac{1}{2} g \Delta t^2$$

The z-velocity acquires a gravitational component:

$$v_z(\Delta t) = v_{z0} - g\,\Delta t$$

Recapture Condition

After the tweezer is re-switched on, an atom at position (x_f, y_f, z_f) with velocity v_f is recaptured if its total energy is negative (bound state):

$$E_{\text{total}} = \tfrac{1}{2} m |\mathbf{v}_f|^2 + U(\mathbf{r}_f) < 0$$

where the Gaussian trap potential is:

$$U(\rho, z) = -U_0 \left(\frac{w_0}{w(z)} \right)^2 \exp\left(\frac{-2\rho^2}{w(z)^2} \right)$$

$$w(z) = w_0 \sqrt{1 + z^2 / z_R^2}$$

So the recapture condition becomes: **½m|v_f|² < U₀·(w₀/w(z_f))²·exp(−2ρ_f²/w(z_f)²)**

Run the Monte Carlo Thermometer

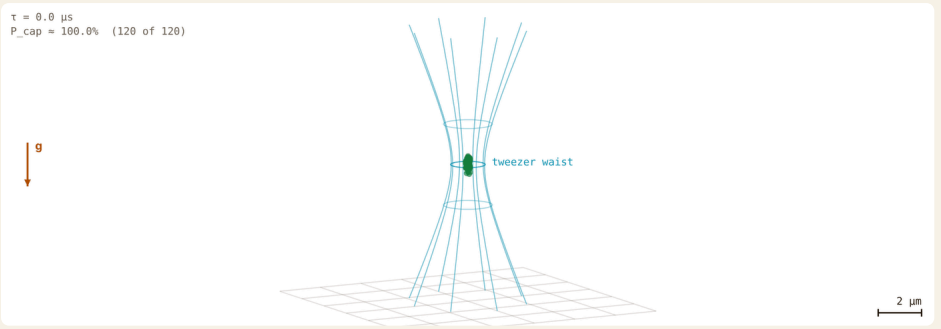
Port of the Python/NumPy vectorised simulation, runs in your browser

Tweezer Parameters

Done, 3 curve(s) · 600 samples each.

ATOM TRAJECTORY SNAPSHOT

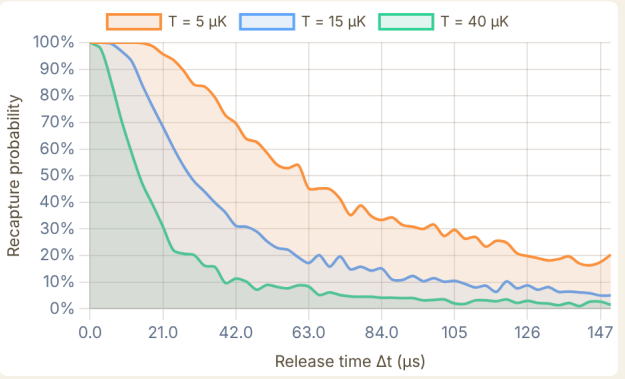
Drag the slider to change release time τ. Green = recaptured, red = escaped. The view is true to scale in 3D (tweezer axis vertical, along gravity, as in the Monte Carlo); drag the figure to turn it.



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

w_0=0.80µm · U_0=351µK · $\omega_r/2\pi$=59.0kHz · P_cap ≈ 100.0%

SURVIVAL PROBABILITY P_CAP(ΔT)



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

TRAP & SYSTEM PARAMETERS

Species: ^{133}Cs
λ: 1064 nm
α: 1163.4 a.u.
I_peak: 1.34 GW/m²
U_0: 351 μK
$\omega_r/2\pi$: 59.0 kHz
ω_z/2π: 17.65 kHz
z_R: 1.89 μm

T (MK)	P_CAP(0)	P_CAP(T_MAX)	½-LIFE (MS)
5	1.000	0.203	61.4 μs
15	1.000	0.050	28.9 μs
40	1.000	0.015	14.2 μs

Species Polarizability Data

D1 + D2 transition data used for dynamic polarizability calculation

SPECIES	ΔE_D1 (E_h)	|D_D1| (EA_o)	ΔE_D2 (E_h)	|D_D2| (EA_o)	A_CORE (A.U.)	MASS (U)
^{133}Cs	0.050932	4.4890	0.053456	6.3238	17.35	133
^{87}Rb	0.057314	4.231	0.058396	5.977	10.54	87
^{23}Na	0.077258	3.5246	0.077336	4.9838	1.86	23
^{6}Li	0.067906	3.317	0.067907	4.689	2.04	6

Polarizability Formula (full)

In atomic units, with ω_au = $E_h/\hbar \approx 4.134\times10^{16}$ rad/s and ω_laser = 2πc/λ:

$$\alpha(\omega) = \sum_{j\in\{D1,D2\}} \frac{\Delta E_j \, |d_j|^2}{3\left(\Delta E_j^2 - (\omega/\omega_{\text{au}})^2\right)} + \alpha_{\text{core}}$$

SI conversion: α_SI = α_au × $4\pi\varepsilon_0 a_0^3$ = α_au × 1.6488×10^{-41} C·m²/V. Note: for trap wavelengths longer than the D1 and D2 lines (red detuning, as in standard tweezers), α > 0 and the trap minimum is at maximum intensity; between or below the lines (blue detuning) α < 0 and the atom is pushed out of the light.

Assumptions & Failure Modes

- ▸**Harmonic initial distribution:** Atom positions are sampled from a Gaussian distribution matching the harmonic approximation of the Gaussian well. Valid when T ≪ U_0/k_B.
- ▸**Ballistic flight:** No inter-atomic collisions (single-atom tweezer). Gravity acts in the −z direction. No light forces during release.
- ▸**Full Gaussian recapture:** The recapture condition uses the exact Gaussian potential (not harmonic approx), including the z-dependence through $w(z) = w_0\sqrt{(1 + z^2/z_R^2)}$.
- ▸**Scalar polarizability:** We use only the scalar component; vector/tensor contributions are neglected (valid for linearly polarized, far-detuned tweezers in the electronic ground state).
- ▸**Single-transition approximation:** Only D1 and D2 lines contribute (core correction α_core included). Higher excited states add <1% for λ > 700 nm.

- ▸**Recapture efficiency:** The simulation does not model atom re-heating during recapture, quantum tunnelling, or motional-state filtering. It predicts the classical recapture probability.
- ▸**Thermometry, not direct trap-frequency metrology:** Release-recapture primarily extracts temperature for a known potential. Trap frequencies enter through the assumed waist, depth, and Gaussian geometry; use parametric heating or sideband spectroscopy for a direct frequency measurement.
- ▸**Statistical noise:** For N_MC = 600, the binomial standard error $\sqrt{(p(1-p)/N)}$ is at most $0.5/\sqrt{600} \approx 2\%$. Increase to 1500 for smoother curves.

Sources

Tuchendler et al. (2008), energy distribution and thermometry of a single atom in a tweezer Kaufman et al. (2012), single-atom tweezer cooling and thermometry context Steck alkali data sheets for line and polarizability inputs Savard et al. (1997), PRA 56 R1095 — laser-noise-induced heating of trapped atoms; heating rate model from intensity fluctuations Gehm et al. (1998), PRA 58 3914 — dynamics of noise-induced heating in atom traps; parametric and linear heating Thompson et al. (2013), PRL 110 133001 — coherent single-atom sideband cooling to the motional ground state in a tweezer

CHAPTER 17 SUMMARY

Release-and-recapture confirms whether single atoms are cold enough for reliable trapping and rearrangement. To reach the quantum motional ground state ($\bar{n} \to 0$) required for high-fidelity Rydberg gates, we next turn to sub-recoil single-atom cooling.

Problems for Chapter 17

1. USE THE TOOL Simulate survival curves for 5, 10 and 20 μK. Record the release time at which survival falls to 50% for each.
2. BY HAND Explain with a simple ballistic argument why the 50% release time scales roughly as $1/\sqrt{T}$, and check whether your data from problem 1 follow it.
3. DESIGN Lower the Strehl ratio from 1 to 0.6. Why does the inferred temperature change if you keep fitting with the ideal trap? What does that mean for a real measurement?

Use the tool: on the interactive page. By hand: from the formulas in this chapter. Design: a decision a lab would make. Worked solutions are available to instructors on request (see Part V).

CHAPTER 18 · EXPERIMENT STEP 06

18

Ground-State Cooling of Trapped Atoms

🧊 Single-atom Cooling · Deep dive

A unified theory of optical cooling in neutral atom tweezers, from the Lamb-Dicke regime to quantum ground state preparation. Based on Phatak et al., *Phys. Rev.* A 110, 043116 (2024) (*arXiv*:2406.19153).

> *When a single atom is tightly confined in the Lamb-Dicke regime (η_LD $\ll 1$), motional sidebands resolve cleanly—allowing optical cooling all the way to the vibrational ground state $\bar{n} \approx 0$.*

IN THIS CHAPTER Prerequisites: Ch. 8, Ch. 13, Ch. 17

- Understand the Lamb-Dicke parameter η_LD and why recoil heating is suppressed inside a stiff harmonic trap.
- Compare resolved sideband cooling (RSC), Raman sideband cooling, Λ-enhanced gray molasses, and EIT cooling.
- Evaluate cooling rates and final mean phonon numbers n̄_f for alkali and alkaline-earth-like species.

WEBSITE COMPANION Live interactive page ↗ · Notebook: Open in Colab · download laser-cooling.ipynb · Week 10

🎯

18.1 Why Cool Atoms to the Ground State?

In a neutral-atom optical tweezer experiment, a single atom sits inside a tightly focused laser beam, confined like a harmonic oscillator. The atom's thermal motion, its occupation of motional energy levels $|n\rangle$, is the primary source of decoherence and gate infidelity in quantum computing experiments. Laser cooling techniques reduce this thermal motion until the atom sits in the motional ground state $|n = 0\rangle$, where quantum coherence is maximized.

🌡️
~150 μK
Doppler limit T_D = ħΓ/(2k_B): Rb 146 μK, Cs 126 μK, Na 235 μK

❄️
~1 μK
Sub-Doppler: gray molasses, PG cooling limit

⚛️
⟨n⟩ < 0.1
Ground state: Raman sideband, resolved SB

📐
η ≪ 1
Lamb-Dicke regime: the key parameter for all cooling

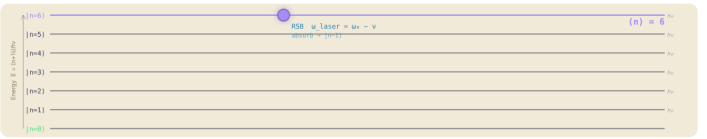
Calculator output at the page's default settings.

The Unified Spin Cooling Framework

All two-photon optical cooling techniques, polarization gradient (PG), gray molasses (GM), Λ-enhanced GM, EIT, and Raman sideband (RSB), share the same fundamental structure. After adiabatically eliminating the excited state and expanding to first order in the Lamb-Dicke parameter, each produces a Hamiltonian and collapse operators with a common form: a differential light shift between spin states equal to the trapping frequency (providing resonance), preferential optical pumping to the lower energy spin state (providing directionality), and coherent coupling between motional states (providing the actual cooling transition).

This unified picture is the central result of Phatak et al. (2024). It reveals three factors that determine the final temperature: (1) large pumping asymmetry γ_+/γ_-, (2) small Lamb-Dicke parameter η, and (3) a maximally dark lower-energy state.

Phonon ladder — watch the atom cool step-by-step to |n=0⟩ via sideband absorption + spontaneous emission



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

📐

18.2 The Lamb-Dicke Regime

Foundation

A tightly trapped atom behaves like a quantum harmonic oscillator with discrete energy levels $|n\rangle$ separated by $\hbar\omega$ ($\omega = 2\pi\nu$), where ν is the trap frequency. The spatial extent of the ground state wavefunction is the zero-point motion x_0. The Lamb-Dicke parameter η compares this motion to the photon wavelength, it is the single most important parameter in laser cooling.

TRAP HAMILTONIAN AND MOTIONAL STATES

The atom in the trap is described by a harmonic oscillator:

TRAP HAMILTONIAN

$$H_{\text{trap}} = \hbar\omega\left(a^\dagger a + \tfrac{1}{2}\right), \quad \omega = 2\pi\nu$$

Calculator output at the page's default settings.

Zero-point motion

$$x_0 = \sqrt{\frac{\hbar}{2m\omega}} = \sqrt{\frac{\hbar}{4\pi m\nu}}$$

Lamb-Dicke parameter

$$\eta = k\,x_0 = k\sqrt{\frac{\hbar}{2m\omega}} = k\sqrt{\frac{\hbar}{4\pi m\nu}}$$

LD EXPANSION OF LIGHT-MATTER COUPLING

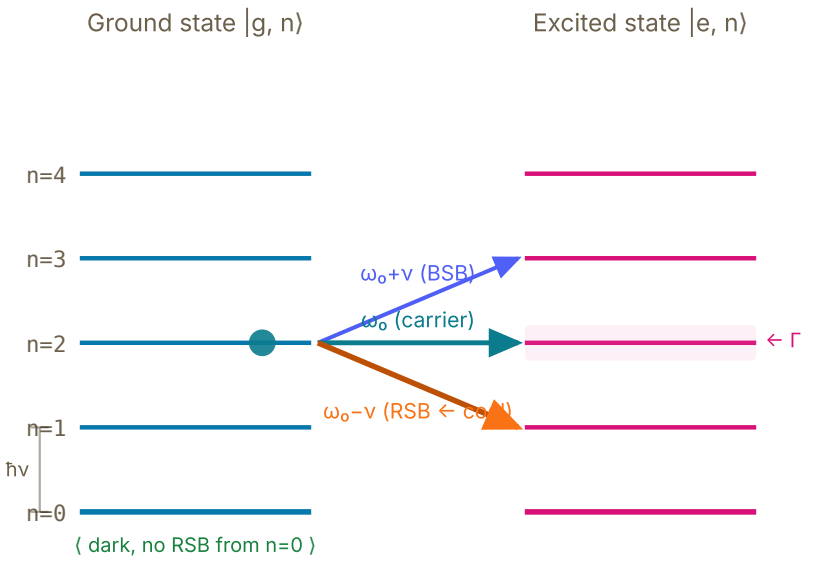
$$e^{ik\hat{x}} \approx 1 + i\eta(a + a^\dagger) \quad (\eta \ll 1)$$

Physical Meaning of η

When $\eta \ll 1$, the atom's position uncertainty during a photon interaction is much smaller than the photon wavelength. A photon absorption is unlikely to change the motional quantum number ($\Delta n \neq 0$ transitions are suppressed by η). This suppression is what makes sideband cooling efficient: the carrier transition ($\Delta n = 0$) dominates, and both sidebands are weaker than the carrier by η^2 (red: $\eta^2 n$, blue: $\eta^2(n+1)$). Cooling comes from tuning the laser to the red sideband ($\Delta n = -1$) so that, with resolved sidebands, the blue sideband ($\Delta n = +1$, heating) is far off resonance.

Typical values: $\eta \approx 0.14$–0.27 for ^{87}Rb in optical tweezers with $\nu \approx 200$–$50\,\text{kHz}$ and $\lambda \approx 780\,\text{nm}$ ($\eta \approx 0.19$ at 100 kHz).

Key insight: The Lamb-Dicke parameter sets the fundamental limit for all cooling techniques. Tighter traps (larger ν) give smaller η and colder final temperatures. This is why tweezer experiments push for MHz-level trap frequencies.



Harmonic oscillator energy levels. The atom in the tweezer trap has discrete motional states spaced by ħν. Three transitions are possible from any level: carrier (ω_0, Δn=0), blue sideband ($\omega_0+\nu$, Δn=+1, heating), and red sideband ($\omega_0-\nu$, Δn=−1, cooling). For resolved sideband cooling: $\nu \gg \Gamma$ so the RSB can be driven selectively. The ground state n=0 has no RSB → population accumulates there.

18.3 Resolved Sideband Cooling

Single-photon

Resolved sideband (RSB) cooling is the simplest route to the motional ground state, applicable when the trap frequency exceeds the photon linewidth ($\nu \gg \Gamma$). The atomic spectrum splits into a carrier (at ω_0) and discrete sidebands at $\omega_0 \pm n\nu$. By tuning a laser to the red sideband at $\omega_0 - \nu$, every absorption removes one quantum of motional energy.

Cooling Cycle

Step 1: Drive the red sideband, a photon at $\omega_0 - \nu$ takes $|g, n\rangle \to |e, n-1\rangle$. Step 2: Spontaneous emission returns the atom to the ground electronic state. In the Lamb-Dicke regime, spontaneous emission almost always leaves n unchanged (it changes n only with probability of order η^2), so the atom lands in $|g, n-1\rangle$. The net effect per cycle: one motional quantum removed. The state $|g, 0\rangle$ is dark (no red sideband to drive), so the population accumulates there.

Resolved sideband steady-state (Eq. 38)

$$\langle n \rangle_{ss} = \frac{1}{4}\left(\frac{\Gamma}{\nu}\right)^2 \left(\alpha + \frac{1}{4}\right)$$

α is the emission-direction factor: $\alpha = 1$ for emission into 1D (used in the paper's simulations), $\alpha = 1/3$ averaged over all directions. The $\frac{1}{4}$ comes from off-resonant carrier absorption, the α term from the recoil of the emitted photon.

CONDITION FOR RESOLVED SIDEBANDS

$$\nu \gg \Gamma \quad \Leftrightarrow \quad \frac{\Gamma}{\nu} \ll 1$$

Physical Mechanism

The cooling rate is set by the photon scattering rate on the red sideband: $R_- = \eta^2 \Gamma$. Heating comes from off-resonant excitation: the blue sideband, and the carrier followed by a spontaneous photon whose recoil changes n, each suppressed by $(\Gamma/\nu)^2$. In steady state these balance to give $\langle n \rangle_{ss} = \frac{1}{4}(\Gamma/\nu)^2(\alpha + \frac{1}{4})$.

The formula is striking: the final temperature depends only on the ratio Γ/ν, not on the laser intensity or the Lamb-Dicke parameter. Making the trap tighter (larger ν) directly lowers the temperature floor.

Example: For ^{87}Rb ($\Gamma/2\pi = 6$ MHz) in a 200 kHz trap the formula gives $\langle n \rangle \approx \frac{1}{4}(30)^2 \times 1.25 \approx 280$ (with $\alpha = 1$): meaningless, because the sidebands are not resolved. Resolved single-photon SB cooling needs narrow lines such as Sr $^1S_0 \to {}^3P_1$ (689 nm, $\Gamma/2\pi = 7.4$ kHz), or two-photon Raman transitions (section below).

Sideband Cooling Limit Calculator

Sideband regime

Unresolved ×

⟨n⟩ resolved ¼(Γ/ν)²(α+¼)

n/a (unresolved)

Doppler limit ⟨n⟩ (Δ = −Γ/2)

29.825

Resolved = $\nu > \Gamma$ (both as ordinary frequencies). The resolved formula only applies in that regime. Doppler limit (unresolved, $\nu \ll \Gamma$, at the optimal $\Delta = -\Gamma/2$): $\langle n \rangle = (1+\alpha)\Gamma/(4\nu) - \frac{1}{2}$, i.e. $k_B T_D = \hbar\Gamma/2$ for $\alpha = 1$.



Resolved sideband cooling animation

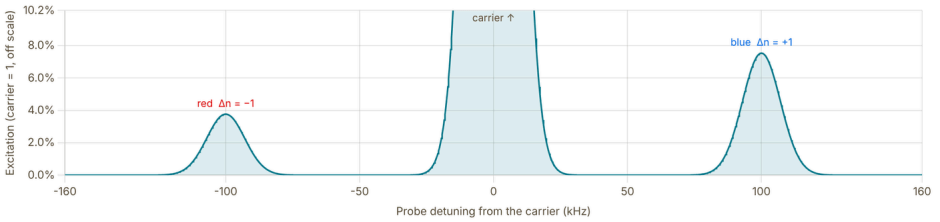
Resolved sideband cooling cycle. The population is driven from a higher motional state to the excited-state red sideband, then decays into the next-lower motional state.

Interactive

How do you know it worked? Reading ⟨n⟩ off the sidebands

Probe the cooled atom weakly and scan the probe across the carrier. The red sideband ($\Delta n = -1$) needs a quantum to take away, so its height is $\propto \eta^2\langle n\rangle$; the blue sideband ($\Delta n = +1$) can always add one, so it is $\propto \eta^2(\langle n\rangle + 1)$. Their ratio is the thermometer, and it needs no calibration of η or of the probe power:

$$\frac{P_{\text{red}}}{P_{\text{blue}}} = \frac{\langle n\rangle}{\langle n\rangle + 1} \quad \Longrightarrow \quad \langle n\rangle = \frac{R}{1-R}$$



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

η, ν (from the Lamb-Dicke calculator)

η = 0.194, ν = 100 kHz

Red / blue ratio R

0.500

Ground-state population (thermal) 1/(⟨n⟩+1)

50.0 %

Red sideband 3.77 %, blue 7.54 % of the carrier. Inferring ⟨n⟩ = R/(1−R) from these heights gives back 1.00.

Model: weak (perturbative) probe in the Lamb-Dicke regime, sideband strengths $\eta^2\langle n\rangle$ and $\eta^2(\langle n\rangle + 1)$ relative to the carrier, which is drawn with height 1 and runs off the top of the plot. The ratio holds for any phonon distribution; $P(n{=}0) = 1/(\langle n\rangle + 1)$ assumes a thermal one. Each line is drawn as a Gaussian of full width ν/6, a display choice, not a property of your probe. Strong (π-pulse) probing and $\eta \gtrsim 0.3$ need the full Rabi frequencies $\Omega_{n,n\pm1}$. Used this way by Monroe et al., PRL 75, 4011 (1995) (trapped ion) and Kaufman, Lester & Regal, PRX 2, 041014 (2012) (single atom in a tweezer).

18.4 Unresolved Sideband Cooling

Doppler regime

When $\nu \ll \Gamma$, the trap sidebands are not spectrally resolved, the photon linewidth is so broad that a single laser drives both the carrier and nearby sidebands simultaneously. This is the regime most atoms find themselves in after a MOT. Standard Doppler cooling operates here, and the cooling limit is set by competition between the viscous force (cooling) and random recoil kicks (heating).

The Unresolved Sideband Limit

When the sidebands overlap, the atom cannot distinguish cooling transitions from carrier transitions. Expanding the rate equations for $\nu \ll \Gamma$ gives (cf. Eq. 39 of Phatak et al.):

Unresolved sideband limit (Eq. 39)

$$\langle n\rangle_{ss} = \frac{1+\alpha}{8}\frac{\Gamma}{\nu}\left(\frac{\Gamma}{2|\Delta|}+\frac{2|\Delta|}{\Gamma}\right)-\frac{1}{2} \xrightarrow{\Delta=-\Gamma/2} \frac{(1+\alpha)\Gamma}{4\nu}-\frac{1}{2} \gg 1$$

Equivalent temperature

$$k_B T_D = \frac{\hbar\Gamma}{2} \quad \text{(Doppler temperature, angular-frequency convention)}$$

Example: For Rb ($\Gamma/2\pi = 6\,\text{MHz}$) at $\nu = 100\,\text{kHz}$ and $\alpha = 1$: $\langle n\rangle \approx 6\,\text{MHz}/(2\times 100\,\text{kHz}) - \frac{1}{2} \approx 30$, matching $k_B T_D/h\nu$. Still very far from the ground state.

The Doppler Temperature

The Doppler limit arises from the competition between two processes: the laser's viscous force damps atomic momentum (cooling), while random recoil from spontaneous emission heats the atom. At the optimal detuning $\Delta = -\Gamma/2$, the minimum achievable temperature is $T_D = \hbar\Gamma/(2k_B)$.

For ^{87}Rb: $T_D = 146\,\mu\text{K}$. For the ground state in a 100 kHz trap, we need $T \lesssim h\nu/k_B \approx 5\,\mu\text{K}$. Doppler cooling alone does not reach the ground state for broad-line atoms.

Crossing to Sub-Doppler

The solution is two-photon cooling, the techniques that follow. By using spin-dependent light shifts and coherent dark states, these methods surpass the Doppler limit while still using the same broad transitions. The Lamb-Dicke parameter η then controls how efficiently the motional sidebands are addressed.

18.5 Dark-State Sideband Cooling

Standing wave

A standing wave creates a spatially varying light-matter coupling. At the nodes of the standing wave, the electric field vanishes and the atom does not couple to light, it becomes "dark" to the laser. An atom cooled to a motional state whose wavefunction concentrates near the node has its scattering heavily suppressed, enabling very deep cooling.

Position-Dependent Coupling

In a standing wave, the interaction Hamiltonian contains $\sin(k\hat{x})$ or $\cos(k\hat{x})$ instead of $e^{ik\hat{x}}$. Near a node (at $\hat{x} = 0$), expanding to first order:

NEAR A NODE: $\sin(KX) \approx KX$

$$\hat{H}_I \propto \Omega\, k\hat{x}\, |\text{e}\rangle\langle\text{g}|$$

Matrix element for motional transition

$$\langle n-1|\, k\hat{x}\, |n\rangle = kx_0\sqrt{n} = \eta\sqrt{n}$$

Enhanced Cooling at the Node

At the node, the $n \to n-1$ coupling scales as $\eta\sqrt{n}$, which goes to zero as $n \to 0$. The state $|n=0\rangle$ is therefore dark at the node, it does not scatter photons. This creates a natural dark state for cooling and allows the motional ground state population to accumulate without being further heated.

Different Trap Frequencies

When the cooling laser and the trap have different length scales (e.g., a red-detuned tweezer and a near-resonant cooling laser), the motional states seen by each can differ. The paper treats this case using a squeezing transformation, showing that an effective Lamb-Dicke parameter

Effective LD parameter (Eq. 55 regime)

$$\eta_{\text{eff}} = \eta\sqrt{\frac{\nu_{\text{cool}}}{\nu_{\text{trap}}}}$$

applies when the cooling and trapping frequencies differ. A tighter trap ($\nu_{\text{trap}} > \nu_{\text{cool}}$) reduces η_{eff}, pushing the cooling limit lower.

> **Practical note:** In optical tweezer experiments, the tweezer itself provides the tight trap, while a separate near-resonant beam provides the cooling sideband drive. This geometry is exactly the "different trap frequencies" case and allows MHz-trap tweezers to cool using MHz-linewidth transitions.

🔄

18.6 Polarization Gradient Cooling

Sub-Doppler

Polarization gradient (PG) cooling uses two counter-propagating laser beams with different polarizations to create a spatially varying light shift that depends on the atom's internal spin state. The atom's spin follows the local light polarization adiabatically, and the resulting Sisyphus mechanism extracts kinetic energy efficiently. Temperatures well below the Doppler limit are achievable.

The Sisyphus Mechanism

In the lin $\perp$ lin geometry, two counter-propagating beams create a polarization gradient (the intensity stays uniform). The light shift of the m_F sublevels varies sinusoidally in space, and the optical pumping rate between sublevels also varies. An atom in the higher-m_F state climbing a potential hill is pumped by optical pumping to the lower-m_F state at the hill top, like Sisyphus rolling a boulder up forever. Each pump cycle dissipates kinetic energy equal to the light shift U_0.

Lin ⊥ Lin vs σ⁺–σ⁻

In lin ⊥ lin, the polarization goes from linear to circular and back over half a wavelength, which is what creates the light-shift hills above. In σ^+-σ^- the light is linearly polarized everywhere, with the axis rotating along the beam: there are no hills, and cooling comes instead from a motion-induced population imbalance between m_F sublevels (Dalibard & Cohen-Tannoudji 1989). Both reach comparable sub-Doppler temperatures.

Temperature Limit

PG cooling does not have a simple closed-form temperature limit, it depends on the spin F, the laser intensity, and detuning. The temperature scales approximately as:

PG cooling temperature (approximate)

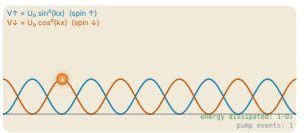
$$k_B T \sim \frac{\hbar \Omega^2}{|\Delta|} \quad (\sim U_0, \text{ the light shift})$$

where Ω is the Rabi frequency and Δ the detuning. Lowering the laser power reduces U_0 and thus the temperature, but eventually the atom is not cooled fast enough and the technique breaks down. Typical results: $T \sim 1\text{–}10\,\mu\text{K}$ for alkalis.

> **Key difference from sideband cooling:** PG cooling does not resolve individual motional sidebands. It reaches sub-Doppler temperatures through the Sisyphus mechanism but does not prepare the motional ground state. It is usually the first step before sideband or GM cooling in tweezer experiments.

Spin Cooling Model for PG

Phatak et al. show that PG cooling is captured by the same unified spin model as GM and EIT. After adiabatically eliminating the excited state, the effective ground-state dynamics has: (1) a position-independent Hamiltonian giving spin-dependent light shifts, (2) position-independent collapse operators for optical pumping, and (3) position-dependent collapse operators (recoil heating in the dark state). The ratio γ_-/γ_+, how asymmetric the pumping is toward the cooling spin state, controls the final temperature.



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Sisyphus cooling animation. The blue and orange curves show the sinusoidal light-shift potentials for the two spin states (↑ and ↓), which are π out of phase. The atom (cyan dot) climbs a potential hill, then is optically pumped (green arrow) to the bottom of the other spin state's potential. Energy equal to the light shift U_0 is dissipated at each pump event.

18.7 Gray Molasses Cooling

Dark states

Gray molasses (GM) cooling uses transitions where coherent superpositions of ground spin states can be completely dark to the light. Unlike Sisyphus cooling where all spin states scatter, in GM the atom spends most of its time in a non-scattering dark state. Cooling occurs when the dark state becomes bright only at the turning points of the classical motion, and the atom is then pumped back to a dark state with less energy.

Dark States and the $F \to F' = F$ Transition

For a $J \to J' = J$ (or $F \to F' = F$) transition driven with blue-detuned light, coherent superpositions of ground m_F states can be constructed that have zero dipole matrix element to all excited states. These are the "dark states." An atom pumped into a dark state stops scattering. The bright state (orthogonal superposition) has a positive light shift, sitting higher in energy. With blue detuning, bright states are trapped in regions of high intensity and dark states live in low-intensity regions.

The Unified Spin Model (Eq. 72)

After adiabatic elimination and expanding to first order in η, Phatak et al. derive the effective spin-cooling Hamiltonian:

GM spin Hamiltonian

$$H = \nu a^\dagger a + \Delta_s |\uparrow\rangle\langle\uparrow| + \Omega_R k\hat{x}\, \hat{F}_x$$

Collapse operators (pumping)

$$L_+ = \sqrt{\gamma_+}\,\hat{F}_+(1 + ik\hat{x}), \quad L_- = \sqrt{\gamma_-}\,\hat{F}_-(1 + ik\hat{x})$$

COOLING LIMIT (EQ. 72)

The steady-state mean phonon number in the spin model is:

Gray molasses steady state (Eq. 72)

$$\langle n \rangle = \frac{s}{1 - s}, \quad s = \frac{4\eta^2\gamma_h + \gamma_-}{\eta^2\gamma_h + \gamma_+}$$

where γ_h is the heating rate from the position-dependent operators, γ_+ is the pumping rate toward the dark state (cooling), and γ_- is the residual rate back toward the bright state (limiting factor).

> **Fundamental limit:** For $\gamma_- \to 0$ (perfectly dark lower state), $s \to 4\eta^2\gamma_h/\gamma_+$. The temperature scales as η^2, the same functional dependence as all two-photon cooling techniques. Achieving $\langle n \rangle \ll 1$ requires $\eta^2\gamma_h/\gamma_+ \ll 1$.

Gray Molasses Cooling Limit Calculator

Enter the spin model parameters to compute $\langle n \rangle_{ss}$ from Eq. 72. The ratios γ_-/γ_+ and γ_h/γ_+ characterize how asymmetric the pumping is and how much recoil heating the dark state experiences.

s parameter

0.1765

⟨n⟩ steady state

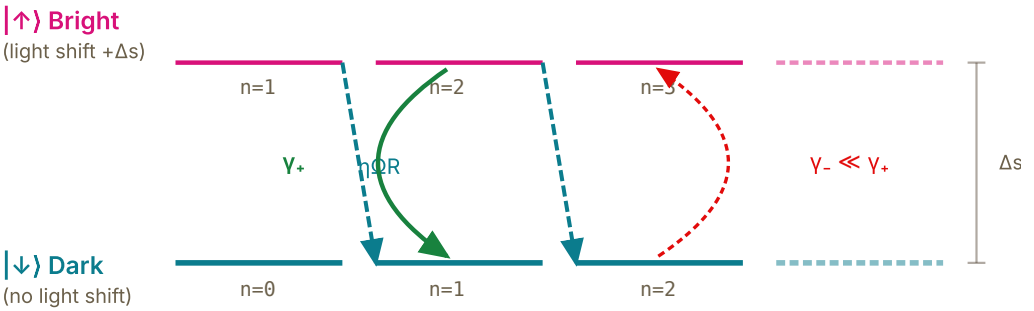
0.2143

Ground state fraction

82.4%

Cooling possible?

Yes (s < 1)



Spin cooling model for gray molasses. Blue-detuned light creates a differential light shift Δs between the bright $|\uparrow\rangle$ and dark $|\downarrow\rangle$ states. Optical pumping γ_+ transfers atoms to the dark state. The coherent Raman coupling $\eta\Omega_R$ (dashed) links $|\uparrow,n\rangle$ to $|\downarrow,n-1\rangle$, removing one motional quantum per Raman+pump cycle. The residual back-pumping rate $\gamma_- \ll \gamma_+$ sets the temperature floor.



Gray molasses dark-state cooling animation

Dark-state cooling loop. Raman coupling transfers population from a dark motional state into the neighboring bright state, excitation promotes it to the excited state, and spontaneous decay returns it to the lower dark motional state.

✨

18.8 EIT and Λ-Enhanced Gray Molasses

Two-photon dark states

Electromagnetically induced transparency (EIT) and Λ-enhanced gray molasses (Λ-GM) improve over standard GM by using two coherent laser beams to create dark states from superpositions of different ground hyperfine manifolds. This "Λ-configuration" makes the dark state darker than a single-photon GM dark state, reducing γ_-/γ_+ and enabling lower temperatures.

The Λ-Configuration

Two ground states $|g_1\rangle$ and $|g_2\rangle$ (from different hyperfine levels) are each coupled to a common excited state $|e\rangle$ by laser fields Ω_1 and Ω_2 with detunings Δ_1 and Δ_2. The key is to transform to the bright/dark basis:

Bright and dark states (Eq. 76)

$$|g_B\rangle = \frac{\Omega_1|g_1\rangle + \Omega_2|g_2\rangle}{\Omega_{rms}}$$
$$|g_D\rangle = \frac{-\Omega_2|g_1\rangle + \Omega_1|g_2\rangle}{\Omega_{rms}}$$

RAMAN COUPLING BRIGHT ↔ DARK (EQ. 85)

$$\Omega_r = \frac{\Omega_1\Omega_2}{2\Delta}$$

EIT RESONANCE CONDITION

$\Delta_1 = \Delta_2$ (two-photon resonance)

Calculator output at the page's default settings.

Why EIT/Λ-GM is Better than GM

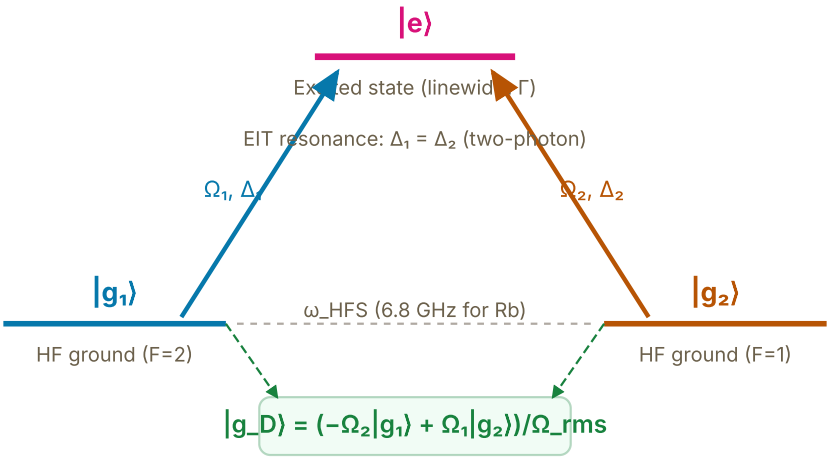
In standard GM, the dark state is a single ground-state spin superposition that is dark only at the leading order in the coupling. In the Λ-configuration, the dark state $|g_D\rangle$ is dark to both laser fields simultaneously due to destructive interference, the contribution from $|g_1\rangle$ and $|g_2\rangle$ cancel exactly. This makes $\gamma_- \to 0$ much more effectively, directly reducing the steady-state $\langle n \rangle$.

The same spin-cooling model (Eq. 72) applies, but with significantly smaller γ_-/γ_+, extracted from simulations of ^{87}Rb: $\gamma_-/\gamma_+ \approx 0.006$ for Λ-GM vs. $\gamma_-/\gamma_+ \approx 0.11$ for standard GM.

Coherence requirement: EIT/Λ-GM requires the two lasers to be phase-coherent to within the two-photon linewidth. Any relative phase noise $\delta\phi$ with frequency ω_ϕ drives coupling between $|g_D\rangle$ and $|g_B\rangle$ at rate $\sim (\Delta_1 - \Delta_2)$. For 100 kHz traps, coherence within ~1 kHz is required, naturally satisfied if both frequencies come from the same laser via an AOM or EOM.

EIT Heating Resonance

When $\Delta_1 - \Delta_2 = 2\nu$ (the blue sideband two-photon condition), a heating resonance appears. The optimal cooling occurs at $\Delta_1 = \Delta_2$ when the bright-state light shift equals the trap frequency, $\Omega^2/\Delta = \nu$.



The Λ-level configuration for EIT / Λ-GM cooling. Two coherent laser beams with Rabi frequencies Ω_1 and Ω_2 address two hyperfine ground states $|g_1\rangle$ and $|g_2\rangle$ to a common excited state $|e\rangle$. When $\Delta_1 = \Delta_2$ (two-photon resonance), destructive interference creates the dark superposition $|g_D\rangle = (-\Omega_2|g_1\rangle + \Omega_1|g_2\rangle)/\Omega_{rms}$ that does not couple to either laser. This dark state is much darker than a single-beam GM dark state, leading to lower γ_-/γ_+ and deeper cooling.

18.9 Raman Sideband Cooling

Ground state prep

Raman sideband (RSB) cooling is one of the highest-fidelity routes to motional ground-state preparation for neutral atoms in tweezers. Unlike PG, GM, and EIT, which all use the same beams for both cooling and optical pumping, RSB separates these roles. Far-detuned Raman lasers drive the motional transition while a carefully designed near-resonant optical pump-

ing beam uses angular momentum selection rules to suppress recoil heating from the dark state.

The RSB Cooling Cycle (Fig. 9)

Step 1 (Raman): Two far-detuned counter-propagating lasers at Rabi frequencies Ω_1, Ω_2 and detuning Δ drive the two-photon transition $|g_1, n\rangle \to |g_2, n-1\rangle$ via stimulated Raman. The Raman coupling:

Effective Raman coupling (Eq. 86–87)

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

$$\langle g_1, n-1 | \, H_{\text{eff}} \, | g_2, n \rangle = -\Omega_R \, \eta \sqrt{n}$$

(Optical pumping):

A separate near-resonant beam pumps $|g_2\rangle \to |g_1\rangle$ via spontaneous emission. The key is that angular momentum selection rules ($\Delta m_F = 0, \pm 1$) can be exploited so that $|g_1\rangle$ remains dark to the pump beam. Spontaneous emission in the Lamb-Dicke regime keeps the atom at $n-1$ rather than n.

Why RSB Beats EIT/Λ-GM

In EIT and Λ-GM, the counter-propagating lasers that provide cooling also do optical pumping. This means position-dependent collapse operators act on the dark state: the atom in $|g_D\rangle$ scatters photons from the cooling beams at rate $\propto \eta^2 \gamma_h$. This heating rate is fundamental and sets the η^2 floor in Eq. 72.

In RSB, the optical pump beam is completely separate and designed to be dark to $|g_1\rangle$. The far-detuned Raman beams have negligible scattering ($\Gamma_R \propto \Omega^2/\Delta^2 \approx 0$). So the position-dependent collapse operator from pumping is absent, removing the $\eta^2 \gamma_h$ floor entirely.

RSB spin model (Eq. 88)

$$H = \nu a^\dagger a + (\Delta_1 - \Delta_2)|g_2\rangle\langle g_2| + \Omega_R k \hat{x} \, \hat{F}_x$$

$$L_+ = \sqrt{\gamma_+} \, \hat{F}_+ (1 + ik\hat{x})$$

Practical limits: RSB cooling is limited not by the fundamental theory but by technical factors: (1) coherence time between $|g_1\rangle$ and $|g_2\rangle$ is typically a few ms due to magnetic field fluctuations; (2) the carrier transition ($\Delta n = 0$) is driven with Rabi frequency Ω_R, a factor

$1/(\eta\sqrt{n})$

larger than the red sideband, and can populate higher n if not carefully timed; (3) phase coherence of both Raman lasers is required.

Comparison: PG / GM / EIT / RSB at a Glance

The progression from PG to RSB represents successive improvements in how dark the dark state is to recoil heating:

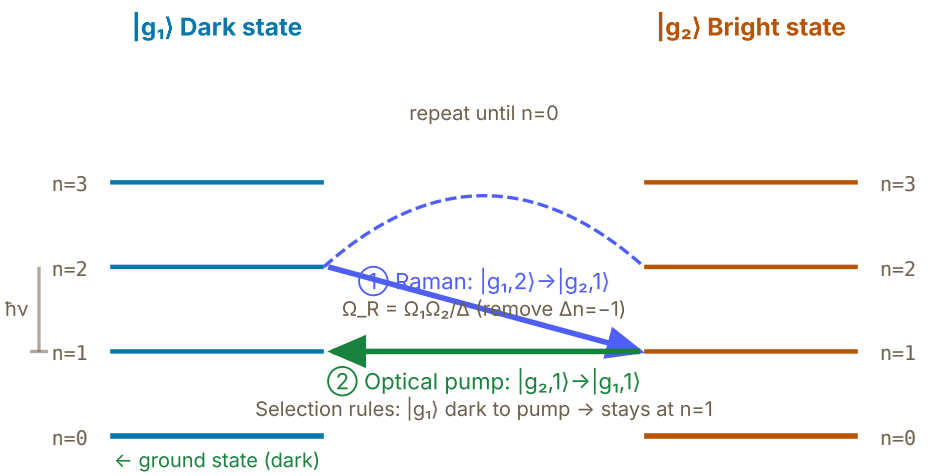
- **PG:** No dark state; all spin states scatter. Sisyphus mechanism provides cooling. Limit: $T \sim U_0 \sim \hbar\Omega^2/|\Delta|$.
- **GM:** Single-laser dark state. Position-dependent pumping heats the dark state at rate $\eta^2\gamma_h$. Limit: $\langle n \rangle \propto \eta^2\gamma_h/\gamma_+$.
- **EIT/Λ-GM:** Two-laser dark state via interference. γ_-/γ_+ reduced by $\sim 20\times$. Same $\eta^2\gamma_h$ floor from counter-propagating pump.
- **RSB:** Separate pump beam, dark to $|g_1\rangle$. No $\eta^2\gamma_h$ floor. Limit: technical coherence and off-resonant carrier driving.

Phase Coherence Requirement

Both RSB and Λ-GM require the two Raman lasers to maintain phase coherence. Any relative frequency noise at timescales faster than the coherence time τ_c drives unwanted bright-dark mixing. The requirement is:

$$\delta\nu_{\text{laser}} \lesssim \eta\,\Omega_R \lesssim \eta\,\nu$$

For $\eta = 0.1$ and $\nu = 100\,\text{kHz}$, this requires laser coherence at the $\sim 10\,\text{kHz}$ level, readily satisfied by generating both frequencies from the same laser using an AOM or microwave EOM.



RSB two-step cooling cycle (Fig. 9 of Phatak et al.). Step ①: two far-detuned Raman beams drive $|g_1,n\rangle \to |g_2,n-1\rangle$ via stimulated two-photon transition, reducing n by 1. Step ②: a near-resonant optical pump returns $|g_2\rangle \to |g_1\rangle$. Because $|g_1\rangle$ is dark to the pump beam by angular momentum selection rules, no recoil heating enters the dark state, eliminating the $\eta^2\gamma_h$ heating floor present in GM and EIT.

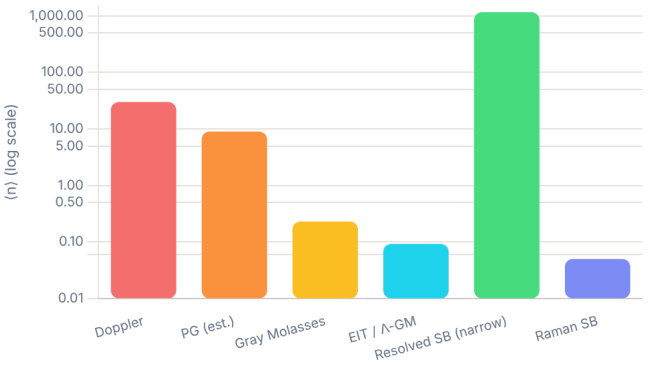
📊

18.10 Technique Comparison

Summary

All techniques are benchmarked for the same atom and trap. The chart and table below show theoretical $\langle n \rangle$ limits using parameters representative of an atom in a 100 kHz optical tweezer, with η = 0.1 by default (^{87}Rb at 780 nm actually has $\eta \approx 0.19$; move the slider). Use the calculator to explore how the limits shift for your atom and trap.

Steady-state ⟨n⟩ for different cooling techniques (default η = 0.1, Γ of Rb D2)



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

TECHNIQUE	⟨N⟩ LIMIT	REGIME	BEAMS	DARK STATE	KEY REQUIREMENT	TYPICAL ⟨N⟩ (RB, 100 KHZ)
Doppler	$(1+\alpha)\Gamma/(4\nu) - \frac{1}{2}$	UNRESOLVED	1 color	None	$\Delta = -\Gamma/2$	~30
Resolved SB	$\frac{1}{4}(\Gamma/\nu)^2(\alpha+\frac{1}{4})$	RESOLVED $N \gg \Gamma$	1 (narrow line)	$\lvert n=0\rangle$ dark	$\nu \gg \Gamma$	~0.01 (Sr/Yb)
Polarization Gradient	$\sim U_0/k_B$	SUB-DOPPLER	2 (counter-prop.)	Partial	Power & detuning tuning	2–5
Gray Molasses	$s/(1-s)$, $s \approx 4\eta^2\gamma_h/\gamma_+$	SUB-DOPPLER	2 (counter-prop., blue det.)	Spin dark state	$F \to F' = F$ transition	0.3–1
EIT / Λ-GM	Same as GM, smaller γ_-	DEEP SUB-DOPPLER	2 colors, coherent	Two-photon dark state	Phase coherence; $\Delta_1 = \Delta_2$	0.05–0.3
Raman Sideband	Technical (coherence time)	GROUND STATE	2 Raman + 1 pump	Separate pump (sel. rules)	Phase coherence; $\Omega_R < \nu$	<0.05

Insight 1: η² Scaling

All two-photon cooling techniques have a temperature floor scaling as η^2. Tighter traps (higher ν, smaller η) directly improve the final temperature. This is why modern tweezer experiments push toward 1 MHz trap frequencies.

Insight 2: Pumping Asymmetry

The ratio γ_-/γ_+ is the key material parameter. A large asymmetry means the atom quickly returns to the cooling spin state after each cycle. GM achieves $\gamma_-/\gamma_+ \approx 0.11$; Λ-GM achieves ≈ 0.006; RSB achieves ≈ 0 by design.

Insight 3: Dark State Quality

The final temperature is set by how dark the dark state is to recoil heating. Single-laser dark states (GM) have a position-dependent heating rate $\eta^2\gamma_h$. Two-laser dark states (EIT) reduce γ_-. Separate-pump RSB eliminates position-dependent heating from the optical pump entirely.

Primary reference: Phatak, S.S. et al., "A Generalized Theory for Optical Cooling of a Trapped Atom with Spin," *Phys. Rev. A* 110, 043116 (2024) (arXiv:2406.19153). All equations on this page are from that paper. The full master equation derivation, adiabatic elimination, and detailed comparison of all techniques for ^{87}Rb are there.

18.11 Papers, Textbooks & Videos

Go deeper

The resources below are organized by how deep you want to go, from a 10-minute video to graduate-level review articles. Start with the videos if you want intuition first, then move to the textbook chapters, then the original papers once you want to work through the math.

🖥️

Hood Lab, Single-atom Cooling Simulations

github.itap.purdue.edu/hoodlabpurdue/Trapped-atom-cooling-simulations

The companion simulation code for Phatak et al. (2024). Contains master equation solvers for all cooling techniques discussed on this page, resolved and unresolved sideband, PG, GM, Λ-GM, EIT, and RSB, validated against the paper's figures. Requires Python + QuTiP. Run your own simulations, explore parameter sweeps, and reproduce the paper's results.

Purdue GitHub →

Videos, Build the Intuition

🏅

William D. Phillips, Nobel Lecture (1997)

Nobel Prize · "Nearly Absolute Zero: The Story of Laser Cooling"

The clearest physical intuition for laser cooling from one of its inventors. Phillips walks through Doppler cooling, optical molasses, and sub-Doppler mechanisms with hand-drawn diagrams. The written lecture and slides are freely available on the Nobel Prize site. Essential first read.

Foundational · Nobel lecture

🎥

How We Make the Coldest Things in the Universe

Veritasium · ~10 min

Visual, accessible introduction to laser cooling for a general science audience. Good for building the core mental model of photon momentum kicks before diving into the physics. Derek Muller visits NIST to see real laser cooling in action.

Introductory · Visual

🎓

The Coolest Stuff in the Universe — W. D. Phillips

NIST / YouTube · ~60 min public lecture

Phillips delivers his famous public lecture on laser cooling — much more engaging than reading the Nobel text. Covers Doppler cooling, the optical molasses surprise, sub-Doppler temperatures, and BEC. Highly recommended before tackling the technical papers.

Foundational · Full lecture

🎓

Gray Molasses & EIT Cooling — Ye Lab, JILA

JILA / Jun Ye group · Research overview + papers

The Ye Lab at JILA pioneered many sub-Doppler techniques for alkaline-earth atoms. Their research pages link to key papers on gray molasses, EIT cooling, and sideband cooling in optical lattices. Excellent for connecting the theory on this page to experiment.

Graduate level · Sub-Doppler · EIT

Textbooks & Free Notes

📄

Quantum and Atom Optics, Daniel Steck

steck.us/teaching · Free PDF · Updated continuously

The definitive free graduate textbook for AMO physics. Chapters 9–11 cover laser cooling comprehensively: Doppler cooling, optical molasses, the Wigner-Weisskopf theory, and sub-Doppler mechanisms. Also covers the master equation formalism used by Phatak et al. The spin tensor coefficients in the paper's Appendix reference Steck directly.

Free PDF · Most recommended

📋

Quantum Dynamics of Single Trapped Ions

Leibfried, Blatt, Monroe, Wineland · Rev. Mod. Phys. 75, 281 (2003)

The standard review for sideband cooling of trapped ions, and the physics is identical for neutral atoms in tweezers. Covers the Lamb-Dicke regime, motional state preparation, sideband thermometry, and ground-state cooling in rigorous detail. Highly cited (5000+).

Review article · Sideband cooling

📋

The Semiclassical Theory of Laser Cooling

Stenholm · Rev. Mod. Phys. 58, 699 (1986)

The foundational review that established the theoretical framework for laser cooling, before sub-Doppler cooling was even discovered. Covers the semiclassical and quantum treatments of radiation pressure, the Fokker-Planck approach to cooling, and the Doppler limit.

Review article · Doppler / foundations

📋

Laser Cooling Below the Doppler Limit by Polarization Gradients

Dalibard & Cohen-Tannoudji · JOSAB 6, 2023 (1989)

The original paper predicting and explaining sub-Doppler cooling via polarization gradients, the Sisyphus mechanism. Written the same year the effect was discovered experimentally. The theoretical framework here directly underlies PG and GM cooling.

Original paper · PG cooling

Key Experimental Papers

⚛️

A Generalized Theory for Optical Cooling of a Trapped Atom with Spin

Phatak et al. · PRA 110, 043116 (2024) · arXiv:2406.19153

The primary reference for this entire page. Derives the unified spin-cooling framework unifying PG, GM, Λ-GM, EIT, and RSB under a single formalism. All equations on this page are from here.

Theory · Unified framework

🧲

Imaging a ^{6}Li Atom in an Optical Tweezer 2000 Times with Λ-Enhanced Gray Molasses

Blodgett, Peana, Phatak et al. (Hood Lab) · PRL 131, 083001 (2023)

Experimental demonstration of Λ-GM cooling in a single-atom optical tweezer, the experiment at Purdue that motivated the generalized theory paper. Uses ^{6}Li with the $F = 1/2, 3/2 \to F' = 3/2$ (D1) Λ-configuration.

Experiment · Λ-GM · Hood Lab

🔬

Gray-Molasses Optical-Tweezer Loading: Controlling Collisions for Scaling Atom-Array Assembly

Brown, Thiele, Kiehl et al. (Regal Lab) · PRX 9, 011057 (2019)

Uses blue-detuned gray molasses during tweezer loading to control the light-assisted collisions that normally limit single-atom loading to about half the traps, raising the loading probability well above 50% (rubidium, Regal lab).

Experiment · GM in tweezers

🎯

Cooling a Single Atom in an Optical Tweezer to Its Quantum Ground State

Kaufman, Lester, Regal · PRX 2, 041014 (2012)

The landmark paper demonstrating RSB cooling of a single neutral atom in a tweezer to $\langle n \rangle < 0.05$, the first ground-state preparation in a tweezer. Foundational for the neutral-atom quantum computing platform.

Experiment · RSB · Landmark

🎯

Coherence and Raman Sideband Cooling of a Single Atom in an Optical Tweezer

Thompson, Tiecke, Zibrov et al. · PRL 110, 133001 (2013)

Independent demonstration of RSB in a tweezer. Includes detailed characterization of the coherence properties and an analysis of the technical limits discussed in the theory section above.

Experiment · RSB · Coherence

✨

Ground State Laser Cooling Using Electromagnetically Induced Transparency

Morigi, Eschner, Keitel · PRL 85, 4458 (2000)

The theoretical proposal of EIT cooling (demonstrated soon after with a trapped ion by Roos et al., PRL 85, 5547 (2000)). Establishes the bright/dark basis picture and the resonance condition. Directly relevant to EIT cooling of neutral atoms in tweezers.

Theory · EIT cooling

CHAPTER 18 SUMMARY

Ground-state cooling suppresses Doppler dephasing and thermal position spread inside each tweezer. We can now put loading, imaging, rearrangement, and readout together into a complete experimental cycle.

Problems for Chapter 18

1. USE THE TOOL Compute the Lamb–Dicke parameter for ^{87}Rb in a 100 kHz trap with 780 nm light.
2. BY HAND Check problem 1 with $\eta = k\sqrt{\hbar/2m\omega}$ and explain why a lighter atom in the same trap is harder to cool to the ground state.
3. DESIGN Use the technique comparison to choose a ground-state cooling method for ^{6}Li and for ^{133}Cs in 100 kHz tweezers, and justify each with the relevant linewidth and Lamb–Dicke parameter.

Use the tool: on the interactive page. By hand: from the formulas in this chapter. Design: a decision a lab would make. Worked solutions are available to instructors on request (see Part V).

CHAPTER 19 · EXPERIMENT STEP 07

19

End-to-End Tweezer Experimental Cycle

🧩 Virtual Tweezer Lab · Simulator

Every neutral-atom quantum computer repeats the same cycle: load atoms at random, photograph them, drag them into a perfect block, check, entangle, measure. Step through one simulated shot on the camera, then run a thousand and see which part of the cycle spoils them.

> *A real neutral-atom shot chains together stochastic loading, initial imaging, AOD rearrangement, verification imaging, quantum operations, and final state readout—each step carrying finite loss and duration.*

IN THIS CHAPTER Prerequisites: Ch. 8, Ch. 16, Ch. 18

- Trace the full sequence of a tweezer-array shot from ~50% stochastic loading to defect-free array assembly.
- Quantify how vacuum lifetime, imaging loss, and move efficiency compound into the overall shot success probability.
- Identify which stage bottlenecks repetition rate and array yield as atom counts scale.

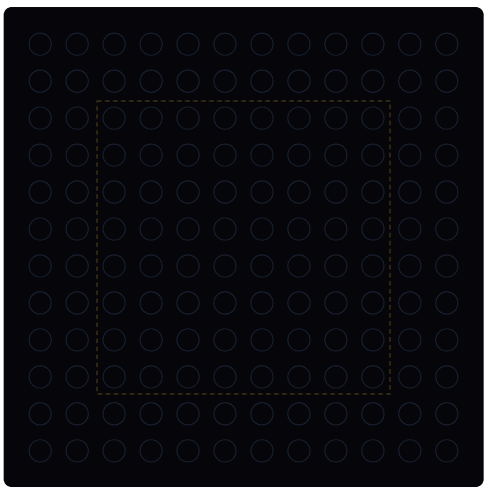
WEBSITE COMPANION Live interactive page ↗ · Week 10

01 · ONE SHOT

19.1 The experimental cycle, on the camera

Interactive

1. Load
2. Image
3. Rearrange
4. Verify
5. Entangle
6. Read out



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Camera · idle

empty tweezer atom (camera / tracked) moving tweezer loss / hole Rydberg pulse target block

Shot log

Tweezers / target
144 / 8×8

Atoms loaded (truth)
—

Detected (≥ 14 photons)
—

False detections / missed
—

Lost while imaging
—

Moves · sequence time
—

Lost in moves / collisions
—

Background-gas losses
—

Holes in target
—

Complete Bell pairs
—

Pairs reading equal bits
—

02 · YOUR APPARATUS

19.2 Settings

Green tags are published numbers; amber tags are illustrative defaults to replace with your own system's.

Threshold ≥ 14 photons · false-detection rate 2.9×10^{-8} · miss rate 1.7×10^{-7} · expected loaded 72.0 of 144, need 64 · P(enough atoms) = 92.2%

03 · A THOUSAND SHOTS

19.3 Where the defects come from

The same simulation without the animation, repeated. Every number below comes from these settings.

Interactive

1000 shots done

DEFECT-FREE SHOTS (MONTE CARLO)	WITHOUT REARRANGING, P^{K^2}	ENOUGH ATOMS LOADED
32.2% ± 1.5% (1000 shots)	**5.4×10^{-20} chosen 8×8 block happens to be full**	**92.2%** P(Bin(144, 0.50) ≥ 64): the ceiling for rearranging
MOVES PER SHOT	HOLES PER SHOT	BELL PAIRS READING EQUAL BITS
31.6 sequence 32 ms at 1 ms per move	**1.34 of 64 target sites**	**99.7%** expected F + (1−F)/2 = 99.8%

Calculator output at the page's default settings.

Shots spoiled by…

Not enough atoms

8.6%

False detection

0.0%

Lost while imaging

46.4%

Lost in a move

26.9%

Collision with a missed atom

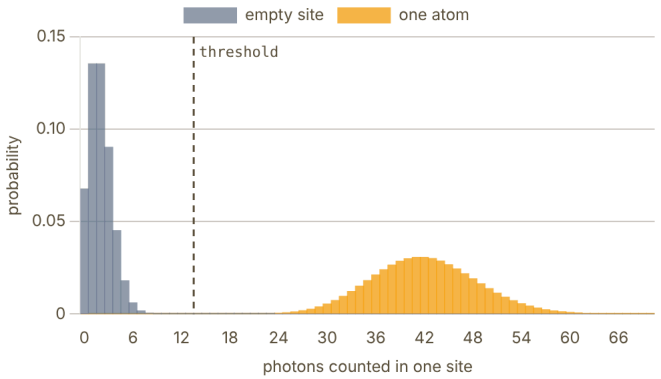
0.0%

Background-gas loss

15.1%

A shot counts under every cause that left at least one hole in the target, so the bars can add to more than 100%.

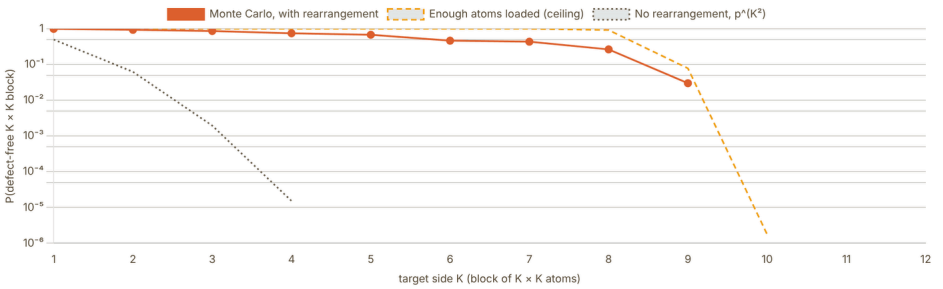
Photon counts per site



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Empty sites are Poisson with mean 2.0; occupied ones with mean 42.0. Bars are weighted by 1 − p and p. Threshold ≥ 14: 2.9×10^{-8} of empty sites read as full, 1.7×10^{-7} of atoms are missed.

Defect-free probability vs target size



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Monte Carlo: 200 shots per size with the settings above (other K values use the same array and losses). Sizes with no defect-free shot (K = 10–12) are below about 1/200 and are not plotted. Curves below 10^{-6} are cut off.

19.4 The exponential wall

Without rearrangement

Each tweezer loads independently with probability p, so a chosen $K \times K$ block is full with probability p^{K^2}. At $p = 0.5$ an 8×8 block is full once in $2^{64} \approx 1.8 \times 10^{19}$ tries. Blockade loading caps p near one half (Schlosser 2001); enhanced loading raises it, but even $p = 0.89$ gives $0.89^{64} \approx 6 \times 10^{-4}$.

Rearrangement turns this around: take an image, then use a steerable tweezer to carry spare atoms into the holes (Barredo et al. 2016; Endres et al. 2016). Now the hard limit is having enough atoms, $P\big(\mathrm{Bin}(M^2, p) \geq K^2\big)$, which is close to 1 once the array holds comfortably more than K^2/p traps.

With rearrangement, a slower wall

Each move risks the atom being carried, and the whole sequence takes time while every atom can be knocked out by background gas. If a shot needs n moves of duration t, then roughly

$$P_{\text{defect-free}} \approx (1-\varepsilon)^{n}\, e^{-K^2 n t/\tau} \times (\text{imaging terms}),$$

and n grows with K^2. The decay is still exponential in the number of atoms, just with a far smaller exponent. Try it: raise K with 1% loss per move and watch the Monte Carlo fall away from the “enough atoms” curve.

Beyond one shot: keep the array full

The newest machines stop treating each shot as a fresh start. Norcia et al. filled 1225-site ^{171}Yb arrays incrementally from a repeatedly refilled reservoir, reaching 99% per-site occupancy (PRX *Quantum* 5, 030316, 2024). Chiu et al. kept an array of over 3,000 atoms running for more than two hours by reloading 300,000 atoms per second into tweezers (*Nature* 646, 1075, 2025). Manetsch et al. trapped 6100 Cs atoms in one array (*Nature*, 2025). This lab models the single-shot cycle those systems are built from.

19.5 Assumptions

- **Loading:** each tweezer holds one atom with probability p, independently; never two.
- **Imaging:** a site's count is Poisson with mean N_{bg} (empty) or $N_{bg} + N_{at}$ (occupied). The threshold is the integer that minimises the sum of the false-positive and false-negative rates. An occupied site survives the image with the set probability.
- **Assignment:** empty target sites, taken from the centre outward, each claim the nearest unused atom detected outside the block. Atoms already detected inside the block stay. Paths that cross other traps are allowed.
- **Moves:** a move from a false-positive or already-lost site delivers nothing. A real atom is lost with probability ε. If it lands on a site that held an undetected atom, light-assisted collisions eject both.
- **Vacuum:** after the moves, every atom survives with $e^{-T/\tau}$, where T = moves × time per move.
- **Defect-free** is judged from the true occupancy. The verification image is drawn for show; its own errors are not counted.
- **Bell pairs:** neighbouring columns in the block are paired. A pair with both atoms is a good Bell pair with probability F, giving equal Z-basis bits; otherwise its bits are random. The expected agreement is $F + (1 - F)/2$. Agreement in one basis is necessary but does not prove the fidelity.

Sources

Schlosser et al. (2001), *Nature* 411, 1024 — sub-Poissonian loading of single atoms in a microscopic dipole trap Brown et al. (2019), *Phys. Rev.* X 9, 011057 — gray-molasses optical-tweezer loading Barredo et al. (2016), *Science* 354, 1021 — atom-by-atom assembler of defect-free 2D arrays Endres et al. (2016), *Science* 354, 1024 — atom-by-atom assembly of defect-free 1D arrays Evered et al. (2023), *Nature* 622, 268 — high-fidelity parallel entangling gates on a neutral-atom quantum computer Norcia et al. (2024), PRX *Quantum* 5, 030316 — iterative assembly of ^{171}Yb atom arrays with cavity-enhanced optical lattices Chiu et al. (2025), *Nature* 646, 1075 — continuous operation of a coherent 3,000-qubit system Manetsch et al. (2025), *Nature* — a tweezer array with 6100 highly coherent atomic qubits

CHAPTER 19 SUMMARY

Thinking in terms of the full experimental cycle reveals how small per-step losses compound at scale. To close Part II, we gather the everyday optical, atomic, and angular-momentum calculators used at the bench.

Problems for Chapter 19

1. USE THE TOOL Run 1,000 shots and sweep the target block size. At what size does the defect-free success probability fall below 50%?
2. BY HAND If each atom survives the sequence with probability p, a block of N atoms succeeds with probability ≈ p^N. Use your data to estimate p, and predict the success probability for N = 100.
3. DESIGN You need a 100-atom defect-free block with at least 50% probability per shot. Which settings (loss per move, survival of one image, vacuum lifetime) would you improve first, and by how much?

Use the tool: on the interactive page. By hand: from the formulas in this chapter. Design: a decision a lab would make. Worked solutions are available to instructors on request (see Part V).

CHAPTER 20 · EXPERIMENT STEP 08

20

Everyday AMO Reference Calculations

🧮 Quick Lab Console · Calculator

A console for the everyday numbers that decide whether an AMO setup is plausible: beams, RF power, recoil, trap frequencies, cavity scales, angular momentum, and unit conversions.

> *At the optical bench, researchers constantly convert between Gaussian beam waists, Rayleigh ranges, dBm/mW RF powers, Zeeman shifts, cavity FSRs, and Clebsch-Gordan transition strengths.*

IN THIS CHAPTER Prerequisites: Ch. 1, Ch. 8, Ch. 11

- Compute Gaussian beam propagation, fiber-coupling mode matching, and focal spot sizes rapidly.
- Calculate recoil velocities, Doppler shifts, Zeeman splittings, optical trap frequencies, and EOM/AOM parameters.
- Evaluate Clebsch-Gordan coefficients and relative dipole transition strengths for multi-level atomic manifolds.

WEBSITE COMPANION Live interactive page ↗ · Notebook: Open in Colab · download lab-calculators.ipynb · Week 4

Quick lab workflows

20.1 Pick the calculation by the lab question.

The console is organized so you can move from a practical experimental task to the right calculator family without remembering where every formula lives.

01 **Design beam optics**

Gaussian beam propagation, fiber coupling, collimation/focusing, NA, and telescope magnification.

02
Check an AOM/RF chain
Convert dBm, account for gain/loss, estimate deflection, and sanity-check drive powers.

03
Estimate atom scales
Recoil, Doppler limits, Zeeman shifts, saturation intensity, and velocity classes.

Calculator output at the page's default settings.

04 **Design a tweezer**

Estimate waist, Rayleigh range, trap frequencies, Lamb-Dicke parameter, and cavity-adjacent scales.

05
Couple angular momenta
Clebsch-Gordan coefficients, selection rules, and the Wigner-Eckart reference.

06
Check modulator sidebands
EOM/AOM sideband depths via Bessel functions, plus double-pass shift and efficiency.

07
Convert lab units fast
Use the unit converter when notebook units, vendor units, and paper units do not agree.

Calculator output at the page's default settings.

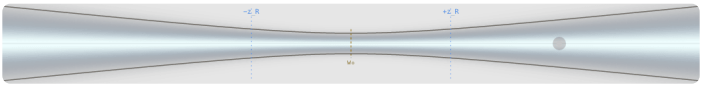
7 chapters on this page. Open one from a card, or from the chapter bar below, which stays pinned while you read.

20.2 Optics & Beam Calculations

Gaussian beam propagation, collimation, focusing, NA, telescope magnification.

Gaussian Beam Profile

Side-view of beam propagation from the waist. Uses λ and w_0 from the Rayleigh Range inputs above.



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

w_0 = 950.00 µm · z_R = 2664.74 mm · λ = 1064 nm

20.3 Power, RF & AOM Calculators

mW / dBm, gain/loss chains, AOM Bragg angles, shot noise.

Shot Noise on Photodetector

Photon shot noise sets the quantum noise floor for a given detection bandwidth.

$$I_{\mathrm{dc}} = \frac{\eta e P}{h\nu} \qquad i_{\mathrm{shot}} = \sqrt{2e\, I_{\mathrm{dc}}\, \mathrm{BW}} \quad [\mathrm{A_{rms}}]$$

$$\mathrm{NEP} = \frac{i_{\mathrm{shot}}}{R} \quad [\mathrm{W}]$$

Photocurrent I_dc 5.841 µA Shot noise i_shot 1368.095 pA_rms NEP (noise equiv. power) 2342.201 pW

20.4 Atomic Physics Calculators

Recoil, Doppler, de Broglie, Zeeman, saturation intensity, Maxwell–Boltzmann velocities.

Photon Recoil

Recoil velocity, single-photon recoil energy, and two-recoil scatter heating scale. Sets the energy scale for laser cooling.

$$v_{\mathrm{rec}} = \frac{\hbar k}{m} = \frac{h}{m\lambda} \qquad E_{\mathrm{rec}} = \frac{(\hbar k)^2}{2m}$$
$$T_{\mathrm{rec}} = \frac{E_{\mathrm{rec}}}{k_{\mathrm{B}}} \qquad T_{\mathrm{scatter}} \approx \frac{2E_{\mathrm{rec}}}{k_{\mathrm{B}}}$$
$$E_{\mathrm{rec}}/h = \frac{\hbar k^2}{4\pi m} \quad [\mathrm{kHz}]$$

Common mistake, recoil factor-of-two. E_{rec} is the kinetic energy from one momentum kick $\hbar k$. A scatter event usually has absorption plus spontaneous-emission recoil, so the heating scale is often near $2E_{\mathrm{rec}}$, but the emitted photon projection depends on geometry and dimensional averaging.

Recoil velocity v_rec 3.5214 mm/s Single recoil E_rec/k_B 99.1788 nK = 0.0992 µK Abs.+em. heating scale 2E_rec/k_B 198.3576 nK Recoil frequency E_rec/h 2.0666 kHz

Thermal de Broglie Wavelength

λ_dB ≫ inter-particle spacing → quantum degeneracy. Also for matter-wave interferometry.

$$\lambda_{\mathrm{dB}} = \frac{h}{\sqrt{2\pi m k_{\mathrm{B}} T}}$$

de Broglie wavelength 15.1382 nm 0.015138 µm

Thermal Velocity Distribution

Maxwell–Boltzmann rms, most-probable, and mean speeds.

$$v_{\mathrm{rms}} = \sqrt{\frac{3k_{\mathrm{B}}T}{m}} \qquad v_{\mathrm{prob}} = \sqrt{\frac{2k_{\mathrm{B}}T}{m}} \quad [\text{most probable}]$$

$$v_{\mathrm{mean}} = \sqrt{\frac{8k_{\mathrm{B}}T}{\pi m}} \quad [\text{mean speed}]$$

v_rms 0.1369 m/s v_prob 0.1118 m/s v_mean 0.1262 m/s

20.5 Trap & Cavity Calculators

Optical tweezer frequencies, Lamb–Dicke parameter, Fabry–Pérot cavity, mode matching.

Optical Tweezer Trap Frequencies

Harmonic approximation near the trap minimum. Radial and axial frequencies from trap depth and beam waist.

$$\omega_r = \sqrt{\frac{4U_0}{mw_0^2}} \quad \text{[radial]} \qquad \omega_z = \sqrt{\frac{2U_0}{mz_R^2}} \quad \text{[axial]}$$

$$z_R = \frac{\pi w_0^2}{\lambda}$$

Radial ω_r/2π 59.24 kHz Axial ω_z/2π 14933.3 Hz = 14.9333 kHz Rayleigh range z_R 2.66 μm

Lamb–Dicke Parameter

η = k x_zpf. Resolved sideband cooling requires η ≪ 1. x_zpf = √(ℏ/2mω) is the zero-point motion.

$$x_{\text{zpf}} = \sqrt{\frac{\hbar}{2m\omega_{\text{trap}}}} \qquad \eta = k\cos\theta \cdot x_{\text{zpf}}$$

$$\langle n \rangle_{\min} \approx \frac{1}{4}\left(\frac{\Gamma}{\omega}\right)^2\left(\alpha + \frac{1}{4}\right) \quad \text{[RSB cooling limit; } (\Gamma/4\omega)^2 \text{ without recoil]}$$

x_zpf (zero-point motion) 19.49 nm Lamb–Dicke parameter η 0.1017 Regime ✅ Lamb–Dicke (η≪1)

Optical Cavity / Fabry–Pérot

FSR, finesse, linewidth, round-trip time. Used for laser locking, ULE cavities, transfer cavities.

$$\text{FSR} = \frac{c}{2nL} \qquad \mathcal{F} = \frac{\pi\sqrt{R}}{1-R}$$

$$\delta\nu = \frac{\text{FSR}}{\mathcal{F}} \quad \text{[linewidth FWHM]}$$

FSR 1.4990 GHz = 1498.96 MHz Finesse F 31414 Linewidth δν 0.0477 MHz = 47.72 kHz Round-trip time 0.667 ns

References

Foot (2005), *Atomic Physics* — trap frequencies, recoil, Zeeman, and laser cooling formulas Grimm, Weidemüller & Ovchinnikov (2000), *Adv. At. Mol. Opt. Phys. 42* — optical dipole trap depth, polarizability, Lamb-Dicke parameter Yariv (1989), *Quantum Electronics* — gaussian beam propagation, cavity mode matching, AOM optics Saleh & Teich (2019), *Fundamentals of Photonics* — fiber coupling, telescope design, beam waist transformations NIST CODATA 2018 — fundamental constants (ℏ, k_B, c, m_e, μ_B) used throughout Steck alkali data sheets — species-specific σ_0, I_sat, Γ, recoil energy for Rb, Cs, Li, Na, K

20.6 Clebsch–Gordan Coefficients

⟨$j_1 m_1$; $j_2 m_2$ | JM⟩ via Racah formula: exact ±√(p/q) results, CG tables, and Wigner–Eckart reference.

$$|J,M\rangle = \sum_{m_1,m_2} \langle j_1 m_1;\, j_2 m_2 \mid JM\rangle \, |j_1,m_1\rangle |j_2,m_2\rangle$$

Selection rules, coefficient is zero unless:

- $M = m_1 + m_2$ (z-component conservation)
- $|j_1 - j_2| \le J \le j_1 + j_2$ (triangle rule)
- $|m_i| \le j_i$ for $i = 1, 2$

$$J \in \{|j_1 - j_2|,\ |j_1 - j_2| + 1,\ \ldots,\ j_1 + j_2\}$$

$$\sum_J (2J+1) = (2j_1+1)(2j_2+1) \quad \text{[dimension check]}$$

Exchange: $\langle j_1 m_1; j_2 m_2 | JM\rangle = (-1)^{j_1+j_2-J} \langle j_2 m_2; j_1 m_1 | JM\rangle$
Time-rev: $\langle j_1 m_1; j_2 m_2 | JM\rangle = (-1)^{j_1+j_2-J} \langle j_1 -m_1; j_2 -m_2 | J -M\rangle$

CG Table for fixed j_1, j_2, J

Rows = m_1, columns = m_2. Cell: $\langle j_1 m_1; j_2 m_2 \mid J, m_1+m_2\rangle$. Keep $j \le 4$ for fast computation.

$M_1 \setminus M_2$	-1/2	1/2
-1	0	−√(2/3)
0	1/√3	−1/√3
1	√(2/3)	0

j_1=1 ⊗ j_2=1/2 → J=1/2 ⊕ J=3/2 Dimension: 6 = 2 + 4 ✓

Wigner–Eckart Theorem

$$\langle \alpha', j', m' \mid T^{(k)}_q \mid \alpha, j, m\rangle = \langle j, m;\, k, q \mid j', m'\rangle \times \frac{\langle \alpha' j' \| T^{(k)} \| \alpha j\rangle}{\sqrt{2j'+1}}$$

The CG coefficient carries ALL m-dependence.
The reduced matrix element is independent of m, m', q.

E1 ($k=1$): $\Delta j = 0, \pm 1$ (no $0 \to 0$); $\Delta m = 0$ (π), ± 1 ($\sigma^{\pm}$)
M1 ($k=1$): Same Δj, Δm; $\Delta l = 0$

E2 ($k=2$): $\Delta j = 0, \pm 1, \pm 2$; $\Delta m = 0, \pm 1, \pm 2$
Hyperfine: All $\langle F, m_F | T | F', m_{F'}\rangle$ reduce to one $\langle F \| T \| F'\rangle$ via Wigner–Eckart

Wigner 3j relation:

$$\langle j_1 m_1;\, j_2 m_2 \mid JM \rangle = (-1)^{j_1 - j_2 + M} \sqrt{2J+1} \begin{pmatrix} j_1 & j_2 & J \\ m_1 & m_2 & -M \end{pmatrix}$$

In Python: `sympy.physics.wigner.wigner_3j(j1, j2, j3, m1, m2, m3)`

20.7 EOM / AOM Sideband Calculator

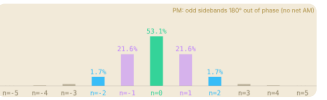
Bessel-function sideband depths for phase modulation (PM), the two-sideband spectrum of pure amplitude modulation (AM), and AOM double-pass shift and efficiency.

EOM Sideband Depths — Bessel Functions $J_n(\beta)$

A phase-modulating EOM with modulation index β splits the carrier into sidebands at ±Ω, ±2Ω, … Each sideband carries power fraction $J_n(\beta)^2$ relative to the input. **AM:** $E_0(1 + m \cos \Omega t)$ has only the ±1 sidebands, each with field m/2 and in phase with each other (here β is read as the depth m, 0–1). **PM/FM:** sidebands have alternating phase (±1 sidebands are π out of phase) — no net AM on a PD.

$$P_n/P_0 = J_n(\beta)^2 \qquad \sum_{n=-\infty}^{\infty} J_n(\beta)^2 = 1$$

Carrier (n=0) J0(β) = 0.7290 53.14% each
Sideband n=±1 J1(β) = 0.4650 21.62% each
Sideband n=±2 J2(β) = 0.1321 1.75% each
Sideband n=±3 J3(β) = 0.0244 0.06% each
Sideband n=±4 J4(β) = 0.0033 0.00% each
Sideband n=±5 J5(β) = 0.0004 0.00% each
β = 1.080 · PDH signal ∝ $J_0 J_1$ = 0.3390 (max at β≈1.08)



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Spectrum — carrier + sidebands (height ∝ field amplitude)

Key β Values

Special modulation depths used in common lab applications.

B	$J_0(B)^2$	$J_1(B)^2$	USE
1.08	0.531	0.216	PDH: max error signal ∝ $J_0 J_1$
2.405	≈ 0	0.269	Carrier suppression (STIRAP, EIT)
5.520	≈ 0	0.116	2nd carrier null
0.5	0.881	0.059	Weak modulation (small sidebands)

AOM Frequency Shift & Double-Pass Setup

AOMs shift laser frequency by ±f_AOM. A **double-pass** configuration (retroreflect through AOM a second time) shifts by 2f_AOM, is polarization-maintaining, and makes the output direction independent of frequency — essential for laser locking scan ranges without beam steering.

$$\Delta f_{\text{single}} = \pm f_{\text{AOM}} \qquad \Delta f_{\text{double}} = 2 f_{\text{AOM}} \qquad \eta_{\text{DP}} = \eta_{\text{SP}}^2$$

Single-pass shift +80 MHz Double-pass shift +160 MHz Single-pass efficiency 75.0 % Double-pass efficiency 56.3 % Double-pass power loss 43.8 %

Double-pass alignment tip: the cat's-eye retroreflector (λ/4 plate + mirror) must send the beam *back through the AOM crystal*, not just the fiber coupler. Align by checking that the retro beam overlaps the AOM input beam on an iris, then optimize the second-order extinction (should be >30 dB).

EOM Half-Wave Voltage V_π

Converts RF power to modulation index β = π V_RF / V_π. Useful for comparing specs across EOMs.

$$\beta = \frac{\pi V_{\text{RF}}}{V_\pi} = \frac{\pi \sqrt{2 R P_{\text{RF}}}}{V_\pi}$$

V_RF (peak) 3.16 V Modulation index β 0.050 Carrier power 99.88 % ±1 sideband power 0.06 % each

🏔️

Hood Lab: We use a resonant EOM at 9.192 GHz to generate Cs hyperfine sidebands for saturated absorption spectroscopy and a non-resonant broadband EOM at 20 MHz for PDH locking to the 685 nm ULE cavity. The 9 GHz resonant EOM is particularly unforgiving — driving off resonance by even 5 MHz drops the modulation depth by >10 dB; always match the drive frequency to the EOM's resonant frequency with a VNA before use.

20.8 Unit Converter

Instant conversions for the units AMO physicists actually use. Enter any value and all equivalents update live.

CHAPTER 20 SUMMARY

With Part II complete, we can load, cool, image, and characterize single trapped atoms with precision. In Part III, we encode quantum information in those atoms and build a programmable quantum processor.

Problems for Chapter 20

1. USE THE TOOL Use the Zeeman calculator for $g_F = 1/2$, $m_F = 2$ at 1 G, then at 10 G.
2. BY HAND Check problem 1 with $\Delta E/h = g_F m_F \mu_B B/h$, using $\mu_B/h \approx 1.4$ MHz/G. Then plan an RF chain: a 10 dBm driver must deliver 2 W (33 dBm) to an AOM through a 2 dB cable loss. How much amplifier gain is needed? Check it with the dBm calculator.
3. DESIGN Using the Clebsch–Gordan table, compare the σ^+ transition strengths from $m_F = 0$ and $m_F = 2$ on the ^{87}Rb D2 line ($F = 2 \rightarrow F' = 3$). Explain why experiments cycle on the stretched transition for imaging.

Use the tool: on the interactive page. By hand: from the formulas in this chapter. Design: a decision a lab would make. Worked solutions are available to instructors on request (see Part V).

PART III

From Atoms to a Quantum Computer

Turning trapped atoms into qubits: how qubits work and lose coherence, how decoupling protects them, how Rydberg blockade entangles them, what limits gate fidelity, how gates are benchmarked, how circuits run on moving atoms, and how processors are linked with photons.

21	Quantum Computing Fundamentals	Quantum Computing Learning Path
22	Decoherence & Optical Bloch Equations	Decoherence Lab
23	Dynamical Decoupling & Noise Spectroscopy	Dynamical Decoupling
24	Rydberg Blockade Physics	Rydberg Blockade Lab
25	Two-Qubit Gate Error Budgets	Rydberg Gate Error Budget
26	Randomized Benchmarking	Randomized Benchmarking
27	Zoned Neutral-Atom Processors	Zoned Atom Processor
28	Remote Entanglement & Quantum Networks	Remote Entanglement

CHAPTER 21 · QUANTUM STEP 01

21

Quantum Computing Fundamentals

🔵 Quantum Computing Learning Path · Learning path

Interactive visualizations from single qubits on the Bloch sphere to Rydberg blockade and Wigner functions. No prior quantum background required.

> *Part III turns our single-atom tweezer array into a quantum computer. We begin by establishing the mathematical and physical language of qubits, gates, entanglement, error correction, and magic states.*

IN THIS CHAPTER Prerequisites: Ch. 1, Ch. 18

- Map single-qubit states and rotations onto the Bloch sphere and connect them to atomic Rabi oscillations and hyperfine clock qubits.
- Understand entanglement, Rydberg blockade two-qubit gates, and core algorithms (Grover, Shor, analog simulation).
- Grasp the foundations of quantum error correction (surface codes, threshold theorem) and non-Clifford magic states.

WEBSITE COMPANION Live interactive page ↗ · Week 1

Course map

21.1 Read it like a short course, then use it as a reference.

The page now has a visible learning path: build a single qubit, learn how gates and measurement work, add entanglement, then connect the abstractions to Rydberg neutral-atom hardware.

01
Qubit language
Bloch sphere, superposition, phase, and Rabi rotations.

02
Operations
Gates, measurement, optical pumping, and hyperfine-state control.

03
Many-qubit physics
Entanglement, two-qubit gates, Rydberg blockade, and decoherence.

04
Algorithms and QEC
Algorithms, analog simulation, error correction, and what hardware must deliver.

Calculator output at the page's default settings.

Quantum Quiz

21.2 Test your understanding

> 0/30 mastered

8 questions pulled from all 15 topics below, weighted toward whatever you've missed before. Progress is saved in this browser.

🔵

21.3 Bloch Sphere

single qubit

Single-qubit state space

Any pure single-qubit state can be written as $|\psi\rangle = \cos(\theta/2)|0\rangle + e^{i\varphi}\sin(\theta/2)|1\rangle$, mapping to a point on the unit sphere (Bloch sphere) with polar angle $\theta \in [0,\pi]$ and azimuthal angle $\varphi \in [0,2\pi]$. The Bloch vector **n** = (sin θ cos φ, sin θ sin φ, cos θ) represents the expectation values ⟨σx⟩, ⟨σy⟩, ⟨σz⟩.

Special states

|0⟩ (north pole): θ=0, Bloch vector = (0,0,+1)
|1⟩ (south pole): θ=π, Bloch vector = (0,0,−1)
|+⟩ = (|0⟩+|1⟩)/√2: θ=π/2, φ=0 → (+x equator)
|+i⟩ = (|0⟩+i|1⟩)/√2: θ=π/2, φ=π/2 → (+y equator)
Mixed state: ρ = I/2 → Bloch vector = origin (not on sphere)

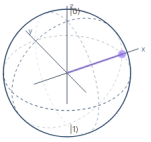
$$|\psi\rangle = \cos(\theta/2)|0\rangle + e^{i\varphi}\sin(\theta/2)|1\rangle$$

Bloch vector:

$$\langle\sigma_x\rangle = \sin\theta\cos\varphi \qquad \langle\sigma_y\rangle = \sin\theta\sin\varphi \qquad \langle\sigma_z\rangle = \cos\theta$$

$$P(|0\rangle) = \cos^2(\theta/2) \qquad P(|1\rangle) = \sin^2(\theta/2)$$

📖 Quantum Country, visual intro 💻 IBM Quantum: single-qubit states 🔧 Quirk, interactive circuit/state



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

|ψ⟩: bx=1.00, by=0.00, bz=0.00 | P(0)=50.0%

〰️

21.4 Superposition & Interference

Mach-Zehnder

Why superposition enables interference

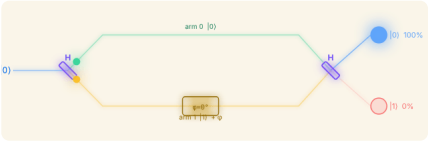
A single photon in state $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$ simultaneously "explores" both arms of the interferometer. At the second beamsplitter the two amplitudes combine — constructively or destructively depending on the relative phase φ. This is quantum interference: not two photons, but one photon's amplitude interfering with itself.

The MZI implements $U_{MZI} = H \cdot R_z(\varphi) \cdot H$. Sweeping φ continuously transfers the photon from one output port to the other — a full, reversible control of quantum probability.

$$U_{MZI} = H \cdot R_z(\varphi) \cdot H$$

$$P(|0\rangle) = \cos^2(\varphi/2) \qquad P(|1\rangle) = \sin^2(\varphi/2)$$

Constructive interference at $\varphi = 0$ (all into $|0\rangle$); destructive at $\varphi = \pi$ (all into $|1\rangle$)



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

🔍 **Which-path information destroys interference.** If any measurement — however subtle — reveals which arm the photon took, the interference fringes vanish completely, even if you never read the result. Knowledge itself, not the act of looking, collapses the superposition. This is the quantum eraser principle.

📖 Quantum Country: superposition & interference 🌐 Mach-Zehnder interferometer (Wikipedia) 📄 Hardy 2001: Quantum theory from five axioms (arXiv)

🌀

21.5 Rabi Oscillations

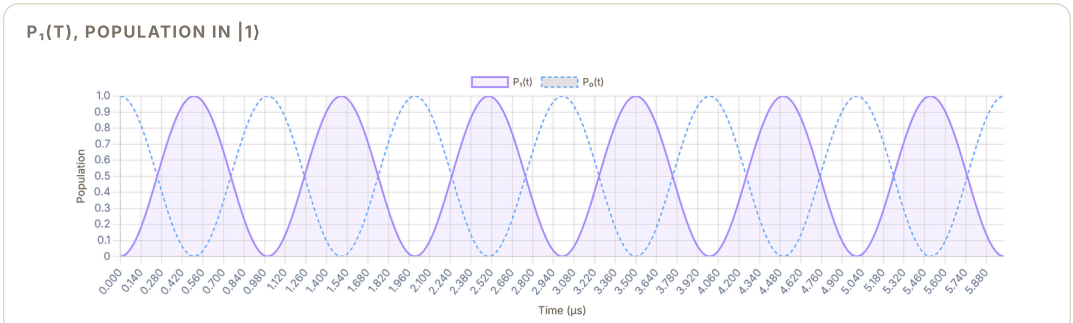
coherent drive

Driven two-level system

A resonant laser drive with Rabi frequency Ω causes coherent oscillations between $|0\rangle$ and $|1\rangle$. Off-resonance (detuning Δ), the generalised Rabi frequency is $\Omega' = \sqrt{(\Omega^2+\Delta^2)}$ and the oscillation amplitude is reduced to $(\Omega/\Omega')^2$.

$$P_1(t) = \left(\frac{\Omega}{\Omega'}\right)^2 \sin^2\left(\frac{\Omega' t}{2}\right) \quad \text{where} \quad \Omega' = \sqrt{\Omega^2 + \Delta^2}$$

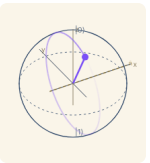
Resonance ($\Delta = 0$): $P_1 = \sin^2(\Omega t/2)$, full population transfer
π-pulse time: $t_\pi = \pi/\Omega$ $\quad[|0\rangle \to |1\rangle$ exactly]



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

$\Omega' = \sqrt{(\Omega^2+\Delta^2)}$ = 1.00 MHz, amplitude = 100.0%, π-pulse t = 500 ns

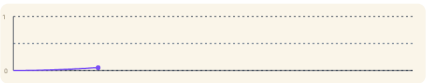
Bloch vector precession



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Ω'=1.00 MHz P_1=5%

$P_1(t)$ live



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

The Bloch vector rotates around the *effective field* $\hat{\mathbf{n}}_\mathbf{eff} = (\Omega, 0, -\Delta)/\Omega'$ at angular frequency Ω'. At resonance (Δ=0) the rotation axis is $\hat{x}$— the vector sweeps from $|0\rangle$ to $|1\rangle$ and back. Off-resonance, the axis tilts toward $\hat{z}$ and the cone shrinks.

📊 Steck: Cs D-line data (Ω reference) 💻 QuTiP: driven qubit dynamics 📄 Steck: Quantum Optics notes, two-level atoms (free PDF)

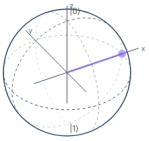
⚡

21.6 Quantum Gates

unitary operations

Unitary transformations on the Bloch sphere

Single-qubit gates are 2×2 unitary matrices acting on $|\psi\rangle$. Geometrically they are rotations of the Bloch sphere. Pauli gates rotate by π around x, y, or z axes. The Hadamard H swaps the $|0\rangle/|1\rangle$ basis with $|\pm\rangle$. Phase gates (S, T) rotate around z.



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

|ψ⟩: bx=1.00, by=0.00, bz=0.00 | P(0)=50.0%

🔧 Quirk, drag-and-drop gate sandbox 💻 IBM Quantum: single-qubit gates 📄 Barenco et al. 1995: elementary gates for quantum computation (universal gate sets)

📏

21.7 Measurement & Born Rule

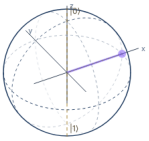
projection, statistics

Born rule & measurement bases

Measuring $|\psi\rangle = \cos(\theta/2)|0\rangle + e^{i\varphi}\sin(\theta/2)|1\rangle$ projects onto an orthonormal basis. In the **Z basis** $P(|0\rangle) = \cos^2(\theta/2)$ — the azimuthal angle φ is *completely undetectable*. Switch to the **X** or **Y basis** and φ becomes visible: $P(|+\rangle) = \frac{1}{2}(1+\sin\theta\cos\varphi)$, $P(|+i\rangle) = \frac{1}{2}(1+\sin\theta\sin\varphi)$. The post-measurement state always collapses to one eigenstate of the chosen basis.

Z basis: $P(|0\rangle) = \cos^2(\theta/2)$ — φ is undetectable

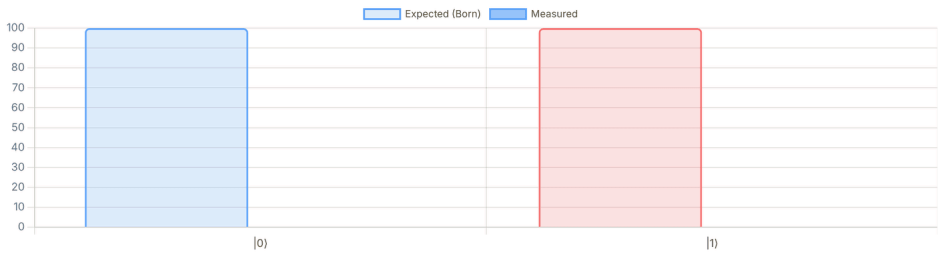
State on Bloch sphere



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

|ψ⟩: bx=1.00, by=0.00, bz=0.00 | P(0)=50.0%

MEASUREMENT HISTOGRAM — Z BASIS



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

📖 Quantum Country: measurement & Born rule 💻 IBM Quantum: measurement basics

💡

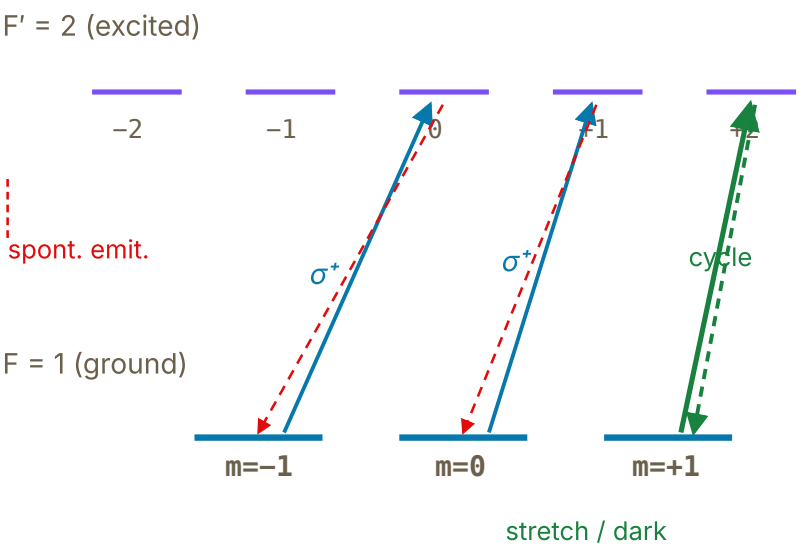
21.8 Optical Pumping & State Preparation

qubit initialization

Driving atoms into a single quantum state with light

Optical pumping uses **circularly polarized resonant light** to redistribute atomic population among magnetic sublevels m_F. A σ^+ beam drives $\Delta m_F = +1$ transitions — atoms absorb photons and reach excited states inaccessible from the highest m_F level. After repeated absorption-emission cycles, population accumulates in the **dark state** (the sublevel that cannot absorb) or a **cycling (stretch) state**. This is how we initialize qubits: prepare all atoms in one specific m_F before coherent manipulation.

Σ⁺ OPTICAL PUMPING, LEVEL DIAGRAM (F=1 EXAMPLE)



POPULATION DYNAMICS, SCATTER PHOTONS INTERACTIVELY

Start: equal population in $m_F = -1, 0, +1$. σ^+ light pumps toward $m_F = +1$.

$m_F = -1$ 33.3%

$m_F = 0$ 33.3%

$m_F = +1$ 33.3%

Scattered: 0

Click "+1 photon" to begin pumping the population toward $m_F = +1$.

Selection rules (electric dipole): σ^+: $\Delta m_F = +1$; σ^-: $\Delta m_F = -1$; π: $\Delta m_F = 0$
Scattering rate from sublevel m_F (σ^+ beam):

$$R_m = \frac{\Gamma}{2} \cdot \frac{s \cdot |C_m|^2}{1 + s|C_m|^2 + (2\Delta/\Gamma)^2}$$

where $|C_m|^2$ = squared Clebsch–Gordan coeff. for $|F, m_F\rangle \to |F', m_F + 1\rangle$
Dark state (σ^+, only if $F' \leq F$): $m_F = +F$ has no allowed σ^+ absorption. For $F' = F + 1$ the stretched state $m_F = +F$ is instead a *cycling* state.
Pumping time: $\tau_{pump} \approx 2F/(R_{scatter} \cdot \varepsilon_{branch})$

Clock-state preparation (|F, m_F=0⟩)

For qubit encoding in clock states (m_F=0), a different strategy is used: pump first to a high-field-seeking state with σ^+, then apply a microwave or RF π-pulse to transfer to $|F=2, m_F=0\rangle$. Alternatively, use a combination of $\sigma^+ + \pi$ light to create a dark state at m_F=0. In practice, Rb-87 clock-state qubits are initialized with >99% fidelity in ~100 μs.

EIT & dark states in Λ systems

Coherent population trapping (CPT) is a related phenomenon: two fields in a Λ configuration drive an atom into a *coherent dark state*, a superposition of two ground states that has zero absorption. This underpins **EIT (electromagnetically induced transparency)** and **STIRAP** transfer. Unlike optical pumping (which is incoherent), CPT/EIT preserves quantum coherence.

📄 Happer (1972), Optical pumping (Rev. Mod. Phys., the canonical review) 📄 Fleischhauer, Imamoglu & Marangos (2005), EIT review (Rev. Mod. Phys.) 📊 Steck alkali data sheets, CG coefficients, branching ratios for all species

🔑

21.9 Hyperfine Qubits & Atomic Clock States

AMO qubit encoding

Why atoms are excellent qubits

Neutral alkali atoms encode qubits in their **hyperfine ground states**: two long-lived magnetic sub-levels of the electronic ground state split by the interaction between nuclear spin I and electron spin J. The interaction is $H = A_{hf}\,\mathbf{I}\cdot\mathbf{J}$, which for a J = ½ ground state splits the levels by $\Delta E = A_{hf}(I + ½)$, where A_{hf} is the magnetic dipole hyperfine constant. **Clock states** ($|F, m_F = 0\rangle$) are first-order insensitive to magnetic field fluctuations, enabling $T_2 > 10$ seconds, orders of magnitude longer than superconducting qubits (~100 μs).

$$\hat{H}_{hf} = A_{hf}\,\vec{I}\cdot\vec{J} + B_{hf}\left[3(\vec{I}\cdot\vec{J})^2 + \ldots\right] \quad \text{(quadrupole term, needs } I \geq 1 \text{ and } J \geq 1)$$

Qubit encoding (e.g. ^{133}Cs): $|0\rangle = |F{=}3, m_F{=}0\rangle$, $|1\rangle = |F{=}4, m_F{=}0\rangle$
Splitting: $\Delta E/h = 9{,}192{,}631{,}770$ Hz (defines the SI second!)
2nd-order Zeeman shift (clock states):

$$\Delta\nu(B) \approx +427\ \text{Hz/G}^2\ \text{(Cs)}, \quad +575\ \text{Hz/G}^2\ \text{(Rb-87)}$$

Single-qubit gate: resonant MW or two-photon Raman pulse; $t_\pi \sim 5\,\mu$s (Raman, $\Omega/2\pi \sim 100$ kHz)

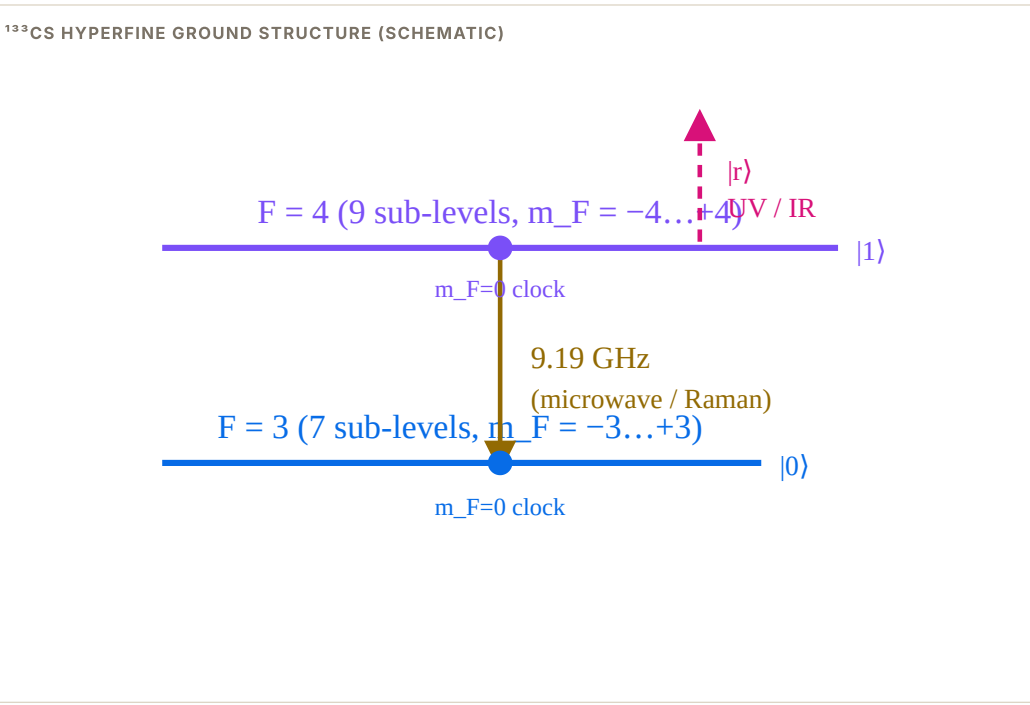
Qubit encoding strategies

Clock qubit (m_F=0 → m_F'=0): best coherence, first-order B-insensitive. Used in most neutral-atom QC experiments.

Stretched-state qubit (m_F=F → m_F'=F'): strongest coupling to MW field, but first-order B-sensitive, requires ≲1 mG stability.

Optical qubit (ground ↔ metastable excited state, e.g. Sr, Yb): very fast gates via optical fields; coherence limited by excited-state lifetime and photon scattering.

For Rydberg gate operation, either qubit state is coupled to a Rydberg level |r⟩ via a two-photon (IR + UV) or single-photon (UV ~318 nm for Cs) transition. The blockade between two atoms in |rr⟩ enforces a conditional phase on |11⟩.



Hyperfine qubit comparison

ATOM	HYPERFINE SPLITTING	CLOCK STATES	T_2 (BEST)	GATE
^{87}Rb	6.835 GHz	|1,0⟩ → |2,0⟩	~1 s (with DD)	Rydberg CZ
^{133}Cs	9.193 GHz (SI clock!)	|3,0⟩ → |4,0⟩	12.6 s (DD, 6100 atoms)	Rydberg CZ
^{6}Li	228.2 MHz	|F=1/2⟩ → |F=3/2⟩ (m_F pair)	~100 ms	Rydberg / tweezer
^{171}Yb	≈0.75 kHz/G × B (nuclear Zeeman)	|m_I=+½⟩ → |m_I=−½⟩ ($^1S_0$, F=½)	>10 s	Rydberg CZ
^{88}Sr (optical)	429.2 THz (clock line)	$^1S_0$ → $^3P_0$	>1 min	Rydberg / optical

💡 **^{171}Yb** is emerging as a leading neutral-atom qubit: its $^1S_0$ ground state has J = 0, so the nuclear spin I = 1/2 (m_I = ±½) is a clean two-level system that couples to magnetic fields only through the small nuclear moment (≈0.75 kHz/G), enabling exceptionally long coherence. (The 12.6 GHz F=0 ↔ F=1 hyperfine qubit belongs to the *ion* $^{171}Yb^+$.)

🏆 **Record T_2**: Neutral atom clock qubits in optical tweezers have demonstrated T_2 ≈ 40 s (spin echo) for ^{87}Sr nuclear-spin qubits (Barnes et al., Nat. Commun. 2022) and T_2 = 12.6 s across a 6100-atom Cs array (Manetsch et al. 2025). For comparison, superconducting transmon qubits: T_2 ~ 100–500 μs.

📊 Steck: Cs D_2 data (hyperfine constants) 📊 Steck: Rb-87 data 📄 Ma et al. 2022: ^{171}Yb qubit array (Princeton) 📄 Saffman 2016: neutral atom QC review 📄 Evered et al. 2023, long T_2 + 99.5% gates

🔗

21.10 Entanglement & Bell States

non-separability

Bell states, maximally entangled two-qubit states

A state is entangled if it cannot be written as $|\psi_A\rangle \otimes |\psi_B\rangle$. The four Bell states form an orthonormal basis for the two-qubit Hilbert space and are maximally entangled: measuring one qubit instantly determines the other's outcome, regardless of separation, "spooky action at a distance" (Einstein, 1935).

$\|\Phi^+\rangle$ `(|00⟩ + |11⟩)/√2`	$\|\Phi^-\rangle$ `(|00⟩ − |11⟩)/√2`	$\|\Psi^+\rangle$ `(|01⟩ + |10⟩)/√2`
$\|\Psi^-\rangle$ `(|01⟩ − |10⟩)/√2`	$\|00\rangle$ (SEPARABLE) `|0⟩⊗|0⟩, not entangled`	$\|++\rangle$ (SEPARABLE) `(|0⟩+|1⟩)/√2 ⊗ same`

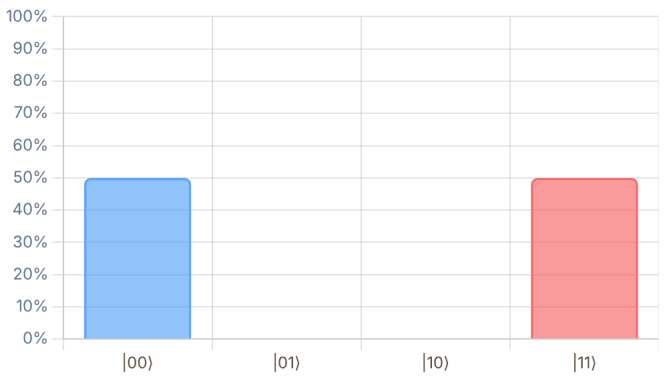
Calculator output at the page's default settings.

$|\Phi^+\rangle$, 🔗 Entangled

Maximally entangled. Measuring qubit A in $|0\rangle$ → qubit B collapses to $|0\rangle$. Measuring $|1\rangle$ → $|1\rangle$.

$(|00\rangle + |11\rangle)/\sqrt{2}$ Ket coefficients: $|00\rangle$: amplitude = +0.707, prob = 50% $|01\rangle$: amplitude = +0.000, prob = 0% $|10\rangle$: amplitude = +0.000, prob = 0% $|11\rangle$: amplitude = +0.707, prob = 50%

MEASUREMENT PROBABILITIES P(OUTCOME)



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Interactive

IS IT JUST CORRELATION? PLAY THE CHSH GAME

A referee sends Alice and Bob random bits x and y. Without talking, they answer bits a and b, and win if $a \oplus b = x \cdot y$ (answers equal, except when both got a 1, when they must differ). Any classical strategy, even with shared random numbers agreed in advance, wins at most **75%** of rounds. Sharing $|\Phi^+\rangle$ and measuring along angles chosen from x and y wins up to $\cos^2(\pi/8) \approx$ **85.4%**: the CHSH violation, $S = 2\sqrt{2}$ instead of 2. Correlated coins cannot do this; entangled qubits can.

Angles are measurement-basis (polarizer) angles in the x–z plane; for $|\Phi^+\rangle$ the answers agree with probability $\cos^2(\theta_A - \theta_B)$. Why this is not faster-than-light messaging: Alice's answers on their own are 50/50 whatever Bob measures; the pattern appears only when the two records are compared, which needs an ordinary signal. (Clauser, Horne, Shimony & Holt 1969; loophole-free tests: Hensen et al. 2015, Giustina et al. 2015, Shalm et al. 2015.)

Interactive

TELEPORTATION: WHAT THE TWO CLASSICAL BITS ARE FOR

Alice holds an unknown qubit $|\psi\rangle$ and half of a shared $|\Phi^+\rangle$; Bob holds the other half. Alice does a CNOT and a Hadamard, then measures her two qubits. Each of the four results comes up **25% of the time whatever $|\psi\rangle$ is**, so Alice learns nothing about it, and Bob's qubit is left in one of four rotated copies of $|\psi\rangle$. Only when Alice's two bits reach him, by an ordinary signal, does he know which correction undoes the rotation. Until then his qubit, averaged over the four cases, is the centre of the sphere: no information at all.



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Alice's $|\psi\rangle$



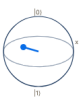
Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

m_1m_2 = 00 · Bob fixes with I



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

m_1m_2 = 01 · Bob fixes with X



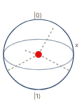
Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

m_1m_2 = 10 · Bob fixes with Z



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

m_1m_2 = 11 · Bob fixes with ZX



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Bob, before the bits arrive

Bloch vector of $|\psi\rangle$ = (0.75, 0.43, 0.50). Average of Bob's four possible states = (0.00, 0.00, 0.00): the maximally mixed state.

Standard circuit (Nielsen & Chuang §1.3.7): with outcomes m_1 (Alice's $|\psi\rangle$ qubit) and m_2 (her half of the pair), Bob holds $X^{m_2}Z^{m_1}|\psi\rangle$ and applies $Z^{m_1}X^{m_2}$. On the Bloch sphere X flips y and z, Z flips x and y, so the four cases average to zero. Proposed by Bennett et al., PRL 70, 1895 (1993); the same move, with a magic state instead of $|\psi\rangle$, is the T-gate gadget in

Magic States.

📄 Horodecki et al. 2009: quantum entanglement (Rev. Mod. Phys. 81, 865) 📄 Clauser, Horne, Shimony & Holt 1969: the CHSH inequality (PRL 23, 880) 📖 Quantum Country: entanglement 🔧 Quirk, build Bell states visually

🔀

21.11 Two-Qubit Gates

entangling operations

Entangling gates

Two-qubit gates act on the 4D Hilbert space spanned by $|00\rangle$, $|01\rangle$, $|10\rangle$, $|11\rangle$. CNOT, CZ, and iSWAP can generate maximal entanglement from product states, they are universal when combined with single-qubit gates.

CNOT matrix (basis: $|00\rangle$,$|01\rangle$,$|10\rangle$,$|11\rangle$)

1	0	0	0
0	1	0	0
0	0	0	1
0	0	1	0

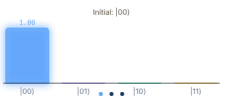
CNOT, Controlled-NOT

Flips the target qubit if and only if the control qubit is $|1\rangle$. Together with single-qubit gates, CNOT is universal for quantum computing. Applied to $|+0\rangle \rightarrow$ Bell state $|\Phi^+\rangle$.

Truth table

INPUT		OUTPUT
$\|00\rangle$	→	$\|00\rangle$
$\|01\rangle$	→	$\|01\rangle$
$\|10\rangle$	→	$\|11\rangle$
$\|11\rangle$	→	$\|10\rangle$

Bell state $|\Phi^+\rangle$ generation: H ⊗ I → CNOT



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

🔧 Quirk, apply CNOT & CZ visually 💻 IBM Quantum: two-qubit systems 📄 Zhang, Vala, Sastry & Whaley 2003: geometric theory of nonlocal two-qubit operations (perfect entanglers)

⚛️

21.12 Rydberg Atoms & Blockade

neutral-atom QC

Rydberg blockade mechanism

Exciting two nearby atoms to Rydberg states shifts one atom's resonance by V(R) = C_6/R^6 due to van der Waals interactions. When V(R) ≫ ħΩ_R, the double excitation |rr⟩ is off-resonant; this is the **Rydberg blockade**. Within the blockade radius R_b, only one atom can be excited at a time, enabling two-qubit entangling gates.

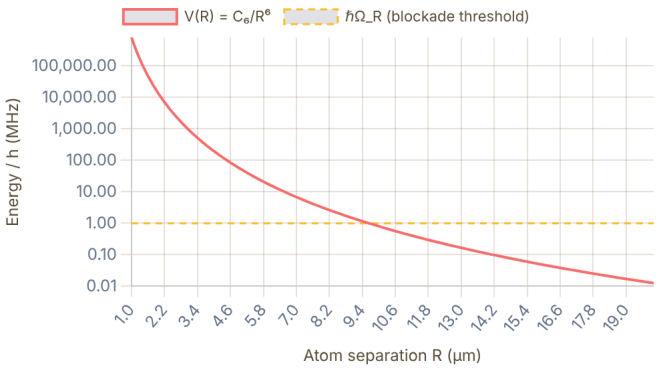
$$C_6 \propto n^{*11} \qquad (n^* = \text{effective principal quantum number})$$

$$R_b = \left(\frac{C_6}{\hbar\Omega_R}\right)^{1/6} \quad [\text{blockade radius}]$$

$$V(R) = \frac{C_6}{R^6} \quad [\text{van der Waals interaction}]$$

Gate fidelity: $\mathcal{F} \approx 1 - (\Gamma/\Omega_R) - \alpha\langle n_{motion}\rangle$

V(R) = C₆/R⁶ INTERACTION POTENTIAL



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

📄 Jaksch et al. 2000, original Rydberg gate proposal (PRL 85, 2208) 📄 Evered et al. 2023, 99.5% CZ gate (Harvard) 📄 Saffman 2016: neutral atom QC review 📄 Bluvstein et al. 2024, 48 logical qubits

📉

21.13 Decoherence: T_1 & T_2

open quantum systems

Bloch sphere collapse

Real qubits interact with their environment. T_1 (relaxation time) is how long the excited state population survives, the qubit decays from |1⟩ to |0⟩. T_2 (coherence/dephasing time) is how long superpositions survive. Always $T_2 \leq 2T_1$. The off-diagonal density matrix element $|\rho_{01}|$ decays as e^{-t/T_2}.

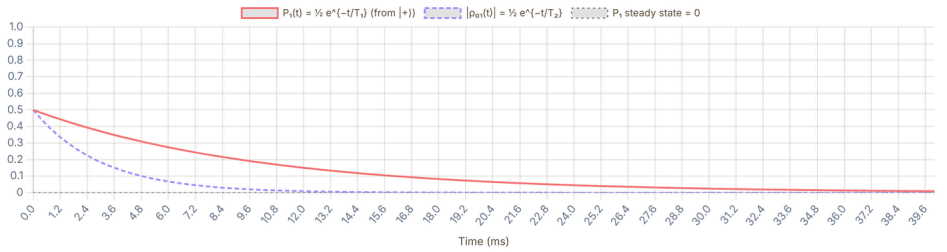
Bloch equations (Lindblad):

$$\dot{s}_z(t) = -\frac{s_z - 1}{T_1} \quad \Rightarrow \quad s_z(t) = 1 - (1 - s_{z0})\, e^{-t/T_1} \quad (|0\rangle = +\hat{z} \text{ is the state } T_1 \text{ decays to})$$

$$|\rho_{01}(t)| = |\rho_{01}(0)|\, e^{-t/T_2}$$

Typical (alkali hyperfine qubits in tweezers): $T_1 \gtrsim 1$ s (trap-light Raman scattering), $T_2^* \approx 1$–5 ms, $T_2 \approx 100$ ms (echo); always $T_2 \leq 2T_1$

POPULATION $P_1(T)$ AND COHERENCE $|P_{o1}(T)|$



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

T_1 = 10 ms, T_2 = 3 ms, T_2/T_1 = 0.30 (max = 2.0, pure dephasing = 1.0). Gate time must satisfy t_gate ≪ T_2.

📄 Zurek 2003: decoherence review 💻 QuTiP: Lindblad master equation 📄 Bluvstein et al. 2024, coherence in 48 logical qubits

⚙️

21.14 Quantum Algorithms

Grover, Deutsch-Jozsa, QAOA

Deutsch-Jozsa Algorithm

The first proof-of-principle quantum speedup. Determines whether a function f:{0,1}^{n} → {0,1} is **constant** (same output always) or **balanced** (half 0, half 1) in a *single* oracle query, versus O(2^{n-1}+1) classically. Uses Hadamard superposition + phase kickback from the oracle, then a final Hadamard layer, all |0⟩ output means constant, any other outcome means balanced.

Grover's Search Algorithm

Finds a marked item in an unsorted database of N entries in O(√N) oracle queries , quadratic speedup over classical O(N). Each Grover iteration applies: (1) the oracle (phase flip on target state), then (2) a diffusion operator (inversion about the mean amplitude). After ≈ (π/4)√N iterations the target state has amplitude ≈ 1. Relevant for neutral-atom QC because each gate cycle is slow, minimising query count is critical.

Optimal iterations: k ≈ (π/4) √N

QAOA, Quantum Approximate Optimisation

A near-term variational algorithm (NISQ-friendly) for combinatorial optimisation. Alternates between a *cost unitary* $e^{-i\gamma H_C}$ encoding the problem and a *mixing unitary* $e^{-i\beta H_B}$, with angles (γ, β) optimised classically. Neutral-atom arrays are a natural hardware match: Rydberg interactions implement the cost Hamiltonian natively, and the array geometry can be programmed to encode graph problems.

Circuit complexity & quantum advantage

Quantum advantage requires circuits deep enough that classical simulation is intractable but shallow enough to survive decoherence. Current neutral-atom processors (Lukin/Greiner/Browaeys groups) achieve ~100-qubit circuits at 99%+ two-qubit gate fidelity. Grover gives proven quadratic speedup; exponential speedups (Shor's factoring) require fault-tolerant hardware, motivating the QEC section below.

Grover's algorithm: amplitude amplification on N=8 states, target = state ★



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

📖 Matuschak & Nielsen: Quantum Computing for the Very Curious (free, interactive) 💻 IBM Quantum Learning, Grover & QAOA tutorials with runnable circuits 📄 Ebadi et al. 2022, quantum optimization of maximum independent set on a 289-atom Rydberg array (Science 376, 1209)

🔬

21.15 Analog Quantum Simulation

Hubbard model, spin models, arrays

Digital vs. Analog quantum computation

Digital QC decomposes any unitary into discrete gates, making it fully programmable but requiring error correction for deep circuits. **Analog simulation** engineers a physical Hamiltonian H_sim that mimics a target model H_target, then lets the system evolve. No gate decomposition is needed; the physics *is* the algorithm. AMO platforms excel here because atom–atom interactions are highly controllable and the Hilbert space scales exponentially with atom number, making classical simulation intractable beyond ~50–60 qubits.

Bose-Hubbard & Fermi-Hubbard models

Cold atoms in optical lattices realise the Hubbard Hamiltonian:

$H = -t\,\Sigma(a_i^\dagger a_j + \text{h.c.}) + (U/2)\,\Sigma\, n_i(n_i-1) - \mu\,\Sigma\, n_i$

where t is the tunnel coupling (laser intensity), U is the on-site interaction (Feshbach resonance tunable), and μ is chemical potential. At $U/t \gg 1$: Mott insulator. At $U/t \ll 1$: superfluid. The Fermi-Hubbard model (fermionic atoms) is believed to describe high-T_c superconductivity, directly accessible in the lab.

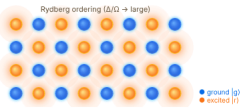
Rydberg spin models & quantum magnetism

Rydberg tweezer arrays naturally implement Ising and XY spin Hamiltonians:

$H = \Sigma_i\,(\Omega/2)\sigma^x_i - \Delta\, n_i + \Sigma_{i<j}\, V(r_{ij})\, n_i\, n_j$

$V(r) = C_6/r^6$ for van der Waals, C_3/r^3 for resonant dipole exchange (e.g. nS–nP pairs). By choosing array geometry, any graph's adjacency can be encoded. Experiments have demonstrated Z_2 symmetry breaking, topological spin liquids (Semeghini et al. 2021), and frustrated magnetism in programmable arrays.

Rydberg array: Z_2 symmetry breaking as Ω/Δ is tuned — atoms (●) and blockade halos shown



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

📄 Ebadi et al. 2021, quantum phases of matter in 256-atom Rydberg array (Harvard) 📄 Semeghini et al. 2021, topological spin liquid in Rydberg array 📄 Gross & Bloch 2017, quantum simulations with ultracold atoms in optical lattices (review)

🛡️

21.16 Quantum Error Correction

fault tolerance

Why QEC is essential

Physical qubits have error rates ~10^{-3}–10^{-2} per gate. Running a useful algorithm (e.g., Shor factoring a 2048-bit number) requires ~10^{8} logical operations, impossibly many without error correction. QEC encodes one **logical qubit** in many physical qubits so that errors can be *detected and corrected* without ever measuring the logical state directly (which would collapse it). The key insight: measuring error *syndromes* (parities) reveals which error occurred without revealing the encoded information.

3-qubit bit-flip code: $|0_L\rangle = |000\rangle$, $|1_L\rangle = |111\rangle$
Syndrome operators: $Z_1Z_2 \otimes I$ and $I \otimes Z_2Z_3$

Surface code (distance-d, rotated): encodes 1 logical qubit in d^2 data + d^2-1 measure qubits $= 2d^2-1$ physical qubits.

$$p_L \approx A \cdot \left(\frac{p}{p_{th}}\right)^{\lceil (d+1)/2 \rceil} \qquad p_{th} \approx 1\%$$

Example: $d = 7$, $p = 0.1\%$ physical $\Rightarrow p_L \approx A\,(0.001/0.01)^4 = A \times 10^{-4} \approx 10^{-5}$ per logical cycle (for $A \approx 0.1$)

Key QEC concepts

Logical qubit: encoded in 2d²−1 physical qubits (d² data + d²−1 ancilla); protected from errors up to weight ⌊(d−1)/2⌋, any combination of that many physical errors is correctable.

Syndrome measurement: ancilla qubits measure weight-4 Z⊗Z⊗Z⊗Z and X⊗X⊗X⊗X stabilizers (weight 2 at the boundary), parities only, never the logical state itself.

Threshold theorem: if physical error rate p < p_{th}, increasing code distance d *exponentially suppresses* logical errors — making arbitrarily reliable computation possible from noisy components.

Transversal gates: some logical gates (Pauli, CNOT, Hadamard for certain codes) can be applied bitwise across physical qubits without spreading errors, critical for fault tolerance.

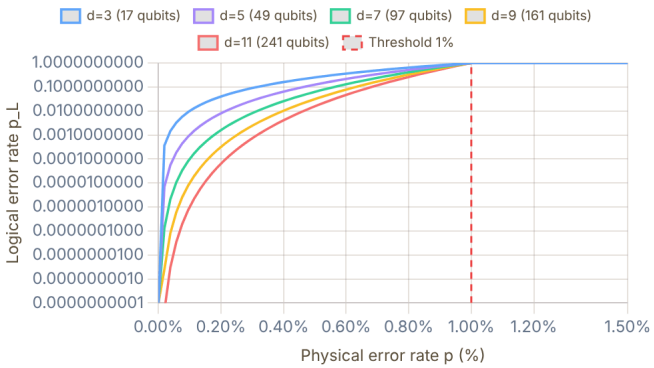
2024–26 State of the Art

Google Willow (2024/25): surface-code memories on a 105-qubit superconductor, demonstrating *exponential* error suppression below threshold as code distance increased up to d=7.

Bluvstein et al. / Harvard (Nature 2024): 48 logical qubits in neutral atom tweezer array; transversal logical gates demonstrated with fault-tolerant circuits.

Quantinuum H2 (2024): 99.9% two-qubit fidelity in trapped ions; 100× error suppression on encoded logical qubit vs raw physical qubit. **Update (2026)**: Quantinuum's Helios (98 qubits, all-to-all) pushed this to a peer-reviewed 99.92% two-qubit fidelity and, separately, 48 fully error-corrected logical qubits.

Surface code: logical vs physical error rate



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Below threshold ($p < 1\%$), larger $d \rightarrow$ exponentially lower p_L

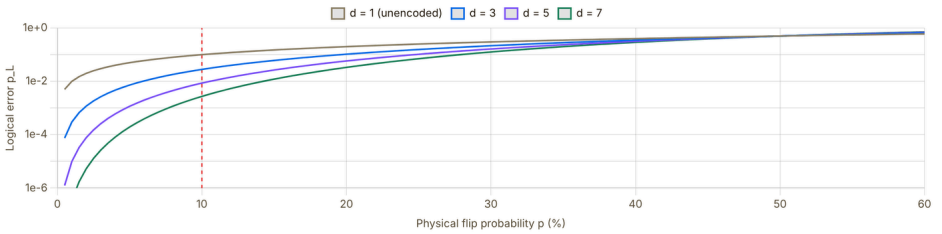
Current neutral-atom gates: ~99.5% fidelity (Evered 2023) $\rightarrow p \approx 0.5\%$, well below the 1% threshold. Target for fault-tolerant useful computation: $p < 0.1\%$.

Interactive

TRY IT: THE REPETITION CODE, AND WHERE A THRESHOLD COMES FROM

Store one bit as d copies. The parity checks between neighbours (the squares) light up wherever two neighbours disagree, which says *where* the flips are without ever reading the stored value. The decoder picks the smallest set of flips that explains the lit checks. Click data qubits to flip them: the code survives any ⌊(d−1)/2⌋ flips, and fails, silently, when more than half flip, because then the smaller explanation is the wrong one.

0 flips · lit checks: none · decoder flips back qubits: none
Stored bit intact, all checks quiet.



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Curves: $p_L = \sum_{k>d/2} \binom{d}{k} p^k (1-p)^{d-k}$ for independent bit flips and perfect parity checks; points: the decoder above run on random errors. Below the crossing a bigger code is exponentially better; above it, worse. Here the crossing sits at 50% because only data qubits err and only bit flips are possible. Real codes must also catch phase flips (the surface code checks both), and every check is itself a noisy circuit, which is why the surface-code threshold in the chart above is about 1%, not 50%.

Physical qubit overhead for one logical qubit

CODE	PHYSICAL / LOGICAL QUBIT	THRESHOLD P_TH	STATUS (2025)
3-qubit bit-flip	`3`	`~50%`	Corrects X errors only, not universal
Steane [[7,1,3]]	`7`	`~1%`	All single-qubit errors; CSS code
Surface code d=5	`49`	`~1%`	Google Willow, below-threshold surface-code memory ✓
Surface code d=7	`97`	`~1%`	Google Willow also reported larger-distance surface-code memory
Neutral-atom logical qubits	`hundreds of atoms`	`code-dependent`	Harvard 2024 used reconfigurable atom arrays and transversal logical gates, not a simple d=7 surface-code row
Surface code d=25	`1249`	`~1%`	Target for Shor's algorithm at scale
Bivariate bicycle (IBM)	`~10–15× fewer`	`~0.5%`	More compact; proposed with simulations (Bravyi et al., Nature 2024)

📄 Fowler et al. 2012, surface code primer 📄 Google Willow 2024 (arXiv) 📄 Bluvstein et al. 2024, 48 logical qubits (Harvard) 🎓 Preskill QEC lecture notes, Caltech (free) 💻 IBM Quantum: QEC course (free)

✨

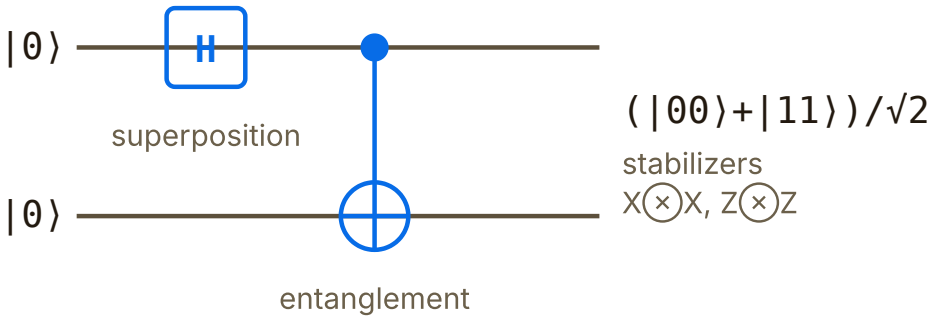
21.17 Magic States & the Clifford Trap

why entanglement is not enough

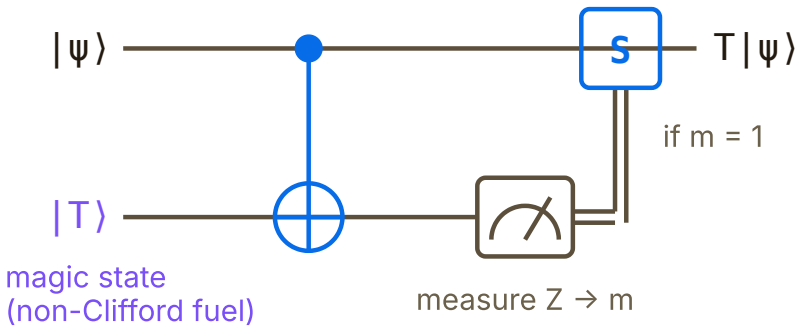
The paradox: superposition + entanglement is not a quantum advantage

A Hadamard makes superposition and a CNOT makes a maximally entangled Bell state, yet any circuit built only from **Clifford gates** (H, S, CNOT or CZ, Paulis, Pauli measurements) on stabilizer inputs can be simulated efficiently on a laptop: the **Gottesman–Knill theorem**. The trick is that such a state is fully described by N stabilizers (for the Bell state, X⊗X and Z⊗Z), a table of about $4N^2$ bits that updates in polynomial time per gate (Aaronson & Gottesman 2004), instead of 2^N complex amplitudes. To outrun classical computers you also need **non-Clifford** operations, and in a fault-tolerant machine those come from consumable **magic states**.

The Clifford trap · classically simulable

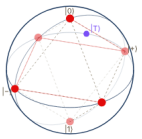


Breaking out · inject a magic state to apply T



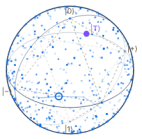
Check of the gadget: with $|\psi\rangle = a|0\rangle + b|1\rangle$ and $|T\rangle = (|0\rangle + e^{i\pi/4}|1\rangle)/\sqrt{2}$, the CNOT gives $a|0\rangle(|0\rangle + e^{i\pi/4}|1\rangle) + b|1\rangle(|1\rangle + e^{i\pi/4}|0\rangle)$. Outcome m = 0 leaves $a|0\rangle + b\, e^{i\pi/4}|1\rangle = T|\psi\rangle$; m = 1 leaves $T^\dagger|\psi\rangle$, and $S\,T^\dagger = T$. Only Clifford operations touch the data qubit.

RANDOM GATE SEQUENCES FROM |0⟩: WHERE CAN THE QUBIT GO?



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Clifford {H, S}: 649 gates · 6 distinct states (always ≤ 6)



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Clifford + T {H, T}: 649 gates · 87% of the sphere visited (200 equal-area cells)

Each step applies a random gate from the set. H and S only make 90° and 180° turns, so the state hops among the six vertices of the *stabilizer octahedron* (dashed), the 24 rotations of the single-qubit Clifford group. T is a 45° turn about z; combined with H it gives a rotation by an angle that is an irrational multiple of π, so the states it reaches fill the sphere densely: that is universality. The purple dot marks |T⟩ = T|+⟩ on the equator at φ = 45°, outside the octahedron. Drag a sphere to rotate it.

Magic states

Clifford gates map the six cardinal states (±x, ±y, ±z on the Bloch sphere) onto each other; their convex hull is the stabilizer octahedron. Any state outside it is **magic**: it cannot be made by Clifford operations on stabilizer states, even with measurements and classical feed-forward. The standard fuel is

$$|T\rangle = T|+\rangle = \tfrac{1}{\sqrt{2}}\left(|0\rangle + e^{i\pi/4}|1\rangle\right).$$

One |T⟩ plus the Clifford gadget above applies one T gate. Bravyi & Kitaev (2005) showed that ideal Clifford operations plus noisy copies of such states are enough for universal computation, provided the copies can be purified. (Naming varies: Bravyi & Kitaev call this equatorial state, up to a Clifford rotation, an |H⟩-type state and reserve "T-type" for the state along $(1,1,1)/\sqrt{3}$.)

Why |T⟩ is the convenient one

It sits exactly 45° along the equator, so teleporting it through the data qubit needs only a CNOT, a Z measurement and, half the time, an S (90°) fix-up. It is also an eigenstate of a Clifford operator:

$$SX\,|T\rangle = e^{i\pi/4}\,|T\rangle,$$

equivalently the +1 eigenstate of $(X + Y)/\sqrt{2}$. Measuring that Clifford checks whether a stored |T⟩ is still a |T⟩ without destroying it, which is what *cultivation* (below) does over and over.

R gates and gate synthesis

The phase-gate family is $R_\phi = \mathrm{diag}(1, e^{i\phi})$: $Z = R_\pi$, $S = R_{\pi/2}$, $T = R_{\pi/4}$ are fixed-angle members (R_k with $\varphi = 2\pi/2^k$ appears in the quantum Fourier transform). R_x, R_y, $R_z(\theta)$ are continuous rotations about the Bloch axes. An error-corrected machine cannot do an arbitrary continuous rotation natively, so it approximates $R_z(\theta)$ by a string of H, S and T: to accuracy ε this takes about $3\log_2(1/\varepsilon)$ T gates (Ross & Selinger), i.e. roughly 30–100 magic states per rotation for ε between 10^{-3} and 10^{-10}.

What the classical simulator pays

state vector: 18 PB
stabilizer tableau: 1.3 kB
Clifford-only: polynomial time

State vector: 2^N amplitudes × 16 bytes. Stabilizer tableau: $2N(2N+1)$ bits. With t T gates, the best stabilizer-based simulators (Bravyi & Gosset 2016) cost ~$2^{0.23t}$ (sampling) to $2^{0.5t}$ (one probability) times polynomial factors: a few T gates are cheap, hundreds are not. Exponential in the number of *magic* gates, not in the entanglement.

Purifying quantum resources: distillation vs cultivation

Noisy resources are cleaned up by spending many imperfect copies (or many checks) to get fewer good ones. Three protocols share the word, but they work at different layers:

	BELL-PAIR (ENTANGLEMENT) DISTILLATION	MAGIC-STATE DISTILLATION	MAGIC-STATE CULTIVATION
Purpose	High-fidelity entangled pairs between two nodes of a network	Clean |T⟩ states for fault-tolerant non-Clifford gates	Clean |T⟩ states, much more cheaply
Needs QEC?	No: runs on raw physical qubits	Yes: runs on logical qubits	Yes: grows inside a small, then larger, code patch
How	Both parties apply local CNOTs between two noisy pairs, measure the sacrificial pair and compare parities over a classical channel; keep the surviving pair only if they agree (Bennett et al. 1996; Deutsch et al. 1996)	Logical Clifford circuits (e.g. the 15-to-1 protocol) merge many noisy |T⟩ copies into one cleaner copy; levels can be stacked (Bravyi & Kitaev 2005)	Inject one physical |T⟩ into a tiny code, repeatedly measure the Clifford check above with stabilizer rounds, discard any run that flags an error, then expand the survivor into a full surface code (Gidney, Shutty & Jones 2024)
Output quality	Fidelity rises each round, limited by local gate errors	Arbitrarily low error by adding levels	Simulated 2×10^{-9} at 10^{-3} circuit noise (4×10^{-11} at 5×10^{-4}); the authors suggest further distillation may never be needed. First hardware demo: fidelity ≥ 0.9999(1), 40× lower error, keeping 8% of attempts (Rosenfeld et al., *Science* 2026)
Cost	At least two pairs in, one out per round, with probabilistic success	Large "factories" of many logical qubits	About the gate count of one lattice-surgery CNOT of equal reliability

Try one round of entanglement distillation in the Remote Entanglement guide. The hardware context for cultivation is on the Google vs IBM page.

📄 Aaronson & Gottesman 2004, stabilizer simulation (PRA 70, 052328) 📄 Bravyi & Kitaev 2005, Clifford gates + noisy ancillas (PRA 71, 022316) 📄 Bravyi & Gosset 2016, Clifford-dominated simulation (PRL 116, 250501) 📄 Ross & Selinger, optimal Clifford+T approximation of z-rotations 📄 Gidney, Shutty & Jones 2024, magic state cultivation (arXiv) 📄 Rosenfeld et al. 2026, cultivation on hardware (*Science*)



21.18 Further Reading & Resources

curated library

Textbooks & Free Courses

Nielsen & Chuang, Quantum Computation and Quantum Information The standard reference. Ch. 1–4 for foundations, Ch. 10 for QEC. ("Mike & Ike")

Quantum Country (Matuschak & Nielsen) Free online "mnemonic medium", read and remember via spaced repetition. Best first read.

IBM Quantum Learning Free structured courses from IBM: basics → algorithms → error correction → hardware.

Preskill Ph229 Lecture Notes (Caltech) Free 9-chapter notes on quantum information, error correction, and fault tolerance.

Aaronson: Quantum Computing Since Democritus (lecture notes) Excellent conceptual treatment, complexity theory and quantum foundations.

AMO Physics References

Steck Alkali Data Sheets (all species) Gold standard: Rb, Cs, Li, Na, K, transition frequencies, linewidths, hyperfine constants.

Saffman 2016: Neutral Atom Quantum Computing Review Comprehensive review of Rydberg-based QC. Essential reading for AMO PhD students.

Kaufman & Ni 2021: Atomic Array Review Optical tweezer arrays, from single-atom loading to many-body entanglement.

QuTiP, Quantum Toolbox in Python Open-source Python: simulate Lindblad dynamics, Rabi oscillations, Wigner functions, and more.

Browaeys & Lahaye 2020: Many-body physics with individually controlled Rydberg atoms State-of-the-art review of programmable Rydberg simulators and quantum gates.

Landmark Experimental Papers

Jaksch et al. 2000, original Rydberg gate proposal (PRL 85, 2208) The foundational theory paper proposing Rydberg blockade for neutral-atom two-qubit gates.

Evered et al. 2023 (Harvard), 99.5% Rydberg CZ gate State-of-the-art neutral-atom gate fidelity with Cs atoms; long T_2 via dynamical decoupling.

Bluvstein et al. 2024 (Harvard), 48 logical qubits Fault-tolerant logical circuits in neutral atom arrays; transversal gates demonstrated.

Google Willow 2024, surface code below threshold First exponential suppression of logical errors with increasing code distance.

Fowler et al. 2012, surface code for quantum computation The definitive introduction to surface codes. Essential for understanding fault tolerance.

Interactive Tools & Simulators

Quirk, Browser Quantum Circuit Simulator Drag-and-drop gates; real-time Bloch sphere, probability amplitudes, entanglement display.

IBM Quantum Composer Circuit composer with free access to real IBM quantum hardware (127–433 qubit devices).

QuTiP (Python) Simulate open systems, master equations, Rabi oscillations, Wigner functions in Python.

Quantum Computing Stack Exchange Q&A forum, expert answers on everything from gate sets to QEC to AMO hardware.

arXiv quant-ph, daily new papers Stay current, essentially all QC papers appear here before journal publication.

CHAPTER 21 SUMMARY

These fifteen core concepts form the shared vocabulary of quantum information processing. In real hardware, however, every qubit interacts with its environment—so we next study decoherence quantitatively.

Problems for Chapter 21

1. USE THE TOOL In the Rabi oscillation explorer set $\Omega/2\pi = 1$ MHz and $\Delta = 0$ and read off the π-pulse time. Then set $\Delta/2\pi = 1$ MHz. How do the maximum excited-state population and the oscillation frequency change?
2. BY HAND Derive $P_{max} = \Omega^2/(\Omega^2 + \Delta^2)$ and the generalised Rabi frequency $\Omega' = \sqrt{(\Omega^2 + \Delta^2)}$ and check both against your readings in problem 1.
3. DESIGN With optimal angles the CHSH game is won with probability $\cos^2(\pi/8) \approx 0.854$, against 0.75 for any classical strategy. Play 1,000 rounds. How many rounds would you need to separate the quantum and classical win rates by five standard deviations?

Use the tool: on the interactive page. By hand: from the formulas in this chapter. Design: a decision a lab would make. Worked solutions are available to instructors on request (see Part V).

CHAPTER 22 · QUANTUM STEP 02

22

Decoherence & Optical Bloch Equations

🌀 Decoherence Lab · Deep dive

A resonantly driven two-level atom, watched two ways at once — as a vector on the Bloch sphere, and as the population trace a detector would actually show. Four regimes, one equation: Steck's optical Bloch equations, integrated live, not pre-baked.

> *Ideal unitary gates never exist in isolation: spontaneous emission, magnetic fluctuations, and trap-induced light shifts continuously shrink the Bloch vector through relaxation (T_1) and dephasing (T_2*, T_2).*

IN THIS CHAPTER Prerequisites: Ch. 21

- Distinguish energy relaxation T_1, inhomogeneous free-induction decay T_2*, and homogeneous coherence time T_2 (with $T_2 \le 2T_1$).
- Solve the optical Bloch equations to see how dissipation and pure dephasing deform Bloch-sphere trajectories.
- Match experimental protocols (population decay, Ramsey fringes, Hahn echo) to the specific noise mechanism each isolates.

WEBSITE COMPANION Live interactive page ↗ · Week 11

22.1 The Bloch sphere, refreshed

State, drive, and what the axes mean

Two levels $|g\rangle, |e\rangle$, driven by a laser with Rabi frequency Ω and detuning $\Delta = \omega - \omega_0$. Package the state as a Bloch vector $\langle \boldsymbol{\sigma} \rangle = (\langle \sigma_x \rangle, \langle \sigma_y \rangle, \langle \sigma_z \rangle)$: **z** is the population inversion $\rho_{ee} - \rho_{gg}$ (−1 = all $|g\rangle$, +1 = all $|e\rangle$); **x, y** are the two quadratures of the atomic dipole, only nonzero for a superposition. The whole vector precesses like a gyroscope, $\partial_t \langle \boldsymbol{\sigma} \rangle = \boldsymbol{\wp} \times \langle \boldsymbol{\sigma} \rangle$, $\boldsymbol{\wp} = \Omega \hat{x} - \Delta \hat{z}$ (Steck, *Quantum and Atom Optics*, §5.4).

$$\partial_t \langle \sigma_x \rangle = \Delta \langle \sigma_y \rangle - \gamma_\perp \langle \sigma_x \rangle$$
$$\partial_t \langle \sigma_y \rangle = -\Delta \langle \sigma_x \rangle - \Omega \langle \sigma_z \rangle - \gamma_\perp \langle \sigma_y \rangle$$
$$\partial_t \langle \sigma_z \rangle = \Omega \langle \sigma_y \rangle - \Gamma(\langle \sigma_z \rangle + 1)$$

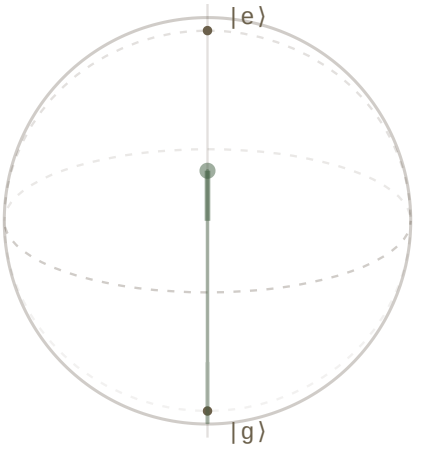
Steck §5.5, "the optical Bloch equations." Every dual-view panel below runs this exact system, integrated numerically (RK4) in real time with different $(\Gamma, \gamma_\perp, \Delta)$ — nothing is a pre-rendered animation. $\Gamma = 1/T_1$ (longitudinal/energy decay), $\gamma_\perp = 1/T_2$ (transverse/coherence decay), $\gamma_\perp = \Gamma/2 + \gamma_c$ with γ_c the "extra" dephasing.

22.2 Coherent driving — the baseline

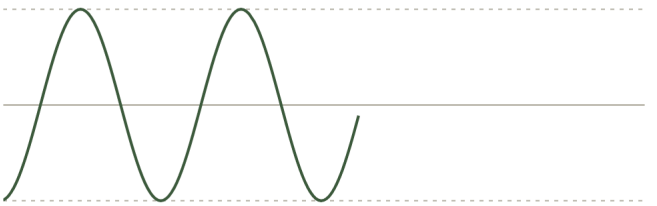
No environment at all: $\Gamma = 0$, $\gamma_\perp = 0$, on resonance. Left: the vector orbits a closed great circle forever. Right: the population inversion it produces — an undamped sine that never loses contrast.

Rabi flopping Γ=0, $\gamma_\perp$=0, unitary

Bloch sphere



Σ
z t



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

σ_z(t) = -0.11

$$\sigma_z(t) = -\cos(\Omega t), \qquad P_e(t) = \sin^2\left(\frac{\Omega t}{2}\right)$$

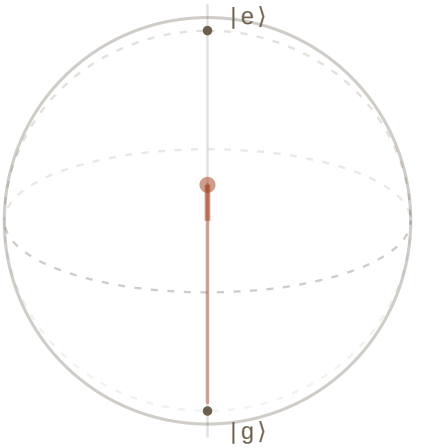
Steck §5.2. Full-contrast Rabi flopping, forever — the trajectory every regime below decays away from.

22.3 Dissipation — T_1

Turn on spontaneous emission: $\Gamma > 0$, and with no extra dephasing $\gamma_\perp = \Gamma/2$ exactly (natural/homogeneous broadening). Real energy now leaves into the vacuum field every time the atom sits in $|e\rangle$. Watch several full oscillations decay before the sweep resets.

Damped Rabi oscillation $1/T_1 = \Gamma$

Bloch sphere



Σ
z t



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

σ_z(t) = -0.18

$$\sigma_z(t) = -1 + \frac{\Omega^2}{\Omega^2 + \Gamma^2/2}\left[1 - e^{-3\Gamma t/4}\left(\cos\Omega_\Gamma t + \frac{3\Gamma}{4\Omega_\Gamma}\sin\Omega_\Gamma t\right)\right], \qquad \Omega_\Gamma = \sqrt{\Omega^2 - (\Gamma/4)^2}$$

Torrey's exact resonance solution, Steck §5.5.2 (starting in $|g\rangle$). Same oscillation as §02, now spiraling into the ball and settling on the pole. Ω_Γ and the $e^{-3\Gamma t/4}$ envelope are exactly what you fit to Rabi-flop data to extract T_1.

What actually happens physically

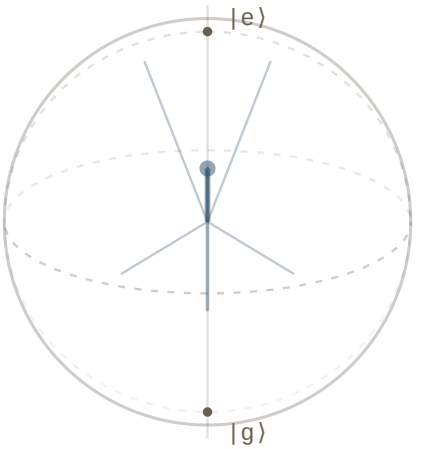
The atom couples to the continuum of vacuum electromagnetic modes. Each time it's in $|e\rangle$ it can emit a real photon into a mode you're not tracking — the energy genuinely leaves the two-level system. That's why **z moves**: population flows one way, $|e\rangle \to |g\rangle$, irreversibly.

22.4 Dephasing — T_2^* / T_φ

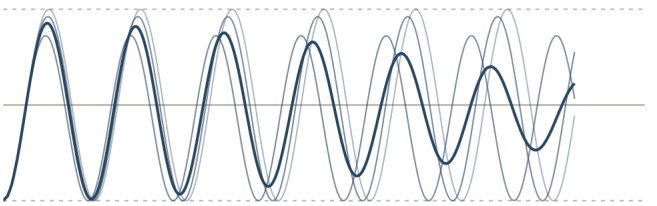
Set $\Gamma = 0$ again — no energy loss anywhere — but now five atoms (or five repeated shots) each see a slightly different detuning Δ_i: laser-linewidth jitter, a field gradient, spatial intensity variation across a tweezer array. Watch the five thin traces: **none of them decay**. What decays is only their average.

Ensemble of 5, detuning spread ±0.4Ω Γ=0, $1/T_\varphi$

Bloch sphere — 5 members + live average



Detector signal — 5 members + live average



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

⟨σ_z⟩ = 0.22 (n=5 ensemble)

The subtlety worth being precise about

Free precession by itself — one atom rotating around $\hat{z}$ at its own detuning rate — is **not** decoherence. It's exactly reversible; a single trajectory's phase is perfectly well-defined the whole time (Steck §5.4: "free precession... about $\hat{z}$ at the detuning rate"). Dephasing only becomes a real decoherence mechanism once you average over a spread of rates you can't track shot-to-shot (Steck §5.4.2).

That's also why it's partly fixable: a π-pulse (spin/photon echo) reverses each atom's phase and lets the faster ones catch up to the slower ones, rephasing at time $2T$. This works for the *static* spread shown here — it defines $\mathbf{T_2^*}$ — but not for dephasing that's genuinely random in time (collisions, laser phase noise), which defines the intrinsic γ_c folded into T_2.

$$\gamma_\perp = \frac{\Gamma}{2} + \gamma_c \quad (\gamma_c\text{: homogeneous/intrinsic — laser phase noise, collisions — not echo-recoverable})$$

$$T_2^* \le T_2 \quad \text{once a static ensemble/shot-to-shot spread is folded in (echo-recoverable back toward } T_2)$$

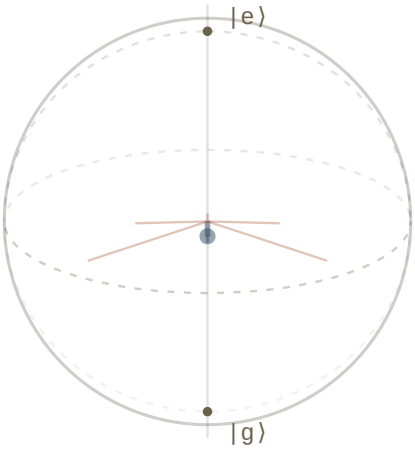
Steck §5.5 (definitions), §5.4.2 (echo). Each member above: exact undamped generalized-Rabi solution $\tilde{\Omega}_i = \sqrt{\Omega^2 + \Delta_i^2}$, Steck §5.2.2 — no approximation.

22.5 Decoherence — T_2, both at once

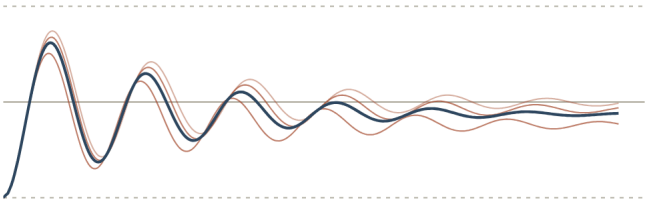
Real experiments have both: $\Gamma > 0$ *and* a Δ_i spread. The same five atoms individually relax toward $|g\rangle$ while also drifting out of phase with each other — the ensemble average spirals inward and down at once, over several visible oscillations before settling.

Both mechanisms, same ensemble $1/T_2$ (net)

Bloch sphere — combined



Detector signal — combined



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

⟨σ_z⟩ = -0.12 (n=5 ensemble)

$$\frac{1}{T_2} = \frac{1}{2T_1} + \frac{1}{T_\phi}$$

Decoherence isn't a third mechanism — it's the name for whatever combination of dissipation and dephasing is acting on $|\rho_{eg}|$. The trace color blends rust→blue as a visual reminder of that sum, not a claim about which one "wins" first.

Same physics, a different textbook's letters

Fox, *Quantum Optics: An Introduction* (Ch. 9, Fig. 9.10), draws this exact pair of Bloch-sphere pictures under the labels T_2' (transverse) and T_1 (longitudinal) instead of T_ϕ and T_1. T_2' is Fox's name for the same pure/intrinsic dephasing time called γ_c/T_ϕ here, and his relation $1/T_2 = 1/(2T_1) + 1/T_2'$ is the identical equation boxed above — just a prime instead of a ϕ subscript. Don't let the different letters imply different physics across books.

Fox's caption states two things worth making explicit here too. First: "the Bloch vector of the relaxed state is only meaningful for the entire ensemble rather than for individual atoms" — true of every spiral and fan on this page. A single atom never smoothly shrinks; it stays a pure state on the Bloch sphere's surface and jumps discontinuously to $|g\rangle$ at one random instant. The smooth curves you're watching are what averaging over many such jumps — or many repeated shots — looks like; a single-trajectory ("quantum jump") picture of §03's decay would look completely different from its smooth panel here. Second: "longitudinal decay inevitably causes transverse relaxation as well" — exactly what $\gamma_\perp = \Gamma/2 + \gamma_c \geq \Gamma/2$ already says in §01: T_1 decay always costs you some coherence too, never only population.

22.6 How you'd actually measure each one

PROTOCOL	ISOLATES	SIGNATURE
Rabi flopping, drive left on	a mix of both	damped oscillation: Ω_Γ and $e^{-3\Gamma t/4}$ envelope (§03)
Population decay, drive off	T_1 only	$\sigma_z(t) \to -1$ exponentially at rate Γ
Ramsey ($\pi/2$ – wait T – $\pi/2$)	T_2*	fringe contrast decays with T (§5.4.1); the atomic-clock protocol
Hahn echo ($\pi/2 - T - \pi - T$ – detect)	T_2 (removes the static T_2* spread)	echo amplitude recovers signal lost to inhomogeneous dephasing; extended by CPMG / DD sequences on the Dynamical Decoupling playground

22.7 Cheat sheet

MECHANISM	MOVES Z?	BLOCH PICTURE	ECHO-RECOVERABLE?	TYPICAL CAUSE
Dissipation (T_1)	Yes	spiral to pole	No	spontaneous emission, off-resonant scattering
Dephasing, intrinsic (γ_c)	No	shrink toward axis, fixed latitude	No	laser phase noise, collisions
Dephasing, inhomogeneous (T_2^*)	No	fan-out, fixed latitude	Yes (spin echo)	detuning / light-shift spread across ensemble or shots
Decoherence (T_2, net)	whatever T_1 does	combination	partially	whatever mix of the above is actually present

All four panels integrate Steck's boxed optical Bloch equations (§01) numerically (RK4) in real time — the "combined" panel in §05 is not a separately-derived closed form, it's the same equations with both $\Gamma>0$ and $\Delta\neq0$ plugged in at once. Reduced-motion viewers see one representative frame per panel instead of the live sweep.

22.8 References & Further Reading

📘 Steck, *Quantum and Atom Optics* — Ch. 5 (the optical Bloch equations, coherent Rabi flopping, Torrey's damped-resonance solution, dephasing): the primary source for every closed-form result on this page 📘 Fox, *Quantum Optics: An Introduction* (Oxford Master Series), Ch. 9, Fig. 9.10 — cross-checked against §04's dephasing/T2' relation 📄 Zurek (2003) — decoherence, einselection, and the quantum origins of the classical 🔵 Quantum Computing Learning Path — Decoherence: T_1 & T_2 (the short-form version of this page) 🛡️ Dynamical Decoupling — Hahn echo, CPMG, and filter functions for extending T_2

CHAPTER 22 SUMMARY

Measuring T_1, T_2^*, and T_2 tells us which noise channels dominate our atomic memory. Once slow dephasing is diagnosed, we can actively suppress it using dynamical decoupling pulse sequences.

Problems for Chapter 22

1. USE THE TOOL Watch Sections 03–05. Sketch the Bloch-vector trajectory and the σ_z signal for pure T_1 decay, pure dephasing and both together, and mark where the vector ends up in each case.
2. BY HAND Show that $T_2 \leq 2T_1$ follows from $1/T_2 = 1/2T_1 + 1/T_\varphi$.
3. DESIGN A Ramsey signal decays as a Gaussian in time, while the echo signal hardly decays. What does this say about the noise, and which measurement would you do next?

Use the tool: on the interactive page. By hand: from the formulas in this chapter. Design: a decision a lab would make. Worked solutions are available to instructors on request (see Part V).

CHAPTER 23 · QUANTUM STEP 03

23

Dynamical Decoupling & Noise Spectroscopy

🛡️ Dynamical Decoupling · Simulator

How π-pulse sequences extend qubit coherence by engineering the spectral filter function — from Hahn echo to XY-16, with live filter function and T_2 enhancement calculators. Physics context: Cs-133 neutral atom arrays (Manetsch 2025, T_2 = 12.6 s).

> *Because magnetic gradients and slow laser drifts fluctuate on millisecond-to-second timescales, applying timed π-pulse trains refocuses accumulated phase errors and extends coherence by orders of magnitude.*

IN THIS CHAPTER

Prerequisites: Ch. 22

- Derive the filter-function formalism $\chi(T) = (1/2\pi) \int S(\omega)|\tilde{Y}(\omega,T)|^2 \, d\omega$ for Ramsey, Hahn echo, and CPMG-N sequences.
- Compare CPMG with error-compensating XY-4, XY-8, and XY-16 sequences under spatially inhomogeneous Rabi frequencies.
- Use multi-pulse coherence decay to reconstruct the environmental noise power spectral density $S(\omega)$.

WEBSITE COMPANION Live interactive page ↗ · Notebook: Open in Colab · download dd-playground.ipynb · Week 11

01 / The Core Problem

23.1 T_2^* vs T_2: Why Coherence Decays

Three timescales govern qubit memory: energy relaxation T_1, true dephasing T_2, and the free-induction-decay time T_2^*. Understanding how they differ is the starting point for dynamical decoupling.

Timescale Hierarchy

T_1 (energy relaxation): spontaneous emission, photon scattering, sets a hard ceiling. Always $T_2 \leq 2T_1$.

T_2 ("true" T_2): dephasing from dynamic noise that changes during the experiment. A spin-echo can refocus the static part; only the time-varying part survives.

T_2^* (FID time): the apparent dephasing seen in free induction decay — includes *both* static inhomogeneity across the ensemble AND dynamic noise. Almost always $T_2^* \ll T_2 \leq 2T_1$.

- Static inhomogeneity: B-field gradient, differential light shifts, laser frequency drift
- Dynamic noise: magnetic field fluctuations, motional dephasing, laser phase noise
- In neutral-atom arrays: T_2^* limited by inter-atom B-field gradient, T_1 can be $\gg T_2$

Free Induction Decay Models

For Gaussian inhomogeneous broadening (dominant in most neutral atom experiments):

$$W_{\mathrm{FID}}(T) = e^{-(T/T_2^*)^2}$$

For Lorentzian broadening (Markovian / exponential noise):

$$W_{\mathrm{FID}}(T) = e^{-T/T_2^*}$$

Both collapse to a simple envelope, but the physical origin differs. The key: T_2^* is a *measured* number. T_2 requires refocusing (spin echo or DD).

Cs-133 T_2^*

14.0 ms

Free Ramsey dephasing, Manetsch 2025

T_2 with XY-16

12.6 s

DD-measured T_2 (pulse error limited)

T_1 (Cs-133)

$\gg T_2$

Spontaneous emission negligible

Fundamental limit

$T_2 \leq 2T_1$

Irreducible energy relaxation

> **Note on Manetsch 2025:** The free Ramsey dephasing time is $T_2^* = 14.0(1)$ ms, while dynamical decoupling extends coherence to $T_2 = 12.6(1)$ s. This is exactly the regime where DD matters: refocusing suppresses slow magnetic-field and light-shift noise, but pulse errors eventually become the new limit.

The General Framework: Phase Accumulation

A qubit in a fluctuating field $\delta\omega(t)$ accumulates random phase:

$$\varphi = \int_0^T y(t)\,\delta\omega(t)\,dt$$

where $y(t) = \pm 1$ is the **switching function** determined by the pulse sequence. Coherence in the Gaussian noise approximation:

$$W(T) = \langle e^{i\varphi} \rangle = e^{-\chi(T)}$$

$$\chi(T) = \frac{1}{2\pi} \int_0^\infty S(\omega)\,|\tilde{Y}(\omega, T)|^2\,d\omega$$

Here $S(\omega)$ is the one-sided noise power spectral density [rad²/s] and $\tilde{Y}(\omega, T) = \int_0^T y(t)\,e^{i\omega t}\,dt$ is the **filter function**. T_2 is defined by $\chi(T_2) = 1$, i.e., $W(T_2) = 1/e$.

02 / Measuring T_2^*

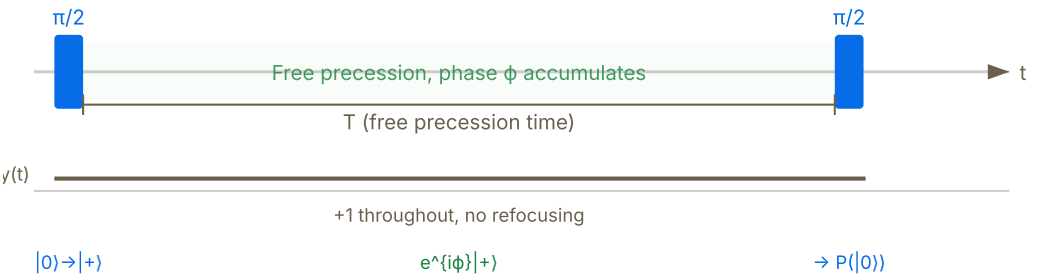
23.2 Ramsey Spectroscopy: The Free-Precession Experiment

Ramsey spectroscopy is the fundamental technique for measuring qubit coherence. Before you can extend T_2, you need to measure T_2*, and Ramsey is how you do it. It also reveals *what* is dephasing your qubit, pointing toward which DD sequence to use.

The Ramsey Sequence

The sequence is: **$\pi/2$ pulse → free evolution T → $\pi/2$ pulse → measure**.

1. **First $\pi/2$ pulse:** rotates qubit from $|0\rangle$ to the equator of the Bloch sphere, creating superposition $|+\rangle = (|0\rangle + |1\rangle)/\sqrt{2}$.
2. **Free evolution T:** the qubit precesses at its Larmor frequency. Any detuning δ from the drive frequency causes it to rotate in the equatorial plane. Noise $\delta\omega(t)$ causes random phase $\varphi = \int_0^T \delta\omega\, dt$.
3. **Second $\pi/2$ pulse:** maps the equatorial phase back to a population difference, the phase is now readable.
4. **Measure:** $P(|0\rangle) = (1 + \cos(\delta T + \varphi))/2$. For a single detuning δ, this oscillates as T increases ("Ramsey fringes"). For an ensemble with many δ values, the fringes wash out at rate $1/T_2$*.



Ramsey sequence. The switching function y(t) = +1 throughout, phase from all fields (static AND dynamic) accumulates. The second $\pi/2$ converts the accumulated phase into a measurable population difference.

What You Measure: Ramsey Fringes & T_2*

If you apply a drive at detuning $\delta = \omega_{\text{drive}} - \omega_{\text{qubit}}$, the qubit precesses at δ during T. Sweeping T gives oscillating fringes:

$$P(|0\rangle) = \frac{1}{2}\left[1 + \cos(\delta T)\, e^{-(T/T_2^*)^n}\right]$$

$n = 2$: Gaussian broadening (static disorder), sharp fringe contrast loss at $T \sim T_2^*$.
$n = 1$: Lorentzian (Markovian dynamic noise), exponential decay envelope.

Reading the Ramsey Signal

- **Fringe frequency → detuning δ** (how far off-resonance you are). Use this to lock to qubit frequency.
- **Fringe contrast decay → T_2*** (dephasing time). The contrast $|C(T)| = e^{-(T/T_2^*)^n}$ gives you T_2* directly from the decay envelope.
- **Fringe shape (Gaussian vs exponential):** Gaussian decay → static inhomogeneity dominates (e.g., B-field gradient across array). Exponential decay → dynamic noise dominates (e.g., laser phase noise). This tells you whether DD will help dramatically (Gaussian case, where static gradient is the enemy) or only modestly (exponential case).
- **Multiple revivals:** If fringe contrast dips then revives, there is a discrete noise source (e.g., a 50 Hz power line pickup).

> **In neutral atom arrays:** Run a Ramsey experiment on all ~100–6100 atoms simultaneously. Each atom is an independent clock. The ensemble-averaged signal washes out at T_2*. But even *single-atom* Ramsey (no ensemble average) shows coherence decay, this is the true single-qubit T_2*. Manetsch 2025 measured free Ramsey dephasing T_2* = 14.0(1) ms and extended the coherence to T_2 = 12.6(1) s in their 6100-atom Cs-133 array using dynamical decoupling.

Ramsey vs Spin Echo: What Each Measures

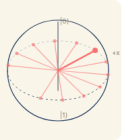
Ramsey (FID): $\varphi = \int_0^T \delta\omega(t)\,dt$, integrates ALL noise (static + dynamic) with $y(t) = +1$. Measures T_2^*.

Spin echo: $\varphi = \int_0^{T/2} \delta\omega\,dt - \int_{T/2}^{T} \delta\omega\,dt$, subtracts the two halves. Static fields cancel; only noise that changes between the two halves survives. Measures a longer $T_2^{\text{echo}} \geq T_2^*$.

The ratio T_2^{echo}/T_2^* tells you how much of the dephasing is static (large ratio → mostly static → gradient problem → DD is very effective) vs dynamic (ratio ≈ 1 → dynamic noise dominates → DD helps less). In alkali tweezer arrays: the ratio can be 10–100× when static B-field gradients dominate.

FID Dephasing vs CPMG Refocusing — side by side

Free Induction Decay (no DD)



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

t=1.22s |C|=0.00

CPMG-4 (4 π-pulses)



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

t=1.22s |C|=0.85

FID: spins fan out and never refocus. CPMG: π-pulses flip the fan, reversing dephasing.

03 / The Simplest Refocusing

23.3 Hahn Echo: One π-Pulse

The spin echo (Hahn 1950) uses a single π-pulse at t = T/2 to refocus all *static* field inhomogeneity. It is the prototype for all dynamical decoupling.

How It Works

The π-pulse at t = T/2 flips the qubit, making $y(t) = +1$ for $t < T/2$ and $y(t) = -1$ for $t > T/2$.

For a *constant* field $\delta\omega = \text{const}$:

$$\varphi = \delta\omega \int_0^{T/2} dt - \delta\omega \int_{T/2}^{T} dt = 0$$

The static contribution cancels exactly. Only noise that *changes* between the two halves (i.e., noise at $\omega \sim 1/T$) survives.

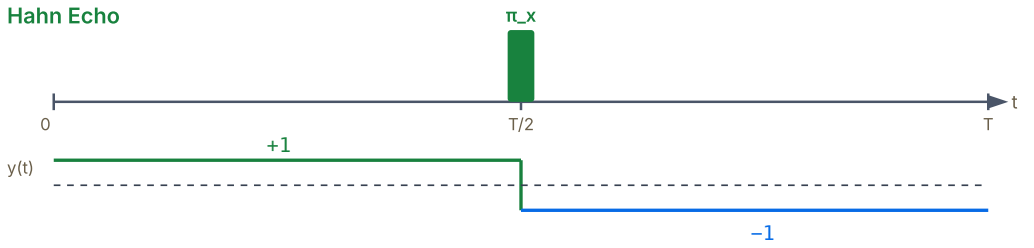
Filter Function

The Hahn echo filter function:

$$|\tilde{Y}_{\mathrm{HE}}(\omega, T)|^2 \propto \frac{\sin^4(\omega T/4)}{\omega^2}$$

Key properties:

- Zero at $\omega = 0$, static fields refocused
- Peak near $\omega \approx 4.7/T$ (where $\tan(\omega T/4) = \omega T/2$), sensitive to noise at the echo timescale
- High-pass character: suppresses $\omega \ll 1/T$ noise



Refocusing intuition: Imagine all atoms in an ensemble processing at slightly different rates due to B-field gradients. At T/2, the π-pulse reverses their relative phases. At T, all atoms realign, the "echo", regardless of their individual offset frequencies. In the ideal static-offset model, this refocusing is exact as long as the offsets do not change during [0,T].

03 / N-Pulse Sequences

23.4 CPMG: N Equally-Spaced π-Pulses

The Carr-Purcell-Meiboom-Gill (CPMG) sequence is the standard multi-pulse dynamical decoupling protocol. N π-pulses, equally spaced, push the filter function peak to higher and higher frequencies.

Pulse Timing

N π-pulses at positions:

$$t_k^{\mathrm{CPMG}} = \frac{(2k-1)\,T}{2N}, \quad k = 1, \ldots, N$$

The inter-pulse spacing is $\tau = T/N$. The first pulse is at $\tau/2$, then every τ.

The switching function $y(t)$ alternates between $+1$ and -1 at each pulse, starting at $+1$.

Filter Function

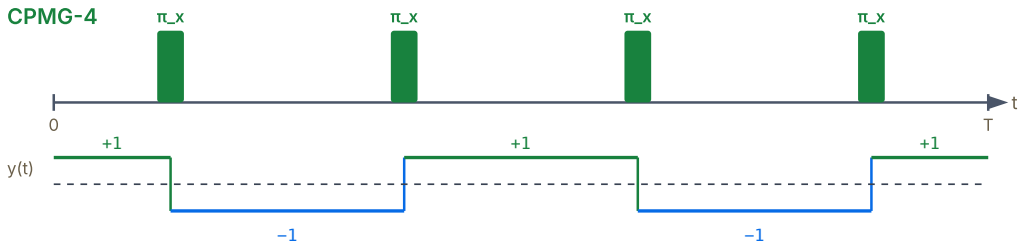
$$|\tilde{Y}_{\mathrm{CPMG-N}}(\omega, T)|^2 = \frac{16\sin^4(\omega T/4N)\,\sin^2(\omega T/2)}{\omega^2\cos^2(\omega T/2N)} \quad (N \text{ even; use } \cos^2(\omega T/2) \text{ for odd } N)$$

Cywiński et al., PRB 77, 174509 (2008). For $N = 1$ this reduces to the Hahn-echo result $16\sin^4(\omega T/4)/\omega^2$.

Peak location (where $\cos(\omega T/2N) = 0$): $\omega_c = \pi/\tau = \pi N/T$. More pulses → peak shifts to *higher* frequency.

For $1/f$ noise ($S(\omega) \propto 1/\omega$): higher ω_c → less noise power → better coherence:

$$T_2^{\text{CPMG}}(N) \propto N^{\alpha/(1+\alpha)}, \quad S(\omega) \propto \omega^{-\alpha}$$



CPMG caveat: All pulses rotate around the same axis (x). If the π rotation angle is slightly off by ϵ, errors accumulate *coherently*: after N pulses the error is $\sim N\epsilon$ (not

$\sqrt{N}\epsilon$

). This is fine for single-qubit experiments but catastrophic for large arrays where different atoms see different pulse areas.

04 / Robust Sequences

23.5 XY-4, XY-8, XY-16: Cancelling Pulse Errors

XY sequences use the same timing as CPMG but alternate the rotation axis between x and y. This causes pulse errors to cancel to progressively higher order.

XY-4

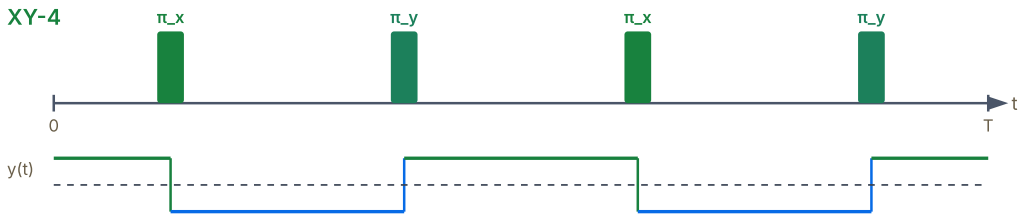
Four pulses in a cycle: $\pi_x\, \pi_y\, \pi_x\, \pi_y$. Same spacing as CPMG-4, the filter function is *identical* for pure z-dephasing.

But: a small over-rotation error on π_x is partially cancelled by the subsequent π_y. First-order pulse errors cancel within one cycle.

XY-8 and XY-16

XY-8: $X\, Y\, X\, Y\, Y\, X\, Y\, X$, an XY-4 cycle followed by its time-reverse (Gullion et al. 1990). Cancels both first- *and* second-order errors (rotation angle AND axis tilt).

XY-16: An XY-8 cycle followed by its phase-inverted copy, $\bar{X}\, \bar{Y}\, \bar{X}\, \bar{Y}\, \bar{Y}\, \bar{X}\, \bar{Y}\, \bar{X}$ (where $\bar{X} = -\pi_x$, a π rotation about $-x$). Third-order error suppression. This is the sequence used in Manetsch 2025 for T_2 = 12.6 s in a 6100-atom Cs-133 array.



Same timing as CPMG-4, different rotation axes

Key distinction: In the ideal perfect-pulse limit, CPMG-N = XY-4 = XY-8 = XY-16 give *identical* coherence decay, the filter function is the same (it only depends on timing). For *imperfect* pulses (real life), XY-16 >> CPMG because pulse errors cancel to higher order.

Sequence Comparison

SEQUENCE	PULSES / CYCLE	ORDERS CORRECTED	FILTER FUNCTION	NOTES
`CPMG-N`	N	None (coherent accumulation)	Same as XY for z-noise	Simple; axis-error builds up linearly
`XY-4`	4	1st order	Identical to CPMG-4	Standard; good for 1D noise axis
`XY-8`	8	2nd order	Identical to CPMG-8	Compensates axis tilt + angle error
`XY-16`	16	3rd order	Identical to CPMG-16	Used in Manetsch 2025 (T_2 = 12.6 s)
`KDD`	16	Higher (multi-axis)	Similar shape	Knill DD; best for >1 noise axis

05 / The Key Insight

23.6 Filter Function Formalism

Every DD sequence acts as a bandpass filter on the noise PSD. Choose the sequence whose filter has minimal overlap with where the noise actually lives.

The Dephasing Integral

All DD sequences obey the same master equation:

$$\chi(T) = \frac{1}{2\pi} \int_0^\infty S(\omega)\, |\tilde{Y}(\omega, T)|^2 \, d\omega$$

T_2 is defined by $\chi(T_2) = 1$. To maximize T_2: engineer $|\tilde{Y}(\omega, T)|^2$ so it has **minimum overlap** with $S(\omega)$.

1/f Noise ($S(\omega) \propto 1/\omega$)

Most AMO systems: B-field fluctuations, laser frequency noise. Noise concentrated at *low* frequency.

Strategy: push filter peak to *higher* ω_c by adding more pulses. Each doubling of N pushes ω_c up by 2×, cutting the overlap with 1/f noise.

Result: $T_2 \propto N^{0.5}$ (for $\alpha = 1$)

White Noise ($S(\omega) = \text{const}$)

Markovian / shot-noise limited. Noise power flat across all frequencies.

DD does nothing for white noise, the integral $\int S\,|\tilde{Y}|^2\, d\omega$ is fixed by the total filter weight, not its peak location.

In this regime T_2 is pulse-error limited. Rabi driving better than DD.

Peaked Noise ($S(\omega) \sim \delta(\omega - \omega_0)$)

Mechanical vibrations, AC power line (60 Hz), laser relaxation oscillation.

Strategy: tune $\tau = \pi/\omega_0$ to place a filter *null* exactly at ω_0. This is "resonance avoidance."

CPMG with specific N and T can null a specific noise frequency exactly.

$$\chi(T) = \frac{1}{2\pi}\int_0^\infty S(\omega)\,|\tilde{Y}(\omega,T)|^2\,d\omega, \qquad T_2 \text{ defined by } \chi(T_2) = 1$$

Physical picture: The switching function $y(t)$ is a square wave. Its Fourier transform $\tilde{Y}(\omega, T)$ has power concentrated around $\omega_c = \pi N/T$. The dephasing $\chi(T)$ is literally a weighted inner product of the noise spectrum and the filter function — minimizing their overlap is the design principle of every DD sequence.

06 / Interactive Calculator

23.7 Filter Function Viewer

Compare filter functions $|\tilde{Y}(\omega, T)|^2$ for FID, Hahn echo, CPMG-4, and CPMG-16 against the noise PSD. Adjust total time T and noise exponent α to see how overlap changes.

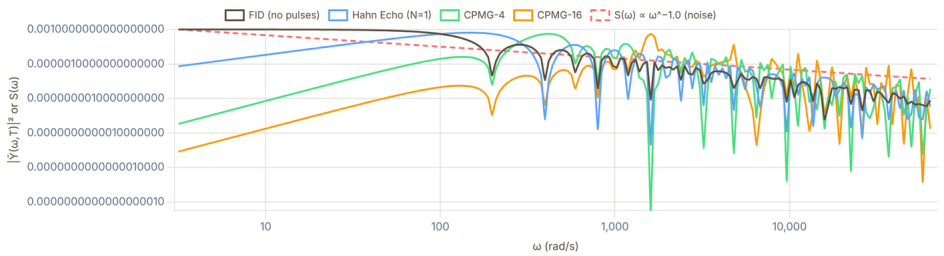
Filter Functions vs Noise PSD LOG-LOG

FID: χ = 1.000

Hahn: χ = 0.104

CPMG-4: χ = 0.031

CPMG-16: χ = 0.008



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Solid lines: filter functions $|\tilde{Y}(\omega,T)|^2$. Dashed line: noise PSD $S(\omega) \propto \omega^{-\alpha}$ (arb. units). Chi values normalized so $\chi_{FID}(T) \approx 1$ at $T = T_2^*$.

07 / T_2 Enhancement

23.8 T_2 vs N: Enhancement Factor Calculator

Compute $T_2(N)/T_2^*$ as a function of the number of CPMG pulses for different noise spectra. Binary search finds T_2 for each N by requiring $\chi(T_2, N) = 1$.

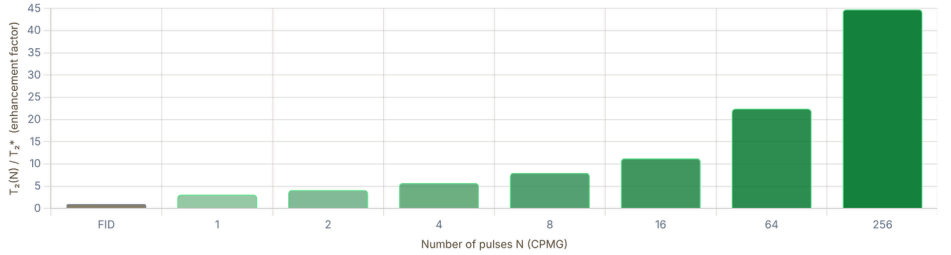
T_2 Enhancement Factor vs Number of Pulses NUMERICAL

N=1: 3.11×

N=16: 11.25×

N=256: 44.79×

Scaling: $T_2 \propto N^{0.49}$



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Enhancement $T_2(N)/T_2^*$ vs N for CPMG sequences. Expected scaling: α=0 → N^0 (no improvement), α=1 → N^0.5, α=2 → N^1.

> **Expected scaling for α = 1.0:** $T_2(N) \propto N^{0.50}$ (theoretical). Measured from fit (N=4→256): N^0.49. 1/f noise (α≈1): square-root improvement, doubling N multiplies T_2 by $\sqrt{2} \approx 1.41$.

08 / Noise Characterization

23.9 Noise Spectroscopy via DD

DD sequences don't just protect qubits, they can also *measure* the noise spectrum that's causing decoherence (Biercuk et al., Nature 2011).

Protocol

1. Fix total time T; sweep N from 1 to ~1000.
2. For each N, measure the coherence W(T,N) and extract χ(T,N) = −ln W.
3. The filter function for CPMG-N has a sharp peak at $\omega_c = \pi N/T$.
4. The noise PSD at ω_c is proportional to χ at that peak:

$$S(\omega_c) \approx \frac{2\pi\,\chi(T,N)}{|\tilde{Y}(\omega_c,T)|^2 \Delta\omega_{\text{peak}}}$$

5. By sweeping T as well, reconstruct $S(\omega)$ over many decades of frequency.

Manetsch 2025 Application

Used XY-16-based sequences on a 6100-atom Cs-133 array to characterize coherence and pulse-error limitations. Identified dominant sources:

- **DC to ~1 Hz:** magnetic field drift (1/f character)
- **~10 Hz–1 kHz:** B-field gradient fluctuations (main T_2^* limitation)
- **>1 kHz:** technical laser/RF noise floor

This spectroscopy directly informs which DD sequence to use and how many pulses N are needed to optimally suppress the noise.

Why this matters for large arrays

In a 6100-atom array, atoms at different positions see slightly different B-fields (gradient). The noise spectroscopy reveals that the gradient fluctuates as a correlated noise source, it's the same noise at every atom, just scaled by position. This means all 6100 atoms can be simultaneously decoupled with the same pulse sequence (XY-16) synchronized across the array.

If the noise were *independent* per atom (e.g., each atom's individual laser scatter), DD would still work but would need to be applied on a per-atom basis.

09 / Experimental Context

23.10 Connection to Neutral Atom Experiments

How dynamical decoupling is used in current state-of-the-art neutral atom qubit experiments.

Manetsch et al., Nature 2025 (arXiv:2403.12021)

System: 6,100-atom Cs-133 array in optical tweezers. Hyperfine qubit: Cs clock-state qubit in the $6S_{1/2}$ ground-state manifold.

DD sequence: XY-16 chosen over CPMG for two reasons: (1) atoms at different positions see different Rabi frequencies (Gaussian laser profile) → spatially varying pulse errors → CPMG errors build up coherently; (2) XY-16 cancels these to 3rd order.

Results: T_2* = 14.0(1) ms from free Ramsey dephasing and T_2(XY-16) = 12.6(1) s under dynamical decoupling. The long-time coherence is therefore a DD-enhanced memory benchmark, not a free-precession T_2*.

Noise source: Primarily B-field gradient fluctuations (dB/dx · x-position of atom). Secondary: differential AC Stark shift from tweezers (light shift noise).

Rb87 vs Yb171 Qubits

Rb87 (Lukin, Greiner, Kaufman groups): Uses $|F=1\rangle \leftrightarrow |F=2\rangle$ hyperfine qubit. First-order Zeeman sensitivity: $\sim 0.7\,\mathrm{MHz/G}$ differential shift → magnetic noise couples directly to qubit → T_2* limited by B-field gradients → DD essential.

Yb171 (Ye, Thompson groups): Uses $|I=1/2, m_I=\pm 1/2\rangle$ nuclear spin qubit. Weak nuclear-spin Zeeman shift ($\sim 0.75\,\mathrm{kHz/G}$ vs $\sim 0.7\,\mathrm{MHz/G}$ for field-sensitive ^{87}Rb transitions, $\sim 933\times$ smaller; note ^{87}Rb's $|1,0\rangle \leftrightarrow |2,0\rangle$ clock state is also first-order insensitive). T_2* already very long → less need for aggressive DD. DD still used but marginal improvement.

General Principles for Large Arrays

As arrays scale to thousands of qubits, two effects make DD increasingly critical:

- **Spatial B-gradient noise:** grows with array size → T_2* shrinks as $\sim 1/\sqrt{N_{\mathrm{atoms}}}$ for random gradient.
- **Pulse inhomogeneity:** atoms at edges of Gaussian beam get systematically different pulse areas → XY-16 (not CPMG) becomes mandatory.
- **Noise spectroscopy:** characterizing the noise spectrum of the full array (not just single atoms) is needed to choose optimal DD parameters.

Array size

6,100

Cs-133 atoms (Manetsch 2025)

DD sequence

XY-16

3rd-order pulse error cancellation

T_2^* (free)

14.0 ms

Free Ramsey dephasing

T_2 (XY-16)

12.6 s

Limited by pulse errors at this scale

Perspective: In most AMO systems (smaller arrays, less-perfect clocks states), $T_2^* \ll T_2$ and DD gives 10–100× improvement. Manetsch 2025 shows how dramatic the gap can be between free dephasing and a carefully engineered DD sequence: milliseconds without refocusing become seconds with XY-16, until pulse errors and technical noise set the new ceiling.

23.11 Sources

Degen, Reinhard & Cappellaro (2017), *Rev. Mod. Phys.* 89 035002 — quantum sensing review; dynamical-decoupling filter functions and noise spectroscopy Álvarez & Suter (2011), *PRL* 107 230501 — measuring the spectrum of colored noise by dynamical decoupling Manetsch et al. 2025, 6100-atom Cs tweezer array and long coherence Viola, Knill & Lloyd (1999), *PRL* 82 2417 — original dynamical decoupling proposal; bang-bang control to suppress decoherence Carr & Purcell (1954), *Phys. Rev.* 94 630 — spin echo sequence; the classical forerunner of all refocusing pulse sequences Biercuk et al. (2009), *Nature* 458 996 — optimized DD sequences for noise suppression in trapped-ion qubits Du et al. (2009), *Nature* 461 1265 — preserving electron spin coherence in solids (irradiated malonic acid) by optimal (Uhrig) dynamical decoupling

CHAPTER 23 SUMMARY

Dynamical decoupling protects idle qubits over seconds-long storage windows. To perform universal quantum computation, we now introduce the entangling mechanism of neutral-atom processors: Rydberg blockade.

Problems for Chapter 23

1. USE THE TOOL Use the enhancement-factor calculator for $\alpha = 1$ and $\alpha = 2$. Fit the slope of $T_2(N)/T_2^*$ against N on a log scale.
2. BY HAND Compare your fitted exponents with $N^{\alpha/(\alpha+1)}$, the scaling expected for $S(\omega) \propto \omega^{-\alpha}$ noise.
3. DESIGN In the filter-function viewer, where is the passband of CPMG with N pulses in total time T? Design a noise-spectroscopy measurement that maps $S(\omega)$ between 1 and 100 kHz.

Use the tool: on the interactive page. By hand: from the formulas in this chapter. Design: a decision a lab would make. Worked solutions are available to instructors on request (see Part V).

CHAPTER 24 · QUANTUM STEP 04

24

Rydberg Blockade Physics

🔮 Rydberg Blockade Lab · Calculator

Choose a Rydberg state, compute single-atom scales, estimate blockade radius, and interpret whether the numbers are useful for neutral-atom gates or quantum simulation.

> *Exciting trapped neutral atoms to high principal quantum numbers ($n \sim 50$–100) makes them enormous and easily polarized: the van der Waals coefficient grows roughly as n^{*11}, so two atoms a few microns apart interact strongly.*

IN THIS CHAPTER Prerequisites: Ch. 1, Ch. 8, Ch. 21

- Estimate quantum defects and n*, the van der Waals coefficient ($C_6 \propto n^{*11}$ within a series), the radiative lifetime of low-l states ($\tau \propto n^{*3}$) and the blockade radius $R_b = (C_6/\hbar\Omega)^{1/6}$.
- Verify the strong-blockade condition $V(R) \gg \Omega$ required to suppress doubly excited $|rr\rangle$ population during two-qubit gates.
- Compare single-photon UV and two-photon excitation schemes across alkali and alkaline-earth-like qubits.

WEBSITE COMPANION Live interactive page ↗ · Notebook: Open in Colab · download rydberg-calculator.ipynb · Week 12

Blockade design path

24.1 Start with the state, then ask whether the interaction is useful.

Rydberg physics is powerful because several quantities scale violently with principal quantum number. The page is now ordered like a lab estimate: pick a state, compute single-atom properties, estimate blockade, then check gate limits and caveats.

01
Choose state
Select species, n, l, j and inspect quantum defect, binding energy, radius, and lifetime.

02
Estimate interaction
Use the quick C_6/blockade model to set a first-pass tweezer spacing and Rabi frequency.

Calculator output at the page's default settings.

03 **Interpret the regime**

Compare blockade, Rabi frequency, decay, and expected gate errors before trusting the number.

04 **Check caveats**

Magnetic sublevels, Förster defects, electric fields, and pair-state mixing require ARC or paper-specific data.

Validity note: binding energies and n* values are reliable quick estimates; C_6 values are order-of-magnitude estimates. Real blockade strengths depend on magnetic sublevel, polarization, quantization axis, Förster defects, pair-state mixing, and electric fields. Use ARC, PairInteraction, or paper-specific C_6 tables before designing an experiment.

EFFECTIVE QUANTUM NUMBER
46.8687
$n^* = n - \delta_0 - \delta_2/(n-\delta_0)^2$

BINDING ENERGY
0.0062 eV 50.0 cm⁻¹

ORBITAL RADIUS ⟨R⟩
174.4 nm
3295 a_0 = 174.4 nm · $\langle r \rangle = [3n^{*2} - l(l+1)]/2 \cdot a_0 \approx 1.5\, n^{*2}\, a_0$

RADIATIVE LIFETIME
147.2 µs
$\tau \approx \tau_0 \times n^{*3}$

Calculator output at the page's default settings.

Quantum defect formula

In alkali atoms the core electrons shift the effective principal quantum number from n to n* = n − δ(n,l,j). The Rydberg–Ritz formula:

$\delta(n, l, j) = \delta_0 + \delta_2 \,/\, (n - \delta_0)^2$

gives the binding energy E_b = R_y / n*² (in eV) with R_y = 13.6057 eV. Mean orbital radius: ⟨r⟩ = [3n*² − l(l+1)]/2 · a_0 ≈ 1.5 n*² a_0 for low-l states. Radiative lifetime: τ ∝ n*³ for low-l states, ∝ n*⁵ for high-l circular states.

Rydberg Blockade Lab

Van der Waals blockade radius R_b for two atoms both driven to the Rydberg state. C_6 values are approximations scaled from published data, use ARC for precision calculations.

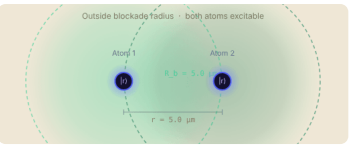
|C_6| APPROX
15.0 GHz · µm⁶

BLOCKADE RADIUS R_B
4.97 µm

INTERACTION AT R
1.0 MHz MHz

Calculator output at the page's default settings.

⚠️ **Not in blockade regime**, r = 5.0 µm > R_b = 4.97 µm. Reduce separation or increase n/Ω.



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Van der Waals blockade

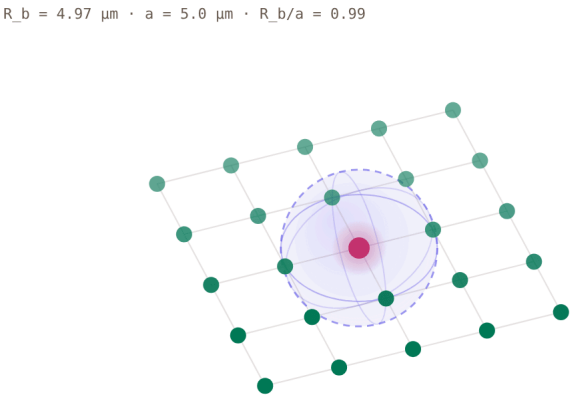
Two atoms in |n,l,j⟩ interact approximately via $U(r) = -C_6/r^6$ in the van der Waals regime and away from near-degenerate pair-state resonances. When |U(r)| ≫ ħΩ, double excitation |rr⟩ is energetically forbidden — the Rydberg blockade. The blockade radius is defined by $|U(R_b)| = \hbar\Omega$:

R_b = (|C_6| / ħΩ)^(1/6)

C_6 often scales approximately as n^{*11} within one series, so doubling n roughly multiplies R_b by $2^{11/6} \approx 3.6$. Treat the sign and magnitude as state-specific: S, P, and D states can differ strongly because of angular factors and Förster resonance structure.

In an array: how many neighbours does one excitation block?

The same R_b and spacing as above, now in a 5 × 5 tweezer array (or two stacked layers, one spacing apart). The pink atom is excited to $|r\rangle$; every atom inside its blockade sphere (indigo) is shifted by $U \geq \hbar\Omega$ and can't follow.



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

excited $|r\rangle$ blockaded ($d < R_b$) free to be excited

$R_b/a = 0.99$: one excitation blocks **0 of 24** other atoms in this array. The spacing is larger than R_b, so neighbours can be excited independently. Lower the atom separation (or raise n) above to bring them inside the sphere.

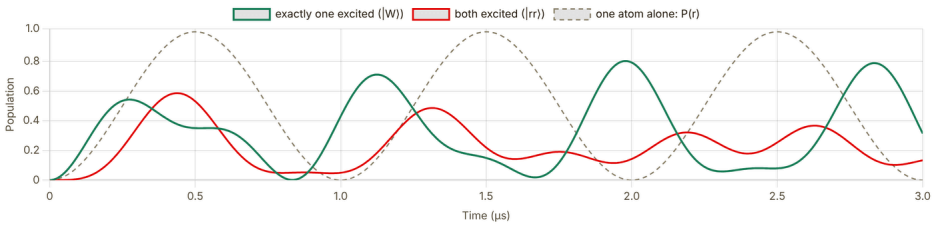
NEIGHBOUR SHELL	ATOMS	U/ħΩ	BLOCKADED	E ≈ $(D/R_B)^{12}$
a = 5.0 µm	4	0.96	no	—
1.414 a = 7.1 µm	4	0.12	no	—
2 a = 10.0 µm	4	0.01	no	—
2.236 a = 11.2 µm	8	7.7e-3	no	—
2.828 a = 14.1 µm	4	1.9e-3	no	—

$U(d)/\hbar\Omega = (R_b/d)^6$ follows from $|U(R_b)| = \hbar\Omega$ and $U \propto 1/d^6$; the last column is the page's own imperfect-blockade estimate, for a pair at that distance. Same caveats as the calculator: isotropic van der Waals interaction with an approximate C_6, no angular dependence or Förster resonances, so treat it as a first-pass layout check.

Interactive

Two atoms, one laser: watch the blockade happen

Drive both atoms with the same Ω (the slider above) and start them in |gg⟩. Without interaction each atom Rabi-flops on its own, so half the time both are excited. With $U \gg \hbar\Omega$ the doubly excited |rr⟩ is pushed off resonance: the pair oscillates between |gg⟩ and the entangled $|W\rangle = (|gr\rangle + |rg\rangle)/\sqrt{2}$ instead, and **$\sqrt{2}$ faster**, because both atoms drive the same single excitation. Move the separation slider across R_b and watch the green curve take over.



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

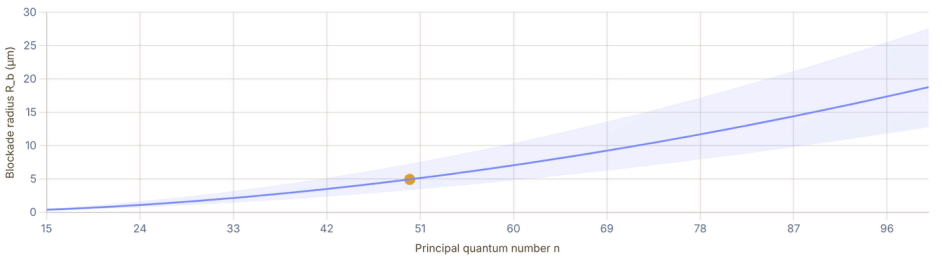
r = 5.0 μm, R_b = 4.97 μm → $U/\hbar\Omega$ = 0.96 · largest P(rr) = 0.586
In between: the pair partly reaches |rr⟩ and the dynamics are not a clean oscillation, which is what an imperfect blockade gate error looks like.

Exact evolution of $H/\hbar = \frac{\Omega}{\sqrt{2}}(|gg\rangle\langle W| + |W\rangle\langle rr| + \text{h.c.}) + U\,|rr\rangle\langle rr|$ on resonance, no decay, with Ω and U(r) = $|C_6|/r^6$ taken from the calculator (same approximate C_6). For $U \gg \hbar\Omega$ the largest |rr⟩ population tends to $\frac{1}{2}(\Omega/U)^2$, so the page's imperfect-blockade estimate $\varepsilon \approx (d/R_b)^{12} = (\Omega/U)^2$ is the conservative side of it. The $\sqrt{2}$ speed-up was first seen with two atoms in tweezers by Gaëtan et al., *Nat. Phys.* 5, 115 (2009); N atoms inside one blockade sphere oscillate at $\sqrt{N}\,\Omega$.

🔬

Rb blockade context: There are multiple experimentally used Rb Rydberg excitation paths. A common textbook/lab route is 780 nm (5S→5P) + 480 nm (5P→nS/nD). The Harvard/QuEra high-fidelity gate in Evered et al. instead used 420 nm + 1013 nm via the 6P manifold to drive ^{87}Rb to $|53S_{1/2}\rangle$. This calculator estimates blockade physics once the chosen Rydberg state and Rabi frequency are specified; the laser wavelengths mainly determine scattering, Doppler, and technical constraints.

BLOCKADE RADIUS R_B VS. PRINCIPAL QUANTUM NUMBER N, RB $NS_{1/2}$, Ω/2Π = 1.00 MHZ



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Shaded region: R_b for Ω/2π = 0.1 – 10 MHz. Dashed line: current Ω. Dot: selected n.

Quantum defects reference, selected species

Values from published literature (Li et al. 2003 for Rb; Goy et al. 1982 / Weber & Sansonetti 1987 for Cs; Lorenzen & Niemax 1983 for Li, Na, K). $n^* = n - \delta_0 - \delta_2/(n-\delta_0)^2$.

SPECIES	SERIES	Δ_0	Δ_2	N*(N=50)
Rb	$nS_{1/2}$	3.1311804	0.1784	46.8687
	$nP_{1/2}$	2.6548849	0.2900	47.3450
	$nP_{3/2}$	2.6416737	0.2950	47.3582
	$nD_{3/2}$	1.3480917	-0.6029	48.6522
	$nD_{5/2}$	1.3464657	-0.5960	48.6538
Cs	$nS_{1/2}$	4.0493532	0.2391	45.9505
	$nP_{1/2}$	3.5914767	0.3500	46.4084
	$nP_{3/2}$	3.5590268	0.3740	46.4408
	$nD_{3/2}$	2.4754562	0.0090	47.5245
	$nD_{5/2}$	2.4663544	0.0120	47.5336
Li	$nS_{1/2}$	0.3995101	0.0290	49.6005
	$nP_{1/2}$	0.0471780	-0.0240	49.9528
	$nP_{3/2}$	0.0471780	-0.0240	49.9528
	$nD_{3/2}$	0.0021290	-0.0120	49.9979
	$nD_{5/2}$	0.0021290	-0.0120	49.9979
Na	$nS_{1/2}$	1.3479690	0.0614	48.6520
	$nP_{1/2}$	0.8554240	0.1119	49.1445
	$nP_{3/2}$	0.8545660	0.1130	49.1454
	$nD_{3/2}$	0.0148670	-0.0410	49.9851
	$nD_{5/2}$	0.0146560	-0.0400	49.9854
K	$nS_{1/2}$	2.1801985	0.1300	47.8197
	$nP_{1/2}$	1.7138920	0.2330	48.2860
	$nP_{3/2}$	1.7201680	0.1930	48.2797
	$nD_{3/2}$	0.2769700	-0.0100	49.7230
	$nD_{5/2}$	0.2770000	-0.0100	49.7230

C_6 scaling and Förster resonances

For like-state pairs $|n\ell j, n\ell j\rangle$, C_6 arises from second-order dipole-dipole coupling to nearly degenerate pair states $|n'l', n''l''\rangle$. C_6 diverges near a **Förster resonance** (exact degeneracy between initial and intermediate pair states), shifting to a C_3/r^3 interaction. Some pair states lie close to a Förster resonance, and a small electric field can tune them onto it for large, controllable interactions (e.g. Ravets et al. 2014, Rb, Browaeys group). Early two-atom blockade experiments used ordinary van der Waals pairs: Gaëtan et al. (2009) used Rb $58D_{3/2}$.

Blackbody radiation limit

At room temperature (T = 300 K, k_BT/h = 6.25 THz), Rydberg states with $\hbar\omega$_n,n' < k_BT are strongly mixed by blackbody photons. Effective lifetime drops below the spontaneous emission limit for n > 30–40 at 300 K. Cryogenic environments (4 K, 77 K) extend T_1 by orders of magnitude. For n = 50 Rb at 300 K: τ_eff ~ 50 μs (vs ~100 μs spontaneous-only).

Rydberg gate fidelity estimate

A Rydberg CZ gate requires: (1) Ω t_gate = π (π-pulse to |r⟩), (2) blockade U ≫ ħΩ, (3) gate time t_gate ≪ τ_Rydberg. Key fidelity-limiting errors:

$$\varepsilon_{\text{se}} \approx \frac{\pi}{\Omega\,\tau_{\text{Ryd}}} \quad \text{[spontaneous emission error]}$$

$$\varepsilon_{\text{blk}} \approx \left(\frac{\hbar\Omega}{U(r)}\right)^2 = \left(\frac{r}{R_b}\right)^{12} \quad \text{[imperfect blockade error]}$$

Typical: $n = 70$, $\Omega = 2\pi \times 1$ MHz, $\tau \approx 400\,\mu$s $\Rightarrow \varepsilon_{\text{se}} \sim 10^{-3}$
$r = 5\,\mu$m, $R_b = 10\,\mu$m $\Rightarrow \varepsilon_{\text{blk}} \sim 10^{-4}$

Two-photon excitation

Direct UV single-photon excitation (e.g., 318 nm for Cs, ~297 nm for Rb) is possible but common experiments use **two-photon excitation** via an intermediate state (e.g., Rb generic route: 780 nm (D2) + 480 nm → nS/nD; Rb Harvard/QuEra gate route: 420 nm + 1013 nm via 6P; Cs: 852 nm + 509 nm → nP). The effective Rabi frequency is Ω_eff = $\Omega_1\Omega_2$ / (2Δ) where Δ is the single-photon detuning. Large Δ reduces photon scattering from the intermediate state.

🔬

Gate fidelity target: State-of-the-art Rb tweezer processors target CZ gate fidelity ≥ 99.5%. Dominant errors include Rydberg-state decay, intermediate-state scattering, Doppler dephasing, blockade leakage, and pulse inhomogeneity. Increasing Ω reduces gate time and Rydberg decay error, but demands higher optical power and tighter technical control in whichever excitation path is used (420/1013 nm, 780/480 nm, or another platform-specific route). Use the Gate Fidelity Budget tool to explore the full error budget.

Key Papers — Rydberg Physics & Neutral-Atom Gates

[1] Jaksch et al., *Fast Quantum Gates for Neutral Atoms*, PRL 85, 2208 (2000) — original Rydberg blockade gate proposal. DOI

[2] Urban et al., *Observation of Rydberg blockade between two atoms*, Nature Physics 5, 110 (2009) — first experimental blockade. DOI

[3] Gaëtan et al., *Observation of collective excitation of two individual atoms in the Rydberg blockade regime*, Nature Physics 5, 115 (2009) — concurrent blockade observation (Browaeys group). DOI

[4] Levine et al., *Parallel implementation of high-fidelity multiqubit gates with neutral atoms*, PRL 123, 170503 (2019) — first >97% Rydberg CZ gate. DOI

[5] Evered et al., *High-fidelity parallel entangling gates on a neutral-atom quantum computer*, Nature 622, 268 (2023) — 99.5% two-qubit gate fidelity (Harvard/QuEra). DOI

[6] Saffman, Walker & Mølmer, *Quantum information with Rydberg atoms*, Rev. Mod. Phys. 82, 2313 (2010) — comprehensive review of blockade physics and gates. DOI

References & Tools

Source confidence note. Quick blockade-radius outputs here are useful order-of-magnitude estimates. For publishable values of C_6, pair potentials, Stark maps, or Förster defects, use ARC or a dedicated pair-state calculation. ROUGH ESTIMATE

💻 ARC, Alkali Rydberg Calculator (Python, Sibalic et al. 2017), precision C_6, pair potentials, Stark maps PEER-REVIEWED 📄 Saffman, Walker & Mølmer, Rev. Mod. Phys. 82, 2313 (2010), definitive Rydberg QC review PEER-REVIEWED 📄 Saffman 2016, Quantum computing with atomic qubits and Rydberg interactions: progress and challenges (J. Phys. B 49, 202001) PEER-REVIEWED 📄 Evered et al. 2023, High-fidelity parallel entangling gates (Harvard, 99.5%) PEER-REVIEWED 📖 Gallagher, Rydberg Atoms (Cambridge, 1994), quantum defects, lifetimes, interactions TEXTBOOK

CHAPTER 24 SUMMARY

Rydberg blockade turns micron-spaced neutral atoms into fast, switchable entangling gates. Reaching fault-tolerant fidelities (>99.5%) requires accounting for every physical error source in a quantitative gate budget.

Problems for Chapter 24

1. USE THE TOOL For ^{87}Rb, compare n = 50 and n = 70: record the C_6 coefficient, the lifetime and the blockade radius at $\Omega/2\pi = 1$ MHz.
2. BY HAND Using $C_6 \propto n^{11}$ and $R_b = (C_6/\hbar\Omega)^{1/6}$, predict the ratio of blockade radii and compare with problem 1.
3. DESIGN Choose an atom spacing and lattice geometry (square or triangular) so that nearest neighbours are blockaded but next-nearest neighbours are not. Check it in the tool.

Use the tool: on the interactive page. By hand: from the formulas in this chapter. Design: a decision a lab would make. Worked solutions are available to instructors on request (see Part V).

CHAPTER 25 · QUANTUM STEP 05

25

Two-Qubit Gate Error Budgets

📊 Rydberg Gate Error Budget · Calculator

Quantify every error source in a Rydberg two-qubit gate, from spontaneous emission and Doppler dephasing to laser noise, finite blockade, and SPAM. Identify your dominant bottleneck and compare to state-of-the-art experimental fidelities.

> *A Rydberg CZ gate takes only a few hundred nanoseconds, yet spontaneous decay, finite blockade strength, Doppler shifts, and laser phase noise each nibble away at its fidelity.*

IN THIS CHAPTER Prerequisites: Ch. 18, Ch. 22, Ch. 24

- Decompose two-qubit CZ infidelity into Rydberg lifetime decay, finite-blockade leakage, Doppler dephasing, and laser intensity/phase noise.
- Run the two-pulse CZ gate of Levine et al. (2019) with finite blockade and Rydberg decay, and compare it with time-optimal pulses that spend less time in the Rydberg state.
- Identify the dominant error lever to optimize when pushing gate fidelity from 99% toward 99.9%.

WEBSITE COMPANION Live interactive page ↗ · Notebook: Open in Colab · download fidelity-budget.ipynb · Week 12

A two-qubit Rydberg CZ gate must achieve raw fidelity F > 99% to meet common fault-tolerance thresholds for practical quantum error correction (surface-code thresholds are order 1%, depending on noise model and decoder). But which error source dominates? This tool builds a complete error budget across 8 contributions, spontaneous emission from the Rydberg state, Doppler dephasing, laser phase noise, blockade leakage, SPAM, atom loss, B-field dephasing, and Rabi inhomogeneity, so you can immediately see which term to attack first. Compare against published Rydberg-gate benchmarks from Rb and Yb platforms.

25.1 Gate Fidelity & Error Budget

Total Gate Fidelity $F = 1 - \Sigma\varepsilon_i$ 80.973%

❌ Below 99%, significant noise reduction needed. Identify and fix the dominant error.

Rydberg lifetime τ

132 μs

n^*=56.9, BBR-corrected (300 K)

Total error $\Sigma\varepsilon_i$

1.903e-1

= 1 − F

Dominant error source

Doppler (thermal)

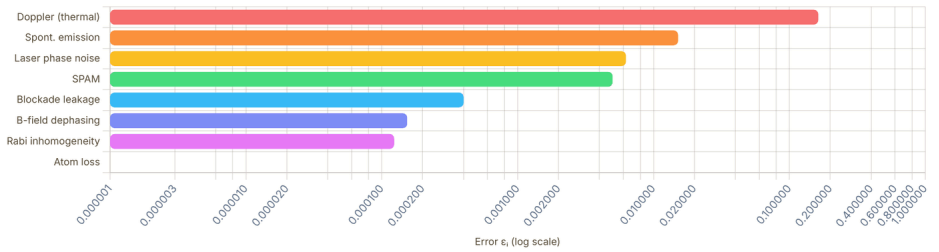
85.7% of total error (ε = 1.63e-1)

PER-SOURCE ERROR BUDGET

ERROR SOURCE	E_i	FRACTION	CONTRIBUTION	WHAT IF?
⚑ Doppler (thermal)	1.63e-1	85.7%		
Spont. emission	1.52e-2	8.0%		
Laser phase noise	6.27e-3	3.3%		
SPAM	5.00e-3	2.6%		
Blockade leakage	4.00e-4	0.2%		
B-field dephasing	1.54e-4	0.1%		
Rabi inhomogeneity	1.23e-4	0.1%		
Atom loss	2.00e-7	0.0%		
Total $\Sigma\varepsilon_i$	1.903e-1	100%		

Press “Set to 0” on any row to see what removing that error source alone would buy; every readout on the page follows.

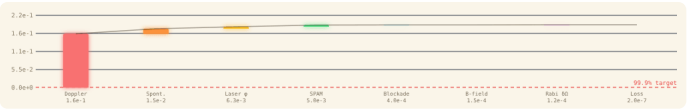
ERROR BREAKDOWN, RELATIVE CONTRIBUTIONS



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

ERROR WATERFALL — CUMULATIVE BUILD-UP

Each bar adds its error to the running total. The dashed line marks a practical high-fidelity target (0.1% error = 99.9% fidelity); surface-code threshold estimates are closer to ~1% for simple noise models.



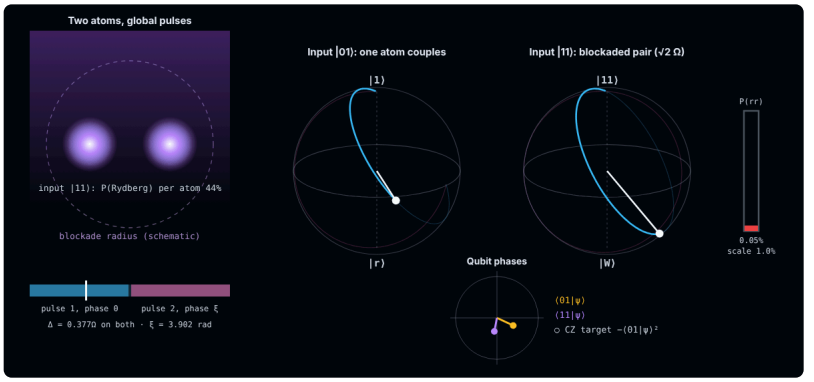
Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

25.2 Inside the CZ Pulse

Interactive

The budget above adds error formulas. Here the gate itself runs: the two-pulse protocol of Levine et al. (2019), integrated exactly for finite blockade and Rydberg decay. Both atoms see two global pulses with detuning $\Delta \approx 0.377\Omega$, the second with its laser phase jumped by ξ. An atom alone ($|01\rangle$) and a blockaded pair ($|11\rangle$) follow different loops, each returns to the ground

state, and the phases they pick up satisfy $\phi_{11} = 2\phi_{01} - \pi$: a CZ gate up to single-qubit phases.



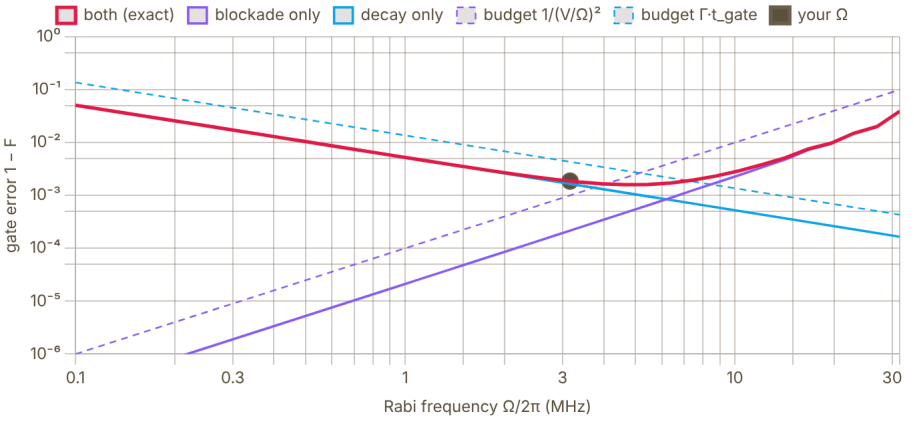
Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

t = 122 ns of 432 ns pulse 1 · laser phase 0

SETTINGS

Pulse: $\Delta/\Omega = 0.377376$, $\tau\Omega = 4.292654$ per pulse ($2\tau = 2.7328\pi/\Omega$), $\xi = 3.902396$ rad, found here by solving for an exact CZ at $V \to \infty$; Levine et al. quote $\Delta \approx 0.377\Omega$ and $2\tau \approx 2.732\pi/\Omega$. Rates are illustrative defaults; "use the budget" copies the page's own numbers.

GATE ERROR VS RABI FREQUENCY (EXACT SIMULATION)



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

With $V/2\pi = 100$ MHz and a 100.00 µs Rydberg lifetime the error is smallest near $\Omega/2\pi \approx 4.6$ MHz ($1 - F \approx 1.6\text{e-}3$): faster pulses leak past the blockade, slower ones decay. The dashed budget lines are the formulas used above; the exact pulse spends only part of the gate in $|r\rangle$, so $\Gamma\cdot t_gate$ overestimates the decay error.

Model: $H = \frac{\Omega}{2}\sum_i (e^{i\varphi}|1\rangle_i\langle r| + \text{h.c.}) - \Delta\sum_i |r\rangle_i\langle r| + V|rr\rangle\langle rr|$, with $\varphi = 0$ then ξ. $|01\rangle$ is a two-level problem; $|11\rangle$ couples to $|W\rangle = (|1r\rangle + |r1\rangle)/\sqrt{2}$ at $\sqrt{2}\,\Omega$ and on to $|rr\rangle$ (shift V). Decay is the no-jump loss $-i\Gamma/2$ per Rydberg excitation, $\Gamma = 1/\tau_r$. Fidelity $F = (|\mathrm{Tr}\, M|^2 + \mathrm{Tr}\, M^\dagger M)/20$ (Pedersen, Møller & Mølmer 2007) with the best single-qubit phase removed. Laser noise, Doppler shifts and leakage back into the qubit are not included. Sources: Levine et al., PRL 123, 170503 (2019) · Pedersen, Møller & Mølmer, Phys. Lett. A 367, 47 (2007).

25.3 State-of-the-Art Comparison

EVERED ET AL. 2023 (HARVARD)
99.5%
Rb-87, $|53S_{1/2}\rangle$ CZ gate with 60 pairs in parallel. Nature 622, 2023.

Evered et al. 2026 (Harvard) PREPRINT

99.854%

Rb-87 CZ gate, raw fidelity averaged over 17 gate sites in a 442-atom array (99.941% with loss post-selection). arXiv:2604.25987, April 2026 — not yet peer-reviewed.

Muniz et al. 2025 (Atom/JILA)

99.72%

^{171}Yb nuclear-spin CZ gate with post-selection; 99.40% without post-selection. PRX Quantum 2025.

MA ET AL. 2023 (PRINCETON)
98.0%
Metastable ^{171}Yb qubit with mid-circuit erasure conversion. Nature 622, 2023.

FTQC THRESHOLD (SURFACE CODE)
99.0%
Typical threshold for fault-tolerant QC with surface codes (~1% physical error rate).

TYPICAL LAB (OPTIMISED)
97–99%
Well-optimised Rydberg gate without extensive noise reduction (Δν~1 kHz, T~5 μK).

YOUR SIMULATION
80.973%
Dominant: Doppler (thermal) (86% of error)

Calculator output at the page's default settings.

25.4 Error Source Formulas

Radiative decay during gate

During a Rydberg gate, atoms spend time in the Rydberg state (lifetime τ_Ryd ∝ n*³). Any spontaneous emission event destroys the coherence and constitutes an error. At room temperature, blackbody radiation (BBR) drives transitions to neighbouring Rydberg states, effectively halving the lifetime for n ~ 50–70.

$$\tau_{\mathrm{Ryd}} \approx \tau_0 \times n^{*3} \qquad (\tau_0 \sim 1\ \mathrm{ns\ for\ alkali}\ nS)$$

With BBR (300 K): $\Gamma_{\mathrm{total}} = \Gamma_{\mathrm{rad}} + \Gamma_{\mathrm{BBR}} \approx 2\Gamma_{\mathrm{rad}}$ ($n \sim 60$)
Cryogenic (< 10 K): $\Gamma_{\mathrm{total}} \approx \Gamma_{\mathrm{rad}}$ (BBR suppressed)
Spontaneous emission error:

$$\varepsilon_{se} = \Gamma_{total} \times t_{gate} = t_{gate}/\tau_{Ryd}$$

Random Doppler phase accumulated during the gate

A thermal atom moving at speed v accumulates a random phase δφ = k_eff × v × t_gate from the Rydberg excitation laser. For a Gaussian velocity distribution, the gate fidelity is reduced by the dephasing factor.

$$v_{\text{rms}} = \sqrt{\frac{k_B T}{m}} \quad \text{[1D thermal velocity]}$$

Phase uncertainty: $\delta\phi = k_{\text{eff}} \cdot v_{\text{rms}} \cdot t_{\text{gate}}$
Dephasing error:

$$\varepsilon_D = \frac{1}{2}(k_{\text{eff}} \cdot v_{\text{rms}} \cdot t_{\text{gate}})^2$$

Single-photon excitation ($\lambda \sim 300$–480 nm): $k_{eff} = 2\pi/\lambda_{Ryd}$
Example: Rb, $T = 5\,\mu$K, $\lambda = 480$ nm, $t_{gate} = 2\,\mu$s $\Rightarrow$

$v_{rms} = \sqrt{k_B T/m} = 21.9$

mm/s, $\delta\phi = 0.57$ rad, $\varepsilon \approx 0.16$ (the calculator's default; colder atoms or a shorter gate cut this quadratically)

Laser coherence time vs gate duration

A laser with Lorentzian linewidth Δν has coherence time T_coh = 1/(π Δν). If the gate time t_gate ≳ T_coh, phase noise significantly degrades the fidelity. Narrow-linewidth lasers (<1 Hz) are needed for high-fidelity Rydberg gates.

$$T_{\text{coh}} = \frac{1}{\pi\,\Delta\nu}$$

Phase noise error (white frequency noise):

$$\varepsilon_\phi \approx \pi\,\Delta\nu \cdot t_{\text{gate}}$$

For $t_{\text{gate}} \ll T_{\text{coh}}$: $\varepsilon_\phi \ll 1$ ✓; for $t_{\text{gate}} = T_{\text{coh}}$: $\varepsilon_\phi \approx 1$ ×
State-of-the-art: $\Delta\nu < 1$ Hz (ULE cavity) $\Rightarrow T_{coh} > 300$ ms $\Rightarrow \varepsilon_\phi < 10^{-5}$ for $t_{gate} = 2\,\mu$s

Leakage into doubly-excited |rr⟩ state

The Rydberg blockade assumes the interaction U ≫ Ω. When this is not satisfied, the doubly-excited |rr⟩ state acquires a small amplitude ~ Ω/U, causing gate errors. For the standard CZ protocol:

$$\varepsilon_{\text{blockade}} \approx \left(\frac{\Omega}{U}\right)^2$$

$U/\Omega \geq 10$: $\varepsilon < 1\%$; $U/\Omega \geq 50$: $\varepsilon < 0.04\%$; $U/\Omega \geq 100$: $\varepsilon < 0.01\%$
C_6/h for alkali nS states near $n \sim 60$–70 is typically order 10^2–10^3 GHz·μm^6 and is strongly state- and Förster-defect-dependent. At $r = 5\,\mu$m ($r^6 \approx 1.6 \times 10^4\,\mu$m^6) this gives interaction shifts of roughly 6–60 MHz, so $U/\Omega \gg 1$ is achievable for $\Omega/2\pi \sim 1$ MHz.

Magnetic field noise ($g_F m_F \sim 1$):

$$\delta\phi_B = \frac{\mu_B \cdot \Delta B \cdot t_{\text{gate}}}{\hbar} \qquad \varepsilon_B = \frac{1}{2}\left(\frac{\mu_B \cdot \Delta B \cdot t_{\text{gate}}}{\hbar}\right)^2$$

Use clock states ($m_F = 0$) to suppress by $\sim 10^3$
Rabi inhomogeneity:

$$\varepsilon_\Omega = \frac{1}{2}\left(\frac{\pi\,\delta\Omega/\Omega}{2}\right)^2 \quad \text{[suppress with composite pulses/DRAG]}$$

SPAM: $\varepsilon_{\text{SPAM}} = p_{\text{prep}} + p_{\text{meas}}$
Atom loss: $\varepsilon_{\text{loss}} = t_{\text{gate}}/\tau_{\text{trap}}$ ($\tau_{\text{trap}} \sim 10$–$100$ s in UHV tweezers)

25.5 Key References

📄 Evered et al. (2023), High-fidelity parallel entangling gates on a neutral-atom QC (Nature) 📄 Evered et al. (2026), High-fidelity entangling gates and nonlocal circuits with neutral atoms (arXiv preprint) 📄 Muniz et al. (2025), High-fidelity universal gates in the ^{171}Yb nuclear-spin qubit (PRX Quantum) 📄 Ma et al. (2023), High-fidelity gates and mid-circuit erasure conversion in ^{171}Yb (Nature) 📄 Jaksch et al. (2000), Fast quantum gates for neutral atoms (original Rydberg gate proposal) 📄 Saffman, Walker & Mølmer (2010), Quantum information with Rydberg atoms (Rev. Mod. Phys.) 📄 Ebadi, Wang, Levine et al. (2021), Quantum phases of matter on a 256-atom programmable quantum simulator (Nature)

CHAPTER 25 SUMMARY

An itemized error budget turns fidelity engineering from guesswork into targeted upgrades. Once a gate is built, we need a SPAM-insensitive protocol to measure its true error rate in the lab.

Problems for Chapter 25

1. USE THE TOOL With the default settings, identify the largest error source. Double the Rabi frequency and record what changes.
2. BY HAND The Doppler dephasing error grows with temperature. Halve the atom temperature and check that the Doppler term changes as you would predict from the velocity spread.
3. DESIGN Starting from the defaults, reach a 99.5% CZ fidelity by changing at most three parameters. Defend each change as achievable in a real lab.

Use the tool: on the interactive page. By hand: from the formulas in this chapter. Design: a decision a lab would make. Worked solutions are available to instructors on request (see Part V).

CHAPTER 26 · QUANTUM STEP 06

26

Randomized Benchmarking

📈 Randomized Benchmarking · Deep dive

A thorough interactive reference for RB, from the SPAM problem and Clifford twirling, through the decay model and IRB, to practical design choices and state-of-the-art neutral atom results. Includes working calculators for standard RB and interleaved RB.

> *In neutral-atom arrays, state-preparation and measurement (SPAM) errors (typically 0.5–2% per shot) are often comparable to or larger than the gate errors we want to measure—making direct single-gate tomography unreliable.*

IN THIS CHAPTER Prerequisites: Ch. 21, Ch. 25

- Explain how Clifford twirling converts general gate-independent noise into an exponential decay F(m) = A p^m + B that separates gate error from SPAM.
- Extract error per Clifford r_C and isolate individual target gates (such as CZ) using Interleaved Randomized Benchmarking (IRB).
- Design statistically sound RB experiments (depths, sequences, shots) and recognize breakdown regimes such as leakage and crosstalk.

WEBSITE COMPANION Live interactive page ↗ · Notebook: Open in Colab · download rb-explorer.ipynb · Week 13

Section 01

26.1 Why Randomized Benchmarking? The SPAM Problem

Before RB existed, measuring gate fidelity seemed simple: prepare $|0\rangle$, apply gate G, measure. But this conflates three separate error sources that you cannot disentangle with a single experiment.

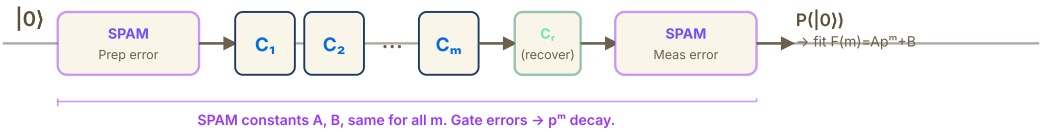
State Preparation and Measurement (SPAM) Errors

Every qubit experiment involves three steps: (1) prepare a known state, (2) apply your gate or circuit, (3) measure the output state. Steps 1 and 3 both have imperfections called **state preparation errors** and **measurement errors**, collectively SPAM. In neutral atom experiments, SPAM typically contributes 0.5–2% error per shot, often comparable to or *larger* than the gate error you're trying to measure.

If you apply gate G once and measure, you see: $P(|0\rangle) = F_{\text{gate}} \cdot F_{\text{SPAM}}$. You cannot extract F_{gate} without knowing F_{SPAM} independently. Even worse, SPAM errors fluctuate with laser power drifts, alignment, and atom temperature — making them hard to characterize reliably.

The Key Insight: SPAM Errors Don't Scale with Circuit Depth

Randomized Benchmarking's crucial observation: apply a *random sequence* of m Clifford gates followed by the recovery gate C_r that inverts the whole sequence. The ideal output is always $|0\rangle$. As you increase m, gate errors accumulate, but SPAM errors are the same no matter how long the circuit is. This means SPAM appears as constants A and B in the fit model, while gate errors appear in the exponential decay p^m. You get a clean separation.



Bottom line: By varying the circuit depth m and fitting the survival probability decay, you extract the decay constant p which depends *only* on gate errors. SPAM gets absorbed into the pre-factors A and B and disappears from the gate error estimate. This is what makes RB so powerful, it's SPAM-insensitive by construction.

Section 02

26.2 The Clifford Group

RB uses random Clifford gates, not arbitrary unitaries. The choice is deliberate: Clifford operations have a unique property that makes the RB decay model exact.

Definition: Normalizer of the Pauli Group

The Clifford group $\mathcal{C}_n$ consists of all n-qubit unitaries U that *map Pauli operators to Pauli operators under conjugation*:

$$U \in \mathcal{C}_n \iff UPU^\dagger \in \mathcal{P}_n \text{ for all } P \in \mathcal{P}_n$$

In other words, Clifford gates **normalize the Pauli group**. This means if you conjugate any Pauli by a Clifford, you get another Pauli (up to a phase). Paulis get "shuffled around", they never become non-Pauli.

Group Sizes

- **1-qubit:** 24 elements, generated by H and S, or equivalently $X_{\pi/2}$ and $Z_{\pi/2}$. These correspond to the 24 orientation-preserving symmetries of a cube (the chiral octahedral group O).
- **2-qubit:** 11,520 elements, generated by 1Q Cliffords plus CNOT (or CZ).
- **n-qubit:** grows as $\sim 2^{O(n^2)}$, exponentially large, but still efficiently simulable via the Gottesman-Knill theorem (since Cliffords map stabilizer states to stabilizer states).

Why Cliffords? The Twirling Argument

The key insight behind RB is **Clifford twirling**. Given any noise channel Λ (which could be completely arbitrary, coherent errors, non-Markovian, anything), averaging over the Clifford group converts it into a *depolarizing channel*:

$$\frac{1}{|\mathcal{C}_n|}\sum_{C\in\mathcal{C}_n} C^\dagger \Lambda(C\rho C^\dagger)C = p\rho + (1-p)\frac{I}{d}$$

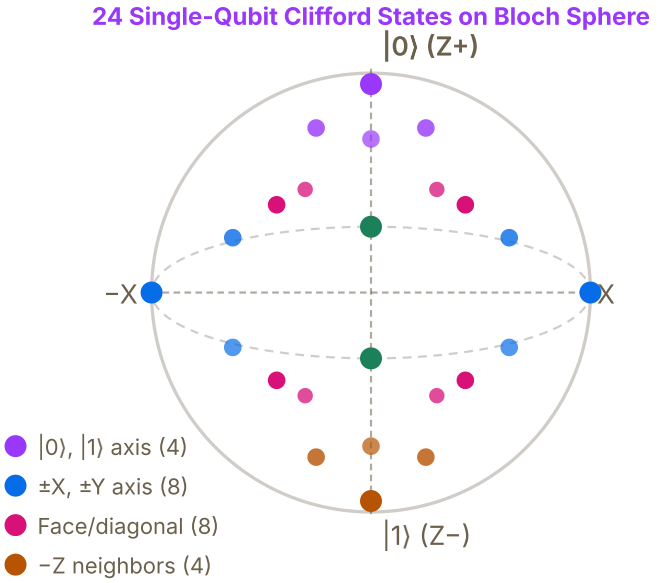
This is the twirl. It means: no matter what your noise actually is, coherent Z rotations, amplitude damping, cross-talk, laser phase noise, after Clifford twirling it *looks like* depolarizing noise parameterized by a single number p. The RB decay $F(m) = Ap^m + B$ is therefore exact (not approximate) for any noise that is gate-independent.

1-Qubit Clifford Decompositions in $\{X_{\pi/2}, Z_{\pi/2}\}$ Basis

Modern neutral atom hardware implements 1Q gates natively as $X_{\pi/2}$ and $Z_{\pi/2}$ pulses (virtual Z via phase shift). Every 1Q Clifford can be decomposed into at most 3 such pulses, giving an average of ≈ 1.875 native gates per Clifford.

GATE / AXIS	NATIVE GATE SEQUENCE	# NATIVE GATES	COMMENT
Identity I	— (none)	0	Virtual only
X_π (X gate)	$X_{\pi/2} \cdot X_{\pi/2}$	2	Two half-X pulses
Y_π (Y gate)	$Z_{\pi/2} \cdot X_\pi \cdot Z_{-\pi/2}$	3	Axis rotation

Z_π (Z gate)	$Z_{\pi/2} \cdot Z_{\pi/2}$	0 (virtual)	Phase update
H (Hadamard)	$Z_{\pi/2} \cdot X_{\pi/2} \cdot Z_{\pi/2}$	1 + 2 virtual	1 physical pulse
S ($Z_{\pi/2}$)	$Z_{\pi/2}$	0 (virtual)	Frame rotation
T_3 (120° face rotations)	1–3 pulse combos	1–3	8 elements, avg 2
T_6 (180° edge rotations)	$X_{\pi/2}$ or $Z_{\pi/2}$	1–2	6 elements
Avg over all 24	—	1.875	Theoretical average



Section 03

26.3 The RB Decay Model

From first principles: why the survival probability decays exponentially, what the fit parameters mean, and how to extract gate error and useful circuit depth.

Derivation: From Twirling to Exponential Decay

1. **Start with gate-independent depolarizing noise.** Each Clifford gate C_i is followed by a noise channel: $\tilde{C}_i(\rho) = \Lambda(C_i \rho C_i^\dagger)$ where $\Lambda(\rho) = p\rho + (1-p)\frac{I}{d}$ is depolarizing with parameter $p \in [0, 1]$.

2. **After m gates, average over all random Clifford sequences.** By the twirling argument, the noise from m sequential gates compounds as p^m (since each application of Λ shrinks the Bloch vector by factor p). The averaged state before measurement is: $\overline{\rho_m} = p^m |\psi_{\text{prep}}\rangle\langle\psi_{\text{prep}}| + (1-p^m)\frac{I}{d}$

3. **Apply the recovery gate and measure.** The measured survival probability (probability of recovering $|0\rangle$) is: $F(m) = \langle 0|\overline{\rho_m}|0\rangle = A \cdot p^m + B$ where $A = \langle\psi_{\text{prep}}|0\rangle^2 - 1/d$ absorbs SPAM, and $B = 1/d$ is the long-depth asymptote (fully mixed state for $d = 2$ gives $B = 1/2$).

4. **Extract error per Clifford from p.** For a d-dimensional system, depolarizing parameter p relates to error per Clifford: $r_C = \frac{(d-1)(1-p)}{d}$. For 1-qubit ($d = 2$): $r_C = (1-p)/2$, so $p = 1 - 2r_C$. For 2-qubit ($d = 4$): $r_C = 3(1-p)/4$, so $p = 1 - (4/3)r_C$.

5. **Convert to error per native gate.** Since each Clifford is compiled from $\bar{n}$ native gates on average, the error per native gate is $\varepsilon_{\text{gate}} \approx r_C/\bar{n}$. For 1Q: $\bar{n} \approx 1.875$. For 2Q, counting CZ gates only (they dominate the error): $\bar{n} \approx 1.5$.

$$F(m) = A \cdot p^m + B$$

$$r_C = \frac{(d-1)(1-p)}{d} \quad [\text{1Q: } r_C = \tfrac{1-p}{2},\ \text{2Q: } r_C = \tfrac{3(1-p)}{4}]$$

$$\varepsilon_{\text{gate}} \approx \frac{r_C}{\bar{n}} \quad \text{where } \bar{n} \approx 1.875\ (\text{1Q}),\ \approx 1.5\ \text{CZ}\ (\text{2Q})$$

$$N_{\text{useful}} = \frac{1}{2r_C}\ (\text{1Q}),\ \frac{3}{4r_C}\ (\text{2Q}) \quad [\text{depth at which } p^m \approx 1/e, \text{ i.e. the decaying part } A\,p^m \text{ has fallen to } 1/e]$$

Physical Meaning of Parameters

- p **(decay constant):** Proximity to identity channel. $p = 1$ means perfect gates, $p = 0$ means fully depolarizing. Extracted directly from exponential fit.
- A **(amplitude):** Absorbs SPAM: $A = F(0) - 1/d$, so $A \leq 1/2$ for one qubit. With ~5% SPAM error in neutral-atom experiments, $F(0) \approx 0.95$ and $A \approx 0.45$. Does *not* affect r_C.
- B **(asymptote):** Long-depth limit where the state is fully mixed. For 1Q: $B \to 1/2$ (chance). For 2Q: $B \to 1/4$. In practice, B is a fit parameter that captures imperfect long-depth behavior.
- r_C **(error per Clifford):** The quantity you report. Benchmark comparing platforms. Equivalent to infidelity per Clifford averaged uniformly over the Clifford group.

Average Gate Fidelity

The average gate fidelity of a noisy channel $\tilde{C}$ relative to ideal C is:

$F_{\text{avg}}(C, \tilde{C}) = \int_\psi \langle\psi|C^\dagger \tilde{C}(|\psi\rangle\langle\psi|)C|\psi\rangle\, d\psi$

Under the depolarizing model, $F_{\text{avg}} = p + (1-p)/d$. For a 1Q gate: $F_{\text{avg}} = p + (1-p)/2 = 1 - r_C$. For a native gate: $F_{\text{avg,gate}} \approx 1 - \varepsilon_{\text{gate}} = 1 - r_C/\bar{n}$.

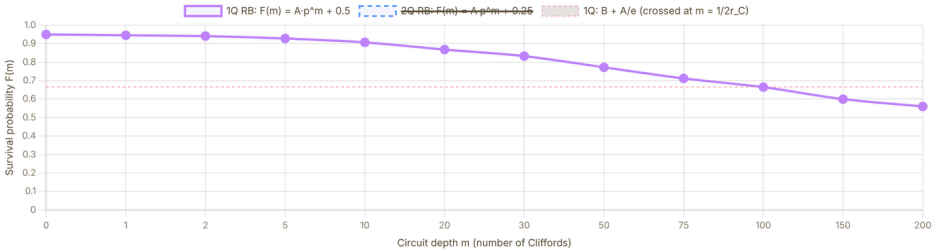
Surface code fault-tolerance requires $F_{\text{avg}} \gtrsim 99\%$ per gate, corresponding to $r_C \lesssim 0.01$ per Clifford. Leading neutral atom systems achieve $r_C \sim 0.003$–0.005 for 2Q gates.

Section 04

26.4 Interactive Standard RB Calculator

Adjust r_C and SPAM amplitude to see how the RB decay curve changes. Compare 1Q and 2Q benchmarks.

RB DECAY CURVE, F(M) = A · Pᵐ + B



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

p (1Q decay) 0.99000

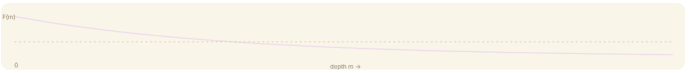
r_C (err/Clifford) 0.500%

ε_gate ≈ r_C/1.875 0.267%

N_useful = 1/(2r_C) 100

F_avg gate 99.733%

Live RB Simulation



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Each dot = one random Clifford circuit run. Dots accumulate depth-by-depth, tracing the exponential decay.

Reading the chart: The dashed line sits at $B + A/e$ for the 1Q curve: it crosses that curve at the characteristic depth $N_{\text{useful}} = 1/(2r_C)$, where $p^m = 1/e$ (for 2Q the same point is $3/(4r_C)$). The curves level off at $B = 1/d$, not at zero. The 2Q curve decays faster because $p_{2Q} = 1 - (4/3)r_C < p_{1Q} = 1 - 2r_C$ for the same r_C. SPAM amplitude A shifts the whole curve vertically but does *not* change r_C.

Section 05

26.5 Interleaved RB (IRB)

Standard RB gives the average error per Clifford. To measure the fidelity of a *specific gate*, use Interleaved RB.

Why IRB?

The standard RB decay constant p_{ref} reports an average over all 24 (or 11520) Clifford gates. But for building a quantum computer, you care about individual gates: how well does your CZ gate perform? Your $X_{\pi/2}$? IRB isolates a single target gate G.

Protocol

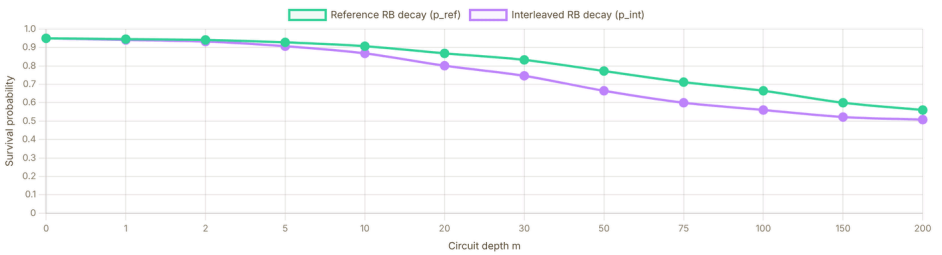
1. **Standard RB:** Run regular RB with random Clifford sequences of depths $m = 1, 2, 5, \ldots$ Fit $F_{\text{ref}}(m) = A_{\text{ref}} \cdot p_{\text{ref}}^m + B$ to get p_{ref}.

2. **Interleaved RB:** Insert the target gate G after *every* random Clifford: $\ldots C_j \cdot G \cdot C_{j-1} \cdot G \cdot C_{j-2} \cdot G \ldots$ The combined sequence is still m pairs of (Clifford, target gate) followed by recovery. Fit $F_{\text{int}}(m) = A_{\text{int}} \cdot p_{\text{int}}^m + B$ to get p_{int}.

3. **Extract gate error:** The ratio $p_{\text{int}}/p_{\text{ref}}$ cancels the average Clifford error, leaving only the contribution of G: $r_{\text{gate}} = \frac{d-1}{d}\left(1 - \frac{p_{\text{int}}}{p_{\text{ref}}}\right)$

$$r_{\text{gate}} = \frac{d-1}{d}\left(1 - \frac{p_{\text{int}}}{p_{\text{ref}}}\right)$$

IRB, REFERENCE VS INTERLEAVED DECAY CURVES



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

p_ref (standard RB) 0.99000

p_int (interleaved) 0.98010

p_int / p_ref 0.9900

r_gate extracted 0.500%

r_gate set 0.500%

F_gate = 1 - r_gate 99.500%

IRB Systematic Errors: The formula $r_{\text{gate}} = \frac{d-1}{d}(1 - p_{\text{int}}/p_{\text{ref}})$ is a lower bound. If the noise of G is highly coherent or gate-dependent, the extracted r_{gate} can underestimate the true error. The bounds are tight only when $r_C \ll r_{\text{gate}}$ (their width grows with the reference error r_C).

Practical note: $p_{\text{int}} \leq p_{\text{ref}}$ always (the interleaved gate adds more error). If $r_{\text{gate}} \gg r_C$, the interleaved curve decays much faster and the ratio is sensitive to fitting noise. Use more sequences (100+) per depth in that regime.

Section 06

26.6 Two-Qubit RB Specifics

Two-qubit benchmarking involves a much larger Clifford group and requires accounting for the structure of 2Q gate decompositions and cross-talk.

2Q Clifford Group Statistics

- **Group size:** 11,520 elements (vs 24 for 1Q). Sampling is still efficient since random 2Q Cliffords can be generated from a small circuit gate set.
- **CZ gate count per 2Q Clifford:** On average ~ 1.5 CZ gates + ~ 8.4 single-qubit gates (in $\{H, S, X_{\pi/2}\}$ basis).
- **Dimension:** $d = 4$, so $B \to 1/4$ and the formula becomes $r_C = \frac{3(1-p)}{4}$.
- **Error budget:** 2Q error per Clifford includes contributions from CZ gate fidelity AND single-qubit gates within the Clifford decomposition. To isolate CZ: use IRB with $G = \text{CZ}$.

Simultaneous RB (SimRB)

Run single-qubit RB on both qubits *simultaneously* while the other qubit is active. Compare to individual 1Q RB on each qubit in isolation.

Crosstalk signature: If $r_C^{\text{sim}} > r_C^{\text{iso}}$, the difference $\Delta r_C = r_C^{\text{sim}} - r_C^{\text{iso}}$ quantifies crosstalk-induced error. In Rydberg arrays, crosstalk arises from: residual Rydberg admixture, off-resonant blockade from neighboring atoms, and laser spillover.

Parallel RB: Running many qubit pairs simultaneously (Evered 2023: 60 pairs) tests that gate fidelity doesn't degrade under realistic many-body load, crucial for demonstrating scalable performance.

PROPERTY	1Q STANDARD RB	2Q STANDARD RB
Clifford group size	`24`	`11,520`
Hilbert space dim d	`2`	`4`
Asymptote B	`0.5`	`0.25`
r_C from p	$(1-p)/2$	$3(1-p)/4$
Avg native gates / Clifford	`~1.875`	`~1.5 CZ + ~8.4 SQ`
Typical r_C neutral atoms	<0.001	0.003–0.007
What limits fidelity?	Laser phase noise, off-resonance, Doppler	Spontaneous emission during Rydberg excitation, Doppler, blockade errors

> **Generators of the 2Q Clifford Group:** The 2Q Clifford group is generated by $\{H \otimes I,\ S \otimes I,\ I \otimes H,\ I \otimes S,\ \text{CNOT}\}$. In a Rydberg processor, you replace CNOT with CZ (related by single-qubit Hadamards). Random 2Q Cliffords can be synthesized as random circuits of depth ∼12–20 in this gate set.

Section 07

26.7 Practical Considerations

How to actually run RB in the lab, choosing sequence counts, depths, shots, and understanding when the RB model breaks down.

Experimental Design Parameters

- **Number of random sequences per depth (K):** 30–100 is typical. More sequences reduce variance in the fit. For high-fidelity systems ($r_C \lesssim 0.003$), use $K \gtrsim 75$ to resolve the small decay.
- **Number of depth values (M):** 8–15 points, log-spaced up to $m_{\max} \sim 3/r_C$ (where survival drops to $\sim e^{-3} \approx 5\%$ above B). Too few depth points → unreliable fit. Too many → waste of time on fully decohered sequences.
- **Shots per circuit (n_{shots}):** 100–1000. More shots reduce shot noise. With atom-resolved imaging, even 200 shots per circuit is often sufficient.
- **Total experiments:** $K \times M \times n_{\text{shots}} \sim 30 \times 10 \times 200 = 60{,}000$ shots minimum.

Statistical Error on r_C

The statistical uncertainty on the fitted error per Clifford scales approximately as:

$$\delta r_C \approx \frac{1}{A \cdot \sqrt{K \cdot M \cdot n_{\text{shots}}}}$$

With $A \approx 0.9$, $K = 50$, $M = 10$, $n_{\text{shots}} = 200$: $\delta r_C \approx \frac{1}{0.9\sqrt{100{,}000}} \approx 0.0035$, about 3.5×10^{-3}. This is the floor on how precisely you can report r_C. For sub-0.1% gate errors, you need $\gtrsim 10^6$ total shots.

Bootstrapping: For confidence intervals, resample the set of sequences (with replacement) and refit many times. Report the 95th percentile range.

When Does RB Break Down?

- **Gate-dependent noise:** If different Clifford gates have very different error rates, the twirl doesn't produce a clean single-exponential decay, you may see multi-exponential behavior. Use **character RB** (Helsen et al.) to handle this.
- **Leakage out of the qubit subspace:** If population leaks into $|2\rangle$, $|3\rangle$ (e.g., second Rydberg state, hyperfine levels), the decay model fails. Use **leakage RB** (Wood & Gambetta) which explicitly tracks leakage as an additional decay channel.
- **Non-Markovian noise:** If noise has temporal correlations (e.g., slow laser frequency drifts, magnetic field fluctuations), different sequences in the same batch see different error rates. The ensemble average is still approximately valid but the fit variance increases.
- **Crosstalk for multi-qubit benchmarking:** If benchmarking qubit i while qubit j is nearby and in a different state, the effective noise on qubit i depends on qubit j's state. Simultaneous RB explicitly captures this.

> **Neutral atom highlight (Evered et al., Nature 2023):** Benchmarked 60 qubit pairs *simultaneously*, each pair undergoing independent RB sequences. Result: 99.5% 2Q CZ gate fidelity holding globally across the array, not just for isolated qubit pairs. This parallel RB approach is essential for proving that cross-atom interactions don't degrade performance at scale. The fact that the isolated-pair and simultaneous-array results agreed confirmed that crosstalk is negligible at typical array spacings ($\sim 3\,\mu\text{m}$ inter-site distance in the tweezer array).

PARAMETER	CONSERVATIVE	STANDARD	HIGH-PRECISION
Sequences per depth K	20	50	100+
Depth points M	6	10	15
Shots per circuit	100	200	1000
δr_C (approx)	~0.010	~0.0035	~0.0009
Total shots	12,000	100,000	1,500,000

Section 08

26.8 Results in Neutral Atom Experiments

State-of-the-art RB benchmarks for neutral atom platforms (Rb, Cs, Yb tweezer arrays), with comparison to fault-tolerance thresholds.

PLATFORM / SPECIES	GATE	R_C (PER CLIFFORD)	E_{GATE} (PER NATIVE GATE)	NOTES / PAPER
^{87}Rb tweezer (Harvard)	2Q CZ (Rydberg)	0.0050	0.50%	Evered et al., *Nature* 622, 268 (2023). 60-pair parallel RB, 99.5% fidelity.
^{133}Cs tweezer (Caltech)	memory / imaging	context	adjacent	Manetsch et al. (2025). 6100-atom Cs array; adjacent alkali scaling benchmark, not an Rb87 single-qubit fidelity number.
^{171}Yb tweezer (JILA)	2Q CZ (Rydberg)	0.0060 raw / 0.0028 post-sel.	0.60% raw / 0.28% post-sel.	Muniz et al., *PRX Quantum* 6, 020334 (2025). 99.40(3)% raw and 99.72(3)% with post-selection.
^{133}Cs tweezer (CsQ / various)	2Q	~0.010	~0.5%	Multiple groups; heavier atom, different blockade physics.
^{87}Rb tweezer (Lukin group)	1Q + 2Q (Rydberg)	0.005 (2Q)	0.50%	Bluvstein et al., *Nature* 604, 451 (2022). Logical qubit experiments with RB characterization.
^{87}Sr tweezer (Atom Computing)	1Q (nuclear spin)	0.0002	0.01%	Barnes et al., *Nature Commun.* 13, 2779 (2022). Nuclear spin qubit in ^{87}Sr with laser-free drives.
Surface code threshold	—	~0.01 (<0.001 target)	~1% (<0.1% target)	Theoretical surface-code threshold is ~1% (Fowler et al., 2012); <0.1% is the practical target well below threshold for low qubit

overhead.

Progress timeline: Rydberg 2Q gate fidelity has improved from $\sim 97\%$ (2019) to $\sim 99.5\%$ (2023) to $\gtrsim 99.7\%$ (2025) in under 6 years. The rapid improvement has been driven by: better atom cooling (reducing Doppler error), shorter gate times (reducing spontaneous emission), and improved laser stability (phase noise).

Why Yb is promising: ^{171}Yb has a nuclear spin qubit ($I = 1/2$, zero electronic spin in $^1S_0$), giving $T_2 \sim$ seconds in a tweezer and nearly zero sensitivity to magnetic field noise. SPAM is also improved: $^1S_0 \to {}^3P_0$ clock transition has near-unity detection fidelity via fluorescence imaging. These factors combine to push r_C well below 10^{-3}.

26.9 References & Further Reading

Foundational RB theory, scalable variants, experimental benchmarks, and QEC context.

Randomized Benchmarking — Original & Theory

Knill et al. (2008), PRA 77 012307 — original randomized benchmarking proposal for Clifford groups Magesan, Gambetta & Emerson (2011), PRL 106 180504 — scalable and robust RB; proves the exponential decay model Magesan et al. (2012), PRL 109 080505 — interleaved RB (IRB): extracting individual gate error rates Emerson, Alicki & Życzkowski (2005), J. Opt. B 7 S347 — scalable noise estimation with random unitary operators; early RB precursor

Experimental Benchmarks (Neutral Atoms)

Evered et al. (2023), *Nature* 622 — 99.5% two-qubit CZ gate fidelity in Rb tweezer arrays via RB Bluvstein et al. (2024), *Nature* 626 58 — logical quantum processor based on reconfigurable atom arrays Saffman, Walker & Mølmer (2010), *Rev. Mod. Phys.* — Rydberg quantum-information review and gate-error context

Error Correction Context

Fowler et al. (2012), PRA 86 032324 — surface code threshold analysis; why ~99% gate fidelity matters Proctor et al. (2022), *Nature Physics* 18 75 — measuring the capabilities of quantum computers with randomized benchmarking Knill, Laflamme & Zurek (1998), *Science* 279 342 — resilient quantum computation, the threshold theorem: how good gates need to be for fault tolerance

CHAPTER 26 SUMMARY

Randomized benchmarking provides the standard, platform-agnostic scorecard for 1Q and 2Q gates. With high-fidelity local gates verified, we can scale up by physically shuttling atoms across a zoned processor.

Problems for Chapter 26

1. USE THE TOOL Run standard RB with $r_C = 0.01$ and the default SPAM fidelity. Read off the decay constant p.
2. BY HAND Show that SPAM errors change the amplitude A and offset B of $F(m) = Ap^m + B$ but not p, and convert p to an average error per Clifford for one qubit.
3. DESIGN Use interleaved RB to estimate the error of one gate from $r_{C,ref}$ and $r_{C,int}$. What assumptions make this estimate unreliable?

Use the tool: on the interactive page. By hand: from the formulas in this chapter. Design: a decision a lab would make. Worked solutions are available to instructors on request (see Part V).

CHAPTER 27 · QUANTUM STEP 07

27

Zoned Neutral-Atom Processors

🎛️ Zoned Atom Processor · Simulator

Pick an algorithm and watch it run on a schematic neutral-atom processor: qubits wait in a storage zone, ride crossed acousto-optic tweezers into an entangling zone for one global Rydberg pulse, and fly up to be read out. A real state-vector simulation runs underneath, and a small compiler decides where every atom should live.

> *Unlike fixed-wiring superconducting chips, neutral atoms in optical tweezers can be coherently transported between dedicated storage, entangling, and readout zones—giving dynamic all-to-all connectivity.*

IN THIS CHAPTER Prerequisites: Ch. 8, Ch. 24, Ch. 26

- Map quantum circuits (GHZ, Grover, Bernstein-Vazirani) onto a zoned architecture separating storage, Rydberg gate, and readout regions.
- Evaluate the trade-off between AOD transport duration, heating/coherence loss, and parallel gate scheduling.
- See how reconfigurable connectivity lowers SWAP overhead for both NISQ circuits and qLDPC/surface-code error correction.

WEBSITE COMPANION Live interactive page ↗ · Week 13

01 · THE DEVICE

27.1 Run an algorithm on a zoned atom array

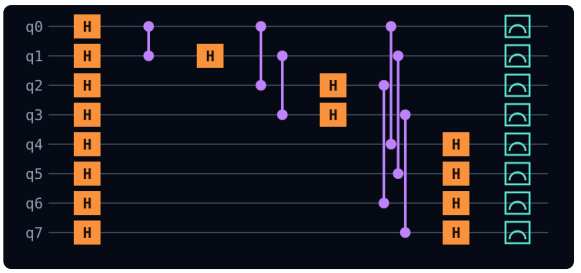
Atoms are coloured by their own qubit state; a pink halo means the atom is entangled with the others.

Interactive



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

8 qubits · ready transport 0 µs



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

qubit in |0⟩ qubit in |1⟩ entangled (halo) AOD tweezer beams local single-qubit pulse global Rydberg pulse

Now

GHZ state · 8 qubits. Equal weight on 00000000 and 11111111. The CNOT tree doubles the entangled set each layer. Press Run or Step.

State vector · probabilities



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

- |00000000⟩ 100.0%

This compilation

Qubits
8

Gates after compiling
15 1q · 7 CZ

Gates cancelled
0

Device steps
8

Rydberg pulses (trips)
3

Atom moves
28

Transport time
524 µs

Estimated fidelity
86.2%

02 · THE COMPILER

27.2 Where an atom sits changes how long the circuit takes

Same circuit, two placements of qubits in the storage row.

Placement comparison

	IN ORDER	OPTIMISED
Storage order (slot → qubit)	q0 q1 q2 q3 q4 q5 q6 q7	q6 q2 q0 q4 q1 q5 q3 q7
Rydberg pulses (trips)	7	3
Atom moves	28	28
Transport distance	1073 µm	995 µm
Transport time	1172 µs	524 µs
Estimated fidelity	86.0%	86.2%

Placing the qubits in a different order saves 4 trips: pairs that interleave in index order sit next to each other once reordered.

Why the order of atoms matters

The moving traps come from a crossed pair of acousto-optic deflectors (AODs): one RF tone per column, one per row, and a trap wherever a row beam crosses a column beam (Bluvstein et al. 2022). In this model two column tones may never pass through each other, since equal frequencies would put both traps in one spot. The grid can only stretch, squeeze and slide, so atoms picked up in one row keep their left-to-right order.

Every pair has to land side by side at a gate site. If two pairs interleave in storage (a, c, b, d with pairs ab and cd), no order-preserving move brings both pairs together. The compiler then splits the layer into separate trips, and each trip costs another round of transport and another Rydberg pulse. The optimiser tries every placement (up to 8! = 40,320) and keeps the one with the fewest trips, then the shortest pair spans.

03 · RUN IT 1,000 TIMES

27.3 What the hardware actually returns

Each shot is a separate noisy trajectory through the same compiled program, read out with imperfect detection.

Interactive

1000 shots

OVERLAP WITH IDEAL DISTRIBUTION
89.4% 1 − TVD, 1000 shots

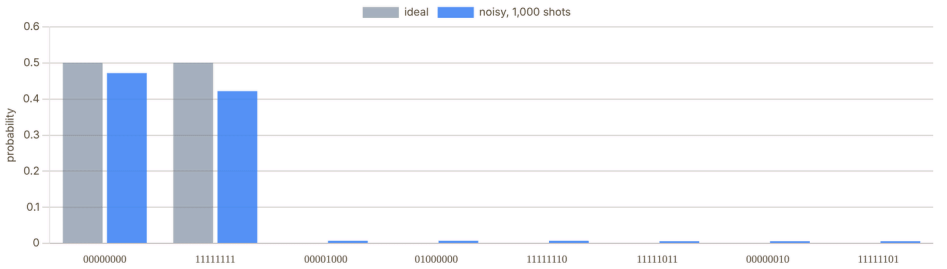
ESTIMATED CIRCUIT FIDELITY
86.2%
$F_1^{15} \cdot F_2^{7}$ · move, idle & readout terms

RIGHT ANSWER
89.4% all-0 or all-1

DISTINCT BIT STRINGS SEEN
34 of 256 possible

Calculator output at the page's default settings.

Measured bit strings (q0 first)



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Noise model: after each single-qubit gate a random Pauli with probability $3(1–F_1)/2$, after each CZ a random two-qubit Pauli with probability $5(1–F_2)/4$ (both reproduce the average gate fidelity of a depolarizing channel), a Z error per carried atom per move, idle dephasing of every qubit during transport (Z with probability $(1 – e^{-t/T_2})/2$ per trip), and a bit flip per qubit at readout. Overlap is $\Sigma \min(p_{ideal}, p_{measured}) = 1 –$ total-variation distance.

04 · WHY ZONES

27.4 Separate storage, entangling and readout

One laser, many gates

A Rydberg CZ needs the two atoms within the blockade radius and away from every other atom. Instead of steering a laser to each pair, the zoned architecture brings the pairs to the laser: Bluvstein et al. split their array into storage, entangling and readout zones and pulse a single global Rydberg laser on every pair in the entangling zone at once, running logical circuits with up to 280 physical qubits (*Nature* 626, 58, 2024). Atoms waiting in storage never see the Rydberg light, and readout happens away from the qubits still in use.

The price is transport. Moving entangled atoms preserves their coherence (Bluvstein et al., *Nature* 604, 451, 2022), but every trip takes time, which is why a compiler that keeps trips few and short matters.

Reading the state panel

Every basis state $|q_0 q_1 \ldots\rangle$ is a bar; its height is $|\langle x|\psi\rangle|^2$ (scaled so the tallest bar fills the panel) and its colour is the phase $\arg\langle x|\psi\rangle$. An atom's colour is its own $P(1)$, and its halo is $1 - |\vec{r}|$, where $\vec{r}$ is the Bloch vector of its reduced density matrix. For a pure global state $|\vec{r}| < 1$ means the qubit is entangled with the rest. In a GHZ state every atom is half blue, half amber, and fully haloed.

05 · WHAT THE MODEL DOES

27.5 Assumptions

- **Gates:** H, X, Z, S, T as local single-qubit pulses; CZ by one global Rydberg pulse on pairs in the entangling zone. CNOT is compiled as H·CZ·H on the target, and adjacent self-inverse gates on the same qubit (H·H, X·X, Z·Z) are cancelled.
- **Scheduling:** as-soon-as-possible layers; each layer's single-qubit gates run together, then its CZ pairs.
- **Transport:** one storage row, one AOD row; atoms keep their left-to-right order; four gate sites. Pairs whose storage intervals overlap need separate trips (greedy interval scheduling). Time = longest move ÷ speed, in and out; gate durations are not included.
- **Geometry is schematic:** zone sizes, spacings and the blockade circle are drawn for clarity, not to scale.
- **Noise:** independent Pauli errors as listed under the histogram; no leakage, atom loss or crosstalk. The estimated fidelity is the product of the per-operation success probabilities.

Sources

Bluvstein et al. (2024), *Nature* 626, 58 — logical quantum processor based on reconfigurable atom arrays (storage, entangling and readout zones) Bluvstein et al. (2022), *Nature* 604, 451 — a quantum processor based on coherent transport of entangled atom arrays (crossed 2D AOD) Evered et al. (2023), *Nature* 622, 268 — high-fidelity parallel entangling gates on a neutral-atom quantum computer Levine et al. (2019), *Phys. Rev. Lett.* 123, 170503 — parallel implementation of high-fidelity multi-qubit gates with neutral atoms

CHAPTER 27 SUMMARY

Zoned shuttling unlocks thousands of dynamically connected qubits inside a single vacuum chamber. Scaling beyond a single chamber requires linking separate processor modules with flying photons.

Problems for Chapter 27

1. USE THE TOOL Run the same circuit with “In order” and “Optimised” placement. Compare the total transport time and the circuit success probability.
2. BY HAND Estimate the circuit success probability by hand as the product of gate, move and readout fidelities, and compare with 1,000 simulated shots.
3. DESIGN Why do zoned architectures separate storage, entangling and readout zones? Give one advantage and one cost of each zone.

Use the tool: on the interactive page. By hand: from the formulas in this chapter. Design: a decision a lab would make. Worked solutions are available to instructors on request (see Part V).

CHAPTER 28 · QUANTUM STEP 08

28

Remote Entanglement & Quantum Networks

🔗 Remote Entanglement · Deep dive

The single qubit is not the bottleneck. The *connection* between qubits is. Remote entanglement generation uses photons to link atoms in separate quantum processors, turning independent tweezer modules into a scalable quantum computer.

> *Because Rydberg blockade acts over ~10 µm, linking separate neutral-atom modules—or distant quantum labs—requires photonic interconnects that herald entanglement between remote atoms.*

IN THIS CHAPTER

Prerequisites: Ch. 11, Ch. 21, Ch. 27

- Follow the two-photon heralded entanglement protocol from atom-photon spin-polarization entanglement through Hong-Ou-Mandel interference and Bell-state measurement.
- Calculate entanglement generation rates from collection NA or cavity cooperativity, fiber attenuation, quantum frequency conversion, and round-trip latency.
- Understand entanglement distillation and quantum repeater chains for overcoming exponential fiber loss.

WEBSITE COMPANION Live interactive page ↗ · Notebook: Open in Colab · download remote-entanglement.ipynb · Week 14

28.1 Why Remote Entanglement?

The scaling wall is real — and photons are the only way through it.

6,100

coherent Cs atom array reported in 2025

<10^6

noisy qubits to factor RSA-2048 (Gidney 2025, preprint)

~15 µm

Max Rydberg blockade radius

∞

Photon range (with fiber)

The single-array bottleneck

The Rydberg blockade — the mechanism behind every neutral-atom two-qubit gate — operates over a range of 5–15 µm. This means that only atoms within a few micrometres of each other can be entangled locally. Single tweezer arrays have reached the several-thousand-atom scale, including the 6,100-atom Cs benchmark reported in 2025; pushing beyond ~10,000 useful, coherent atoms in one module remains a major engineering challenge. Useful fault-tolerant algorithms still need far more: a 2025 resource estimate puts factoring a 2048-bit RSA key at fewer than a million noisy qubits running for under a week (Gidney, arXiv:2505.15917, preprint).

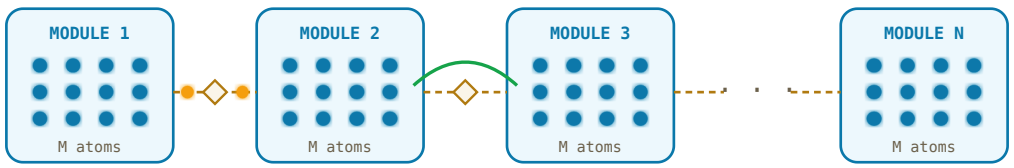
The leading scaling route is **modular architecture**: build many smaller, higher-quality arrays and connect them with quantum links. The link between distant modules is naturally a *photonic interconnect*, because photons can travel macroscopic distances without the motional decoherence that affects material qubits.

> 💡 **Classical analogy:** You don't solve the chip-size problem in classical computing by making one giant chip. You connect many CPUs with a high-speed bus. Modular quantum computing is the same idea — with entanglement replacing the electrical bus.

What remote entanglement enables

- **Modular scaling:** Chain N tweezer arrays, each with M atoms. The effective system size grows as N × M, but each array can be individually optimised.
- **Teleportation-based gates:** A shared Bell pair between modules enables a deterministic quantum gate between one qubit in each module via gate teleportation — even without direct interaction.
- **Quantum networks:** The same photonic link that connects two modules in a lab can connect two labs across a city — enabling distributed quantum computing and quantum key distribution.
- **Quantum repeaters:** Long-distance quantum communication requires intermediate "repeater" nodes with quantum memories. Neutral-atom tweezers are ideal quantum memories.
- **Heterogeneous architectures:** Photonic links can connect different qubit platforms (atoms + superconductors) allowing each to do what it does best.

> 🔗 **Key insight:** Within a module, use Rydberg gates (fast, high-fidelity, 0.2–5 µs). Between modules, use photonic links (slower, heralded, ~ms timescale). The architecture separates the two problems — you don't need remote entanglement to be as fast as local gates, just reliable enough.

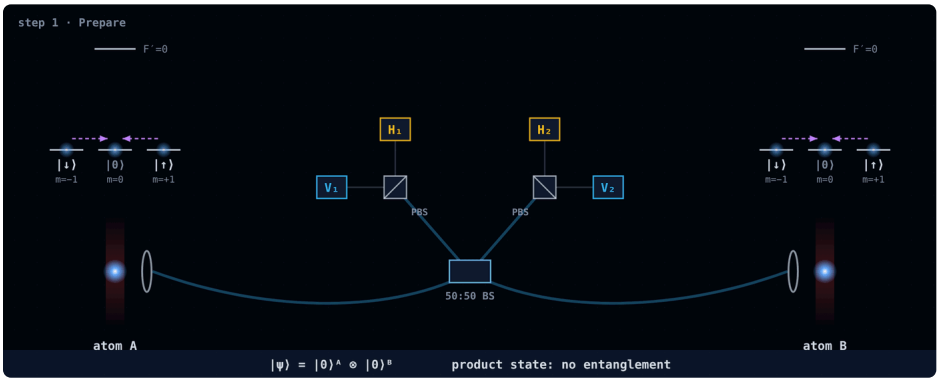


A photon from each neighbouring module meets the other at a Bell-state measurement (◆) in the middle of the link; a coincidence heralds a shared Bell pair (green arc).
N modules × M atoms each: the machine grows by adding modules and links, not by enlarging one array.

28.2 The Protocol — Step by Step

How two atoms that have never interacted become maximally entangled — one photon at a time.

Two-photon heralded entanglement with polarization qubits (the scheme of Hofmann et al. 2012)



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Optical pumping puts each atom in the same state, $|F{=}1, m_F{=}0\rangle$ (call it $|0\rangle$). The two atoms are in a product state $|0\rangle_A|0\rangle_B$: nothing connects them yet, and the qubit states $|\downarrow\rangle = m_F{=}{-}1$ and $|\uparrow\rangle = m_F{=}{+}1$ are empty.

Success probability

Each attempt succeeds with probability:

$$p_{\text{success}} = \frac{\eta_A \cdot \eta_B}{2}$$

where η_A, η_B are the total photon detection efficiencies (collection × fiber × detector) and the ½ is the Bell-state measurement keeping only the two Ψ outcomes. Because both photons must arrive, the rate falls as η^2. For scale: in the first atom–atom experiment each node detected about one photon in a thousand ($\eta \approx 10^{-3}$; Hofmann et al. 2012), so a herald arrived roughly every two minutes (about 3,600 in 107 h of data). Illustrative numbers for a future link: $\eta = 5\%$ gives $p \approx 1.25 \times 10^{-3}$, or about 1,250 pairs/s at 1 MHz attempts; a cavity-level $\eta = 50\%$ gives about 125,000 pairs/s at the same attempt rate. The calculator in section 07 lets you vary all of these.

Why the protocol is heralded (not post-selected)

A key feature: you *know* whether the protocol succeeded without measuring the atoms. A coincidence between two of the four detectors is the herald signal: it tells you the atoms are entangled, and in which Bell state, without revealing which atom is $|\uparrow\rangle$ and which is $|\downarrow\rangle$. This is fundamentally different from post-selection, which would collapse the atomic state. Failed attempts (no coincidence) simply mean you try again, with the atoms reset.

$$|\Psi_{\text{after BSM}}\rangle = \frac{1}{\sqrt{2}}\left(|\uparrow_A\downarrow_B\rangle \pm e^{i\theta}|\downarrow_A\uparrow_B\rangle\right) = |\Psi^{\pm}\rangle$$

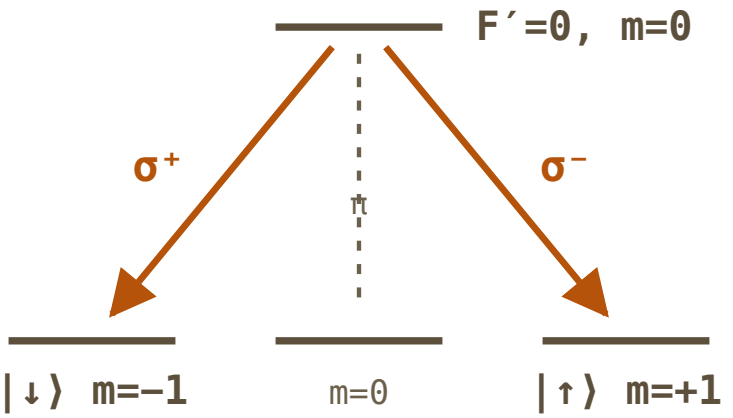
A heralding coincidence means the two photons had orthogonal polarizations, so the atoms are anti-correlated: H_1V_2 or V_1H_2 leaves them in $|\Psi^-\rangle$, H_1V_1 or H_2V_2 in $|\Psi^+\rangle$ (Hofmann et al. 2012). The phase θ can be tracked and corrected with single-qubit rotations. The state of the photons is irrelevant after detection — entanglement is established in the atoms.

28.3 Atom–Photon Entanglement

The fundamental resource: a single atom emitting a photon whose polarization is entangled with the atom's spin.

How it works: selection rules as the quantum interface

Optically pump the atom into $|F{=}1, m_F{=}0\rangle$, then excite it with a short π pulse to $|F'{=}0, m_{F'}{=}0\rangle$. That excited state has three ways down: to $m_F = -1$ by emitting σ^+ light, to $m_F = +1$ by emitting σ^-, or back to $m_F = 0$ with π light, which is not emitted along the quantization axis and so is never collected. A lens on that axis therefore only ever sees the first two, and does not know which one happened:



[87]Rb, as in Volz et al. (2006): one excitation, two collected decay paths.

$$|\Psi\rangle = \frac{1}{\sqrt{2}}\left(|\downarrow\rangle \otimes |\sigma^+\rangle + |\uparrow\rangle \otimes |\sigma^-\rangle\right), \qquad |\downarrow\rangle = |m_F{=}{-}1\rangle,\ |\uparrow\rangle = |m_F{=}{+}1\rangle$$

This is a maximally entangled state between the atom spin and the photon polarization; Volz et al. measured a fidelity of 0.87 with a single trapped [87]Rb atom. The photon then travels to the BSM station, carrying quantum information about the atom without measuring it.

Time-bin encoding (more practical)

Polarization encoding over optical fibers degrades due to birefringence (polarization rotation by stress, temperature). A more robust scheme uses **time-bin encoding**:

- $|\uparrow\rangle \rightarrow$ photon in early time window (E)
- $|\downarrow\rangle \rightarrow$ photon in late time window (L)
- Entangled state:
  $(1/\sqrt{2})(|\uparrow, E\rangle + |\downarrow, L\rangle)$
- Early/late defined by the timing of the excitation pulse on each atom

Time-bin qubits are insensitive to fiber birefringence and are the standard encoding for long-distance quantum links.

Photon indistinguishability — the hidden requirement

For the Bell state measurement to work, the photon from atom A and the photon from atom B must be *perfectly indistinguishable* — identical in frequency, bandwidth, temporal profile, and polarization. If they differ even slightly, the HOM interference is imperfect and entanglement fidelity drops.

The indistinguishability M is related to the visibility of Hong-Ou-Mandel interference. The resulting Bell state fidelity is bounded by:

$$\mathcal{F} \leq \frac{1+M}{2}$$

What degrades indistinguishability for neutral atoms:

- **Motional broadening:** A warm atom Doppler-shifts the emitted photon. Ground-state cooling narrows the emission spectrum dramatically — another reason cooling matters for QC.
- **Inhomogeneous light shifts:** Trapping laser shifts the transition frequency by an amount that varies across the array. Requires magic wavelengths or careful trap design.
- **Spectral diffusion:** Laser noise, magnetic field fluctuations, and phonon coupling all cause shot-to-shot frequency jitter.

✓ **Why ground-state cooling matters here:** Zero-point motion alone gives a Doppler spread $kv_{\mathrm{zpf}}/2\pi$ with

$v_{\mathrm{zpf}} = \sqrt{\hbar\omega/2m}$

: about 19 kHz for ^{87}Rb at 780 nm in a 100 kHz trap, far below the 6 MHz natural linewidth. A hot atom spreads its photons over the thermal velocity distribution instead, so cooling toward $\bar{n} \approx 0$ keeps the two photons spectrally identical.

28.4 Bell State Measurement & Hong-Ou-Mandel Effect

The quantum magic at the beamsplitter — and how imperfect photons degrade entanglement fidelity.

Hong-Ou-Mandel (HOM) effect

When two ideal, indistinguishable photons enter a 50-50 beamsplitter from different input ports, quantum interference forces both photons to exit the same port — the perfect-HOM limit has no coincidence counts. This is the HOM effect (Hong, Ou & Mandel 1987).

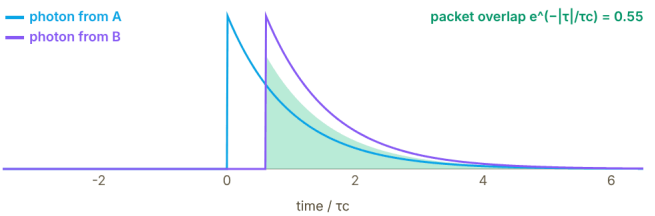
The Bell-state measurement uses exactly this interference, applied to the joint polarization state: the antisymmetric $|\Psi^-\rangle$ always leaves by different ports, the symmetric states leave together. That only works if the two photons are identical in every other degree of freedom (arrival time, frequency, spatial mode), and the HOM dip is how that is checked. For two spontaneously emitted photons offset by a delay τ, the probability of a coincidence behind the beam splitter is:

$$P_{\mathrm{coinc}}(\tau) = \frac{1}{2}\left[1 - \mathcal{M} \cdot e^{-|\tau|/\tau_c}\right]$$

where $\mathcal{M}$ is the overlap of everything except arrival time (0 = distinguishable, 1 = identical) and $\tau_c = 1/\Gamma = 1/(2\pi\Delta\nu)$ is the radiative lifetime of the emitting state ($\Delta\nu$ = natural linewidth, FWHM). Each photon is a wavepacket that switches on at emission and decays as e^{-t/τ_c}; the factor $e^{-|\tau|/\tau_c}$ is the squared overlap of two such packets offset by τ. The deeper the dip, the more identical the photons and the higher the possible entanglement fidelity.

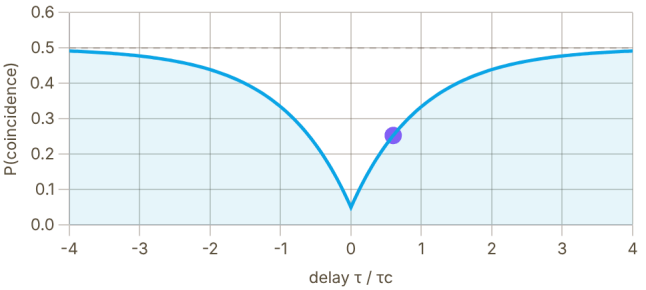
Interactive

Two photons, one beam splitter



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

P(coincidence) **0.253** dip visibility **0.90** max Bell fidelity **0.747**



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Top: the two wavepackets (intensity e^{-t/τ_c} after emission) for the chosen delay; their overlap sets how strongly they interfere. Bottom: the coincidence probability $\frac{1}{2}[1 - \mathcal{M}e^{-|\tau|/\tau_c}]$; the dot is the current delay. The heralded atom pair can be no better than $F = \frac{1}{2}(1 + \mathcal{M}e^{-|\tau|/\tau_c})$.

Linear optics BSM and its 50% limitation

A Bell-state measurement built from beam splitters and polarizers identifies only two of the four Bell states (here $|\Psi^{\pm}\rangle$); the two $|\Phi^{\pm}\rangle$ outcomes put both photons into one detector and look the same. This imposes a 50% efficiency ceiling on the BSM — meaning only half of all entanglement attempts that produce photons result in an unambiguous herald. The remaining 50% are discarded (not errors — just failures that are retried). With beam splitters, phase shifters and detectors alone and no extra photons, 50% is the most any Bell-state analyzer can reach (Calsamiglia & Lütkenhaus, *Appl. Phys.* B 72, 67 (2001)); adding ancilla photons can raise it, for example to 75% (Ewert & van Loock, arXiv:1403.4841). In practice the 50% is accepted and made up for with a higher repetition rate.

28.5 Challenges Specific to Neutral Atoms

Neutral atoms are excellent quantum computers — but photon emission is not their natural language.

Collection efficiency

Ions in Paul traps can be engineered to sit at the focus of a high-NA mirror. Neutral atoms in tweezers face a fundamental geometric challenge: the lens collecting the photon is on one side, and the trapping laser also comes from that side. Without a cavity, collection efficiency is:

$$\eta_{\text{coll}} = \frac{1 - \sqrt{1 - \text{NA}^2}}{2} \approx \frac{\text{NA}^2}{4}$$

For NA = 0.5 (typical tweezer): $\eta_{\text{coll}} \approx 7\%$. Most photons escape sideways. Cavities can boost this to 50–80% but add considerable complexity.

Telecom wavelength mismatch

Alkali atoms (Rb at 780 nm, Cs at 852 nm) emit at wavelengths with several dB/km fiber loss. At 10 km, that leaves roughly 0.001–0.1% transmission before connector and conversion losses. Long-distance links require **quantum frequency conversion** (QFC) to telecom wavelengths: about 0.3–0.35 dB/km near 1310 nm and about 0.18–0.25 dB/km near 1550 nm in low-loss single-mode fiber. Alkaline-earth atoms (Sr at 461/689 nm, Yb at 399/578 nm) face even worse fiber transmission without QFC. State-of-the-art QFC can reach high internal efficiency, but end-to-end efficiency and added noise remain experiment-specific.

Memory coherence vs communication time

After the atom emits a photon, it must wait for:

- Photon transit time: L/c (10 km → 50 µs)
- Classical herald to return: another L/c

- Total latency: 2L/c — the "round-trip" classical communication time

At L = 10 km, this is ~100 µs. The atom must remain coherent for this long — easy for nuclear-spin qubits ($T_2 \sim$ seconds), challenging for microwave-dressed states. The separation between *communication qubit* and *memory qubit* addresses this.

Spectral indistinguishability

Atoms A and B emit photons at slightly different frequencies due to different trapping potentials, residual magnetic fields, and AC Stark shifts. Two photons with spectral overlap less than 1 produce a reduced HOM visibility. Requires active stabilisation: matched trap depths, common-mode field cancellation, optical cavities to narrow linewidth, or real-time spectral demultiplexing.

Waiting for the herald

Remote entanglement is generated probabilistically. While waiting for a herald, the communication atoms are busy and the qubits that need the pair must wait. Attempts are independent, so the mean time to one Bell pair is:

$$\langle t \rangle = \frac{1}{p_{\text{success}}\, f_{\text{rep}}}$$

With $p \approx 0.1\%$ and $f_{\text{rep}} = 1$ MHz (illustrative), that is about 1 ms per Bell pair — hundreds of local Rydberg gate times, during which the computation that needs the pair is stalled. For long fibers f_{rep} itself is capped near $c/(nL)$ by the herald's return trip.

No native strong coupling to photons

Unlike NV centers (spin-selective emission) or quantum dots (Purcell-enhanced emission into a mode), neutral alkali atoms in free space are optically weak — the atom-photon cooperativity is $C = g^2/(\kappa\gamma) \ll 1$ without a cavity. This means each excitation has a low probability of producing a useful photon. Alkaline-earth atoms with narrow "clock" transitions have different challenges: long excited-state lifetimes mean slower photon production rates.

28.6 Approaches to Photon Collection & Emission

The field has converged on several complementary strategies, each with distinct engineering trade-offs.

🔭

High-NA Lens / Parabolic Mirror (Free Space) The simplest approach — collect as many photons as possible with a large aperture

η_{coll} ~ 2–15%

The most direct approach: use the highest available NA objective (0.5–0.9) to collect emitted photons and couple them into single-mode fiber. Requires careful mode matching between the objective's focal mode and the fiber's LP01 mode. A parabolic mirror around the atom can nearly double the collection efficiency by capturing photons from both sides.

- No moving parts, no cavity alignment, integrates naturally with tweezer control optics.
- Collection efficiency limited by NA. For NA = 0.65: $\eta_{\text{coll}} \approx 12\%$. Single-mode fiber coupling reduces this to ~6–8% total.
- **Weinfurter group (LMU):** a high-NA objective collects photons from a single Rb atom in a dipole trap into single-mode fiber; this free-space setup is behind both Hofmann et al. (2012) and van Leent et al. (2022). End-to-end, each node detected roughly one photon per thousand attempts in 2012, because coupling, fiber and detector losses multiply the collection efficiency.
- **Key challenge:** With $\eta \sim 8\%$, the two-photon success probability is $\eta^2/2 \approx 0.3\%$ — about 3,000 Bell pairs/second if attempts could run at 1 MHz (illustrative). Adequate for demonstrations; marginal for linking FTQC modules.

η_{coll} **2–15%** Rate **~kHz at 1 MHz attempts (illustrative)** Complexity **Low** Example **LMU (Weinfurter)**

💎

Cavity QED (Fiber Fabry-Pérot / Macroscopic Cavity) Purcell enhancement channels emission into a single mode — dramatically higher efficiency

η ~ 50–80%

A high-finesse optical cavity surrounding the atom enhances emission into the cavity mode via the Purcell effect. The fraction of photons emitted into the cavity mode (the "beta factor") is:

β = C / (1 + C)

where $C = g^2/(\kappa\gamma)$ is the cooperativity. For $C = 10$: $\beta = 91\%$. The photon exits the cavity mirror and can be coupled directly into single-mode fiber with 60–80% efficiency, giving total $\eta \sim 55$–72%.

- **Rempe group (MPQ):** Pioneered this approach with Rb in macroscopic Fabry-Pérot cavities. 2012 (Ritter et al., Nature): an elementary quantum network linking single Rb atoms in two cavities 21 m apart. 2021 (Daiss et al., Science): photonic quantum-logic gate between two atoms in separate cavities linked by a 60 m fiber.
- **Nanophotonic cavities (Lukin and Vuletić groups, Harvard/MIT):** a single Rb atom trapped next to a photonic-crystal cavity, strongly coupled to its mode (Tiecke et al., *Nature* 508, 241 (2014)).
- **Fiber Fabry-Pérot cavities (FFP):** Micro-machined tips on optical fibers that form a cavity around a single atom. Miniature (tens to hundreds of μm long) with small mode volumes, which helps reach high cooperativity; first used with cold atoms by the Reichel group (ENS Paris).
- **Key challenge:** Cavity alignment stability, acoustic/thermal vibration of cavity mirrors, and integration with the tweezer control optics all require careful engineering. Dielectric mirror surfaces close to the atoms can also collect charge and shift Rydberg levels, which complicates combining a cavity with Rydberg gates.

β **50–95%** η_total **~55–75%** Rate **~10⁵/s at 1 MHz attempts (illustrative)** Groups **MPQ, Harvard/MIT, ENS**

🌈

Quantum Frequency Conversion (QFC) Translate atom-native photons to telecom wavelengths for km-scale fiber links

η_QFC ~ 70–90%

Alkali atoms emit photons at 780–852 nm where standard SMF-28 fiber has 3–5 dB/km loss. At 1550 nm telecom C-band, loss is only 0.18–0.20 dB/km. Quantum frequency conversion (QFC) maps the atom's photon to telecom wavelengths while preserving all quantum information, including polarization and time-bin encoding.

Implemented via:

- **Periodically poled lithium niobate (PPLN) waveguides:** Difference-frequency generation (DFG) with a strong pump laser. For Rb, a 1607 nm pump converts 780 nm photons to 1517 nm (van Leent et al. 2022). Demonstrated with >70% internal efficiency and photon indistinguishability preserved.
- **Nanophotonic AlGaAs and GaP waveguides:** Higher nonlinearity and smaller footprint. Emerging platform.
- **Key demonstration:** van Leent et al. (2022, Nature): two Rb atoms in buildings 400 m apart entangled through up to 33 km of telecom fiber (extra length on spools), with QFC from 780 to 1517 nm.
- **Memory qubit + communication qubit architecture:** In ^{171}Yb ($I = 1/2$) and ^{87}Sr ($I = 9/2$), the nuclear spin serves as the memory qubit (coherence time seconds to minutes), while an optical transition handles photon emission (see the next card). QFC from Yb 399/556 nm to telecom is an active area.

η_QFC **60–92%** Fiber loss at 1550 nm **0.18 dB/km** Demonstrated **33 km (Rb)** Example **Weinfurter (LMU)**

Communication Qubit + Memory Qubit Architecture (AEL atoms) Nuclear-spin memories + electron-spin communication — the alkaline-earth advantage

T_2 > 1 s memory

Alkaline-earth-like (AEL) atoms (^{171}Yb, ^{87}Sr, ^{173}Yb) have a useful two-qubit architecture within a single atom. The **nuclear spin** ($I = 1/2$ for ^{171}Yb) forms an ultra-long-lived memory qubit whose magnetic sensitivity is suppressed to the nuclear-magneton scale because the $^1S_0$ ground state has no electronic angular momentum. The **electronic state** interacts with photons via the $^3P_1$ or $^3P_2$ transitions and serves as the communication qubit.

- **Protocol (schematic):** (1) Drive an optical transition so the atom emits a photon whose polarization or time bin is entangled with its qubit state. (2) A heralded BSM entangles two such communication qubits in different modules. (3) A local gate swaps the link onto nuclear-spin memory qubits. (4) Memories hold the pair while the next link attempts run.
- **Key advantage:** Memory qubits are coherent for seconds while communication is happening (100s of μs latency), so link attempts and local computation can overlap.
- **Groups:** Atom Computing (^{171}Yb), Thompson group (Princeton, ^{171}Yb), Covey group (UIUC, ^{171}Yb), Ye group (JILA, ^{87}Sr), among others.

Memory T_2 **>1 s** Comm. **optical transition** Memory qubit **nuclear spin** Atoms **^{171}Yb, ^{87}Sr**

✓ **Why this matters:** The communication qubit can be "sacrificed" in the photon emission process (which partially destroys coherence) while the memory qubit stays pristine. After heralded success, entanglement is swapped back to the memory. This decoupling is unavailable to alkali atoms with a single qubit register.

DLCZ Quantum Repeater Protocol (Atomic Ensembles) Collective excitation of an ensemble — the foundational quantum repeater scheme

~2001 (landmark proposal)

The DLCZ protocol (Duan, Lukin, Cirac, Zoller, 2001) was the first practically feasible quantum repeater scheme and remains conceptually foundational. Instead of single atoms, it uses *atomic ensembles* where a write laser pulse creates a probabilistic spin wave:

|vacuum⟩ → |0⟩_ensemble|0⟩_photon + ε|1-excitation spin wave⟩|1-photon⟩ + O(ε²)

Detecting the photon heralds the spin wave. Two such nodes interfere their photons at a BSM, creating entanglement between the ensembles. A read pulse then converts the spin wave deterministically into a photon — enabling entanglement swapping across repeater nodes.

- **Key advantage over single-atom schemes:** Collective enhancement — the stored excitation is read out into a well-defined spatial mode with an efficiency that grows with the ensemble's optical depth, so no cavity is needed for good collection.
- **Experiments:** Chou et al. (Kimble group, Caltech; *Nature* 438, 828 (2005)) heralded entanglement between two cold-atom ensembles about 3 m apart. Kuzmich's group (Georgia Tech) demonstrated matter–light state transfer (2004) and remote atomic qubits (2006). Many quantum-memory experiments followed.
- **Limitation for computing:** Ensemble-based qubits are harder to use as computational qubits compared to single-atom registers. DLCZ is most relevant for quantum communication networks rather than modular quantum computing.

Proposal **Duan, Lukin, Cirac, Zoller 2001** First remote ensembles **2005 (Chou et al.)** Best for **quantum networks**

28.7 Entanglement Rate Calculator

Estimate Bell pair generation rates for different hardware configurations.

Photonic Link Parameters

COLLECTION EFFICIENCY	FIBER TRANSMISSION	TOTAL H PER PHOTON
6.70% from lens + Purcell	99.96% 10^(−αL/10), α in dB/km	3.70% coll × FC × fiber × det

SUCCESS PROB / ATTEMPT	BELL PAIR RATE	TIME PER BELL PAIR
0.07% per BSM attempt	684.3 Hz pairs / second	1.5 ms avg latency

Calculator output at the page's default settings.

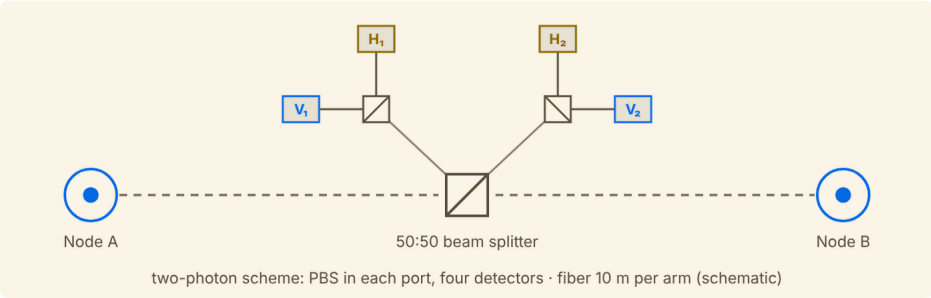
⚠ 100 Hz–1 kHz: fine for demonstrations, slow for linking modules (rough rule of thumb, not a published threshold)

Attempt rate 1.0 MHz (the 2L/v = 0.1 µs herald round trip is not the limit here).

Model: η = collection × fiber coupling × fiber transmission × detector, per photon. The 1310/1550 nm options assume the atom's photon is frequency-converted first; conversion loss is not included. With a cavity, collection = β × 0.85 outcoupling (illustrative). Dark counts, multi-photon events and HOM visibility are left out.

Watch the link run: attempts, losses and heralds

A Monte Carlo of the link above, with the same η, scheme and repetition rate. In every attempt of the two-photon scheme both atoms emit a photon; each survives collection, fiber coupling, the fiber and its detector with probability η. When both arrive, the beam splitter and polarizing beam splitters herald a Bell pair half the time (H_1V_2 or V_1H_2 for $|\Psi^-\rangle$, H_1V_1 or H_2V_2 for $|\Psi^+\rangle$); the other half both photons end up in the same detector and the attempt is discarded, so $p = \eta^2/2$. In the single-photon scheme each atom is excited with probability $\varepsilon = 0.05$ and exactly one detector click heralds, $p \approx 2\varepsilon\eta$. **Slow motion** plays one attempt at a time so you can see where photons are lost; **Real time** runs attempts at the lab repetition rate (one second here is one second of lab time, up to 2×10^5 attempts per frame) and compares the heralded fraction with the calculator.



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

attempts**1**lab time 1.0 µs

photons detected**0** · **0**A · B (η = 3.70%)

heralded pairs**0**none yet

p per attempt—calculator 6.84e-4

Schematic layout (not to scale). Detector dark counts, timing jitter, multi-photon emission and imperfect HOM visibility are left out, as in the calculator.

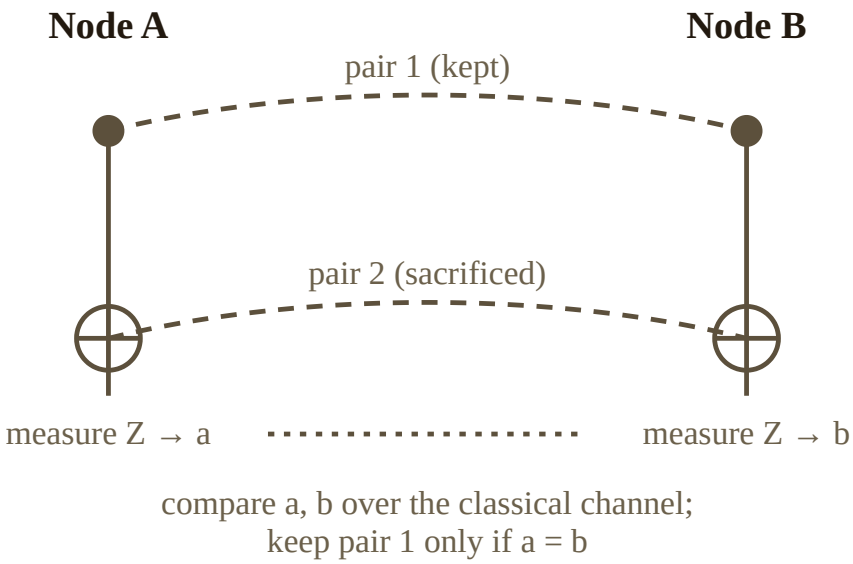
📌 **FTQC requirement:** For surface-code error correction in a modular architecture with physical error rate p_{link} and code distance d, the link needs to generate Bell pairs faster than the error-correction cycle time $\sim d \times t_{\text{gate}}$. The rate actually needed depends on the code, the gate times and how many links run in parallel, so there is no single number; as a rough rule of thumb, rates of order 10^3 pairs/s or more per link are what modular proposals aim for. Cavity-enhanced collection is the most direct way to get there, since the rate scales as η^2.

28.8 Fixing Noisy Pairs: Entanglement Distillation

Real heralded pairs never have fidelity 1. Before pairs are chained or used for gates, several noisy ones can be traded for fewer, better ones.

The recurrence protocol

Heralded pairs arrive with fidelity below 1 (imperfect HOM visibility, dephasing during the wait, detector noise). **Entanglement distillation** (also called purification) turns several noisy pairs into fewer, better ones using only local operations and classical communication, directly on the physical qubits, with no error-correcting code. In the recurrence protocol both nodes hold two pairs; each applies a CNOT from its half of pair 1 to its half of pair 2 (a *bilateral* CNOT), measures its half of pair 2 in Z, and the two nodes compare the results over the classical channel. Equal parities keep pair 1, now with higher fidelity; unequal parities discard it.



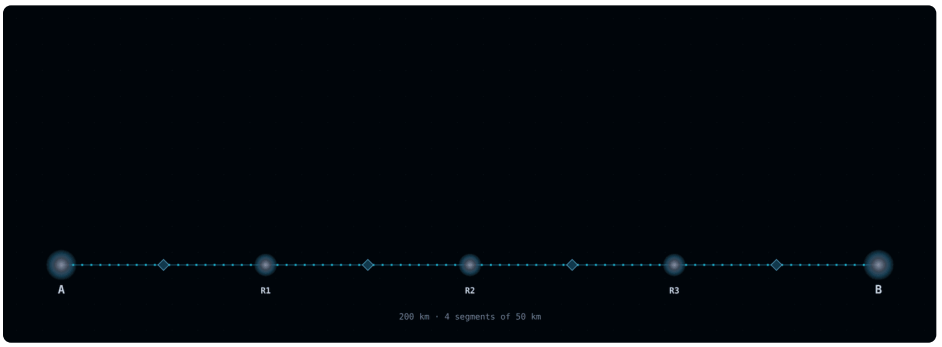
ROUND	FIDELITY	P(KEEP)	RAW PAIRS / OUTPUT
0 (raw)	0.8000	—	1
1	0.8382	76.9%	2.6
2	0.8736	80.7%	6.4
3	0.9045	84.6%	15.2

One round of the Bennett et al. (1996) recurrence for Werner states, with the pairs re-twirled to Werner form between rounds: $F' = \dfrac{F^2 + \left(\frac{1-F}{3}\right)^2}{F^2 + \frac{2F(1-F)}{3} + 5\left(\frac{1-F}{3}\right)^2}$, kept with probability equal to the denominator; it only helps for F > 1/2. Perfect local gates are assumed (in practice local gate errors set a ceiling), and the Deutsch et al. (1996) variant converges faster. Each round consumes two pairs per kept pair, so the raw-pair cost grows like ∏ 2/P(keep). Compare magic-state *distillation* and *cultivation*, which need error correction: Learn Quantum → Magic States.

28.9 Quantum Repeater Chain

One link gets exponentially worse with distance. Cut the fiber into segments, herald each one, store the pairs in atoms, and join them with entanglement swapping. Watch a chain build end-to-end Bell pairs, then see where it beats direct transmission.

Interactive



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

1 s on screen = 249 ms of lab time lab clock 265 ms

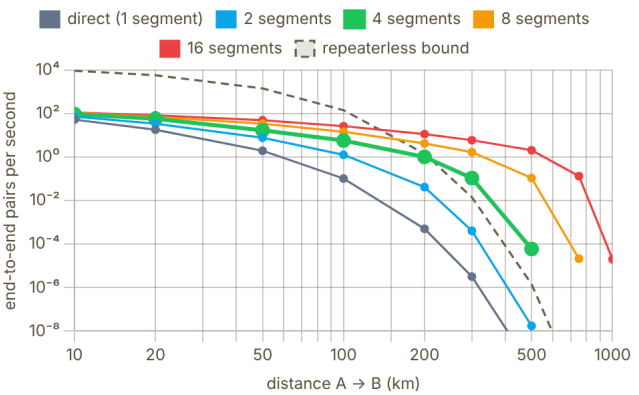
photon attempts on an empty segment stored pair, coloured by fidelity (0.5 → 1) swap (Bell measurement at a repeater) end-to-end pair delivered

Chain parameters

Heralding scheme Two-photon (Barrett & Kok 2005): at most 1/2, quadratic in η. Single-photon: one click from either node; the state keeps an $\alpha|\uparrow\uparrow\rangle$ admixture, so $F \approx 1 - \alpha$ (Humphreys et al. 2018).

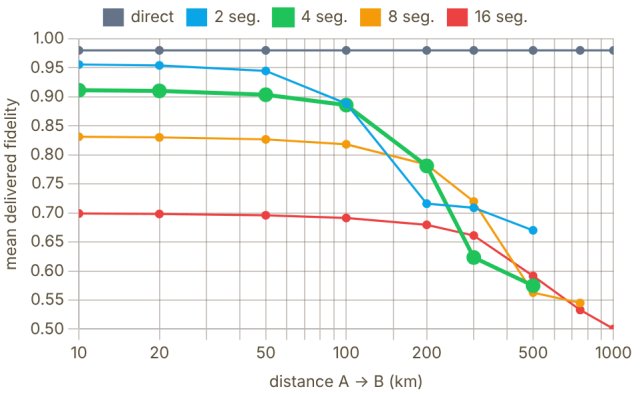
per photon η = 3.17% · attempt p = 5.03e-4 · attempt time 255 µs · mean wait per segment 506 ms · F_0 = 0.980
repeaterless bound over 200 km: 1.5 /s (η = 1.0e-5, one use per 10 µs)

End-to-end rate vs distance



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Delivered fidelity vs distance



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Each point: event-driven Monte Carlo of up to 150 delivered pairs with the parameters above (thick line = your segment count). 5 points delivered nothing within the simulated budget and are left out. The repeaterless bound uses the same node efficiency and one channel use per local overhead time.

What the model does

- **Elementary link:** segment length $L_0 = D/N$, station in the middle, so each photon crosses $L_0/2$: $\eta = \eta_{\text{node}}\, 10^{-\alpha L_0/20}$. Two-photon: $p = \eta^2/2$. Single-photon: $p \approx 2\alpha\eta$ for $\eta \ll 1$, $F_0 = 1 - \alpha$.
- **Attempt time:** $t = L_0/c_f + t_{\text{local}}$ (photon to the station and the herald back, $c_f = c/1.47$). Each empty segment waits a geometric number of attempts.
- **Pairs are Werner states**, $F = (3w + 1)/4$. Storage: $w \to w\, e^{-2t/T_{\text{mem}}}$ (two memories). Swap: $w = w_1 w_2\, w_{\text{swap}}$, done as soon as two neighbouring pairs exist; classical signalling time for the swap results is not counted.
- **Repeaterless bound** (Pirandola et al. 2017): a lossy channel of transmissivity η carries at most $-\log_2(1 - \eta) \approx 1.44\,\eta$ ebits per use. Plotted with $\eta = \eta_{\text{node}}\, 10^{-\alpha D/10}$ at one use per t_{local}.
- No distillation, multiplexing or dark counts. Without distillation, fidelity falls with every swap; the distillation section above shows how pairs are cleaned up.

Sources: Briegel, Dür, Cirac & Zoller, PRL 81, 5932 (1998) · Barrett & Kok, PRA 71, 060310 (2005) · Humphreys et al., *Nature* 558, 268 (2018) · Pirandola, Laurenza, Ottaviani & Banchi, *Nat. Commun.* 8, 15043 (2017) · Sangouard et al., *Rev. Mod. Phys.* 83, 33 (2011)

28.10 Where Remote Entanglement Fits in the Computing Landscape

Not a replacement for the Rydberg gate — a complement that removes the scaling ceiling.

Today

~10^3–10^4

Atoms per array (6,100 Cs atoms reported in 2025). Rydberg gates within one array; atom–atom photonic links shown only in physics experiments.

Next step (illustrative)

~10,000

First inter-module photonic links. 2–5 modules. Error-corrected logical qubits across modules.

LONG TERM (ILLUSTRATIVE)

~10^6

Many-module cluster. Photonic links as the "quantum bus". Fault-tolerant computation.

The two lower cards sketch a plausible trajectory, not a forecast or any company's roadmap.

The two-layer architecture

The mature vision for neutral-atom quantum computing is a two-layer architecture that separates timescales:

- **Layer 1 — Intra-module:** Rydberg gates at 0.2–5 μs. High fidelity (99.5–99.9%). Deterministic. Operates within a single tweezer array. All the work in this site's other QC tools addresses this layer.
- **Layer 2 — Inter-module:** Photonic entanglement at ~10 μs–10 ms per Bell pair. Probabilistic but heralded. Operates between tweezer arrays in different vacuum chambers. Error-corrected logical operations are teleported across this layer.

The layers are complementary. Fault-tolerant quantum computing requires both: logical qubits encoded across many physical qubits (using Rydberg gates within modules) and logical operations on qubits in different modules (using entanglement-based gate teleportation).

Who's working on it?

- **Neutral-atom network experiments:** Weinfurter (LMU, Rb atoms: Hofmann et al. 2012, van Leent et al. 2022 over 33 km) and Rempe (MPQ, atoms in cavities: Daiss et al. 2021).
- **Tweezer arrays with photonic interfaces:** Lukin (Harvard), Thompson (Princeton) and Covey (UIUC), among others, are coupling tweezer-held atoms to cavities or fibers.
- **Companies:** Atom Computing, QuEra, Pasqal and Infleqtion build neutral-atom processors. Atom Computing uses ^{171}Yb nuclear-spin qubits, whose long coherence suits holding a qubit while link attempts run. In March 2026 Google Quantum AI announced a neutral-atom effort alongside its superconducting work, with a Boulder team led by Adam Kaufman (Google, JILA).
- **The trapped-ion benchmark:** Oxford's two-node ion network reached 182 heralded Bell pairs per second at 94% fidelity (Stephenson et al., PRL 124, 110501 (2020)) and later ran a distributed quantum algorithm across the link (Main et al., arXiv:2407.00835, published in *Nature* 2025).

✓ **The optimistic case:** Photonic interconnects between tweezer modules can be added modularly — start with 2 modules (a proof-of-principle), scale to 10 (early FTQC), then hundreds. Each module can be optimized locally before the link is added. The link is additive, not a redesign.

⚠️ **The honest challenge:** Free-space collection with $\eta \approx 5$–10% per node gives $p = \eta^2/2 \approx 0.1$–0.5%, so about 1–5 kHz even at an optimistic 1 MHz attempt rate — likely too slow for large-scale FTQC without cavity enhancement or multiplexing. The cavity integration challenge (alignment stability, compatibility with tweezer control optics, scalability) remains an open engineering problem.

See also: To understand what remote entanglement is *for*, see the Gate Fidelity Budget (what fidelity each local gate contributes), Rydberg Calculator (blockade physics within a module), QC Industry Map (where each company sits), and Quantum Error Correction (why you need millions of qubits).

28.11 Historical Milestones

From the EPR paradox to metropolitan quantum networks — a 90-year arc toward practical quantum interconnects.

1935

Einstein-Podolsky-Rosen (EPR) Paradox

Einstein, Podolsky, and Rosen argue that the perfect correlations of an entangled pair make quantum mechanics "incomplete" if no influence can travel faster than light. Bell (1964) provides the mathematical test; Aspect (1982) demonstrates violations experimentally with photons. Theory

1993

Quantum Teleportation Protocol (Bennett et al.)

Bennett and colleagues propose using a shared Bell pair + classical communication to transfer an arbitrary quantum state between two parties without physically transporting the qubit. Laid the theoretical foundation for all remote entanglement applications. Theory

1997

First quantum teleportation (Zeilinger group)

Bouwmeester et al. teleport polarization states between photons over 1 m. The Nobel Prize in Physics 2022 recognised Clauser, Aspect, and Zeilinger for this line of work.

1999

Cabrillo et al. — Single-photon entanglement scheme

First proposal to entangle two distant atoms using detection of a *single* photon from a superposition of emissions. Success probability $\propto \eta$ (linear in efficiency — a significant theoretical improvement). Theory

2001

DLCZ Quantum Repeater Protocol

Duan, Lukin, Cirac, and Zoller propose a practical quantum repeater using atomic ensembles. This landmark paper provided the first route to long-distance quantum communication, spawning hundreds of experiments worldwide. Ensemble Theory

2005

First heralded entanglement of remote atomic ensembles (Kimble group)

Chou et al. (Caltech; *Nature* 438, 828) detect one photon behind a beam splitter, heralding a single excitation shared between two cold-atom ensembles about 3 m apart: the DLCZ building block. Kuzmich's group (Georgia Tech) showed matter–light state transfer (2004) and entangled remote atomic qubits (2006). Ensemble

2007

Remote entanglement of single ions (Monroe group)

Moehring et al. (Monroe group, then at Michigan; Nature 449, 68) demonstrate heralded entanglement between two $^{171}Yb^{+}$ ions in separate traps about 1 m apart. First entanglement between two individually addressable single qubits. Trapped Ion

2012

Remote entanglement of single neutral atoms (Weinfurter and Rempe groups)

Hofmann et al. (Weinfurter group, LMU Munich; Science 337, 72) herald entanglement between two single Rb atoms in free-space dipole traps 20 m apart. The same year Ritter et al. (Rempe group, MPQ; Nature 484, 195) link single Rb atoms in two optical cavities 21 m apart. Neutral Atom

2015

Loophole-free Bell test (Hensen et al., Delft)

The Hanson group uses NV center remote entanglement over 1.3 km to close all loopholes simultaneously in a Bell inequality test. (Nature 526, 682). NV Center

2020

High-rate remote entanglement of ions (Oxford)

Stephenson et al. (Lucas group, Oxford; PRL 124, 110501) entangle two $^{88}Sr^{+}$ ions in separate traps via a photonic Bell-state measurement at 182 pairs/s with 94% fidelity. Trapped Ion

2021

Photonic quantum gate between distant atoms (Rempe group)

Daiss et al. (MPQ, Science 2021) realise a quantum-logic gate between two Rb atoms each in separate cavity QED nodes, mediated by a shared photon reflected off each cavity. First demonstration of a remote two-qubit gate between neutral atoms. Neutral Atom

2021–22

Three-node quantum network (Delft)

Pompili et al. (*Science* 372, 259, 2021) build a three-node network of NV centers and distribute entanglement between the outer nodes by swapping at the middle one; Hermans et al. (*Nature* 605, 663, 2022) teleport a qubit between non-neighbouring nodes. NV Center

2022

Rb atoms entangled through 33 km of fiber (Weinfurter group)

van Leent et al. (LMU, *Nature* 607, 69) entangle two single Rb atoms in labs 400 m apart through up to 33 km of telecom fiber, using quantum frequency conversion to 1517 nm. Neutral Atom

2024+

Neutral atoms: race to tweezer-array interconnects

Many groups (Lukin/Harvard, Covey/UIUC, Thompson/Princeton, Rempe/MPQ, Weinfurter/LMU and others) are pushing toward photonic interconnects for tweezer-array quantum computers. Modular, photonically linked processors are widely discussed as the route past one array. Neutral Atom

28.12 References & Further Reading

Seminal theory papers, key experimental demonstrations, and review articles on remote entanglement and quantum networking.

Foundational Theory

Kimble (2008), *Nature* 453 1023 — "The quantum internet": the seminal vision paper defining protocols and requirements Duan, Lukin, Cirac & Zoller (2001), *Nature* 414 413 — DLCZ protocol: quantum repeater using atomic ensembles and single photons Barrett & Kok (2005), *PRA* 71 060310 — double-heralded two-photon protocol that is robust to photon loss; used in the Delft NV-centre experiments Wehner, Elkouss & Hanson (2018), *Science* 362 — "Quantum internet: A vision for the road ahead": stage-by-stage roadmap

Key Experimental Demonstrations

Moehring et al. (2007), *Nature* 449 68 — first remote entanglement between two trapped ion qubits via photons Hofmann et al. (2012), *Science* 337 72 — heralded atom–atom entanglement between two labs 20 m apart Daiss et al. (2021), *Science* 371 614 — quantum-logic gate between two cavity-QED modules linked by a 60 m fiber Pompili et al. (2021), *Science* 372 259 — three-node quantum network with NV centers; end-to-end entanglement van Leent et al. (2022), *Nature* 607 69 — remote entanglement of two neutral atoms over 33 km fiber Evered et al. (2023), *Nature* 622 — 99.5% fidelity Rydberg two-qubit gate; sets the bar for local gate quality needed before networking

Reviews & Online Resources

Sangouard et al. (2011), *Rev. Mod. Phys.* 83 33 — comprehensive review of quantum repeater architectures and performance
Azuma et al. (2023), *Rev. Mod. Phys.* 95 045006 (arXiv:2212.10820) — quantum repeaters: from quantum networks to the quantum internet (review)
Quantum Internet Alliance — EU roadmap project; stage definitions and technology readiness levels

CHAPTER 28 SUMMARY

Photonic links remove the single-chamber scaling ceiling, completing the technical arc from a single atom to modular fault-tolerant networks. Part IV serves as a field guide to the literature, research groups, and quantum industry.

Problems for Chapter 28

1. USE THE TOOL In the rate calculator, increase the NA from 0.5 to 0.7. By what factor does the two-photon heralded rate change?
2. BY HAND Explain why the two-photon scheme scales as η^2 while the single-photon scheme scales as η, and what the single-photon scheme gives up in fidelity.
3. DESIGN For a 10 km link, find the memory coherence time needed for a three-segment repeater chain to beat direct transmission, using the repeater tool.

Use the tool: on the interactive page. By hand: from the formulas in this chapter. Design: a decision a lab would make. Worked solutions are available to instructors on request (see Part V).

PART IV

Into the Field

Reading the literature that built the field, finding a research group, and understanding the companies and platforms you may join. Facts about companies change quickly: every number here carries the date it was checked.

29	Foundational AMO Literature	AMO Paper Roadmap
30	Global AMO Research Directory	AMO Group Finder
31	Quantum Computing Industry Landscape	Quantum Industry Map
32	Neutral-Atom Qubits: ^{87}Rb vs. ^{171}Yb	Rb vs Yb Qubit Comparison
33	Trapped-Ion Architectures: Quantinuum vs. IonQ	Quantinuum vs IonQ
34	Superconducting Processors: Google vs. IBM	Google vs IBM
35	Photonic Architectures: PsiQuantum vs. Xanadu	PsiQuantum vs Xanadu

CHAPTER 29 · CAREER STEP 01

29

Foundational AMO Literature

📚 AMO Paper Roadmap · Syllabus

55 papers that build an AMO career, sorted by the stage at which they pay off: from your first week reading about laser cooling to the reviews you cite in a postdoc fellowship application. Pick a track, tick papers off as you read them, and jump from a paper to the toolkit page that puts its physics to work.

> *Part IV transitions from evergreen lab physics to navigating the AMO research community. This chapter organizes 55 landmark papers into structured reading tracks from first-week primers to modern logical-qubit milestones.*

IN THIS CHAPTER — No prerequisites

- Follow curated reading sequences across laser cooling, optical tweezers, cavity QED, Rydberg gates, clocks, and quantum error correction.
- Distinguish foundational theory papers, pedagogical reviews, and landmark experimental demonstrations.
- Connect each canonical paper back to the interactive calculators and chapters in Parts I–III.

WEBSITE COMPANION Live interactive page ↗ · Week 14

29.1 Choose how to read

Do not read all 55 the same way. Some papers are for orientation, some for equations you will reuse, and some for the judgement you need in interviews and proposals. Pick a track to see only its papers; the stages below still set the depth.

Showing all 55 papers

Ticks stay in this browser only. The copied list contains the papers that match the current track and search, with their DOI links.

29.2 Read by career stage

1 / 4

🎓

Undergrad / Entering AMO

These 10 papers form the conceptual skeleton of the field. Read them in order. You do not need to understand every equation — you need to understand *what was achieved* and *why it was surprising*. Each one is short enough to read in an afternoon.

Nobel lectures — How laser cooling began

1. Laser Cooling and Trapping of Neutral Atoms Keystone

W. D. Phillips · Rev. Mod. Phys. 70, 721 (1998)

Phillips explains—with genuine surprise—how his group found atoms colder than the Doppler limit they had calculated. The candid narrative makes the physics memorable: polarization-gradient Sisyphus cooling, optical molasses, and the lesson that experiments sometimes know more than theory.

doi:10.1103/RevModPhys.70.721 →

2. The Manipulation of Neutral Particles

S. Chu · Rev. Mod. Phys. 70, 685 (1998)

Covers the invention of optical molasses, the first optical trap for atoms, atomic fountains and atom interferometry, and the use of optical tweezers on single DNA molecules. Chu's lecture is uniquely interdisciplinary and conveys the spirit of opportunistic exploration across physics and biology that defines creative AMO work.

doi:10.1103/RevModPhys.70.685 →

3. Manipulating Atoms with Photons

C. N. Cohen-Tannoudji · Rev. Mod. Phys. 70, 707 (1998)

The most theoretical of the three, with dressed-state picture, quantum jumps, and subrecoil cooling via VSCPT. Cohen-Tannoudji built the theoretical framework that the whole field still uses — reading this gives you the conceptual tools, not just the results.

doi:10.1103/RevModPhys.70.707 →

Bose–Einstein condensation — The 1995 breakthrough

4. Observation of Bose-Einstein Condensation in a Dilute Atomic Vapor Keystone

M. H. Anderson et al. · Science 269, 198 (1995)

The first BEC: a rubidium-87 condensate at JILA, appearing near 170 nK. Four pages. The bimodal velocity distribution in Figure 2 is the most reproduced image in cold atoms. Read this before reading anything theoretical about BEC.

doi:10.1126/science.269.5221.198 →

5. Bose-Einstein Condensation in a Gas of Sodium Atoms

K. B. Davis et al. · Phys. Rev. Lett. 75, 3969 (1995)

Ketterle's group achieved BEC in sodium four months after the JILA result, using a different trap geometry. Pairs with Anderson 1995 to show that BEC is robust and generalizable — not a one-atom-species miracle — and its optically plugged magnetic trap produced condensates with far more atoms.

doi:10.1103/PhysRevLett.75.3969 →

Trapping — The tools everything else rests on

6. Acceleration and Trapping of Particles by Radiation Pressure

A. Ashkin · Phys. Rev. Lett. 24, 156 (1970)

The paper that launched optical trapping — written 16 years before the first single-beam trap was demonstrated. Ashkin accelerated micrometre-sized spheres with radiation pressure, trapped them between two opposing beams, and proposed doing the same with atoms. The clarity of the physical argument is a masterclass in how to propose a radical experiment.

doi:10.1103/PhysRevLett.24.156 →

7. Trapping of Neutral Sodium Atoms with Radiation Pressure Experiment

E. L. Raab et al. · Phys. Rev. Lett. 59, 2631 (1987)

The magneto-optical trap (MOT) — still the starting point of nearly every cold-atom experiment today. The combination of Zeeman-shifted σ^+/σ^- beams and a quadrupole field is the single most important technical invention in AMO. Four pages; understand the selection-rule argument for why the force is restoring.

doi:10.1103/PhysRevLett.59.2631 →

8. Optical Dipole Traps for Neutral Atoms Review

R. Grimm, M. Weidemüller, Y. Ovchinnikov · Adv. At. Mol. Opt. Phys. 42, 95 (2000)

The definitive review of optical trapping for neutral atoms. Covers the AC Stark shift, trap depth, heating rates, species-specific polarizability, and beam geometry. Every AMO experimentalist has this bookmarked. Start from §2 and use the rest as a reference.

doi:10.1016/S1049-250X(08)60186-X →

Neutral-atom qubits — Why AMO became a technology platform

9. Fast Quantum Gates for Neutral Atoms Keystone

D. Jaksch et al. · Phys. Rev. Lett. 85, 2208 (2000)

The theoretical proposal that ignited Rydberg quantum computing. Jaksch et al. pointed out that the dipole–dipole interaction between two atoms in high-n Rydberg states is strong enough to produce a fast, high-fidelity two-qubit gate. The blockade mechanism described here is now in production at QuEra, Atom Computing, and Pasqal.

doi:10.1103/PhysRevLett.85.2208 →

10. Sub-Poissonian Loading of Single Atoms in a Microscopic Dipole Trap Keystone

N. Schlosser et al. · Nature 411, 1024 (2001)

The birth of the optical tweezer as a single-atom tool. By tightly focusing a laser to a ~1 μm waist, Schlosser showed that collisional blockade limits occupancy to exactly 0 or 1 — the single-atom regime. This is the foundational experiment behind the entire tweezer-array quantum computing industry.

doi:10.1038/35082512 →

🔬

Grad Year 1 — Building Experimental Fluency

These papers form the technical bedrock. Read the technique papers relevant to your lab first; read the rest over 6 months. By the end of your first year, you should be able to look at any ultracold-atom setup and identify the purpose of every major beam and coil.

Laser cooling theory — Below the Doppler limit

11. Laser Cooling Below the Doppler Limit by Polarization Gradients Theory

J. Dalibard & C. Cohen-Tannoudji · J. Opt. Soc. Am. B 6, 2023 (1989)

The definitive theory of Sisyphus cooling. The dressed-state picture of a moving atom climbing a potential hill, losing energy by optical pumping at the top, and repeating is one of the most beautiful mechanisms in AMO. Essential before touching any sub-Doppler experiment.

doi:10.1364/JOSAB.6.002023 →

12. Observation of Atoms Laser Cooled Below the Doppler Limit Experiment

P. D. Lett et al. · Phys. Rev. Lett. 61, 169 (1988)

The surprise experiment that showed sodium atoms at 43 μK — six times below the Doppler limit — and forced the field to invent Sisyphus cooling as an explanation. The story of how this result was initially disbelieved and then explained is a useful lesson in how experimental discovery works.

doi:10.1103/PhysRevLett.61.169 →

Evaporative cooling — The path to quantum degeneracy

13. Evaporative Cooling of Trapped Atoms Review

W. Ketterle & N. J. van Druten · Adv. At. Mol. Opt. Phys. 37, 181 (1996)

The quantitative theory of evaporation: truncation parameter η, elastic collision rate, and the runaway regime condition. Every student designing an evaporation ramp needs this paper. The phase-space density scaling ρ_PSD ∝ N^α (α ≫ 1 in runaway) gives the key intuition for why forced evaporation works so efficiently.

doi:10.1016/S1049-250X(08)60101-9 →

Optical lattices — Atoms as condensed-matter Hamiltonians

14. Cold Bosonic Atoms in Optical Lattices Theory

D. Jaksch et al. · Phys. Rev. Lett. 81, 3108 (1998)

This four-page paper derived the Bose-Hubbard Hamiltonian for atoms in an optical lattice and predicted the superfluid-to-Mott insulator transition. It opened the entire field of quantum simulation with cold atoms. The mapping between lattice depth, tunneling J, and on-site interaction U is derived here for the first time in this context.

doi:10.1103/PhysRevLett.81.3108 →

15. Quantum Phase Transition from a Superfluid to a Mott Insulator in a Gas of Ultracold Atoms Keystone

M. Greiner et al. · Nature 415, 39 (2002)

The experimental realization of the Bose-Hubbard Mott transition — announced with interference pattern images showing the coherence collapse and restoration across the phase boundary. One of the most influential cold-atom papers ever. Demonstrated that cold atoms can simulate a quantum phase transition that is hard to study in solid-state systems.

doi:10.1038/415039a →

Quantum gas microscopy — Seeing every atom

16. A Quantum Gas Microscope for Detecting Single Atoms in a Hubbard-Regime Optical Lattice Keystone

W. S. Bakr et al. · Nature 462, 74 (2009)

The quantum gas microscope: site-resolved single-atom fluorescence imaging in a 2D optical lattice. This invention transformed quantum simulation from a bulk measurement to a site-by-site snapshot. The imaging technique (high-NA fluorescence in optical molasses while a deep lattice pins the atoms) is still the template for modern microscope experiments.

doi:10.1038/nature08482 →

Fermi degeneracy — Ultracold fermions

17. Onset of Fermi Degeneracy in a Trapped Atomic Gas Keystone

B. DeMarco & D. S. Jin · Science 285, 1703 (1999)

The first degenerate Fermi gas with potassium-40 — a different quantum statistics from BEC, with Pauli exclusion suppressing s-wave collisions between identical fermions. DeMarco and Jin's strategy of evaporating a mixture of two spin states, so the atoms can still thermalize, is still the standard approach.

doi:10.1126/science.285.5434.1703 →

Feshbach resonances — Tuning interactions

18. Feshbach Resonances in Ultracold Gases Review

C. Chin, R. Grimm, P. Julienne, E. Tiesinga · Rev. Mod. Phys. 82, 1225 (2010)

The comprehensive review of magnetically tunable s-wave scattering lengths. The tables of resonance positions for Rb, Cs, Li, Na, K, and more are worth the price of admission on their own. The coupled-channel formalism gives you the language to read theory papers about strongly interacting regimes.

doi:10.1103/RevModPhys.82.1225 →

Laser stabilization — Keeping your lasers on resonance

19. Laser Phase and Frequency Stabilization Using an Optical Resonator

R. W. P. Drever et al. · Appl. Phys. B 31, 97 (1983)

The original PDH paper — a collaboration between a gravitational-wave physicist (Drever) and a laser-spectroscopy group (Hall); the technique became the gold standard for laser frequency stabilization in every precision spectroscopy lab. The error signal derivation is compact and should be understood analytically, not just plugged into a datasheet.

doi:10.1007/BF00702605 →

20. An Introduction to Pound-Drever-Hall Laser Frequency Stabilization Review

E. D. Black · Am. J. Phys. 69, 79 (2001)

The pedagogical companion to Drever 1983 — explicitly designed for students. Black walks through every step of the PDH derivation with physical intuition at each stage. Read this before building your first locking circuit. The error-signal-versus-detuning plots are what to keep in your notebook.

doi:10.1119/1.1286663 →

Optical tweezers — Single-atom control

21. Observation of a Single-Beam Gradient Force Optical Trap for Dielectric Particles

A. Ashkin et al. · Opt. Lett. 11, 288 (1986)

The demonstration that a single tightly focused Gaussian beam can trap a particle in 3D using the gradient force — without a counterpropagating beam. This is the foundational tweezer geometry used for both biological molecules and cold atoms. Nobel Prize (Ashkin) in 2018.

doi:10.1364/OL.11.000288 →

22. Cooling a Single Atom in an Optical Tweezer to Its Quantum Ground State Experiment

A. M. Kaufman, B. J. Lester & C. A. Regal · Phys. Rev. X 2, 041014 (2012)

Ground-state cooling of a single ^{87}Rb atom in a tweezer with three-dimensional Raman sideband cooling. The measurement of the sideband asymmetry to extract temperature is a technique every tweezer experimentalist uses. This is the paper that made tweezers a precision quantum platform, not just a trapping tool.

doi:10.1103/PhysRevX.2.041014 →

BEC theory — Understanding the condensate

23. Theory of Bose-Einstein Condensation in Trapped Gases Review

F. Dalfovo et al. · Rev. Mod. Phys. 71, 463 (1999)

The comprehensive theory review written in the four years after the first BEC experiments. The Thomas-Fermi density profile, collective modes, vortices, and the mean-field breakdown — all derived cleanly from the Gross-Pitaevskii equation. The paper to read before tackling Pitaevskii & Stringari's textbook.

doi:10.1103/RevModPhys.71.463 →

Spectroscopy — Reading the atom's internal structure

24. Optical Pumping Review

W. Happer · Rev. Mod. Phys. 44, 169 (1972)

Still the definitive review of optical pumping physics — selection rules, rate equations, coherences, and density matrix formalism. Long but structured so you can read §I–IV and have the key results. Every lab that does state preparation (which is every AMO lab) builds on the mechanisms described here.

doi:10.1103/RevModPhys.44.169 →

Reviews to keep open on your desk — Lattices, tweezers, Rydberg atoms

25. Ultracold Atomic Gases in Optical Lattices: Mimicking Condensed Matter Physics and Beyond Review

M. Lewenstein et al. · Adv. Phys. 56, 243 (2007)

The broadest review of what optical lattice quantum simulation can do, written at the field's moment of expansion. Covers Hubbard models, spin models, frustrated lattices, disorder, and gauge fields. Read §1–3 in Year 1; return to the specialized sections as your research evolves.

doi:10.1080/00018730701223200 →

26. Quantum Science with Optical Tweezer Arrays of Ultracold Atoms and Molecules Review

A. M. Kaufman & K.-K. Ni · Nat. Phys. 17, 1324 (2021)

A 2021 review by group leaders at JILA and Harvard — covers the full arc from single-atom tweezers to arrays of hundreds of atoms, including Rydberg QC, molecular arrays, and Sr clock tweezers. The best entry point for understanding where the tweezer platform stands today and where it is going. Written accessibly for Year 1 grad students.

doi:10.1038/s41567-021-01357-2 →

27. Rydberg Atom Quantum Technologies Review

C. S. Adams, J. D. Pritchard, J. P. Shaffer · J. Phys. B 53, 012002 (2020)

Concise, up-to-date review of Rydberg atom applications: quantum gates, electrometry, single-photon sources, and microwave sensing. The section on quantum information is the best starting point for understanding how blockade translates to gate fidelity. Compact and highly readable.

doi:10.1088/1361-6455/ab52ef →

28. Many-Body Physics with Individually Controlled Rydberg Atoms Review

A. Browaeys & T. Lahaye · Nat. Phys. 16, 132 (2020)

Browaeys is the pioneer of Rydberg tweezer arrays in Europe; this review traces the path from two-atom blockade gates to programmable quantum simulators with hundreds of atoms. The discussion of how to engineer Ising and XY Hamiltonians from Rydberg interactions is an essential bridge between experiment and theory.

doi:10.1038/s41567-019-0733-z →

Advanced Grad / Thesis Writing

These 17 papers represent the state of the art across the major research directions. By the time you're writing your thesis, you should be able to situate your work in relation to every paper on this list — knowing both what each one achieved and where it fell short.

The canonical review — Read before writing any thesis chapter

29. Many-Body Physics with Ultracold Gases Keystone

I. Bloch, J. Dalibard & W. Zwerger · Rev. Mod. Phys. 80, 885 (2008)

The "Bible" of ultracold quantum matter. 80 pages covering BEC, Fermi gases, optical lattices, low-dimensional physics, and disorder. Essentially every AMO thesis cites this review. If you can work through §IV (many-body effects in optical lattices) and §VIII (BCS–BEC crossover), you have the theoretical foundation to engage with almost any cold-atom result from 2008 onward.

doi:10.1103/RevModPhys.80.885 →

Neutral-atom quantum computing — From blockade to logical qubits

30. Observation of Rydberg Blockade Between Two Atoms Keystone

E. Urban et al. · Nat. Phys. 5, 110 (2009)

The first direct demonstration of Rydberg blockade suppression between two individually trapped atoms 10 μm apart — validating the Jaksch 2000 proposal after nine years. The suppressed Rydberg excitation of the second atom is the key signature. Read alongside Gaëtan 2009 (the collective $\sqrt{2}$ speed-up, seen the same year in a different lab).

doi:10.1038/nphys1178 →

31. Quantum Information with Rydberg Atoms Review

M. Saffman, T. G. Walker & K. Mølmer · Rev. Mod. Phys. 82, 2313 (2010)

The most complete review of Rydberg quantum information: two-qubit gate mechanisms, error sources, single-qubit addressing, scalable architectures, and comparison with ions and superconductors. Its error estimates for each channel are the template for every subsequent gate fidelity budget paper.

doi:10.1103/RevModPhys.82.2313 →

32. A Quantum Processor Based on Coherent Transport of Entangled Atom Arrays Experiment

D. Bluvstein et al. · Nature 604, 451 (2022)

Demonstrated mid-circuit atom shuttling to bring non-neighboring qubits together for entanglement — a key enabler for logical qubits. The reconfigurable connectivity solves the limited-connectivity problem that plagues superconducting architectures. Used to prepare cluster states and surface- and toric-code states of up to 24 qubits.

doi:10.1038/s41586-022-04592-6 →

33. High-Fidelity Parallel Entangling Gates on a Neutral-Atom Quantum Computer Experiment

S. J. Evered et al. · Nature 622, 268 (2023)

The Lukin group's demonstration of two-qubit entangling gates with 99.5% fidelity, performed in parallel on up to 60 rubidium atoms. The key innovations are fast single-pulse optimal-control gates, atomic dark states that reduce scattering, better Rydberg excitation, and improved cooling. This is the benchmark every Rydberg-gate experiment is compared against.

doi:10.1038/s41586-023-06481-y →

34. Logical Quantum Processor Based on Reconfigurable Atom Arrays Experiment

D. Bluvstein et al. · Nature 626, 58 (2024)

The first demonstration of a programmable logical processor with neutral atoms, using reconfigurable arrays with mid-circuit readout, feed-forward, and several logical encodings. The experiment includes distance-scaled surface-code operations, color-code qubits, logical GHZ states, and computational circuits with up to 48 logical qubits. Represents the transition from NISQ to early error-corrected computing in neutral atoms.

doi:10.1038/s41586-023-06927-3 →

35. A Tweezer Array with 6,100 Highly Coherent Atomic Qubits Experiment

H. J. Manetsch et al. (Endres group, Caltech) · Nature 647, 60 (2025)

A 6,100-site ^{133}Cs tweezer array with T_2^* = 14.0(1) ms free Ramsey dephasing extended to T_2 = 12.6(1) s via XY-16 dynamical decoupling — the current coherence record for a hyperfine qubit in an optical tweezer array. Read alongside Bluvstein 2024 above to see how far coherent-array scale and dynamical-decoupling technique have advanced together since.

nature.com →

36. High-Fidelity Entangling Gates and Nonlocal Circuits with Neutral Atoms Experiment Preprint

S. J. Evered et al. (Lukin group) · arXiv:2604.25987 (2026)

The Lukin group's own follow-up to Evered 2023 above: 99.854(4)% raw / 99.941(3)% with loss post-selection CZ gate fidelity, stable across a 10-hour, ~2-million-run benchmark on a 442-atom zoned array, plus nonlocal circuits via coherent atom rearrangement. Not yet peer-reviewed as of late 2026 — read it as the frontier claim to watch, not yet the settled record.

arxiv.org →

Optical lattice clocks — The most precise instruments ever built

37. Quantum State Engineering and Precision Metrology Using State-Insensitive Light Traps Keystone

J. Ye, H. J. Kimble & H. Katori · Science 320, 1734 (2008)

A short review of "magic" (state-insensitive) optical traps: the same idea that makes the strontium and ytterbium lattice clocks work also keeps single atoms coherent in cavity-QED and quantum-information experiments. The bridge between the clock and the tweezer communities, and the clearest explanation of why a magic wavelength matters.

doi:10.1126/science.1148259 →

38. Atomic Clock Performance Enabling Geodesy Below the Centimetre Level Experiment

W. F. McGrew et al. · Nature 564, 87 (2018)

A ^{171}Yb optical lattice clock with systematic uncertainty 1.4×10^{-18} — good enough to detect the gravitational redshift of a 1 cm height difference. The paper that shows optical atomic clocks have crossed from precision measurement into practical geophysical sensing.

doi:10.1038/s41586-018-0738-2 →

Quantum simulation — Large programmable systems

39. Quantum Thermalization Through Entanglement in an Isolated Many-Body System Experiment

A. M. Kaufman et al. · Science 353, 794 (2016)

Using a six-site optical lattice and single-site-resolved measurements, Kaufman et al. directly measured the growth of entanglement entropy during thermalization of a pure quantum state — confirming ETH (Eigenstate Thermalization Hypothesis) with individual-atom resolution. A tour de force of experimental quantum information science in AMO.

doi:10.1126/science.aaf6725 →

40. Quantum Phases of Matter on a 256-Atom Programmable Quantum Simulator Experiment

S. Ebadi et al. · Nature 595, 227 (2021)

The Lukin group's 256-atom tweezer array realizing the checkerboard, striated and star ordered phases of the two-dimensional Rydberg Ising model and probing the quantum phase transitions between them. The programmable geometry is set by an SLM — the paper shows a programmable simulator working at sizes beyond exact classical simulation.

doi:10.1038/s41586-021-03582-4 →

41. Probing Topological Spin Liquids on a Programmable Quantum Simulator Experiment

G. Semeghini et al. · Science 374, 1242 (2021)

Using a Rydberg tweezer array on the ruby lattice, Semeghini et al. detected signatures of a topological Z_2 spin liquid — a phase with long-range entanglement but no symmetry breaking, predicted but never cleanly observed before. Demonstrates that programmable quantum simulators can study phases inaccessible to both classical numerics and solid-state experiments.

doi:10.1126/science.abi8794 →

42. Programmable Quantum Simulations of Spin Systems with Trapped Ions Review

C. Monroe et al. · Rev. Mod. Phys. 93, 025001 (2021)

Even if you work on neutral atoms, understanding what trapped-ion simulators can and cannot do relative to your platform is essential for grant applications, conference talks, and industry interviews. Monroe's review gives the most complete comparison of the two leading programmable quantum simulation platforms.

doi:10.1103/RevModPhys.93.025001 →

Beyond the alkalis — Alkaline-earth and magnetic atoms

43. Two-Orbital SU(N) Magnetism with Ultracold Alkaline-Earth Atoms Theory

A. V. Gorshkov et al. · Nat. Phys. 6, 289 (2010)

Gorshkov et al. showed that alkaline-earth-like atoms (Yb, Sr) with nuclear spin I and two stable electronic orbitals $\{^1S_0, {}^3P_0\}$ realize an exact SU(N) symmetric Hubbard model with N = 2I+1 — inaccessible in any solid-state system. The Kugel-Khomskii-like orbital exchange is derived here. The essential theory paper for any Yb or Sr quantum simulation experiment.

doi:10.1038/nphys1535 →

44. Dipolar Physics: A Review of Experiments with Magnetic Quantum Gases Review

L. Chomaz et al. · Rep. Prog. Phys. 86, 026401 (2023)

Comprehensive review of dysprosium, erbium, and chromium quantum gases — dipole-dipole interactions, quantum droplets, supersolid phases, roton instabilities, and anisotropic collapse. If your research touches any magnetic or dipolar species, this is the reference review. Directly relevant to quantum simulation of frustrated magnets and exotic phases.

doi:10.1088/1361-6633/aca814 →

Quantum error correction — The long game

45. Surface Codes: Towards Practical Large-Scale Quantum Computation Theory

A. G. Fowler et al. · Phys. Rev. A 86, 032324 (2012)

The definitive reference for surface code QEC — threshold theorem, logical error rates as a function of physical error rate and code distance, and resource estimates. The critical threshold of p_th ≈ 1% (physical gate error) is cited by every experimental group as the target. Essential for understanding fault-tolerance roadmaps.

doi:10.1103/PhysRevA.86.032324 →

🚀

Postdoc / Career Transition

These 10 papers are less about technique and more about context: where does AMO sit in the broader landscape of quantum science and technology? Every paper here either defines a major research direction or honestly assesses the state and roadmap of the field. Essential for fellowship applications, faculty job talks, and industry interviews.

Quantum advantage — What can quantum systems actually do better?

46. Quantum Computing in the NISQ Era and Beyond Keystone

J. Preskill · Quantum 2, 79 (2018)

Preskill coined "NISQ" (Noisy Intermediate-Scale Quantum) and gave the clearest honest assessment of what 50–1000 qubit devices can and cannot do. Essential reading before any industry conversation or grant about quantum advantage. The distinction between quantum simulation (near-term) and quantum computing (long-term) is made with unusual clarity.

doi:10.22331/q-2018-08-06-79 →

47. Practical Quantum Advantage in Quantum Simulation Review

A. J. Daley et al. · Nature 607, 667 (2022)

The most rigorous assessment of when and where quantum simulators provide a genuine computational advantage over classical methods. The paper identifies specific physical problems — Hubbard model away from half-filling, frustrated magnetism, non-equilibrium dynamics — where quantum simulation is likely to outperform any classical algorithm. Required reading before writing any "quantum advantage" claim.

doi:10.1038/s41586-022-04940-6 →

48. Quantum Simulators: Architectures and Opportunities Review

E. Altman et al. · PRX Quantum 2, 017003 (2021)

A community-wide assessment of quantum simulation roadmaps across neutral atoms, ions, superconductors, and photonics — written by a committee including both theorists and experimentalists. Honest about what each platform can and cannot achieve in 5–10 years. Indispensable for writing a competitive postdoc fellowship proposal or DOE Early Career application.

doi:10.1103/PRXQuantum.2.017003 →

Quantum sensing — The near-term application with guaranteed demand

49. Quantum Sensing Review

C. L. Degen, F. Reinhard & P. Cappellaro · Rev. Mod. Phys. 89, 035002 (2017)

The comprehensive review of quantum sensing theory: SQL, Heisenberg limit, Fisher information, spin squeezing, and the connections between sensing and computation. Covers NV centers, atom interferometers, and optical atomic clocks under a single framework. Essential if your career touches any precision measurement application, including commercializable atomic clocks and gravimeters.

doi:10.1103/RevModPhys.89.035002 →

Platform landscape — Know your neighbours in the quantum ecosystem

50. Quantum Computing with Neutral Atoms Review

L. Henriet et al. · Quantum 4, 327 (2020)

The Pasqal team's review of neutral-atom quantum computing — written from the perspective of a company actively building hardware. Covers qubit encoding, gate schemes, connectivity, and error rates with an eye toward scalability. Reading this alongside Saffman 2010 shows how much the field advanced in a decade.

doi:10.22331/q-2020-09-21-327 →

51. Superconducting Qubits: Current State of Play Review

M. Kjaergaard et al. · Annu. Rev. Condens. Matter Phys. 11, 369 (2020)

The MIT engineering-quantum-systems group's survey of the platform most quantum roadmaps are measured against: transmon design, single- and two-qubit gates, readout, coherence limits and what it takes to scale. Read it so you can put neutral atoms next to their main competitor in a fellowship proposal, and answer the interview question "why atoms and not transmons?" with specifics.

doi:10.1146/annurev-conmatphys-031119-050605 →

52. Cold and Ultracold Molecules: Science, Technology and Applications Review

L. D. Carr et al. · New J. Phys. 11, 055049 (2009)

Molecules in cold traps offer dipolar interactions, rich internal structure (rotational qubits), and direct chemistry control — all unavailable in monatomic systems. This review covers the production routes (STIRAP from Feshbach molecules, buffer-gas loading, direct laser cooling) and the proposed quantum simulation and computing applications. The molecular platform is a major growth area.

doi:10.1088/1367-2630/11/5/055049 →

Frontier results — What today's machines can already measure

53. Evidence for the Utility of Quantum Computing Before Fault Tolerance Experiment

Y. Kim et al. · Nature 618, 500 (2023)

IBM's 127-qubit Eagle processor simulating a kicked Ising model in regimes inaccessible to classical tensor-network methods — claiming quantum utility before fault tolerance. Highly debated in the community (classical simulation subsequently partially recovered the result). Read the paper and the responses: it models how claims of quantum advantage will be scrutinized in the NISQ era.

doi:10.1038/s41586-023-06096-3 →

54. Experimental Characterization of a Quantum Many-Body System via Higher-Order Correlations Experiment

T. Schweigler et al. · Nature 545, 323 (2017)

Included here not as a landmark experiment, but as an example of the creative breadth of quantum simulation: two tunnel-coupled one-dimensional Bose gases on an atom chip, whose relative phase obeys the sine-Gordon field theory, characterized through its higher-order correlation functions. Postdocs who understand how diverse the "quantum simulation" umbrella is will identify opportunities others overlook.

doi:10.1038/nature22310 →

55. Hartree-Fock on a Superconducting Qubit Quantum Computer Experiment

Google AI Quantum & collaborators · Science 369, 1084 (2020)

Google AI Quantum and collaborators simulating the Hartree-Fock ground-state energy of hydrogen chains and diazene on a superconducting quantum processor with error mitigation. Shows how quantum chemistry is an early application with verifiable classical benchmarks — and demonstrates the VQE hybrid classical-quantum workflow that AMO physicists will be expected to understand in any quantum industry role.

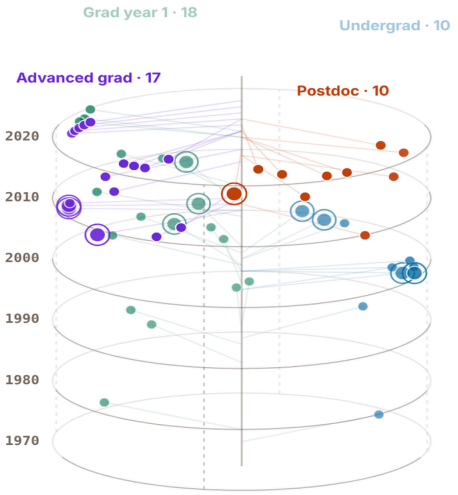
doi:10.1126/science.abb9811 →

29.4 All 55 papers in one view

Height is the year each paper was published. The four faces of the tower are the four career stages in the tabs above, each read left to right in page order. From the keyboard: arrow keys turn it, `]` and `[` step through the papers by year, Enter opens the note.

Undergrad Grad year 1 Advanced grad Postdoc ringed = keystone

Papers per decade: 1970s 2 · 1980s 5 · 1990s 9 · 2000s 11 · 2010s 10 · 2020s 18



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

CHAPTER 29 SUMMARY

Reading primary papers alongside hands-on calculations builds true research fluency. Next, we map the academic research groups around the world advancing these subfields.

Problems for Chapter 29

1. DESIGN Pick one paper from the stage of the roadmap that matches this week. Write a one-page summary: the question, the method, the key figure, and which AMO Toolkit page lets you reproduce a number from it.
2. DESIGN Using one of the platform comparison pages, write a 500-word memo recommending a hardware platform for a specific application (for example, a 100-qubit simulator or a quantum network node). Cite at least three sources.

Use the tool: on the interactive page. By hand: from the formulas in this chapter. Design: a decision a lab would make. Worked solutions are available to instructors on request (see Part V).

CHAPTER 30 · CAREER STEP 02

30

Global AMO Research Directory

🌍 AMO Group Finder · Career map

Where to do a PhD or a postdoc in atomic, molecular and optical physics: a mapped directory of groups you can filter by research area, atom, region and technique, a globe that follows every filter, and a list you can keep.

> *Whether choosing a PhD group, postdoc, or collaborator, knowing which labs pioneer which atomic species and experimental techniques saves months of fragmented searching.*

IN THIS CHAPTER — Prerequisites: Ch. 29

- Survey leading experimental and theoretical AMO groups worldwide by subfield (Rydberg arrays, clocks, BEC/simulation, molecules, ions, cavities).
- Match personal technical interests and target atomic species to active research ecosystems.
- Identify adjacent groups and literature lineages when preparing graduate or postdoctoral applications.

WEBSITE COMPANION Live interactive page ↗

30.1 Start from a research area

Eight areas cover the directory. Pick one to see its groups on the map and in the list below; pick it again to clear it.

30.2 Find groups

USA 40

MANUEL ENDRES
Caltech · USA
Sr and Cs tweezer arrays, Rydberg entanglement, quantum error correction, neutral-atom QC · Rydberg / QC Sr, Cs

Nick Hutzler

Caltech · USA

Laser-cooled polyatomic molecules for precision tests of fundamental symmetries and new-physics searches

Molecules YbOH

KENNETH BROWN
Duke University · USA
Trapped-ion quantum computing, quantum error correction, ion-trap architecture design · Ions Ca+, Yb+

Huanqian Loh

Duke University · USA

Tweezer-trapped atom arrays as building blocks for quantum computing, simulation, and sensing

Rydberg / QC Rb, Cs

Christopher Monroe

Duke University (formerly Univ. Maryland / JQI) · USA

Trapped Yb^{+} ion QC, programmable quantum simulation, quantum networks

Ions Yb+

MARKUS GREINER

Harvard University · USA

Quantum gas microscope, Fermi-Hubbard model, quantum thermodynamics · Quantum Simulation Rb

MIKHAIL LUKIN

Harvard University · USA

Rydberg tweezer arrays, quantum networks, solid-state spin defects · Rydberg / QC Rb, Cs

KANG-KUEN NI

Harvard University · USA

Ultracold polar molecules (KRb, NaCs), molecular tweezers, reaction control · Molecules Na, Cs, K, Rb

Calculator output at the page's default settings.

Adam Kaufman

JILA / Univ. Colorado · USA

Alkaline-earth tweezer arrays, optical-clock qubits, Rydberg gates, quantum simulation and metrology

Rydberg / QC Sr, Yb

Cindy Regal

JILA / Univ. Colorado · USA

Optical tweezer arrays, hybrid quantum systems, single-phonon control in mechanical oscillators

Cavity QED Rb

Ana Maria Rey

JILA / Univ. Colorado · USA

Theory: SU(N) Hubbard model, molecular collisions, metrological squeezing, dipolar physics

Quantum Simulation Theory

James K. Thompson

JILA / Univ. Colorado · USA

Cavity QED, spin squeezing beyond the standard quantum limit, superradiant lasers

Cavity QED Rb

JUN YE

JILA / Univ. Colorado · USA

Sr optical lattice clock, polar molecules, precision spectroscopy, quantum chemistry · Clocks & Metrology Sr, OH

David DeMille

Johns Hopkins University (formerly UChicago) · USA

Laser-cooled polar molecules, electron EDM searches, precision tests of fundamental symmetries

Molecules SrF, TlF

WOLFGANG KETTERLE
MIT · USA
BEC, optical lattices, low-dimensional quantum gases, spinor condensates · BEC / Fermions Na, Li

VLADAN VULETIĆ
MIT · USA
Cavity QED, spin squeezing, optical atomic clocks, quantum metrology · Cavity QED Rb, Yb, Cs

MARTIN ZWIERLEIN
MIT · USA
Fermionic superfluidity, BEC-BCS crossover, quantum gas microscope for fermions · BEC / Fermions Li, Na

Gretchen Campbell

NIST / JQI, Univ. Maryland · USA

Atomtronics, superfluid circuits, quantum sensing with ring-trapped Bose-Einstein condensates

BEC / Fermions Na, Sr

JAMES PORTO
NIST / Univ. Maryland · USA
Optical lattices, Hubbard physics, quantum thermometry, lattice gauge theories · Quantum Simulation Rb

Ian Spielman

NIST / Univ. Maryland · USA

Artificial gauge fields, synthetic spin-orbit coupling, topological band structures

Quantum Simulation Rb

Brian Odom

Northwestern University · USA

Sympathetically-cooled trapped molecular ions, precision spectroscopy, tests of fundamental constants

Ions ions

David Weiss

Penn State University · USA

Optical tweezer arrays for neutral-atom quantum computing, 1D Bose gases, electron EDM tests

Rydberg / QC Cs

BRYCE GADWAY
Penn State University (formerly UIUC) · USA
Synthetic dimensions, quasicrystals, momentum-space lattices · Quantum Simulation Rb

WASEEM BAKR
Princeton University · USA
Quantum gas microscope for fermions, Fermi-Hubbard physics, disorder · Quantum Simulation Li

Lawrence Cheuk

Princeton University · USA

Laser-cooled molecules trapped in optical tweezer arrays for quantum simulation and information

Molecules CaF

Jeff D. Thompson

Princeton University · USA

Yb tweezer arrays, alkaline-earth Rydberg atoms, nanophotonics, erasure conversion for neutral-atom QC

Rydberg / QC Yb

JONATHAN D. HOOD

Purdue University · USA

Li-Cs molecule assembly, optical tweezers, single-atom cooling and imaging · Molecules Li, Cs

RANDALL HULET

Rice University · USA

BEC-BCS crossover, Fermi-Hubbard model, quantum magnetism with fermions · BEC / Fermions Li

TOM KILLIAN

Rice University · USA

Ultracold plasmas, Rydberg atoms in plasmas, strontium BEC · Clocks & Metrology Sr

Mark Kasevich

Stanford University · USA

Atom interferometry, inertial sensing, gravitational wave detection, Sr lattice clock

Clocks & Metrology Rb, Sr

BENJAMIN LEV

Stanford University · USA

Dipolar quantum gases, quantum droplets, Er/Dy mixtures, exotic superfluid phases · Dipolar Dy, Er

Monika Schleier-Smith

Stanford University · USA

Cavity-mediated interactions among laser-cooled atoms for quantum simulation and metrology

Cavity QED Rb

Jonathan Simon

Stanford University (formerly Univ. Chicago) · USA

Photonic and cold-atom quantum matter, synthetic photonic materials, topological photon fluids

Cavity QED Rb

Holger Müller

UC Berkeley · USA

Atom interferometry for tests of gravity and the equivalence principle, precision measurement of constants

Clocks & Metrology Cs

DAN STAMPER-KURN

UC Berkeley · USA

Spinor BEC, cavity optomechanics, quantum gas microscope, orbital magnetism · Cavity QED Rb

David Weld

UC Santa Barbara · USA

Floquet engineering, quantum quasicrystals, non-equilibrium dynamics in quantum degenerate gases

Quantum Simulation Li, Sr

Jacob Covey

Univ. Illinois Urbana-Champaign · USA

Alkaline-earth atom tweezer arrays for distributed quantum computing and atomic-array clocks

Rydberg / QC Yb

Brian DeMarco

Univ. Illinois Urbana-Champaign · USA

Fermi-Hubbard model, disorder, many-body localization, quantum simulation

Quantum Simulation Rb

Subhadeep Gupta

Univ. Washington · USA

Li–Yb quantum mixtures, Yb BEC atom interferometry, Feshbach resonances, heteronuclear molecules

BEC / Fermions Li, Yb

CHENG CHIN

University of Chicago · USA

Feshbach resonances, Efimov physics, superfluid-insulator transition, optical lattice · BEC / Fermions Cs

Canada 4

Lindsay LeBlanc

Univ. Alberta · Canada

Dual-species Rb-K quantum gases, synthetic spin-orbit coupling, non-equilibrium many-body dynamics

Quantum Simulation Rb, K

Kirk Madison

Univ. British Columbia · Canada

Ultracold atom/molecule collisions, precision atomic sensors, quantum pressure metrology

Quantum Simulation Rb, Li

Aephraim Steinberg

Univ. Toronto · Canada

Foundations of quantum mechanics with ultracold atoms and quantum optics, weak measurement, tunneling time

Quantum Simulation Rb

JOSEPH THYWISSEN

Univ. Toronto · Canada

Fermi-Hubbard model, spin dynamics, quantum gas microscope, few-body physics · BEC / Fermions K, Rb

Austria 6

Hannes Bernien

Univ. Innsbruck · Austria

Large-scale optical tweezer arrays, dual-species Rydberg entanglement, quantum simulation atom-by-atom

Rydberg / QC Rb, Cs

FRANCESCA FERLAINO

Univ. Innsbruck · Austria

Dipolar Er/Dy mixtures, supersolid phase, quantum droplets, Feshbach resonances · Dipolar Er, Dy

RUDOLF GRIMM

Univ. Innsbruck · Austria

Feshbach physics, Efimov trimers, ultra-cold mixtures, dipolar Cs BEC · BEC / Fermions Cs, Li

Hanns-Christoph Nägerl

Univ. Innsbruck · Austria

Ultracold molecule formation, low-dimensional quantum gases, precision molecular spectroscopy

Molecules Cs

Rainer Blatt (emeritus) & Christian Roos

Univ. Innsbruck / IQOQI · Austria

Trapped-ion QC, quantum simulation with ions, quantum error correction, metrology

Ions Ca+

Peter Zoller

Univ. Innsbruck / IQOQI · Austria

Theory: AMO quantum information, quantum simulation architectures, open quantum systems

Quantum Simulation Theory

Denmark 2

Jacob Sherson

Aarhus University · Denmark

Quantum control of ultracold atoms, citizen-science-guided optimization (ScienceAtHome), hybrid intelligence

Quantum Simulation Rb

Eugene Polzik

Niels Bohr Institute, Copenhagen · Denmark

Quantum optics with atomic ensembles, quantum memory, teleportation between light and matter

Cavity QED Cs

France 5

Jean Dalibard & Jérôme Beugnon

ENS Paris · France

2D Bose gas, Berezinskii-KT transition, topological defects, quantum turbulence

Quantum Simulation Rb

JAKOB REICHEL

ENS Paris · France

Atom chips, fiber cavity QED, compact atomic clocks, multiparticle entanglement · Cavity QED Rb

Christophe Salomon & Tarik Yefsah

ENS Paris · France

Ultracold Fermi gases, precision atomic clocks, tests of fundamental physics in space (ACES/PHARAO)

BEC / Fermions Li

Alain Aspect & Denis Boiron

Institut d'Optique, Palaiseau · France

He* BEC, atom-atom correlations, Hong-Ou-Mandel for atoms, atom laser, quantum optics

BEC / Fermions He*, Rb

Antoine Browaeys & Thierry Lahaye

Institut d'Optique, Palaiseau · France

Rydberg tweezer arrays, quantum simulation, dipolar atoms, Ising model engineering

Rydberg / QC Rb, Dy

Germany 10

Immanuel Bloch

MPQ / LMU Munich · Germany

Quantum gas microscope, Hubbard model, quantum simulation, optical lattice clocks

Quantum Simulation Rb, Sr

Ignacio Cirac

MPQ Munich · Germany

Theory: quantum information science, tensor networks, AMO platforms for quantum computing

Quantum Simulation Theory

Gerhard Rempe

MPQ Munich · Germany

Cavity QED, atom-photon quantum interface, deterministic single-photon sources, quantum repeaters

Cavity QED Cs, Rb

Andreas Hemmerich

Univ. Hamburg · Germany

Driven optical lattices, non-equilibrium quantum gases, cavity-mediated self-organization

Quantum Simulation Rb

Klaus Sengstock

Univ. Hamburg · Germany

Artificial gauge fields, spinor BEC, non-equilibrium dynamics, frustrated lattice models

Quantum Simulation K, Rb

Selim Jochim

Univ. Heidelberg · Germany

Few-body Fermi systems, deterministic single-atom preparation, quantum simulation with mesoscopic ensembles

BEC / Fermions Li

Markus Oberthaler

Univ. Heidelberg · Germany

Bose-Einstein condensates, matter-wave interferometry, quantum-enhanced metrology, entanglement generation

BEC / Fermions Rb

Matthias Weidemüller

Univ. Heidelberg · Germany

Rydberg gases and spin models, Li–Cs mixtures and Efimov physics, cold chemistry

Rydberg / QC Rb, Li, Cs

TILMAN PFAU

Univ. Stuttgart · Germany

Dipolar quantum gases, quantum droplets, supersolid phases, Rydberg-dressed interactions · Dipolar Dy, Cr, Rb

Christian Gross

Univ. Tübingen (formerly MPQ Garching) · Germany

Quantum gas microscope, quantum simulation, strongly correlated systems

Quantum Simulation Rb

Italy 2

Guglielmo Tino

Univ. Florence / LENS · Italy

Atom interferometry, precision tests of gravity, strontium optical clocks for geodesy

Clocks & Metrology Sr

Philipp Hauke

Univ. Trento · Italy

Theory: quantum simulation of lattice gauge theories, tensor networks, AMO-HEP connection

Quantum Simulation Theory

Netherlands 2

Jook Walraven (emeritus)

Univ. Amsterdam · Netherlands

BEC of hydrogen, quantum reflection, Casimir-Polder interaction, precision hydrogen spectroscopy

BEC / Fermions H, Rb

Florian Schreck

Univ. Amsterdam · Netherlands

Strontium quantum gases, optical lattice clocks, compact ultracold-atom technology for space and metrology

Clocks & Metrology Sr

Spain 1

Leticia Tarruell

ICFO Barcelona · Spain

Ultracold quantum-gas microscopy of strontium (QUIONE), engineered quantum matter, atom-light interfaces

Quantum Simulation K, Sr

Switzerland 2

Tilman Esslinger

ETH Zurich · Switzerland

Fermi-Hubbard model, quantum simulation, optical lattice, non-equilibrium dynamics

Quantum Simulation K, Li

Philipp Treutlein

Univ. Basel · Switzerland

Atom chips, hybrid atom-mechanical systems, entangled atomic ensembles, quantum sensing

Cavity QED Rb

UK 7

CHARLES ADAMS

Durham University · UK

Rydberg atoms, quantum nonlinear optics, microwave-to-optical transduction, EIT · Rydberg / QC Rb

SIMON CORNISH

Durham University · UK

Ultracold Cs molecules (Cs_2, RbCs), tweezer assembly of bialkali molecules, STIRAP · Molecules Cs, Rb

Ed Hinds

Imperial College London · UK

Emeritus professor

Laser-cooled molecules (CaF, YbF), eEDM measurements, atom chips, precision tests of SM

Molecules Rb, CaF

Zoran Hadzibabic

Univ. Cambridge · UK

Homogeneous Bose gases in a box trap, 2D superfluidity, far-from-equilibrium quantum dynamics

BEC / Fermions Rb

Ulrich Schneider

Univ. Cambridge · UK

Quantum gas microscopy, Hubbard physics, engineered disorder and many-body localization

Quantum Simulation K

STEFAN KUHR

Univ. Strathclyde · UK

Quantum gas microscopy, single-atom-resolved imaging of optical lattices · Quantum Simulation Rb

Christopher Foot

University of Oxford · UK

Optical lattices, atom interferometry, superfluid vortices, ultracold mixture physics

Quantum Simulation Rb

China 2

SHI-LIANG ZHU

SCNU Guangzhou · China

Quantum simulation of topological matter, synthetic gauge fields, quantum Hall physics · Quantum Simulation Rb

JIAN-WEI PAN

USTC Hefei · China

Photonic QC, Rydberg tweezer arrays, quantum communication, boson sampling · Rydberg / QC Rb, Sr

India 1

SADIQ A. RANGWALA

RRI Bangalore · India

Hybrid ion-atom traps, cold reactive collisions, atom-ion interactions, laser-cooled Ca^+ · Ions Rb, Ca+

Israel 5

YOAV SAGI

Technion · Israel

Fermi gas spectroscopy, BEC-BCS crossover, Tan contact, strongly interacting fermions · BEC / Fermions K

JEFF STEINHAUER

Technion · Israel

Analog gravity in BEC, Hawking radiation from sonic black holes, quantum correlations · BEC / Fermions Rb

Nir Davidson

Weizmann Institute · Israel

Cold-atom nonlinear dynamics, optical trapping architectures, atom-based quantum simulators

BEC / Fermions Rb

Ofer Firstenberg

Weizmann Institute · Israel

Rydberg EIT, photon-photon interactions via Rydberg blockade, quantum nonlinear optics

Rydberg / QC Rb

Roee Ozeri

Weizmann Institute · Israel

Trapped-ion quantum information and sensing, ion-atom hybrid systems, decoherence-free quantum control

Ions Sr+

Japan 3

Kenji Ohmori

IMS Okazaki · Japan

Rydberg many-body dynamics, ultrafast quantum control, optical-clock-referenced Rydberg spectroscopy

Rydberg / QC Rb

Yoshiro Takahashi

Kyoto University · Japan

Yb quantum simulation, SU(N) Hubbard model, orbital magnetism, two-orbital physics

Quantum Simulation Yb

Hidetoshi Katori

Univ. Tokyo · Japan

Optical lattice clocks (inventor of magic wavelength), Sr and Hg clocks, relativistic geodesy

Clocks & Metrology Sr, Yb

Singapore 2

Wenhui Li

CQT / Nat. Univ. Singapore · Singapore

Ultracold Rydberg gases, Rydberg-atom microwave sensing, dipole-dipole interaction physics

Rydberg / QC Rb

SHAU-YU LAN

Nanyang Tech. Univ. · Singapore

Atom interferometry, compact atomic clocks, quantum sensing, gravimetry · Clocks & Metrology Rb

South Korea 1

Jaewook Ahn

KAIST · South Korea

Rydberg-atom graphs for combinatorial optimization, ultrafast quantum control, atom-based quantum computing

Rydberg / QC Rb

Australia 2

Andrew Truscott

Australian National University · Australia

Metastable helium Bose-Einstein condensates, atom-optics quantum correlation experiments

BEC / Fermions He*

Tyler Neely & Halina Rubinsztein-Dunlop

Univ. Queensland · Australia

BEC, quantum vortex dynamics, 2D quantum turbulence, atom optics, optical tweezers

Quantum Simulation Rb

Brazil 1

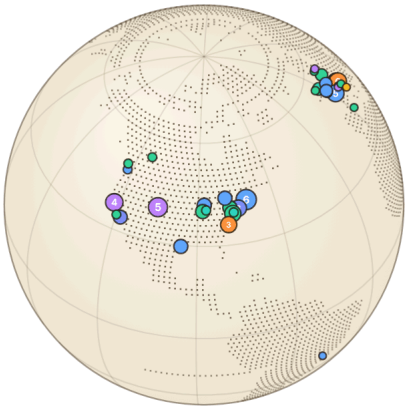
Vanderlei Bagnato

USP São Carlos · Brazil

Bose-Einstein condensation, quantum turbulence in atomic superfluids, ultracold-atom precision sensors

BEC / Fermions Rb

Interactive



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

98 groups on the globe · colours are the research areas · click a marker for details (keyboard: `]` / `[` step through markers, Enter opens)

My list empty

Star a group (☆) to keep it here while you compare. The list stays in this browser.

Technique tags are read from each group's focus line, so they are a guide, not a complete description. The representative paper is one well-known result, not necessarily the group's latest.

30.3 Names to know, by subfield

Curated starting points for four of the busiest subfields. Names with a coloured dot are in the directory and open their profile; the others are on the watchlist below.

30.4 The wider watchlist

More people worth knowing across cold atoms, tweezers, quantum gases, molecules, clocks, interferometry, ions and AMO theory. These entries are names and topics only, so follow up through group pages and recent papers. The search box above filters this list too.

UNITED STATES 15

John Doyle

Harvard University

Cold molecules; precision AMO

Eugene Demler

ETH Zurich (formerly Harvard)

Many-body AMO theory

Norman Yao

Harvard University

AMO many-body theory; Rydberg systems; quantum dynamics

Richard Fletcher

MIT

Quantum gases; many-body dynamics

Soonwon Choi

MIT

Quantum many-body theory; neutral atoms; dynamics

Ronald Fernando Garcia Ruiz

MIT

Precision AMO; exotic atoms and nuclei

David Schuster

Stanford University (formerly University of Chicago)

Hybrid quantum optics and cavity/circuit connections

Deborah Jin legacy / JILA molecule ecosystem

JILA / University of Colorado Boulder

Ultracold molecules

Heather Lewandowski

University of Colorado Boulder

AMO education; molecular beams; precision

Alexey Gorshkov

NIST / University of Maryland

AMO quantum information theory

Kaden Hazzard

Rice University

AMO theory; molecules; Rydberg systems; quantum gases

Paul Kwiat

University of Illinois Urbana-Champaign

Quantum optics; AMO-adjacent quantum information

Mukund Vengalattore

Cornell University

Quantum gases; spinor condensates; nonequilibrium AMO

Erich Mueller

Cornell University

Quantum gas theory; strongly correlated AMO systems

Steven Rolston

University of Maryland

Cold atoms; Rydberg physics; ultracold plasmas

CANADA 1

David Feder / theory ecosystem

University of Calgary

AMO and quantum-gas theory

EUROPE 10

Artur Widera

Kaiserslautern / Mainz, Germany

Ultracold atoms; quantum gases

Sylvain Nascimbene

Sorbonne / LKB, France

Quantum gases; microscopy

Michael Tarbutt / Hinds cold molecule ecosystem

Imperial College London, UK

Cold molecules; precision AMO

Ifan Hughes / Durham AMO ecosystem

Durham University, UK

Cold atoms; Rydberg spectroscopy

Ben Sauer / Ed Hinds lineage

Imperial College London, UK

Cold molecules; precision

Ferdinand Schmidt-Kaler / Tobias Schaetz ecosystem

Germany

Ions and AMO quantum simulation

Dieter Jaksch (University of Hamburg / Oxford) & Florian Mintert (Imperial College London)

AMO quantum theory

Andrew Daley

University of Oxford (formerly Strathclyde)

AMO many-body theory

Christophe Galland / Swiss quantum optics ecosystem

EPFL, Switzerland

Quantum optics; AMO-adjacent

Maciej Lewenstein

ICFO, Spain

AMO many-body theory

ASIA, OCEANIA, ISRAEL AND ELSEWHERE 25

Ben van Linden van den Heuvell / W. Vassen lineage

Amsterdam, Netherlands

Metastable helium; cold atoms

Servaas Kokkelmans

Eindhoven, Netherlands

Ultracold-collision theory

Giovanni Modugno

Florence, Italy

Quantum gases; mixtures

Massimo Inguscio / LENS ecosystem

Florence, Italy

Cold atoms; precision; quantum gases

Leonardo Fallani

Florence, Italy

Optical lattices; quantum simulation

Giacomo Roati

CNR/LENS, Italy

Quantum gases; disorder

Michal Tomza

University of Warsaw, Poland

Ultracold molecules theory

Jakub Zakrzewski

Jagiellonian University, Poland

Quantum chaos; cold-atom theory

Anders Kastberg / Gothenburg-Lund ecosystem

Sweden

Cold atoms; AMO

Barak Dayan

Weizmann Institute, Israel

Quantum optics; atom-photon interfaces

CQT / NUS cold atom ecosystem

National University of Singapore

Quantum gases; AMO

Jochen Geiger / Christian Miniatura ecosystem

CQT, Singapore

AMO theory and experiment

Luming Duan

Tsinghua University, China

Quantum information; neutral atoms/ions theory

Bo Zhao

USTC, China

Ultracold molecules and atoms

USTC cold atom ecosystem

USTC, China

Quantum gases; molecules; optical lattices

Shin Inouye

University of Tokyo, Japan

Quantum gases

Y. Shin

Seoul National University, Korea

Quantum gases

Nick Robins / ANU atom-laser ecosystem

Australian National University

Atom optics; interferometry

Chris Vale

Swinburne University, Australia

Strongly interacting Fermi gases

Matthew Davis / Karen Kheruntsyan theory ecosystem

University of Queensland, Australia

Quantum gas theory

G. S. Agarwal / AMO theory ecosystem

India/US-linked

Quantum optics

Saikat Ghosh

IIT Kanpur, India

Cold atoms; Rydberg/nonlinear optics

IISER/IIT cold atom ecosystem

India

Quantum gases; AMO

Niels Kjaergaard

University of Otago, New Zealand

Ultracold atoms

Ashton Bradley

University of Otago, New Zealand

Quantum gas theory

90 of the 141 watchlist entries are already in the directory and are listed there instead.

CHAPTER 30 SUMMARY

Academic labs drive foundational physics and train the field's talent; increasingly, that same AMO skill set powers commercial quantum computing companies, which we map next.

CHAPTER 31 · CAREER STEP 03

31

Quantum Computing Industry Landscape

💻 Quantum Industry Map · Landscape

A researcher's guide to the hardware race — platforms, companies, metrics, and where neutral atoms fit in. From NISQ to fault tolerance: superconducting, trapped-ion, neutral-atom, silicon-spin, photonic, and topological qubits. Last reviewed **November 2026**.

> *Translating AMO lab skills into industry requires separating peer-reviewed hardware milestones from roadmap projections across superconducting, trapped-ion, neutral-atom, and photonic platforms.*

IN THIS CHAPTER Prerequisites: Ch. 21, Ch. 26, Ch. 28

- Compare physical qubit counts, two-qubit gate fidelities, connectivity, and scaling bottlenecks across all major hardware modalities.
- Distinguish demonstrated experimental metrics from forward-looking commercial roadmaps.
- Map hands-on AMO skills (lasers, vacuum, RF, FPGA control, benchmarking) directly to industry engineering roles.

WEBSITE COMPANION Live interactive page ↗

6. Qubit platforms

compared on this page

27. Companies & labs

in the deep dives and on the globe

6100

Atoms trapped coherently

Manetsch et al., *Nature* 2025

99.92%

Best peer-reviewed 2Q fidelity

Quantinuum Helios, *Nature* 2026

96. Logical qubits, neutral atoms

Harvard/QuEra, *Nature* 2026

How this page is laid out

31.1 Use this as a map, not a press-release feed.

The durable material comes first and the news last, because the news dates fastest. Every number is tagged as peer-reviewed, company claim or preprint where it matters.

01–02
What a qubit platform needs
The DiVincenzo criteria, and analog versus digital machines.

03–06
Six platforms compared
One table, two charts, a map of who builds what, then each platform in depth.

07–09
The road to fault tolerance
Where the field stands, the milestone timeline, and the neutral-atom engineering list.

10–12
News, money and your career
2024–26 flash points, investment, and where an AMO physicist fits.

Calculator output at the page's default settings.

31.2 The DiVincenzo Criteria

Five necessary (and simultaneously required) conditions for a viable qubit platform.

> **DiVincenzo (2000)**, A physical system is a viable qubit platform only if it satisfies all five criteria *simultaneously*, *at scale*, *and with fault-tolerance margins*. No platform has fully satisfied all five yet — this is what makes quantum computing hard.

1️⃣

Scalable Register

Well-defined |0⟩,|1⟩ two-level subspace, replicable in large numbers without performance degradation.

2️⃣

Initialisation

Reliable preparation of the fiducial |000…0⟩ state before each computation: optical pumping (and motional cooling) for atoms and ions, thermalisation and active reset for superconducting circuits.

3️⃣

Long Coherence

$T_2 \gg$ t_gate. Decoherence must be slow vs. gate time. For neutral atoms: $T_2 > 1$ s is achievable; tweezer photon scattering erodes it.

4️⃣

Universal Gates

High-fidelity single-qubit rotations + at least one entangling two-qubit gate (e.g. Rydberg blockade for neutral atoms).

5️⃣

Qubit Readout

Site-specific measurement in the computational basis with near-unity fidelity, without disturbing neighbours.

31.3 Analog vs. Digital Quantum Computing

Two fundamentally different approaches, and why neutral atoms can do both.

Digital (Gate-Based)

Programs decomposed into discrete unitary gates from a universal gate set, applied to qubits, then measured. Mathematically clean and compatible with quantum error correction.

Cost: Error correction needs 100s–1000s of physical qubits per logical qubit. Every gate must be comfortably below the architecture target; practical roadmaps often aim for ~10^{-3} errors even when idealized thresholds are closer to ~10^{-2}.

Who: IBM, Google, IonQ, Quantinuum, QuEra (Rydberg gates)

Analog (Simulation / Annealing)

Engineer a Hamiltonian whose ground state encodes the answer, then adiabatically evolve from an easy initial state. Special-purpose, not universally programmable, but far larger system sizes are accessible in the NISQ era.

Examples: D-Wave's 5000+ superconducting flux qubit annealers; neutral-atom arrays studying quantum magnetism on 2D lattices of hundreds of atoms.

Neutral atoms are unusual: the same array can run Rydberg gates (digital) OR evolve a programmable spin Hamiltonian (analog), just by changing the control protocol. *Pasqal & QuEra exploit this dual capability as a near-term strategy.*

31.4 Six Platforms Side by Side (2025–26)

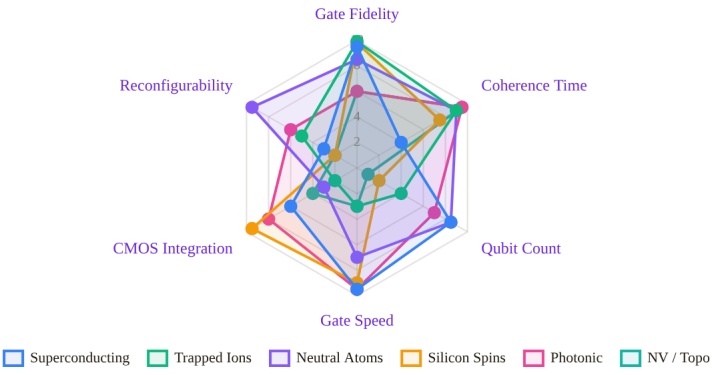
Two-qubit operation values throughout. T_2 = dephasing time; Ops/T_2 = two-qubit gates per coherence window. Highlighted row = neutral atoms (most relevant to AMO physicists).

PLATFORM	ACTIVE QUBITS (2025)	2Q GATE TIME	2Q FIDELITY (BEST)	T_2	OPS / T_2	KEY BOTTLENECK (2025–26)
🔵 Superconducting	105–156 (gate-quality)	50–200 ns	99.5% (Heron/Willow)	100 µs – 0.5 ms	~10^3	Classical wiring at mK; cryo interconnects
🟢 Trapped Ions	56–256 gate-quality (256 = IonQ Superion 256, 2026); #AQ 64 benchmark reported by IonQ Tempo	20–500 µs	99.92% (Quantinuum Helios, 2026)	1 s – minutes	>10^4	Gate speed; modular QCCD scaling

🟣 Neutral Atoms	Hundreds in gate-model processors; 6,100 trapped coherently (2025)	0.2–5 μs	99.5% (Evered 2023); 99.85% raw (2026 preprint)	1–10 s	10^3–10^4	2Q fidelity → 99.9%; mid-circuit measurement scale
🟡 Silicon Spins	6–16	20–100 ns	99.8% (^{28}Si, Mills 2022)	>1 ms (^{28}Si echo)	~10^4–10^5	Device variability; automated tuning
🩷 Photonic	>216 modes (GBS)	Probabilistic	93–97%	N/A	Limited	No deterministic photon–photon gate; source efficiency
💙 NV / Topological	1–few	μs–ms	~99% (NV); undemonstrated (topo)	ms (room temp.) to ~1 s (cryogenic, decoupled) for NV	Varies	Scale (NV); braiding undemonstrated (topo)

31.5 Platform Comparison

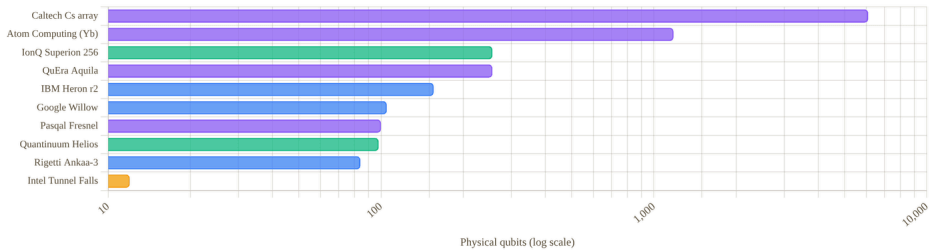
Six dimensions, normalised 0–10. No platform wins on all axes — that's the point.



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

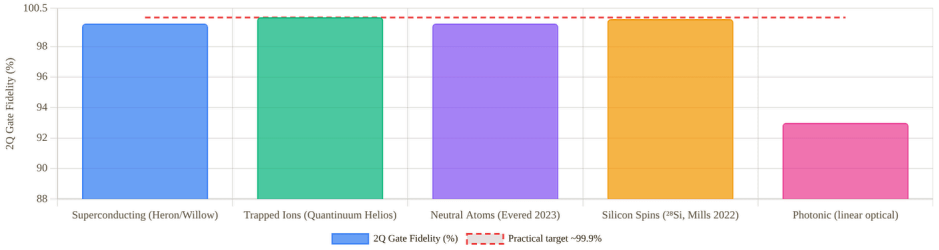
Scores are editorial judgements based on published 2024–26 benchmarks and company roadmaps; roadmap claims are not equivalent to demonstrated hardware. PEER-REVIEWED COMPANY ROADMAP ROUGH ESTIMATE

Qubit Counts & Best Gate Fidelities (≈ 2025–26)



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Atom Computing's 1,180-site ^{171}Yb array is a loading demonstration; gate operations use a subset. IBM Condor (1,121Q) is retired; IBM's active flagships are Heron r2 (156Q) and Nighthawk (120Q, launched Nov 2025) with superior error rates. D-Wave's 4,400+ flux qubits (Advantage2) are quantum annealers — not gate-based and not directly comparable, so they are not plotted. The Caltech array is a research demonstration of trapping and coherence, not a processor running gates. PsiQuantum is not plotted: it has no gate QPU yet. Log scale.

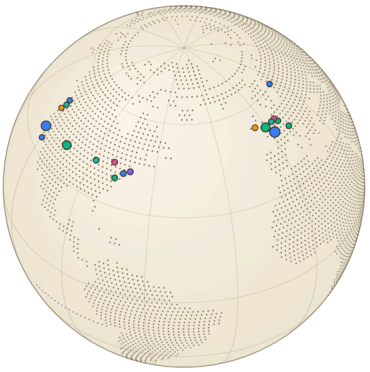


Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Best published 2-qubit gate fidelities. Superconducting: Google Willow / IBM Heron r2 (~99.5%). Trapped ions: Quantinuum Helios 99.92% (Nature 2026, peer-reviewed with Sandia); H2 99.9% (arXiv:2305.03828) shown for context. Neutral atoms: Evered et al., Harvard, *Nature* 2023 (99.5%). Silicon spins: Mills et al. ^{28}Si, *Sci. Adv.* 2022 (99.8%). Red dashed line = practical high-fidelity target ~99.9% (Fowler et al., PRA 2012).

31.6 Where the Hardware Is Built

Every company and lab in the Platform Deep Dives below, placed at its headquarters or main hardware site (city level). Colours are the platforms. Click a marker to open that company in the deep dives; from the keyboard, `]` and `[` step through the markers and Enter opens one.



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

27 companies & labs · HQ or main hardware site, city level

31.7 Platform Deep Dives

How each technology works, key metrics, bottlenecks, leading companies, and key papers.

Superconducting Qubits

How it works

Transmon qubits — Josephson junction circuits cooled to ~15 mK in dilution refrigerators. The Josephson junction makes the LC oscillator anharmonic, isolating the $|0\rangle \leftrightarrow |1\rangle$ transition from the rest of the energy ladder. Microwave pulses (5–8 GHz) drive single- and two-qubit gates. The cross-resonance (IBM) or parametric (Google) interaction is the workhorse two-qubit mechanism.

Key Metrics (2025–26)

Gate time: 10–200 ns 1Q fidelity: >99.9% 2Q fidelity: ~99.5% (Heron/Willow) T_1: 100–500 µs (best Heron) T_2: 50–300 µs Qubit count: 100–500 gate-quality

Principal Bottleneck

Classical wiring at millikelvin: each qubit requires several coaxial lines; heat load and cable density make scaling beyond a few thousand physical qubits extremely hard. Leading solutions: cryo-CMOS multiplexers (Intel Horse Ridge II) and microwave-to-optical transducers for inter-module coupling. IBM's modular Quantum System Two links multiple Heron chips via classical interconnects.

$$E_J = \frac{(\Phi_0/2\pi)^2}{2L_J} \quad \text{[Josephson energy]}$$

$$\omega_q \approx \frac{\sqrt{8E_J E_C}}{\hbar} \qquad E_C = \frac{e^2}{2C_\Sigma}$$

Anharmonicity: $\alpha = \omega_{12} - \omega_{01} \approx -E_C/\hbar \approx -200$ to -350 MHz

Companies (2025–26)

IBM Quantum ↗

Heron r2 (156Q, 2024) and Nighthawk (120Q, launched Nov 2025) are IBM's active flagships — Condor (1,121Q) retired. IBM Quantum System Two is a modular architecture linking Heron chips. Roadmap: Flamingo → Kookaburra for higher qubit quality. Cloud via IBM Quantum Platform.

Google Quantum AI ↗

Willow (105Q, Dec 2024): reported below-threshold error correction — exponential error suppression as code distance increases. A key near-term FTQC demonstration.

Rigetti Computing ↗

Ankaa-3 (84Q, 2024). Publicly traded (RGTI). Cloud hybrid via AWS/Azure. Focus on algorithmic benchmarking over raw qubit count.

IQM Quantum Computers ↗

Finnish startup. IQM Resonance (20Q, Europe-deployed). Co-design QPU architecture. HPC integration focus with CSC Finland and Leibniz Supercomputing Centre.

Alice & Bob ↗

Cat qubits: biased-noise encoding suppresses bit-flips exponentially while accepting higher phase-flip rates, dramatically reducing QEC overhead. Series B raised (2024). Targeting logical qubit demonstration.

D-Wave (QBTS) ↗

Advantage2 (4,400+ flux-qubit annealers, generally available in 2025). Not gate-based — solves QUBO/Ising problems via adiabatic evolution. NASDAQ-listed. Commercial traction in logistics and materials optimisation.

Key Papers

- Google quantum supremacy (Arute et al., *Nature* 2019)
- Surface code logical qubit (Google, *Nature* 2023)
- Willow — below-threshold QEC (Google, *Nature* 2025)
- IBM Heron R2: 156 qubits, heavy-hex layout (IBM blog, Nov 2024)
- Alice & Bob: LDPC-cat codes for low-overhead quantum computing in 2D (Ruiz et al., arXiv:2401.09541)

Trapped Ion Qubits

How it works

Hyperfine "clock" states of $^{171}Yb^{+}$ or $^{40}Ca^{+}$ ions confined in radiofrequency Paul traps. Ions laser-cooled to the motional ground state; shared vibrational modes (phonons) of the ion chain serve as a quantum bus. The Mølmer–Sørensen (MS) gate drives entanglement via off-resonant sideband excitation. Clock-state qubits are insensitive to first-order Zeeman shifts — T_2 can exceed minutes. All-to-all connectivity within a trap zone is native.

Key Metrics (2025–26)

T_2: 1 s – tens of minutes 2Q fidelity: 99.92% (Quantinuum Helios, 2026) 1Q fidelity: $<10^{-4}$ error rate Gate time: 20–500 µs Qubit count: 56–256 (Helios 98; IonQ Superion 256) Ops/T_2: $>10^4$ All-to-all connectivity within zone

Principal Bottleneck & Scaling Path

Mode crowding: N ions share 3N vibrational modes; gates become slower and noisier as the chain grows. **Solution — QCCD:** segment the trap into storage, gate, and readout zones; shuttle ions between them. Quantinuum's H2 (56 qubits) and Helios (98 qubits) implement this with all-to-all connectivity. **Next step:** photonic interconnects linking separate trap modules into a quantum network — the modular ion-trap path to thousands of logical qubits.

$$\hat{H}_{\mathrm{MS}} = \hbar\Omega(\hat{\sigma}_a^{+}\hat{\sigma}_b^{+} + \hat{\sigma}_a^{-}\hat{\sigma}_b^{-})\cos(\delta t)$$

$\delta = \omega_{\mathrm{laser}} - \omega_{\mathrm{qubit}} - \omega_{\mathrm{phonon}}$ (blue sideband detuning)
Gate time: $\tau_g \sim \pi/(\eta^2\Omega)$ where $\eta = kx_{\mathrm{zpf}}$ is the Lamb–Dicke parameter
Infidelity $\propto \bar{n}$ (mean phonon number): requires ground-state cooling $\bar{n} < 0.05$

Companies (2025–26)

IonQ ↗

Forte Enterprise reached #AQ 36 in 2024, and IonQ reported a #AQ 64 Tempo benchmark in 2025. In September 2026, IonQ launched Superion 256 — its first Electronic Qubit Control (EQC)-based commercial system, 256 physical qubits in a standard server-rack footprint, designed to be manufactured "by the hundreds" rather than as one-off lab machines. In October 2026, IonQ announced an all-stock agreement to acquire Zero Trust security firm Elisity for up to $1.3B (incl. up to $150M in earnouts), pairing IonQ's post-quantum cryptography/QKD hardware with Elisity's security software; expected to close Q4 2026, not yet completed. NASDAQ-listed (IONQ). Proprietary #AQ (algorithmic qubit) metric benchmarks circuit performance rather than raw physical qubit count, so do not compare it directly to a hardware qubit count. Partners with AWS, Microsoft Azure, and Google Cloud.

Quantinuum ↗

Helios processor (98 trapped $^{137}Ba^{+}$ ions with $^{171}Yb^{+}$ sympathetic cooling, all-to-all connectivity, launched Nov 2025): best published 2Q gate fidelity (99.92%, peer-reviewed with Sandia, June 2026) among broadly discussed commercial QPUs, succeeding the 56-qubit H2-1 (99.9%, 2024) and demonstrating 48 fully error-corrected logical qubits. Joint venture: Honeywell + Cambridge Quantum. Microsoft made a strategic investment in 2025. Apollo is the next roadmap system.

Oxford Ionics ↗

Part of IonQ since 17 September 2025 (acquisition completed, per IonQ). Oxford spin-out whose Electronic Qubit Control (EQC) drives gates with on-chip electrodes (microwave currents and magnetic-field gradients) instead of steered laser beams, on traps made with standard semiconductor processes. Reported 99.97% two-qubit and 99.9992% single-qubit gate fidelity (announcement, July 2024) COMPANY, and >99.99% two-qubit fidelity without ground-state cooling (arXiv:2510.17286, 2025) PREPRINT. EQC is now the basis of IonQ's Superion systems (see the IonQ card).

AQT (Alpine Quantum Technologies) ↗

$^{40}Ca^{+}$ ion trap systems. European quantum flagship partner. Deployed systems at CERN (HPC integration pilot) and Vienna University of Technology.

eleQtron ↗

German startup. Microwave-driven $^{25}Mg^{+}$ trapped ions. Room-temperature-compatible control electronics. BMBF-funded.

Key Papers

- Mølmer–Sørensen gate (Sørensen & Mølmer, PRL 1999)
- QCCD architecture (Kielpinski et al., *Nature* 2002)
- Quantinuum H2: 99.9% 2Q fidelity (arXiv 2023)
- Quantinuum H2-1: 56-qubit all-to-all connected (arXiv 2024)
- Quantinuum Helios: 98-qubit, 99.92% 2Q fidelity, all-to-all (Nature 2026)
- Fault-tolerant control of an error-corrected qubit (Egan et al., *Nature* 2021)

Neutral Atom Qubits — Optical Tweezers + Rydberg Gates

How it works

Individual neutral atoms (^{87}Rb, ^{133}Cs, ^{171}Yb, ^{88}Sr) trapped in tightly focused laser beams (optical tweezers; ~1 μm waist). Intensity gradient provides a conservative trapping potential. Hyperfine or nuclear-spin ground states form the qubit. Two-qubit gates use the **Rydberg blockade**: transient excitation to n ~ 60–100 Rydberg states, where the van der Waals interaction V = C_6/r^6 ($C_6 \propto n^{11}$) prevents simultaneous excitation of two atoms within the blockade radius r_b. Zone-based architectures (Gemini-class) separate storage, entangling, and readout zones — atoms are transported by AOD-controlled tweezers.

Key Metrics (2025–26)

Array size: hundreds in gate-model processors; 6,100 atoms trapped (2025) Coherence T_2: 1–10 s (hyperfine) 2Q CZ fidelity: 99.5% peer-reviewed (Evered 2023); 99.85% raw in a 2026 preprint Gate time: 0.2–5 μs Identical qubits: same species, locally tunable shifts Reconfigurable 2D/3D geometry in real time Dual-mode: digital gates OR analog Hamiltonian Mid-circuit readout: demonstrated at 48 LQ scale

Engineering Frontiers (2025–26)

① **Gate fidelity → 99.9%+:** Evered et al. 2023 demonstrated 99.5% native CZ on ^{87}Rb. ^{171}Yb and ^{88}Sr nuclear-spin qubits (I = ½, 0) offer longer T_2 and clock-transition magic wavelengths. Optimal-control pulse shaping + deep ground-state cooling ($\langle n \rangle < 0.05$) is the route to fault-tolerance-grade gates.

② **Fault-tolerant logical qubits:** Bluvstein et al. (*Nature* 2024) demonstrated 48 error-corrected logical qubits and 228 logical gates, and Microsoft with Atom Computing entangled 24 logical qubits on ^{171}Yb hardware (Reichardt et al., arXiv:2411.11822, 2024). [[7,1,3]] and surface-code QEC are both being explored.

③ **Scale to 10^4 atoms:** Multiplexed AOD arrays + parallel imaging. QuEra's Gemini-class system targets thousands of atoms. Atom Computing loaded 1,180 ^{171}Yb sites (2023), and Manetsch et al. trapped 6,100 Cs atoms with seconds-scale coherence (*Nature* 2025); running high-fidelity gates across arrays that large is the next step.

④ **Photonic interconnects:** Entanglement between atom arrays via photon-mediated Bell pairs is in early research — essential for modular scale-out beyond a single vacuum chamber.

$$V_{\mathrm{dd}}(r) = \frac{C_6}{r^6}, \quad C_6 \propto n^{11} \quad (\text{van der Waals}, n \gg 1)$$

$$r_b = \left(\frac{C_6}{\hbar \Omega_{\mathrm{Ryd}}}\right)^{1/6} \sim 5\text{–}15\ \mu\mathrm{m} \text{ for } n = 60\text{–}100$$

$$\varepsilon_{\mathrm{motion}} \approx \alpha \langle n \rangle, \quad \langle n \rangle = k_{\mathrm{B}} T / (\hbar \omega_{\mathrm{trap}})$$

Published benchmark: $\langle n \rangle < 0.05$ via resolved-sideband cooling on ^{171}Yb clock transition

Companies (2025–26)

QuEra Computing ↗

Harvard/MIT spin-out (Lukin/Greiner groups). Aquila (256 ^{87}Rb atoms) on AWS Braket — a publicly accessible neutral-atom analog QPU. Gemini-class next-gen targets thousands of atoms in zone-based architecture. Digital gate mode + analog Hamiltonian simulation in one platform. Raised $230M in a financing led by Google and SoftBank (Feb 2025).

Pasqal ↗

French startup (Browaeys/Lahaye, IOGS). Fresnel QPU (100 ^{87}Rb atoms). Hybrid analog-digital. HPC cloud integration (EDF, BASF, Crédit Agricole pilots). Merged with Qu&Co for quantum software stack. European Quantum Flagship partner. Publicly traded (PSQL, Nasdaq, since Aug 28, 2026).

Atom Computing ↗

Headquartered in Berkeley, CA, with its main R&D site in Boulder, CO. Nuclear-spin qubits in ^{171}Yb (I = ½), whose small nuclear magnetic moment makes them weakly sensitive to magnetic noise, so T_2 is long. Loaded a 1,180-site array in 2023; with Microsoft, entangled 24 logical qubits on its hardware (Reichardt et al., arXiv:2411.11822, 2024).

Infleqtion ↗

Formerly ColdQuanta. Sqale cloud QPU (Rb/Cs-based). Quantum networking programme. Acquired SuperTech; contracts with DARPA and US Air Force for quantum sensing + computing. Listed on the NYSE since Feb 2026 (ticker INFQ, Feb 17) after merging with Churchill Capital Corp X.

Key Papers

- Tweezer arrays review (Kaufman & Ni, *Nat. Phys.* 2021)
- QuEra 256-atom analog processor (Ebadi et al., *Nature* 2021)
- 99.5% Rydberg CZ gate (Evered et al., Harvard, *Nature* 2023)
- 48 logical qubits, 228 gates (Bluvstein et al., *Nature* 2024)
- 1,180-qubit ^{171}Yb array (Atom Computing, company announcement, 2023)

Silicon Spin Qubits

How it works

Electron (or hole) spins confined in Si or Si/SiGe quantum dots defined by electrostatic gates. Spin-up $|\uparrow\rangle$ and spin-down $|\downarrow\rangle$ form the computational basis. Single-qubit gates via electron spin resonance (ESR) driven by microwave pulses. Two-qubit gates: exchange interaction J tuned by gate voltage V_ex — the CPHASE gate follows. In isotopically purified ^{28}Si (0.01% residual ^{29}Si) the dominant nuclear spin bath dephasing is removed, giving T_2^* of order 100 µs and T_2 of tens of ms with dynamical decoupling (120 µs and 28 ms in Veldhorst et al., *Nat. Nanotechnol.* 2014).

Key Metrics (2025–26)

Gate time: 20–100 ns 2Q fidelity: 99.8% (^{28}Si best) T_2 (echo): >1 ms in ^{28}Si Ops/T_2: ~10^4–10^5 Qubit count: 6–16 (gate-quality, 2025) Variability: no two dots are identical

The Long Game

Integration density is silicon's superpower: existing 300 mm CMOS fabs could in principle place millions of quantum dots on a single chip — the same infrastructure that makes CPUs. If device-to-device variability can be solved (through automated ML-based tuning or foundry-level uniformity), silicon spins represent the most credible path to millions of physical qubits. Timeline: ~5–10 years from demonstration-scale to hundreds of qubits.

Companies (2025–26)

Intel ↗

Tunnel Falls (12Q, 2023). Built on Intel's 300 mm fab process. Horse Ridge II cryo-CMOS controller. Automotive-grade variability reduction is a core research focus.

Quantum Motion ↗

UK startup. CMOS-compatible Si/SiO_2 platform. Series B raised (2023). Target: foundry-compatible qubit arrays with in-situ tuning. Partnership with GlobalFoundries.

Silicon Quantum Computing ↗

Australian national initiative (UNSW Sydney). Precision atom-by-atom Si:P doping — individual phosphorus donors as qubits. Sub-nm positioning control via STM lithography. 10-qubit demonstrated (2023).

Equal1 ↗

Irish startup. Full-stack QPU on a single CMOS chip — qubits + classical control co-integrated at 4K, eliminating the dilution fridge for the control circuitry.

Key Papers

- Six-qubit Si/SiGe processor (Philips et al., *Nature* 2022)
- 99.8% CNOT in ^{28}Si (Mills et al., *Sci. Adv.* 2022)
- 99.5% 2Q fidelity in ^{28}Si (Noiri et al., *Nature* 2022)
- Intel Tunnel Falls 12Q (arXiv 2024)
- ML-assisted quantum dot tuning (Zwolak et al., *PRApplied* 2020)

Photonic Qubits

How it works

Qubits encoded in polarisation, path, or time-bin modes of single photons in integrated silicon photonics waveguides. Decoherence during transmission is negligible — photons are flying qubits. Single-qubit gates: beam splitters (50:50) and electro-optic phase shifters. The fundamental problem: *photons don't interact* — probabilistic entangling gates (KLM protocol) or matter-mediated nonlinearities are required. Both scale poorly without further architecture innovation.

Fusion-Based QC — PsiQuantum's Bet

Avoid probabilistic two-qubit gates entirely. Pre-generate small entangled resource states (4–6 photons). Fuse them via Bell measurements (fusion gates). Individual fusion failures are heralded and corrected architecturally through the fault-tolerant resource state network. This approach, if single-photon sources and detectors reach ~99%+ efficiency, allows scaling without each gate needing to succeed. PsiQuantum's target: 1M+ photonic modes on GlobalFoundries 300 mm silicon photonics.

Near-Term Value

Quantum key distribution (QKD) networks are commercially deployed today (Toshiba, ID Quantique). Photonic interconnects are essential for linking any qubit module over long distances. Gaussian boson sampling (GBS) on squeezed-light networks offers provable quantum speedup on specific sampling problems.

Companies (2025–26)

PsiQuantum ↗

Major photonic QC bet: >$700M private + AUD $940M Queensland government + US DARPA/DoE contracts (>$1.5B total). Fusion-based architecture on GlobalFoundries 300 mm silicon photonics. No gate QPU yet — building the fab-compatible supply chain first.

Xanadu ↗

Borealis (216 squeezed modes, 2022): demonstrated beyond-classical speed on GBS. PennyLane: open-source quantum ML framework (60k+ GitHub stars). Amazon Braket partner. Photonic fault-tolerance roadmap via GKP encoding.

QuiX Quantum ↗

Si_3N_4 waveguide processors for boson sampling. 20-mode photonic chip (2021). European Quantum Flagship member. Focus on linear-optical circuits for near-term experiments.

Quandela ↗

InGaAs quantum dot single-photon sources with >99% purity. Muse cloud QPU. Focus on high-efficiency, high-indistinguishability photon generation — the rate-limiting hardware for photonic QC.

Key Papers

- KLM protocol (Knill, Laflamme, Milburn, *Nature* 2001)
- GBS photonic advantage (Zhong et al., *Science* 2020)
- Borealis programmable GBS (Madsen et al., *Nature* 2022)
- Fusion-based QC architecture (Bartolucci et al., *Nat. Commun.* 2023)

Solid-State Defects & Topological Qubits

NV Centres & Diamond Defects

Nitrogen-vacancy (NV) centres — a nitrogen atom adjacent to a lattice vacancy — are electron-spin qubits with millisecond coherence at *room temperature*, which is rare among qubit platforms, and about one second at cryogenic temperature with dynamical decoupling (Abobeih et al., *Nat. Commun.* 2018); nearby ^{13}C nuclear spins serve as longer-lived memories. Optically addressable: spin state initialised and read out via spin-selective fluorescence (ZPL at 637 nm). Related platforms: SiV and GeV centres in diamond (Lukin group, Harvard), defects in SiC (neutral divacancy, V_1). Principal near-term roles: quantum repeater network nodes, nanoscale magnetometry and NMR, and as memory qubits in hybrid quantum networks.

Topological Qubits — Majorana Zero Modes

Majorana fermions at the ends of 1D topological superconductors (semiconductor nanowire + s-wave SC in a magnetic field) encode a qubit non-locally — meaning local perturbations cannot flip it. Theoretically this gives hardware-level error protection far below 10^{-6}, dramatically reducing QEC overhead.

Feb 2025 — Microsoft Majorana 1: Microsoft's Station Q team announced the Majorana 1 chip: an InAs/Al heterostructure device designed to host Majorana zero modes, with a topological gap measured via interferometric parity readout. This is the most concrete hardware demonstration yet, though independent verification of non-Abelian braiding remains an open experimental challenge.

⚠️ **Topological qubits: promising physics, unproven computation.** The Feb 2025 Majorana 1 announcement demonstrates a topological gap and parity readout; *braiding* (the operation needed for fault-tolerant gates) has not yet been demonstrated. Peer review and independent replication are ongoing. Treat commercial timelines with skepticism — this remains a physics frontier.

Companies & Groups (2025–26)

Microsoft Azure Quantum ↗

Majorana 1 chip (InAs/Al, Feb 2025). Station Q research lab. Azure Quantum also partners with IonQ, Quantinuum, Rigetti, and Pasqal as a multi-hardware cloud provider; distinguish those cloud partnerships from Microsoft's separate topological-qubit research program.

Quantum Brilliance ↗

Room-temperature NV-centre quantum accelerators. Diamond QPU at 20–25°C — no dilution fridge. Partnered with Oak Ridge National Lab (ORNL). Focus: edge-deployable quantum co-processors for HPC nodes.

Q-NEXT / Argonne ↗

US DOE quantum center. NV centres for quantum repeater nodes. Operating a Chicago–Argonne quantum network testbed for metropolitan quantum-network experiments. Near-term goal: entanglement distribution across 50 km.

QuTech (TU Delft) ↗

Pioneered NV-centre quantum networking (Hensen 2015 loophole-free Bell test). Built a three-node NV network in Delft (Pompili et al. 2021) and entangled nodes in Delft and The Hague over deployed fiber (Stolk et al., *Sci. Adv.* 2024). Topological Josephson junction research also active.

Key Papers

- NV centre review (Doherty et al., *Phys. Rep.* 2013)
- Loophole-free Bell test / NV network (Hensen et al., *Nature* 2015)
- Three-node quantum network (Pompili et al., *Science* 2021)
- Topological quantum wire theory (Kitaev 2001)
- Microsoft Majorana 1 chip (arXiv 2025)

31.8 Where the Field Stands

The goal, the status in 2025–26, and the gap between them.

> 🎯 The End Goal

Fault-tolerant quantum computers with **logical** qubits running Shor's algorithm on RSA-2048 (~20 million noisy qubits at 10^{-3} error rate in the Gidney & Ekerå 2021 estimate; under a million in Gidney's 2025 preprint, at the cost of a week-long run) and simulating industrial molecules such as FeMoco, which published resource estimates put at roughly a hundred to a few thousand logical qubits and billions of logical gates.

📍 Status (2025–2026)

The field is moving from NISQ demonstrations to early fault tolerance: a below-threshold surface-code memory on superconducting qubits (Google Willow), 48 error-corrected logical qubits on 98 trapped ions (Quantinuum Helios), and a 96-logical-qubit fault-tolerant architecture on neutral atoms (Harvard/QuEra). Physical qubit counts run from tens to thousands, depending on whether one counts gate-operated qubits, loaded atoms, or annealer flux qubits. The newest results, including preprints, are in section 10.

🚧 The Gap

Useful surface-code architectures generally target $<10^{-3}$ two-qubit error rates, although threshold estimates are order 10^{-2} for simple noise models — the best published physical two-qubit fidelities are now around the 99.9% level in trapped-ion systems. Published estimates for breaking RSA-2048 range from ~20 million noisy qubits in 8 hours (Gidney & Ekerå 2021) to under a million in under a week (Gidney 2025, preprint), still roughly **3–4 orders of magnitude** beyond today's largest gate-model systems. There is no agreed date for a cryptographically relevant machine; expert estimates spread across the 2030s and later.

31.9 Road to Fault-Tolerant Quantum Computing

Milestone timeline across the three leading platforms (2019–2032). Solid dot = achieved; open dot = projected.

🛡️ Surface Code — Key Numbers

Encodes 1 logical qubit in ~$2d^2$ physical qubits (code distance d). Corrects any error on $<\lfloor d/2 \rfloor$ qubits per syndrome round. At physical error rate 10^{-3}, distance-21 code (~882 physical/logical qubit) gives logical error rate ~10^{-10}/cycle. Google Willow (Dec 2024) demonstrated exponential error suppression *below* threshold: adding more qubits reduced the logical error rate. Breaking RSA-2048 was estimated at **~20 million noisy qubits** (8 hours; Gidney & Ekerå, *Quantum* 5, 433 (2021)), revised to **under a million** (under a week, at 0.1% gate error and a 1 µs surface-code cycle; Gidney, arXiv:2505.15917, 2025 preprint).

🔵 Superconducting

2019 Google Sycamore 53Q quantum supremacy

2023 Google: surface code logical qubit demonstrated

2024 Willow 105Q: below-threshold QEC — a first

2024 IBM shifts to Heron r2; retires Condor; utility-scale focus

2027 ~10 high-quality logical qubits (Google/IBM roadmaps)

2030+ 100–1000 logical qubits; cryo-CMOS interconnects

🟢 Trapped Ions

2021 IonQ 32-qubit system on AWS/Azure

2023 Quantinuum H2: 99.9% 2Q gate fidelity

2024 Quantinuum H2-1: 56Q all-to-all connected

2025 H2 reports quantum volume 33,554,432

2025 Helios launches: 98 qubits, all-to-all, 48 logical qubits

2026 Helios fidelity peer-reviewed: 99.9975% 1Q / 99.921% 2Q

2027+ Apollo roadmap system; modular QCCD network

2032 Fault-tolerant logical qubit processor

🟣 Neutral Atoms

2021 QuEra 256-atom analog processor (Ebadi et al.)

2023 99.5% Rydberg CZ gate (Evered et al., Harvard)

2023 1,180-qubit ^{171}Yb loading (Atom Computing)

2024 48 logical qubits + logical circuits (Bluvstein et al.)

2024 24 entangled logical qubits on ^{171}Yb (Microsoft + Atom Computing, preprint)

2025 6,100 coherent Cs qubits; 96-logical-qubit FT architecture

2026 99.85% raw Rydberg CZ in a 442-atom array (Evered et al., preprint)

2026–28 Repeated active QEC cycles and larger logical processors

Open circles = projected. Pulsing dots = in progress (2025–26). Timelines are roadmap estimates, not guarantees.

31.10 Neutral Atom Roadmap: Engineering Milestones

Six milestones on the path from today's 96-logical-qubit demonstrations to larger fault-tolerant neutral-atom processors.

✓ 48 → 96 Logical Qubits (2024–2025)

Bluvstein et al. (*Nature* 2024) demonstrated a 48-logical-qubit reconfigurable Rb processor. The 2025 architecture paper used up to 448 atoms and 96 logical qubits to demonstrate the core elements of universal fault-tolerant neutral-atom computation: zone-based transport, mid-circuit measurement, feed-forward, and logical operations.

Done ✓

① Gate Fidelity → 99.9% Without Post-Selection

Evered et al. reached 99.5% native CZ in Rb (2023), and a 2026 preprint from the same group reports 99.854% raw, averaged over 17 gate sites in a 442-atom array (arXiv:2604.25987, not yet peer-reviewed). Muniz et al. reported 99.72(3)% for ^{171}Yb with post-selection and 99.40(3)% without. The next milestone is reproducible ≥99.9% two-qubit fidelity without post-selection, at scale and with mid-circuit operations enabled.

Active

② Scale to 10,000+ Atoms

Multiplexed AOD/SLM arrays + parallel imaging. Atom Computing demonstrated 1,180 loaded ^{171}Yb atoms in 2023; Manetsch et al. reported 6,100 highly coherent ^{133}Cs atoms across 12,000 tweezer sites in 2025. The remaining challenge is not loading alone, but maintaining gate fidelity, coherence, transport, and low-loss readout across the full array.

Near-term

③ Full Mid-Circuit Measurement + Feed-Forward at Scale

Demonstrated at the 48-logical-qubit level (Bluvstein 2024). Extending to larger arrays with fast FPGA-based real-time decision making is required for active QEC. The bottleneck is fluorescence imaging speed vs. coherence time — cameras and optics must read out ancilla qubits without disturbing data qubits.

Near-term

④ Repeated Active QEC Below Threshold

Neutral atoms now have logical processors and a universal fault-tolerant architecture demonstration. The next harder milestone is repeated active syndrome extraction where increasing code distance reliably lowers logical error over many cycles, while preserving atom number and avoiding measurement cross-talk.

2027–28

⑤ Photonic Interconnects — Modular Scale-Out

A single vacuum chamber cannot hold millions of atoms. Modular architecture links separate tweezer modules via photon-mediated Bell pairs (atom → photon → fibre → atom entanglement). NV-centre quantum repeater technology provides a blueprint. Relevant for neutral-atom quantum networks and distributed quantum computing.

Research

31.11 2024–2026 Flash Points

The developments that moved the field, newest first — from lab breakthroughs to billion-dollar bets. This is the part of the page that dates fastest.

Jun 2026

🗺️

QuEra Gigaquop Roadmap

QuEra published a fault-tolerance roadmap: Libra, a "megaquop-class" (~10^6 reliable logical operations) system targeted for Amazon Braket around 2028, followed by a "gigaquop-class" (~10^9 reliable logical operations) system with 1,000+ logical qubits at a 10^{-9} logical error rate on 20,000+ physical qubits in a single processing core. QuEra 2026 COMPANY ROADMAP.

Nov 2025

🏆

Quantinuum Helios — 98-Qubit All-to-All

Quantinuum launched Helios (98 fully connected trapped-ion qubits) in November 2025 and demonstrated 48 error-corrected logical qubits. With Sandia it published peer-reviewed gate fidelities of 99.9975% (one qubit) and 99.921% (two qubits), Nature 2026. 2026 preprints on the same hardware report up to 94 error-detected logical qubits and logical memory errors below the physical error rate, down to 5×10^{-6} per cycle with a colour code ($\Lambda\approx12$): arXiv:2602.22211, arXiv:2609.03194, arXiv:2610.24077 PREPRINTS.

2025

💰

Quantinuum H2 — 56 Trapped-Ion Qubits

Quantinuum launched the 56-qubit H2-1 trapped-ion processor in 2024 and reported quantum volume 33,554,432 in 2025. Its strength remains high-fidelity all-to-all gates, with scaling limited by gate speed and modular ion transport. Quantinuum 2024.

Feb 2025

⚡

Microsoft Majorana 1 — Topological Qubit Chip

Microsoft announced the *Majorana* 1 chip, an InAs/Al heterostructure device aimed at topological qubits. Independent confirmation of topological qubit operation and braiding remains open, so treat this as a high-risk frontier rather than a settled platform. In September 2026 Microsoft introduced a second-generation chip (lead instead of aluminium) and gave DARPA on-site access to test it; no qubit-count or fidelity numbers were disclosed. COMPANY

Dec 2024

🌟

Google Willow — Below-Threshold QEC

Google's Willow processors demonstrated below-threshold surface-code memories up to distance 7: logical error per cycle decreased as code distance increased, and the distance-7 memory exceeded the lifetime of its best physical qubit. Nature 2025.

Nov 2024

🔗

Microsoft + Atom Computing — 24 Entangled Logical Qubits

On Atom Computing's ^{171}Yb tweezer hardware, Microsoft's qubit-virtualization stack created and entangled 24 logical qubits, and ran a Bernstein–Vazirani circuit on 28 logical qubits while detecting and correcting atom loss. Reichardt et al., arXiv:2411.11822 PREPRINT.

2024–25

🔧

IBM Roadmap — Utility to FTQC

IBM's public roadmap now emphasizes quality, modularity, and fault tolerance over raw qubit count: Heron-class 156-qubit processors now, Nighthawk/Kookaburra milestones in 2026, and Starling targeted for 2029 with 200 logical qubits and 100 million gates. IBM roadmap. COMPANY ROADMAP

2024–26

⚛️

Harvard/QuEra — 48 to 96 Logical Qubits

Bluvstein et al. demonstrated a 48-logical-qubit reconfigurable neutral-atom processor in Nature 2024. A follow-up published online in 2025 and appearing in Nature 2026 used up to 448 atoms to demonstrate key pieces of a universal fault-tolerant neutral-atom architecture with 96 logical qubits. Nature 2024; Nature 2026.

2023–26

🏛️

PsiQuantum — Government Mega-Funding

Secured AUD $940 million from the Queensland government plus US DARPA and DoE contracts, betting that silicon photonics and wafer-scale fabrication can make fault-tolerant photonic QC manufacturable. In September 2026 the US Department of Commerce finalized a $100M CHIPS Act R&D award targeting electro-optic materials, high-temperature single-photon detectors, and ultra-low-loss photonic packaging. Total private + public backing now exceeds $1.5 B. NIST/Commerce 2026.

31.12 Investment Landscape (2023–2026)

Capital is concentrating around a handful of hardware bets — and governments are piling in.

$1.5B+

PsiQuantum

Private rounds + AUD $940M Queensland government + US DARPA/DoE contracts. Silicon photonic fusion-based QC.

$1.68B

Quantinuum

Nasdaq IPO (QNT, June 2026), following Microsoft's strategic investment and Helios (98 qubits, Nov 2025) achieving peer-reviewed 99.92% 2Q fidelity in June 2026.

$230M

QuEra Computing

Financing led by Google and SoftBank (Feb 2025). Harvard/MIT spin-out. AWS cloud partner.

$3B+

IBM Quantum (cumulative)

Internal R&D since 2016. Utility-scale QPUs, Qiskit ecosystem, cloud access to 100+ organisations worldwide.

$300M+

Atom Computing

Cumulative funding (including its June 2026 financing and a $100M U.S. Commerce Dept. letter of intent) to accelerate deployment of fault-tolerant neutral-atom systems. ^{171}Yb nuclear-spin qubit platform.

$2B

Anderon (IBM)

Standalone IBM quantum-foundry subsidiary; finalized a $1B CHIPS Act award (Sep 2026) matched by $1B of IBM investment to build a 300mm pure-play quantum wafer foundry in Albany, NY, for superconducting qubit arrays and readout electronics.

Pattern: Governments (US, UK, EU, Australia, Japan and others) have announced national programmes worth tens of billions of dollars combined. The bet is: whoever reaches fault tolerance first will have a decisive advantage in cryptography, drug discovery, and materials science.

31.13 AMO Physics → Quantum Industry

How your experimental skills translate — and what to add to be competitive in 2025–26 hiring.

Your AMO Skills → QC Industry Value

Optical tweezers → Neutral atom QC hardware (QuEra, Pasqal, Atom Computing)

Rydberg atoms / blockade → Rydberg gate development; two-qubit gate optimisation

Laser locking (PDH, SAS) → Qubit drive laser stability; photon-source engineering

Ultra-high vacuum systems → All hardware platforms; QPU system engineering

Fluorescence imaging / SNR → Qubit readout fidelity; mid-circuit measurement

MOT / evaporative cooling → State preparation, cooling to ground state ($\langle n \rangle < 0.05$)

RF / microwave spectroscopy → Hyperfine qubit drives; superconducting gate pulses

FPGA / Python data analysis → Quantum control software; real-time QEC feed-forward

Precision measurement / statistics → Gate benchmarking (RB, XEB, process tomography)

Skills to Add for Industry

Quantum Software Frameworks

Qiskit (IBM) Bloqade (QuEra, Julia) PennyLane (Xanadu) Cirq (Google)

Theory to Know

Surface codes & QEC Randomised benchmarking VQE / QAOA basics Quantum error models

Job Roles for AMO PhDs

Quantum Hardware Scientist Systems Engineer Quantum Control Engineer Applications Physicist Error Correction Researcher

Where AMO PhDs Are Most Competitive — by Platform

QuEra / Pasqal / Infleqtion / Atom Computing — Direct fit. Your tweezer + Rydberg + laser experience is the core of their hardware stack. These are the most natural landing spots for an AMO postdoc transitioning to industry.

Quantinuum / IonQ (incl. Oxford Ionics) — Strong overlap: laser cooling, vacuum, spectroscopy, microwave control. Trapped-ion teams heavily recruit from AMO backgrounds.

IBM Research / Google Quantum AI — Research scientist roles for AMO physicists with strong QEC or quantum optics background. Typically require demonstrated software fluency (Qiskit/Cirq) + familiarity with superconducting physics.

National Labs (JILA, NIST, Sandia, Argonne, Oak Ridge) — AMO + QC intersection is well-funded. NIST Ion Storage Group, Argonne Q-NEXT, Sandia Quantum Network all hire AMO PhDs directly. Strong bridge if you want to stay research-adjacent.

PsiQuantum / Xanadu / Quandela — Photonics-adjacent. Relevant if you have quantum optics, cavity QED, or single-photon source experience. Harder cold fit from pure tweezer background unless you build quantum optics skills.

Quantum Software / Cloud (1QBit, Algorithmiq, QC Ware) — Open to strong physicists who learn Python, QEC, and algorithm basics. Less hardware, more classical-quantum interface. Good pivot if you want industry pace over lab life.

31.14 References & Further Reading

Neutral Atoms (Most Relevant)

- Kaufman & Ni, *Nat. Phys.* 2021 — tweezer arrays review PEER-REVIEWED
- Ebadi et al. (QuEra), *Nature* 2021 — 256-atom analog PEER-REVIEWED
- Evered et al., *Nature* 2023 — 99.5% Rydberg CZ gate PEER-REVIEWED

- Bluvstein et al., *Nature* 2024 — 48 logical qubits PEER-REVIEWED
- Atom Computing 1,180-qubit ^{171}Yb array (company announcement, 2023) COMPANY

Superconducting, Ions & Other Platforms

- Arute et al. (Google), *Nature* 2019 — quantum supremacy
- Google Willow — below-threshold QEC, *Nature* 2025
- Google surface code logical qubit, *Nature* 2023
- Quantinuum H2, 99.9% 2Q fidelity (arXiv 2023)
- Quantinuum Helios, 98-qubit all-to-all, 99.92% 2Q fidelity (*Nature* 2026) PEER-REVIEWED
- Philips et al. (Si/SiGe 6Q), *Nature* 2022
- Madsen et al. (Borealis GBS), *Nature* 2022
- Microsoft Majorana 1 chip (arXiv 2025)

Review & Theory

- DiVincenzo 2000 — Five criteria
- Preskill 2018 — NISQ era
- Gidney & Ekerå, *Quantum* 2021 — RSA-2048 needs ~20M qubits
- Gidney 2025 (preprint) — RSA-2048 with <1M noisy qubits in <1 week
- Fowler et al. 2012 — surface code threshold & overhead
- Sørensen & Mølmer, *PRL* 1999 — MS gate
- Bartolucci et al. — fusion-based QC, *Nat. Commun.* 2023

CHAPTER 31 SUMMARY

Every quantum hardware platform faces a distinct engineering trade-off between gate speed, fidelity, and scaling. The final four chapters compare head-to-head matchups within each major modality.

CHAPTER 32 · CAREER STEP 04

32

Neutral-Atom Qubits: ^{87}Rb vs. ^{171}Yb

⚖️ Rb vs Yb Qubit Comparison · Deep dive

Rubidium-87 or ytterbium-171? The two leading neutral-atom qubits compared from the atomic structure up: hyperfine versus nuclear-spin encoding, cooling, coherence, gate fidelity, readout, scaling and erasure conversion, ending in a row-by-row scoreboard you can weight yourself.

> *Within neutral-atom quantum computing, the field has largely converged on two contrasting workhorses: alkali hyperfine qubits (^{87}Rb) and alkaline-earth-like nuclear-spin qubits (^{171}Yb).*

IN THIS CHAPTER

Prerequisites: Ch. 1, Ch. 24, Ch. 31

- Compare ^{87}Rb and ^{171}Yb across laser complexity, cooling, coherence times, erasure conversion, and Rydberg excitation pathways.
- Evaluate published logical-qubit and two-qubit gate milestones from Harvard/QuEra (^{87}Rb) and JILA/Princeton/Atom Computing (^{171}Yb).
- Assess how species choice shapes mid-circuit readout, shuttling, and photonic networking architectures.

WEBSITE COMPANION Live interactive page ↗

32.1 Why These Two Atoms?

The two leading neutral-atom qubit platforms, and why the choice of atom matters deeply.

Neutral atom quantum computers trap individual atoms in optical tweezers and entangle them via Rydberg blockade. The atom species determines everything from qubit coherence to gate speed to the complexity of the laser system. Two atoms have emerged as the leading candidates for scalable fault-tolerant neutral-atom QC, **^{87}Rb** (an alkali, I=3/2, used by the Harvard/QuEra group) and **^{171}Yb** (an alkaline-earth-like rare earth, I=1/2, J=0 ground state, used by Atom Computing). Their physics are fundamentally different, leading to different strengths.

The core trade-off: ^{87}Rb is the mature, well-characterised platform with the strongest *published* Rydberg-gate and logical-processor record: 99.5% two-qubit CZ gates on up to 60 atom pairs operated in parallel, 48 logical qubits encoded in 280 physical atoms (2024), and up to 96 logical qubits in 448 atoms (Bluvstein et al., Nature 649, 2026). ^{171}Yb has a different structural advantage: a J=0 ground state, nuclear-spin qubit, clock-state shelving, and alkaline-earth erasure-conversion protocols. Its best published gate number is 99.40(3)% raw and 99.72(3)% with post-selection. The 6100-atom, 12.6 s coherence result is a **Cs-133** benchmark, not an Rb-87 benchmark; it is included below only as adjacent alkali/tweezer scaling context. **2026 update:** the same Harvard/QuEra collaboration reported a new Rb87 CZ-gate result in an April 2026 preprint — 99.854(4)% raw, 99.941(3)% with loss post-selection, averaged over many gate sites in a 442-atom array and stable across a 10-hour, ~2-million-run benchmark. That would exceed Yb171's best published number on both raw and post-selected fidelity, but it has not yet been peer-reviewed or independently reproduced; see Section 06 for the full picture.

99.5% **Rb87 2Q fidelity (Evered et al. 2023)**	96 **Rb87 logical qubits (Bluvstein et al. 2026; 48 in 2024)**	99.72% **Yb171 2Q fidelity (Muniz et al. 2025, post-sel)**
J = 0 **Yb171 ground state electronic protection**	99.85% **Rb87 2Q fidelity (raw) Evered et al. 2026 — preprint**	

Calculator output at the page's default settings.

> **Source-confidence rule for this page:** PEER-REVIEWED numbers come directly from papers; COMPANY ROADMAP numbers are useful but not equivalent to peer-reviewed benchmarks; ROUGH ESTIMATE entries are order-of-magnitude design guidance. The Cs-133 6100-atom result is explicitly marked as adjacent alkali context, not as an Rb87 panel result.

Who leads where, at a glance

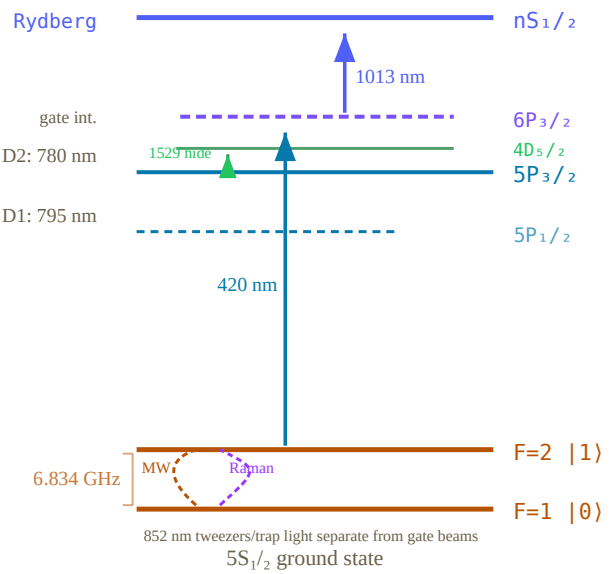
Over the whole table the calls lean **Yb171** (Rb87 8, Yb171 12, 3 even or different), but the categories matter more than the total. Each card counts the *Edge* calls in one block of the full comparison table; click one to see its rows.

Conditional calls ("if it scales", "pending review") count half. Rows are not equally important: read the rows before quoting a winner.

32.2 Atomic Physics Deep Dive

Electronic structure, relevant transitions, and why the physics dictates platform choices.

🔴 ^{87}Rb, Alkali Atom

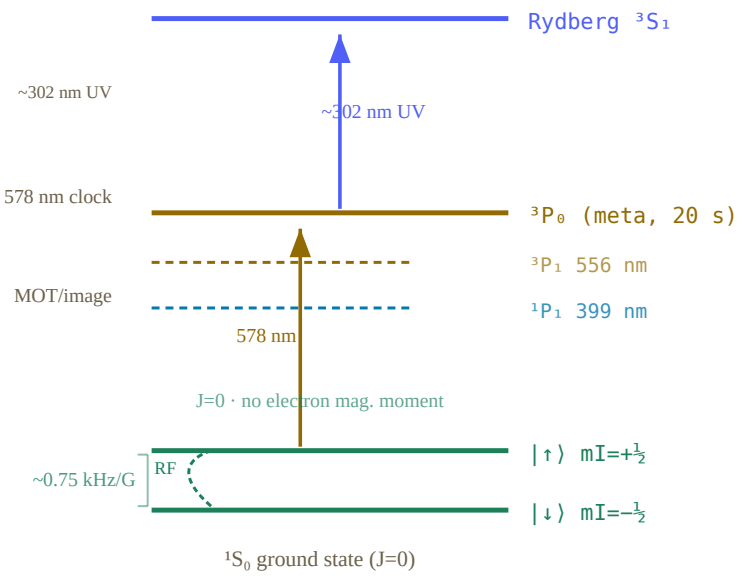


Rb87 level structure. Qubit: hyperfine states of $5S_{1/2}$. Harvard/QuEra high-fidelity gates use two-photon Rydberg excitation through $6P_{3/2}$ with 420 nm + 1013 nm light; state-selective hiding can use the $5P_{3/2} \rightarrow 4D_{5/2}$ transition near 1529 nm. The common 780 nm + 480 nm route via $5P_{3/2}$ is a different Rb Rydberg path, not the Evered 2023 gate path.

- **Z = 37, ground config:** [Kr] $5s^1$, single valence electron
- **Nuclear spin:** I = 3/2
- **Ground state:** $5S_{1/2}$, electronic spin J = 1/2, so F = 1 or 2
- **Cooling transitions:** D1 at 795 nm ($5S \rightarrow 5P_{1/2}$), D2 at 780 nm ($5S \rightarrow 5P_{3/2}$)
- **Rydberg excitation:** Harvard/QuEra gate path via $6P_{3/2}$ (420 nm + 1013 nm); common alternative route via $5P_{3/2}$ (780 nm + 480 nm)

- **State-selective hiding:** Some Harvard/QuEra-style circuits use a 1529 nm beam coupling $5P_{3/2}$ to $4D_{5/2}$ for hiding/shelving operations
- **Natural linewidth:** $\Gamma/2\pi$ = 6.07 MHz (D2)
- **Rydberg interaction scale:** Strongly state- and field-dependent; use ARC/pair-state calculations for exact C_6. Near n≈70 S-states, C_6 is order 10^5–$10^6 \times 2\pi$ MHz·μm^6 ROUGH ESTIMATE

🟢 ^{171}Yb, Alkaline-Earth-Like



Yb171 level structure. Qubit: nuclear spin mI=±½ in $^1S_0$ ground state (J=0). Clock transition to $^3P_0$ at 578 nm. Atom Computing/Muniz-style Rydberg gates use a single 302 nm UV photon from $^3P_0$ to a Rydberg $^3S_1$ state. MOT on $^1P_1$ at 399 nm; narrow-line cooling on $^3P_1$ at 556 nm.

- **Z = 70, ground config:** [Xe] $4f^{14}$ $6s^2$, two valence electrons (closed shell)
- **Nuclear spin:** I = 1/2 (simplest possible)
- **Ground state:** $^1S_0$ (J = 0, critical for qubit protection)
- **Cooling transitions:** $^1P_1$ at 399 nm (broad), $^3P_1$ at 556 nm (intercombination)
- **Clock transition:** $^1S_0 \rightarrow {}^3P_0$ at 578 nm (< 10 mHz linewidth)
- **Rydberg excitation:** Single UV photon (~302 nm) from metastable $^3P_0$ to a Rydberg $^3S_1$ state in the Atom Computing/Muniz gate implementation; overall path from $^1S_0$ is sequential clock shelving plus UV excitation (578 nm + 302 nm)
- **Clock magic wavelength:** near 759 nm for the $^1S_0$–$^3P_0$ optical clock transition; not required to cancel the two $^1S_0$ nuclear-spin qubit light shifts ATOMIC STRUCTURE

> **The J=0 advantage (correctly stated):** In Yb171, the $^1S_0$ ground state has J=0, so the qubit (nuclear spin mI=±½) couples only at the nuclear magneton scale (~0.75 kHz/G, linear in B; $\mu = +0.49\ \mu_N$). Rb87 *clock* states |mF=0⟩ have **zero first-order Zeeman** — their sensitivity is quadratic (575.15 Hz/G^2 × B^2 for the clock transition). At a typical bias field B_0≈1 G and noise δB=1 mG, Rb87 has ≈1.15 Hz sensitivity while Yb171 has ≈0.75 Hz sensitivity; the crossover is near B_0≈0.65 G. The *primary* structural advantages of Yb171's J=0 ground state are: (1) **collisional immunity** — no electronic spin means spin-exchange relaxation is highly suppressed; (2) **spin-independent trapping** of the $^1S_0$ nuclear-spin qubit to leading scalar order; and (3) **erasure conversion** via $^3P_0$.

32.3 Qubit Encoding

How each system defines the two-level qubit, and the physics governing its frequency and noise sensitivity.

Rb87, Hyperfine Qubit

The qubit is encoded in two *hyperfine* ground states of the electronic $5S_{1/2}$ manifold, separated by the nuclear-electron hyperfine coupling A_{hf}:

$$|0\rangle = |F{=}1, m_F{=}0\rangle, \quad |1\rangle = |F{=}2, m_F{=}0\rangle$$
$$\nu_{\text{qubit}} = 6{,}834{,}682{,}610.904\ \text{Hz} \approx 6.834\ \text{GHz}$$

Clock states ($m_F{=}0$): first-order Zeeman is exactly zero.
Quadratic (second-order) Zeeman: $\delta\nu \approx +575.15\ \text{Hz/G}^2 \times B^2$
Sensitivity to fluctuations δB at bias field B_0: $\delta\nu \approx 1150.3\ \text{Hz/G}^2 \times B_0 \cdot \delta B$

- Driven by microwave pulses at 6.834 GHz (or two-photon Raman beams)
- **Clock states |mF=0⟩: zero first-order Zeeman** — symmetry cancels linear Zeeman exactly for mF=0 ↔ mF=0 transition
- Residual quadratic Zeeman: 575.15 Hz/G² × B²; at B=1 G → 575.15 Hz total shift, sensitivity dν/dB = 1150.3 Hz/G
  PEER-REVIEWED
- At typical bias field B_0=1 G, noise δB=1 mG: δν ≈ 1.15 Hz, about 1.5× larger than Yb171's nuclear-spin shift (0.75 Hz); Rb wins only below B_0 ≈ 0.65 G
- Differential light shift from tweezers must be calibrated or nulled

Yb171, Nuclear Spin Qubit

The qubit is encoded in the two nuclear spin states of the $^1S_0$ ground state (I=1/2, J=0). With no electronic angular momentum, the only B-field coupling is through the nuclear magneton:

$$|0\rangle = |m_I{=}-\tfrac{1}{2}\rangle, \quad |1\rangle = |m_I{=}+\tfrac{1}{2}\rangle$$
$$\nu_{\text{qubit}}(B) = -g_I\,\mu_N\,B/h \approx -0.750\ \text{kHz/G} \times B \qquad (\mu = +0.49\,\mu_N,\ g_I = \mu/(I\mu_N) \approx 0.98)$$

Zeeman shift: ~0.75 kHz/G (nuclear magneton scale; qubit frequency ~7.5 kHz at $B = 10$ G)

- Driven by RF pulses (kHz–MHz range in typical lab fields) or two-photon Raman
- First-order Zeeman shift at nuclear magneton scale (~0.75 kHz/G); Rb87 clock states have zero first-order Zeeman and quadratic 575.15 Hz/G²
- For the $^1S_0$ nuclear-spin qubit, scalar light shifts are spin-independent to leading order; 759 nm is the $^1S_0$–$^3P_0$ clock magic wavelength, not a special cancellation between |mI=±1/2⟩ qubit states
- No electronic magnetic moment in ground state → intrinsic noise immunity

B-field sensitivity comparison

$$\underbrace{\delta\nu_{\text{Rb87}}^{(\text{clock})}}_{\text{quadratic Zeeman}} \approx 2 \times 575.15\ \text{Hz/G}^2 \times B_0 \cdot \delta B \qquad \text{vs} \qquad \underbrace{\delta\nu_{\text{Yb171}}}_{\text{nuclear spin, linear}} \approx 0.75\ \text{kHz/G} \times \delta B$$

At $B_0 = 1$ G, $\delta B = 1$ mG: $\delta\nu_{\text{Rb87}}^{(\text{clock})} \approx 1.15$ Hz vs $\delta\nu_{\text{Yb171}} \approx 0.75$ Hz — **Yb171 is ~1.5× lower-noise at 1 G.**
At $B_0 = 10$ G: $\delta\nu_{\text{Rb87}} \approx 11.5$ Hz vs $\delta\nu_{\text{Yb171}} \approx 0.75$ Hz — Yb171 is ~15× better.
At $B_0 = 0.1$ G: $\delta\nu_{\text{Rb87}} \approx 0.115$ Hz vs $\delta\nu_{\text{Yb171}} \approx 0.75$ Hz — Rb87 clock is ~6.5× better at low fields.
Crossover: $B_0 \approx 0.65$ G for any δB (both shifts are linear in δB).

32.4 Laser Cooling & State Preparation

From room temperature to single-atom quantum control, the cooling chain for each platform.

Rb87 Cooling Chain

1. **2D/3D MOT**, Doppler cooling on D2 line (780 nm) to ~200 μK, loading 10^8–10^9 atoms.
2. **Gray Molasses (GM)**, Sub-Doppler cooling on D1 or D2 line using coherent dark states (Λ-enhanced). Achieves (4.0 ± 0.3) μK without a bias field.
3. **Tweezer loading**, Single atoms loaded stochastically into 820–850 nm tweezers. Single-atom detection via fluorescence. ~50% initial loading per site.
4. **Sideband cooling**, Raman or resolved-sideband cooling to motional ground state $\langle n \rangle < 0.1$ per axis. Required to minimize Doppler dephasing during Rydberg gates.
5. **Optical pumping**, Spin initialization to $|F{=}1, mF{=}0\rangle$ with π-polarised light on $F{=}1 \to F'{=}1$ (the $|1,0\rangle \to |1',0\rangle$ transition is forbidden, so $|1,0\rangle$ is dark) + repumper. Fidelity >99.5%.

Key challenge: An 850 nm tweezer is only ~31 THz detuned from the 384 THz D2 line ($5S_{1/2} \to 5P_{3/2}$). Tweezer photon scattering causes qubit dephasing unless the trap is far off-resonance. Typical trap depth ~1 mK.

Yb171 Cooling Chain

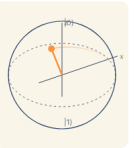
1. **Zeeman slower + 2D MOT**, Broad $^1P_1$ transition at 399 nm ($\Gamma/2\pi$ = 29 MHz) for initial deceleration to ~1 m/s, ~10^9 atoms.
2. **Intercombination MOT**, $^3P_1$ transition at 556 nm ($\Gamma/2\pi$ = 182 kHz, narrow linewidth) for cooling to ~5 μK, close to its Doppler limit $\hbar\Gamma/2k_B \approx 4.4$ μK, without a dedicated sub-Doppler stage.
3. **Optical tweezer loading**, commonly chosen to balance trap depth, scattering, and clock-state needs. 759 nm is magic for the $^1S_0$–$^3P_0$ clock transition; the ground-state nuclear-spin qubit is already nearly spin-independent in scalar traps.
4. **Sideband cooling**, Narrow $^3P_1$ line enables direct sideband-resolved cooling. Achieves ground state $\langle n \rangle < 0.1$.
5. **Nuclear spin initialization**, Optical pumping via $^1S_0 \to {}^3P_1$ with circularly polarized light. Fidelity >99%.

Key advantage: The $^1S_0$ ground-state nuclear-spin qubit has almost no electronic magnetic moment and a mostly spin-independent scalar Stark shift. The 759 nm magic condition is most relevant when coherently using $^3P_0$ clock/shelving states.

32.5 Coherence & Decoherence

How long quantum information survives, and the dominant dephasing mechanisms for each platform.

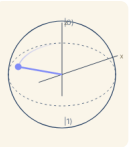
^{87}Rb — schematic dephasing



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

t=0.6 s |r|=0.852

^{171}Yb — schematic dephasing

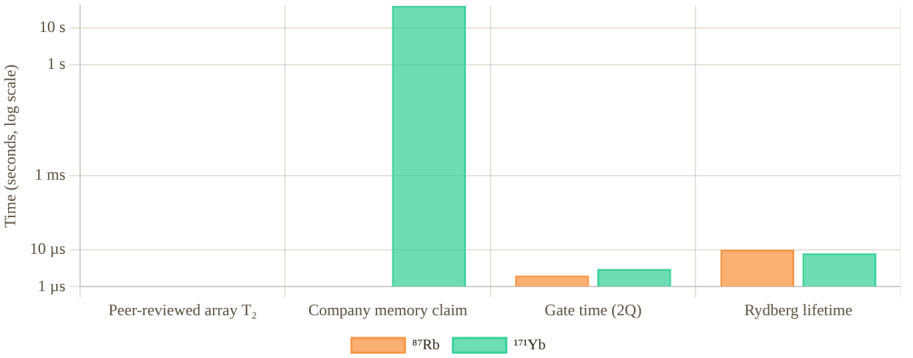


Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

t=0.6 s |r|=0.852

Schematic only: exact T_2 is platform-, encoding-, and sequence-dependent.

Coherence Times (log scale)



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Evidence audit: the 12.6(1) s, 6100-atom benchmark is Cs-133 from Manetsch et al. (Nature 2025), not Rb-87. Rb and Yb coherence numbers below are therefore treated cautiously and tagged by source confidence.

Rb87 Coherence

- **Published Rb array benchmark:** 99.5% CZ gates on up to 60 atom pairs in parallel; coherence during algorithms is sufficient but not the 12.6 s Cs number PEER-REVIEWED
- **Adjacent alkali benchmark:** Cs-133, not Rb-87, reached T_2 = 12.6(1) s in a 6100-atom array NOT RB
- **T_1 (energy relaxation):** >1 s (thermal photon limited)
- **Dominant dephasing:** Magnetic field gradient noise, laser phase noise, differential light shift in tweezer
- **Second-order Zeeman:** 575.15 Hz/G² × B² for the 87Rb clock transition
- **Tweezer scattering:** Off-resonant photon scattering at rate $\Gamma_sc \propto I/\Delta^2$ limits T_1

Rb87 published gate benchmark: $\mathcal{F}_{CZ} = 99.5\%$ on up to 60 parallel atom pairs

Rb87 logical benchmark: 48 logical qubits from 280 physical atoms

The Cs-133 6100-atom coherence result is valuable alkali context, but should not be used as an Rb87 T_2 number.

Yb171 Coherence

- **T_2* (free precession, optical clock):** >100 s demonstrated in optical lattice clocks
- **T_2 / memory:** no peer-reviewed large-array ^{171}Yb T_2 is quoted here. The closest published number is Atom Computing's earlier ^{87}Sr nuclear-spin register: T_2* = 21(7) s, T_2 echo = 42(6) s (Barnes et al., *Nat. Commun.* 13, 2022) PEER-REVIEWED, a sister alkaline-earth species, not a ^{171}Yb number.
- **T_1 (energy relaxation):** Limited by photon scattering from trap, expected comparable to Rb87
- **Dominant dephasing:** Residual differential light shift (trap not perfectly magic), nuclear spin flip from electric quadrupole (suppressed by I=1/2)

- **B-field dephasing:** linear nuclear-spin slope ≈0.75 kHz/G; compare against Rb clock-state quadratic sensitivity at the actual bias field
- **Clock-state advantage:** $^1S_0$–$^3P_0$ magic trapping near 759 nm supports clock/shelving operations; $^1S_0$ nuclear-spin qubit shifts are scalar to leading order

$$\nu_{\mathrm{Yb171}} \approx g_I \mu_N B/h \approx 0.75\ \mathrm{Hz/mG} \times B[\mathrm{mG}]$$

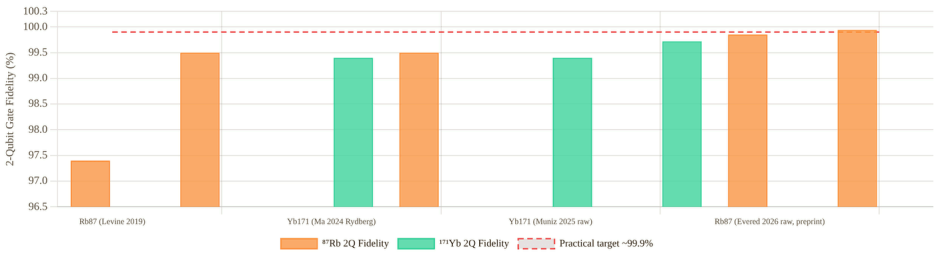
At $B = 1$ G: $\nu_{\mathrm{Yb171}} \approx 0.75$ kHz; Rb87 clock qubit quadratic shift is ≈ 575 Hz
Clock magic wavelength: $\Delta\alpha_{1S0-3P0}(\lambda_{\mathrm{magic}} \approx 759\ \mathrm{nm}) = 0$ for clock/shelving transitions

Bottom line on coherence: Do not compare Rb and Yb using the 6100-atom, 12.6 s number as an Rb result; it is Cs-133. For Rb87, the strongest peer-reviewed evidence is gate/logical-processor maturity. For Yb171, the strongest evidence is the nuclear-spin/clock-state structure and the 2025 gate benchmark. For both, present-day two-qubit gate errors dominate over idle decoherence in most short circuits.

32.6 Gate Performance

Single-qubit and two-qubit gate fidelities, Rydberg excitation physics, and error sources.

Gate Fidelity Comparison (2019–2025)



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Two-qubit gate fidelities, best published values. Rb87: Evered et al. (Nature 2023, peer-reviewed) and Evered et al. (arXiv:2604.25987, 2026, preprint). Yb171: Muniz et al. (PRX Quantum 2025, post-selection and raw). A practical 99.9% target line is shown as a dashed guide.

Rb87 Gate Results

- **Best 2Q CZ fidelity:** 99.5% on 60 qubit pairs simultaneously Evered 2023
- **Parallel entanglement:** 60 Bell pairs in parallel, same 99.5% per pair
- **1Q gate fidelity:** >99.9% (randomized benchmarking, 2024–25)
- **Gate mechanism:** Rydberg blockade CZ gate via two-photon 420+1013 nm excitation to an nS Rydberg state (Evered et al. used $|53S_{1/2}\rangle$)
- **Gate speed:** 0.5–2 μs per two-qubit gate (limited by Rabi frequency)
- **Key innovation:** Optimal-control pulse shaping + atomic dark-state scattering suppression
- **Logical qubit demo:** 48 logical qubits (Bluvstein 2024) and up to 96 logical qubits in 448 atoms (Bluvstein et al., Nature 649, 2026)
- **Array size for gate demos:** Up to 280 physical qubits (Bluvstein 2024)
- **2026 update:** 99.854(4)% raw / 99.941(3)% with loss post-selection, global echo randomized benchmarking averaged over 17 gate sites in a 442-atom zoned array, stable across 10 hours and ~2 million RB runs (Evered et al., arXiv:2604.25987) PREPRINT

$$\mathcal{F}_{\mathrm{CZ}} = 1 - \varepsilon_{\mathrm{Ryd}} - \varepsilon_{\mathrm{scatt}} - \varepsilon_{\mathrm{D}} - \varepsilon_{\Omega}$$

$\varepsilon_{\mathrm{Ryd}} \sim (\Gamma_{\mathrm{Ryd}}\tau_g)$, Rydberg decay error
$\varepsilon_D = \frac{1}{2}(k\, v_{\mathrm{rms}}\, \tau_g)^2$, Doppler dephasing (mitigated by sideband cooling)
Best: 1 − 99.5% = 0.5% total error budget (Evered 2023, peer-reviewed)
2026 preprint: 1 − 99.854% ≈ 0.146% raw error (Evered et al., arXiv:2604.25987, not yet peer-reviewed)

Yb171 Gate Results

- **Best 2Q CZ fidelity:** 99.72(3)% with post-selection (Muniz et al. PRX Quantum 2025) Muniz 2025
- **Without post-selection:** 99.40(3)%, still competitive with Rb87
- **Rydberg gate fidelity (F=1/2 Rydberg):** 99.4(1)% (Ma et al. 2024), 3.3× improvement over earlier results
- **1Q gate fidelity:** >99.5% (Muniz 2025)
- **Gate mechanism:** Rydberg blockade via single UV photon (~302 nm) from $^3P_0$ (metastable) to a Rydberg $^3S_1$ state
- **Gate speed:** ~1–5 μs per two-qubit gate (similar to Rb87)
- **Unique feature:** Can encode qubit in both nuclear spin AND clock state ($^3P_0$), dual encoding protocols
- **Array size for gate demos:** Atom Computing's 1,225-site array (1,180 atoms loaded, 2023 announcement); gate-fidelity demos use smaller subsets

$$\mathcal{F}_{\mathrm{CZ}}^{\mathrm{Yb}} = 99.72\%\ (\text{post-sel}), \quad 99.40\%\ (\text{raw})$$

Nuclear spin qubit: $\varepsilon_B \propto (g_I \mu_N \delta B\, \tau_g/\hbar)^2$, 270× smaller than Rb87
Magic wavelength: $\varepsilon_{\text{light-shift}} \approx 0$, entire error term eliminated

Gate fidelity status (2025–2026): Through 2025, Yb171 held the highest published per-pair two-qubit fidelity (99.72% with post-selection, Muniz et al., peer-reviewed), while Rb87 led on *parallel* gate demonstrations (60 simultaneous pairs at 99.5%, Evered et al. 2023, peer-reviewed) and on logical-processor maturity. In April 2026 the same Rb87 group reported a new record in a preprint — 99.854% raw, 99.941% with loss post-selection — which would exceed Yb171's peer-reviewed number on both raw and post-selected fidelity. That result has not yet been peer-reviewed or independently reproduced, so treat the ranking as provisional until it clears review; Yb171 retains its structural advantages in erasure conversion and non-destructive readout regardless (Sections 07 and 09). The race to consistent 99.9%+ without post-selection, confirmed by independent groups, is the key target for both platforms.

32.7 Readout & SPAM

State preparation and measurement fidelity, the bookend errors of every circuit.

Rb87 Readout

- **Method:** Fluorescence imaging, atoms in |F=2⟩ fluoresce on D2 cycling transition; |F=1⟩ atoms are dark (shelved or blown away)
- **SPAM/readout:** high-fidelity fluorescence readout is well established in Rb arrays; exact numbers depend strongly on protocol and whether loss is included (PROTOCOL-DEPENDENT)
- **Adjacent alkali readout benchmark:** Cs-133 Manetsch 2025 reports 99.98952(1)% imaging survival and >99.99% imaging fidelity; not an Rb number
- **Mid-circuit measurement:** Demonstrated with atom transport (Bluvstein 2022), move ancilla atoms to readout zone, measure, return data atoms
- **Limitation:** Reading destroys the measured atom unless shelving to F=1 is performed before fluorescence
- **Non-destructive readout:** Active area of research; clock-shift measurements explored

Yb171 Readout

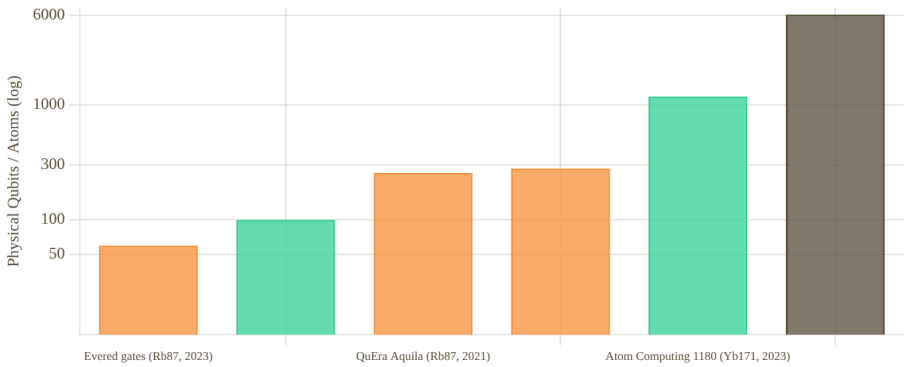
- **Method:** Shelve one spin state to $^3P_0$ (dark), image the other with broad $^1P_1$ transition at 399 nm, scattered photons distinguish $|\uparrow\rangle$ from $|\downarrow\rangle$
- **Repetitive readout:** state discrimination fidelity 0.993(4) with state-averaged survival 0.994(3) in Huie et al. PRX Quantum 2023 PEER-REVIEWED
- **Advantage:** $^3P_0$ metastable state (20 s lifetime) enables truly *non-destructive* qubit storage during readout, image one logical basis without disturbing atoms in the other
- **Clock-state readout:** State-selective fluorescence on the narrow $^3P_1$ transition (556 nm) enables high-SNR readout
- **Mid-circuit measurement:** Structurally enabled by the $^3P_0$ shelving protocol, single-shot QND-like readout of subsets of qubits
- **Dual encoding advantage:** Can store qubits in $^3P_0$ (protected from imaging light) while others are measured

Yb171 structural readout advantage: The $^3P_0$ metastable state (lifetime ~20 s) acts as a long-lived memory state that is far off resonance from the $^1S_0$ imaging transition. By shelving qubits to $^3P_0$ before imaging, Yb171 enables site-selective, low-crosstalk readout that is naturally suited for quantum error correction cycles where only syndrome qubits are measured per round. This is significantly harder to achieve non-destructively in Rb87.

32.8 Array Size & Scalability

Current demonstrations and the path to fault-tolerant arrays of 10,000+ qubits.

Array Size Milestones (physical qubits/atoms)



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Demonstrated physical qubit/atom counts. Rb87 milestones emphasize published gate and logical-processor results. Atom Computing's 1180-qubit Yb result is company-reported. The 6100-atom Manetsch benchmark is Cs-133, shown only as adjacent alkali context.

Rb87 Scaling

- **2021:** QuEra Aquila, 256 atoms (analog mode)
- **2022:** Bluvstein et al., 24-qubit toric code, surface code 19 qubits
- **2023:** Evered et al., 60 qubits with 99.5% 2Q gates
- **2024:** Bluvstein et al., 280 physical qubits; 48 logical qubits in [[8,3,2]] codes (plus surface- and color-code logic)
- **2026:** Bluvstein et al., *Nature* 649, 448 physical atoms; up to 96 logical qubits in a fault-tolerant architecture
- **2025 adjacent alkali:** Manetsch et al. demonstrated 6100 coherent Cs-133 atoms, not Rb87; useful scaling context but not counted as an Rb milestone

- **2026 (preprint):** Evered et al. benchmarked the new CZ-gate record on a 442-atom zoned array (17 gate sites) — a fidelity/stability result, not a new array-size record
- **Rearrangement:** AOD-based atom sorting achieves >99% defect-free loading
- **Roadmap:** 10,000+ qubit arrays, parallel Rydberg gates, photonic interconnects

Yb171 Scaling

- **2022:** Jenkins et al., 10×10 array, 92.73% filling efficiency
- **2023:** Atom Computing, 1,180 atoms loaded in a 1,225-site array (company announcement)
- **2024 (preprint):** Microsoft + Atom Computing, 24 entangled logical qubits; Bernstein–Vazirani on 28 logical qubits (arXiv:2411.11822)
- **2025:** Muniz et al. (Atom Computing collab.), PRX Quantum, 99.72% gate fidelity (post-sel)
- **Gate fidelity at scale:** Not yet demonstrated at 1000+ qubits simultaneously
- **Rearrangement:** Similar AOD-based sorting; magic-wavelength tweezer enables defect-free loading
- **Roadmap:** Scale to 10,000+ qubits, demonstrate surface code, integrate photonic network

> **Current gap:** Rb87 leads this comparison in peer-reviewed logical-processor maturity (48 logical qubits in 2024; up to 96 logical qubits in 448 atoms in Bluvstein et al., Nature 649, 2026) and parallel gate demonstrations (60 pairs at 99.5%). Yb171 leads in the best reported per-pair CZ number when post-selection is allowed (99.72(3)%; raw 99.40(3)%) and has structural erasure/shelving advantages. A major published coherent neutral-atom array benchmark is Cs-133 (6100 atoms), not Rb87 or Yb171.

32.9 Erasure Conversion, Yb171's Structural Advantage

A new paradigm for quantum error correction that alkaline-earth atoms enable especially naturally.

In standard quantum error correction, errors are *Pauli errors* (X, Y, Z) that the code must detect and correct. Erasure errors, errors where we know *which qubit* failed, are much cheaper to correct: an [[n,k,d]] code corrects d-1 erasures but only ⌊(d-1)/2⌋ Pauli errors. If we can *convert* Pauli errors into known erasures, the overhead for fault tolerance drops dramatically.

Wu, Kolkowitz, Puri et al. (Nature Communications 2022) showed that for Yb171 and Sr87, **~98% of gate errors can be converted to erasures** using the $^3P_0$ metastable state. When a qubit undergoes an error during a Rydberg gate, it often ends up in the $^3P_0$ state (detectable as "atom is in wrong sublevel") rather than a Pauli error. Detecting whether an atom is in $^3P_0$ vs $^1S_0$ after each gate round , which is straightforward with site-selective imaging, heralds the error location without collapsing the logical qubit.

$$\text{Standard QEC:} \quad d_{\text{eff}} = d \quad \rightarrow \quad \varepsilon_{\text{logical}} \propto \left(\frac{p}{p_{\text{th}}}\right)^{\lceil d/2 \rceil}$$

$$\text{Erasure QEC:} \quad \text{corrects up to } d-1 \text{ erasures} \quad \rightarrow \quad \varepsilon_{\text{logical}} \propto \left(\frac{p}{p_{\text{th, erasure}}}\right)^{d}$$

With 98% error-to-erasure conversion in the Wu et al. Yb/Sr proposal: simulated surface-code threshold increases from 0.937% to 4.15%.
Practical consequence: a distance-5 surface code with erasures ≈ distance-9 surface code with Pauli errors.

Yb171 Erasure Protocol

- After each gate round, check if atom is in $^1S_0$ (ground) or $^3P_0$ (metastable) via state-selective fluorescence
- Atom in $^3P_0$ = detected error → erasure flag. Atom in $^1S_0$ = passed gate correctly
- The $^3P_0$ state is stable on the timescale of gate operations (τ ~ 20 s), no spontaneous decay during check
- 98% of dominant gate errors (Rydberg decay, Doppler) map to $^3P_0$ state
- Only 2% are "dark" Pauli errors that the erasure protocol misses

- Result: up to ~10× reduction in physical qubit overhead for fault-tolerant computation

Rb87, Limited Erasure

- In Rb87, gate errors are mostly Rydberg decay or atomic motion errors, they deposit the atom back in the ground state without a detectable "flag"
- No equivalent long-lived metastable state accessible from the qubit manifold to flag errors
- Some error conversion possible via Rydberg state monitoring (atom loss detection), but less efficient than Yb171
- Atom loss from the trap is detectable as an erasure (~99% efficiency for atom-loss erasure detection)
- But Pauli errors from coherent control noise are not detectable → remain as standard errors
- **Net erasure fraction:** Architecture-dependent and not a universal published Rb number; atom loss is heraldable, but most coherent Pauli errors are not naturally converted to erasures ROUGH ESTIMATE

2025–2026 update: erasure conversion demonstrated end to end. Zhang, Liu, Bornet, Horvath, Peng, Ma, Huang, Puri & Thompson (Princeton, the same group as Ma et al. 2023) published the first experimental chain showing erasure information helping at every stage, from physical gate to logical teleportation (*Nature Physics*, 2026; arXiv:2506.13724, 2025):

Anti-refocusing (AR) two-qubit gate: $\mathcal{F} = 98.4(1)\%$, vs. $98.9(1)\%$ for a conventional time-optimal gate
AR gate error scales as $(\Delta I/I_0)^4$ instead of $(\Delta I/I_0)^2$ (far less sensitive to laser-intensity noise); detects **38(6)%** of its own errors as erasures mid-circuit
Coherent transport: <6×10⁻⁴ dephasing+bit-flip probability per one-way trip (~2.5 ms); atom loss 1.1(4)%, about half caught as a detectable erasure
[[4,2,2]] logical state-prep fidelity: 87.4(4)% unconditional → 90.2(3)% using erasure info → 98.1(2)% with flag post-selection → **99.5(1)%** with flag *and* erasure post-selection (logical error rate falls 3.6(1)× faster with erasure information)
Logical-qubit teleportation fidelity: 77.1(9)% (random ancilla, syndrome post-selection) → 87(2)% (erasure post-selection) → 80.2(8)% (adaptive, erasure-informed ancilla selection)

Companion result, same issue: Senoo et al. (JILA/Kaufman group), also *Nature Physics* June 2026, report a 99.78(4)% error-detected two-qubit gate fidelity using metastable-encoded qubits — a different qubit encoding from the ground-state hyperfine number quoted elsewhere on this page (Muniz et al., 99.72%), so the two are not directly interchangeable, but both point the same direction: erasure/metastable-state tricks are paying off for Yb171 across multiple groups.

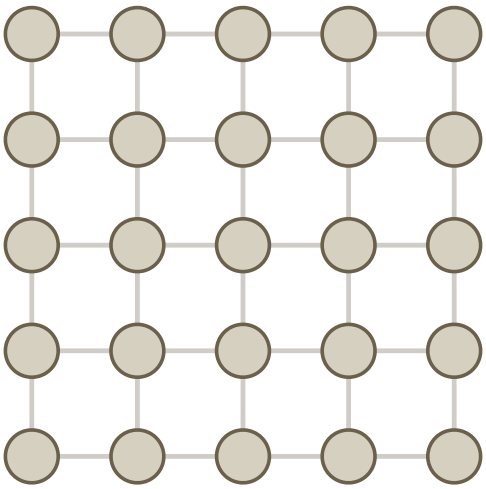
This is the experimental version of the Wu et al. 2022 proposal above: not just a theoretical error-to-erasure conversion rate, but a demonstrated pipeline where knowing *which* qubit failed measurably improves every downstream step — gate, transport, logical encoding, and teleportation — using the same $^3P_0$ metastable-state shelving trick that gives Yb171 its readout advantage (Section 07).

Long-term significance: Erasure conversion may be the most important structural advantage Yb171 has over Rb87 for large-scale fault-tolerant computation. A 10× reduction in qubit overhead means the difference between needing 10,000 physical qubits per logical qubit vs 1,000. Over millions of physical qubits, this could represent years of roadmap advantage.

Throw errors at a code

A distance-d code is guaranteed to fix any mix of **e** erasures (errors whose location is flagged) and **t** hidden errors while **e** + **2t ≤ d − 1**. Each error you throw is flagged with 98% probability, as in the Yb171/Sr87 erasure-conversion proposal cited above. The same errors are then scored with and without the flags.

Code distance



With erasure flags Yb171-style: 98% of gate errors land in $^3P_0$ and are seen

e + 2t = 0 + 2·0 = **0** of 4 → guaranteed correctable

No flagsevery error is a hidden Pauli error

e + 2t = 0 + 2·0 = **0** of 4 → guaranteed correctable

Legend: flagged erasure (location known) · hidden error (location unknown to the decoder). The no-flag case is not labelled Rb87 on purpose: Rb87 atom loss is heraldable, but most coherent errors are not, and there is no single published Rb87 erasure fraction. Beyond the guarantee, correction *may* fail; it doesn't always.

32.10 Full Head-to-Head Comparison

Every major dimension, side by side. Advantage marked in bold color.

Tug of war

Every row of the table below ends with an *Edge* call. Pick the categories you care about and watch the rope move.

Rb87 8

Yb171 12

Across these categories the table leans **Yb171** (12 vs 8, 3 even or different).

Rb87 edge · 9

Even ⁄ different · 3

Yb171 edge · 12

A tally of this page's own *Edge* column, not a score: rows are not equally important, which is exactly why you choose. Conditional calls ("if it scales", "pending review") count half and are hatched. Click any row to jump to it in the table.

CATEGORY	METRIC	^{87}RB	^{171}YB	EDGE

Atomic Physics	Electronic structure	Alkali, 1 valence e^-, $5S_{1/2}$ ground	Alkaline-earth-like, 2 val. e^-, ^1S_o ground	Yb J=0 protection
Nuclear spin	I = 3/2	I = 1/2	Yb simpler	
Ground-state J	J = 1/2 (electronic moment)	J = 0 (no electronic moment)	Yb	
Clock transition	None (microwave HF, 6.834 GHz)	578 nm ($^1S_o \rightarrow {}^3P_o$, <10 mHz)	Yb	
Qubit Encoding	Qubit type	Hyperfine (F=1↔F=2, 6.834 GHz)	Nuclear spin (ml=±½, ~kHz scale)	Tie
B-field sensitivity	~0 first-order; quadratic 575.15 Hz/G² × B² (clock states mF=0)	~0.75 Hz/mG (linear, nuclear magneton scale)	Yb at B > 0.65 G; Rb at lower fields	
Light-shift cancellation	Requires active calibration	^1S_o nuclear-spin scalar shift nearly spin-independent; 759 nm is clock magic for ^1S_o–^3P_o	Yb	
Coherence	T_2 (best published)	No 6100-atom Rb T_2 claim; use Rb gate/logical benchmarks instead	No large-array ^{171}Yb T_2 quoted; sister species ^{87}Sr: T_2 echo 42(6) s (Barnes 2022) PEER-REVIEWED	Yb (structural)
T_2* (free precession)	Not quoted here without an Rb87 array source	Not fully characterized in tweezers	No direct comparison	
Primary dephasing	B-field noise, light shift	Magnetic noise plus residual tensor/vector/scattering effects; clock-state operations benefit from magic trapping	Yb	
Gate Performance	Best 2Q fidelity, peer-reviewed	99.5% raw, parallel (Evered 2023)	99.72% with post-sel (Muniz 2025)	Yb (slightly)
Best 2Q fidelity, incl. preprints PREPRINT	99.941% with post-sel, 99.854% raw (Evered 2026)	99.72% with post-sel (Muniz 2025, peer-reviewed)	Rb if the 2026 preprint holds up under review	
2Q fidelity raw	99.5% (2023, peer-reviewed, 60 pairs); 99.854% (2026, preprint)	99.40% (no post-selection)	Rb (widening, pending peer review)	
Parallel gates	60 qubit pairs simultaneously at 99.5%	Not yet demonstrated at scale	Rb	
Gate speed	0.5–2 µs (2Q Rydberg)	1–5 µs (2Q Rydberg)	Rb (slightly)	
Readout	SPAM/readout fidelity	High; protocol-dependent. Cs 6100-array imaging benchmark is not Rb	0.993(4) state discrimination; 0.994(3) survival (Huie 2023)	Protocol-dependent
Non-destructive readout	Limited, active research	^3P_o shelving enables partial QND readout	Yb	
Scalability	Largest Rb/Yb array benchmark	280 physical atoms in logical-processor demo; 256-atom analog arrays	1,180 qubits (Atom Computing 2023)	Rb peer-reviewed logic; Yb company count
Logical qubit demos	Up to 96 logical qubits (Bluvstein 2026; 48 in 2024, peer-reviewed)	24 entangled logical qubits (48 atoms) and Bernstein–Vazirani on 28 logical qubits (112 atoms), Microsoft + Atom Computing (Reichardt et al., arXiv:2411.11822) PREPRINT	Rb	
Error Correction	Erasure conversion fraction	Architecture-dependent; no universal Rb fraction	~98% (Wu et al. 2022)	Yb major
Qubit overhead reduction	Standard surface code overhead	Surface-code threshold 0.937% → 4.15% in Wu et al. simulation	Yb	

Laser System	Complexity	Simpler: D2 at 780 nm (diode laser)	More complex: 399, 556, 578, 759 nm lasers	Rb
Rydberg excitation lasers	420 nm + 1013 nm for Harvard/QuEra gates; 780 + 480 nm is a common alternative Rb route	~302 nm UV (single photon from $^3P_0$); harder than Rb blue/IR two-photon systems	Rb	
Maturity	Platform maturity	10+ years of tweezer work; most-cited demos	~5 years in tweezers; rapid progress 2022–25	Rb

32.11 Companies & Ecosystem

Who is building what, and how the two platforms map to commercial quantum computing ventures.

^{87}Rb Ecosystem

QuEra Computing ↗

Harvard/MIT spin-out (Lukin/Greiner groups). Aquila (256 atoms, 2021). Cloud access via AWS Braket. Dual analog+digital capability. Key partner for the Harvard group's gate fidelity demonstrations. Leading neutral-atom company by publication record and cloud reach.

Pasqal ↗

French spin-out from Browaeys/Lahaye group (IOGS, Paris). Fresnel processor (100 Rb atoms, 2D programmable). Hybrid analog-digital. Enterprise partnerships: EDF, BASF, Thales. Merged with Qu&Co (2022). Focus: optimization, combinatorial problems, HPC integration.

Infleqtion (ColdQuanta) ↗

Uses Rb (and Cs). Acquired SuperTech. Cloud quantum and quantum networking products. Boulder, Colorado-based.

Harvard/MIT Academic Groups

Lukin group (Harvard): All major Rb gate fidelity/logical qubit demonstrations (Evered 2023, Bluvstein 2024). Greiner group (Harvard): Quantum simulation. Vuletic group (MIT): Tweezer array methods. Most of the landmark papers originate here and license to QuEra.

Endres Group, Caltech

Manuel Endres group at Caltech. Published the 6100-atom **Cs-133** array with T_2=12.6(1) s (Manetsch et al., Nature 2025), a major coherent-array benchmark. Included here as adjacent alkali scaling context, not as an Rb87 result.

^{171}Yb Ecosystem

Atom Computing ↗

Headquartered in Berkeley, CA, with its main R&D site in Boulder, CO. First generation (Phoenix) used ^{87}Sr nuclear-spin qubits (T_2 echo 42 s, Barnes et al. 2022); the 2023 system moved to ^{171}Yb with 1,180 atoms loaded in 1,225 sites (company announcement). With Microsoft: 24 entangled logical qubits and a 28-logical-qubit Bernstein–Vazirani run (arXiv:2411.11822, preprint). Focus: fault-tolerant digital QC with mid-circuit measurement and scalable arrays.

Kaufman Group (JILA/CU Boulder)

Adam Kaufman's group at JILA was the first to demonstrate nuclear-spin Yb171 qubits in optical tweezers (Jenkins et al. PRX Quantum 2022). Key academic partner for Atom Computing. Also works on Sr87 and other alkaline-earth species. The 99.72% gate fidelity result (Muniz et al. PRX Quantum 2025) is an Atom Computing collaboration with JILA and other groups. The group's own Senoo et al. (Nature Physics, June 2026) separately reported 99.78(4)% error-detected fidelity on metastable, erasure-encoded qubits.

NIST / Optical Clock Groups

Several NIST groups (Jun Ye, Andrew Ludlow) have deep Yb171 expertise from optical lattice clock work. The ultranarrow clock transition knowledge transfers directly to tweezer QC. Coherence >100 s demonstrated in optical lattice clocks informs Atom Computing's coherence roadmap.

Also: Sr87 (similar physics)

Fermi National Lab (Jonathan Simon group), Stanford (M. Endres group). Sr87 has I=9/2 (more complex spin than Yb171's I=1/2) but the same J=0 ground state protection and erasure conversion advantage. Nuclear clock groups exploring qubit applications.

32.12 Key Papers & References

Essential reading, all major experimental demonstrations, theory papers, and review articles.

^{87}Rb, Landmark Papers

High-fidelity entangling gates and nonlocal circuits with neutral atoms PREPRINT

Evered, Xu, Li, Geim, Bonilla Ataides, Kalinowski, Bluvstein, Maskara, Kokail, Greiner, Vuletić, Lukin, arXiv:2604.25987 (2026)

99.854(4)% raw / 99.941(3)% with loss post-selection CZ gate fidelity, global echo randomized benchmarking averaged over 17 gate sites in a 442-atom zoned array, stable across a 10-hour, ~2-million-run benchmark. Also demonstrates nonlocal circuits via coherent atom rearrangement. Not yet peer-reviewed; treat as provisional pending journal publication and independent reproduction.

arXiv:2604.25987 ↗

High-fidelity parallel entangling gates on a neutral-atom quantum computer

Evered, Bluvstein, Kalinowski et al., Nature 622, 268–272 (2023)

99.5% CZ gate fidelity on 60 qubit pairs simultaneously using Rydberg blockade with dark-state-assisted scattering reduction. Crossed the NISQ/fault-tolerance boundary for neutral atoms.

nature.com ↗ · arXiv:2304.05420 ↗

Logical quantum processor based on reconfigurable atom arrays

Bluvstein, Evered, Geim et al., Nature 626, 58–65 (2024)

280 physical qubits, 48 logical qubits encoded in transversal color codes. Fault-tolerant operations with mid-circuit measurement and feed-forward. First demonstration of logical gate advantage in neutral atoms.

nature.com ↗ · arXiv:2312.03982 ↗

A large-scale quantum processor with fast, full-coherence operations

Manetsch, Nomura, Bataille et al., Nature 647, 60–67 (2025)

6100-atom **Cs-133** array on a 12,000-site grid. T_2 = 12.6(1) s, 99.98952(1)% imaging survival, and >99.99% imaging fidelity. This is not an Rb87 result; it is listed as adjacent alkali/tweezer scaling context.

nature.com ↗ · arXiv:2403.12021 ↗

Coherent transport and entanglement generation using rearrangeable atom arrays

Bluvstein et al., Nature 604, 451–456 (2022)

Mid-circuit measurement and feed-forward via atom transport. Surface code on 19 qubits, toric code on 24 qubits. First programmable neutral-atom quantum error correction demonstration.

nature.com ↗ · arXiv:2112.03923 ↗

Parallel implementation of high-fidelity multiqubit gates with neutral atoms

Levine, Keesling, Semeghini et al., PRL 123, 170503 (2019)

First high-fidelity (97.4%) parallel Rydberg gates on Rb87 qubits. Introduced the basis for the Harvard group's Rydberg gate protocol.

PRL ↗ · arXiv:1908.06101 ↗

256-atom quantum simulator (QuEra Aquila)

Ebadi, Wang, Levine et al., Nature 595, 227–232 (2021)

256-qubit analog Rb87 quantum simulator studying frustrated magnetism and quantum phase transitions. First 200+ qubit neutral-atom demonstration. Basis for QuEra Aquila product.

nature.com ↗

^{171}Yb, Landmark Papers

Ytterbium nuclear-spin qubits in an optical tweezer array

Jenkins, Lis, Senoo, McGrew, Kaufman, Phys. Rev. X 12, 021027 (2022)

First demonstration of Yb171 nuclear spin qubits in optical tweezers. 10×10 array, 92.73% filling, sub-100 ns single-qubit manipulation. Founded the Atom Computing roadmap. Kaufman group, JILA/CU Boulder.

PRX ↗ · arXiv:2112.06732 ↗

High-fidelity universal gates for ^{171}Yb nuclear-spin qubits

Muniz, Stone, Stack et al. (Atom Computing collaboration, 53 authors), PRX Quantum 6, 020334 (2025)

99.72(3)% two-qubit CZ gate fidelity with post-selection, 99.40(3)% raw. Universal gate set with individually controlled and parallel single- and two-qubit gates. Demonstrates >200 CZ gates on atom pairs. Establishes Yb171 ground-state nuclear spin qubit as a top-tier gate platform.

PRX Quantum ↗ · arXiv:2411.11708 ↗

Rydberg excitation spectroscopy of ^{171}Yb and improved Rydberg gate fidelity

Ma, Peper, Li, Knapp, Liu, Peng, Zhang, Horvath, Burgers, Thompson (Princeton), arXiv:2406.01482 (2024)

CZ fidelity 99.4(1)% using F=1/2 Rydberg series (3.3× better than F=3/2). MQDT spectroscopy of ^{171}Yb Rydberg states. Single-photon UV (~302 nm) excitation from $^3P_0$ to Rydberg states. Thompson group, Princeton.

arXiv:2406.01482 ↗

Erasure conversion for fault-tolerant quantum computing with alkaline-earth atoms

Wu, Kolkowitz, Puri, Thompson, Nature Commun. 13, 4657 (2022)

Theoretical framework showing that 98% of gate errors in Yb171/Sr87 can be converted to detectable erasures using the $^3P_0$ metastable state. ~10× reduction in fault-tolerance overhead. Highly cited theory paper.

Nature Commun. ↗

Logical qubits with erasure conversion using metastable neutral atoms

Zhang, Liu, Bornet, Horvath, Peng, Ma, Huang, Puri, Thompson, Nature Physics (2026); arXiv:2506.13724 (2025)

Experimental follow-up to Wu et al. 2022: an anti-refocusing two-qubit gate at 98.4(1)% fidelity detects 38(6)% of its own errors as erasures; coherent transport loses $<6\times10^{-4}$ per trip with about half of atom loss caught as erasure. A [[4,2,2]] logical code reaches 99.5(1)% state-prep fidelity with flag-and-erasure post-selection (87.4(4)% unconditional), and logical teleportation reaches 87(2)% fidelity with erasure post-selection. First demonstration that erasure information measurably helps end to end, from physical gate to logical teleportation.

nature.com ↗ · arXiv:2506.13724 ↗

High-fidelity entanglement and coherent multi-qubit mapping in an atom array

Senoo, Baumgärtner, Lis, Vaidya, Zeng, Giudici, Pichler, Kaufman, Nature Physics (2026); arXiv:2506.13632 (2025)

Companion result from the same June 2026 *Nature Physics* issue as Zhang et al. above, from the JILA/Kaufman group (same group as Jenkins et al. 2022, first row of this list). Reports an **error-detected** two-qubit gate fidelity of 99.78(4)% using metastable, erasure-encoded qubits — a different encoding from Muniz et al.'s ground-state hyperfine result below, so treat it as a complementary erasure-conversion data point, not a direct replacement for the 99.72% headline number.

nature.com ↗ · arXiv:2506.13632 ↗

Atom Computing: 1,225-site ^{171}Yb nuclear-spin array

Atom Computing, Technical whitepaper + announcement (2023)

Industry announcement: 1,225 tweezer sites, 1,180 atoms loaded, nuclear-spin qubits. A company announcement, not a peer-reviewed benchmark. (Phoenix, the earlier ~100-qubit machine, used ^{87}Sr; its 42 s T_2 echo is in Barnes et al., *Nat. Commun.* 13 (2022).)

atom-computing.com ↗

Two-qubit encoding with ^{171}Yb: clock + nuclear spin dual encoding

Multiple authors, npj Quantum Information (arXiv:2402.13134, 2024)

Proposes encoding logical qubits in both the nuclear spin (fast gates) and clock-state ($^3P_0$, long coherence) of Yb171. Dual-encoding enables gate operations followed by coherent storage, relevant for fault-tolerant architectures.

arXiv:2402.13134 ↗

Reviews & Theory

Tweezer arrays of individual atoms for quantum applications

Kaufman & Ni, Nature Physics 17, 1324–1333 (2021)

Comprehensive review of optical tweezer approaches to neutral-atom quantum computing. Covers both alkali and alkaline-earth platforms. Excellent entry point for the field.

nature.com ↗

Quantum computing with neutral atoms

Henriet et al., Quantum 4, 327 (2020)

Review of neutral atom quantum computing platforms, algorithms, and near-term applications. Covers DiVincenzo criteria, Rydberg blockade physics, and error correction prospects for neutral atoms.

quantum-journal.org ↗

CHAPTER 32 SUMMARY

^{87}Rb has the longest record of published logical-processor demonstrations; ^{171}Yb adds nuclear-spin coherence and erasure conversion, with logical-qubit results of its own (Microsoft and Atom Computing). Next, we examine the analogous architectural split in trapped ions.

CHAPTER 33 · CAREER STEP 05

33

Trapped-Ion Architectures: Quantinuum vs. IonQ

🔗 Quantinuum vs IonQ · Deep dive

The trapped-ion showdown between the two listed trapped-ion companies and their two bets on scale: Quantinuum's laser-driven QCCD racetrack versus IonQ's chip-integrated Electronic Qubit Control. Architecture, gate fidelity, logical qubits and roadmaps first; markets and money last.

> *Trapped ions share neutral atoms' atomic precision and long coherence, but two leading companies—Quantinuum and IonQ—pursue sharply different scaling and control architectures.*

IN THIS CHAPTER

Prerequisites: Ch. 26, Ch. 31

- Compare QCCD ion-shuttling architectures (Quantinuum H-series / Helios) with multi-core photonic-interconnect approaches (IonQ).
- Evaluate peer-reviewed gate fidelities, logical-qubit demonstrations, and commercial deployment metrics.
- Understand how motional mode crowding, optical delivery, and photonic link rates govern ion-trap scaling.

WEBSITE COMPANION Live interactive page ↗

33.1 Why These Two Companies?

The two dominant, publicly-traded trapped-ion quantum computing companies, and why they are worth comparing directly.

Trapped-ion qubits are individually charged atoms confined by electromagnetic fields and manipulated with laser or microwave/RF pulses. The approach has long held the industry's best gate fidelities and full qubit-to-qubit connectivity, at the cost of gate speed and, historically, a very complex laser system for every additional qubit. Two companies dominate the commercial trapped-ion field: **Quantinuum** (formed in 2021 from the merger of Honeywell Quantum Solutions and Cambridge Quantum, majority-owned by Honeywell, IPO'd on Nasdaq in June 2026 under ticker **QNT**) and **IonQ** (spun out of Duke University and the University of Maryland in 2015, publicly traded on the NYSE since a 2021 SPAC merger under ticker **IONQ**). Both are racing toward fault-tolerant quantum computing this decade, but via structurally different bets on how ion-trap hardware should scale.

The core trade-off: Quantinuum has, so far, consistently held the best *peer-reviewed* trapped-ion gate fidelities: its Helios system reported 99.921% two-qubit gate fidelity and 99.9975% single-qubit fidelity in a June 2026 *Nature* paper with over 200 co-authors including Sandia National Laboratories, on a 98-qubit, all-to-all-connected QCCD ("quantum charge-coupled device") processor that still uses laser-driven gates and physical ion shuttling. IonQ, meanwhile, is betting on a fundamentally different manufacturing path: replacing laser-array beam steering with **Electronic Qubit Control (EQC)** — control electronics integrated directly onto the ion-trap chip using standard semiconductor fabrication — acquired via its 2025 purchase of Oxford Ionics and now built at its own SkyWater Technology fab. A component-level EQC demonstration reported >99.99% two-qubit fidelity on a single ion pair in an October 2025 preprint, and in September 2026 IonQ launched Superion 256, a 256-qubit EQC-based system designed to be manufactured "by the hundreds" rather than built one at a time in a lab.
The tension this page is built around: Quantinuum currently leads on demonstrated, peer-reviewed, full-system fidelity; IonQ is betting that chip-based manufacturing will win the race to scale, even if its highest fidelity numbers so far are component-level or roadmap claims rather than full-system, peer-reviewed results.

99.921%
Quantinuum Helios 2Q fidelity (Nature 2026, peer-reviewed)

48
Error-corrected logical qubits (Helios, 2:1 ratio; company / preprint)

99.9%
IonQ Tempo 2Q fidelity (company spec, full system)

256
IonQ Superion 256 physical qubits (Sep 2026)

99.99%
IonQ EQC 2Q fidelity (pair) arXiv:2510.17286 — preprint

Calculator output at the page's default settings.

> **Source-confidence rule for this page:** PEER-REVIEWED numbers come directly from journal papers; PREPRINT numbers are posted but not yet peer-reviewed; COMPANY ROADMAP numbers are product specs, press releases, or future targets, not independent physics results. IonQ's headline 99.99% fidelity is a *two-ion, component-level* preprint result, not a full multi-qubit system number — that distinction is kept explicit everywhere it appears on this page.

Who leads where, at a glance

Over the whole table the calls are **roughly even** (Quantinuum 6, IonQ 6, 10 even or different), but the categories matter more than the total. Each card counts the *Edge* calls in one block of the full comparison table; click one to see its rows.

Conditional calls ("if it scales", "pending review") count half. Rows are not equally important: read the rows before quoting a winner.

33.2 Architecture & Qubit Control

Both companies trap ions in a QCCD architecture — but they address and gate those ions in very different ways.

🟣 Quantinuum: Laser-Driven QCCD

Quantinuum's H-series uses a **quantum charge-coupled device (QCCD)** architecture: ions are held in a linear or racetrack-shaped trap with distinct zones for storage, gates, and readout, and are physically *shuttled* between zones using time-varying electric fields. Two-qubit gates are driven optically, with tightly focused, individually-steered laser beams addressing pairs of ions once they are transported into a gate zone.

- **H1/H2 generation:** $^{171}Yb^{+}$ hyperfine qubits, linear/racetrack QCCD trap COMPANY
- **Helios generation (2025–):** switched the qubit species to $^{137}Ba^{+}$ hyperfine qubits, with co-trapped $^{171}Yb^{+}$ ions used purely as a sympathetic-cooling species PEER-REVIEWED
- **Helios trap geometry:** a rotatable ion-storage ring connected to two gate/operation regions by a four-way junction, giving full all-to-all connectivity without a linear-chain qubit-count penalty
- Gate control hardware: acousto-optic and electro-optic beam-steering systems must scale roughly with qubit count and gate-zone count — the laser system is the dominant engineering bottleneck as qubit number grows
- Advantage: laser gates are extremely well-characterized after 15+ years of AMO ion-trap physics; very high fidelities are routine ESTABLISHED TECHNIQUE

🔵 IonQ: Electronic Qubit Control (EQC)

IonQ's first four hardware generations (Harmony, Aria, Forte, Tempo) also used QCCD-style trapping with laser gates, steered by acousto-optic deflectors (AODs), on $^{171}Yb^{+}$ ions with all-to-all or many-to-many connectivity. Starting with its sixth-generation **Superion** platform, IonQ replaced laser beam-steering with control electronics fabricated directly onto the trap chip.

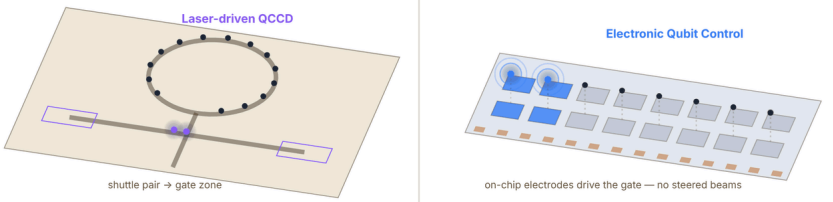
- **Forte / Forte Enterprise / Tempo (2023–2026):** $^{171}Yb^{+}$ qubits, AOD laser gate control, all-to-all (Forte) or many-to-many (Tempo) connectivity COMPANY SPEC

- **Electronic Qubit Control (EQC):** semiconductor-fabricated electrodes drive individual-ion gates directly, eliminating most of the free-space laser-steering optics that otherwise must grow with qubit count; acquired via the ~$1.075B purchase of Oxford Ionics, closed September 2025
- **Manufacturing:** EQC chips are fabricated at SkyWater Technology, a Minnesota semiconductor foundry IonQ agreed to acquire for ~$1.8B in January 2026 (shareholder-approved May 2026); IonQ reports design-cycle time fell from 9 months to 2 months and wafer-lot output rose 12× over six months after the acquisition COMPANY
- **Superion 256 (Sept 2026):** first EQC-based commercial system, 256 qubits, a standard server-rack footprint, and power draw below a typical GPU rack — explicitly designed to be "built by the hundreds," i.e. as a manufacturable product rather than a one-off lab machine

> **Why this is the real story of this comparison:** lasers scale linearly-to-superlinearly in complexity with qubit count — more ions and more parallel gate zones require more individually-steered beams, more optical power, and more alignment stability. Electronic control scales more like a semiconductor chip: once a design is proven, it can in principle be stamped out on a wafer. Quantinuum is still the fidelity leader with the mature, laser-based approach; IonQ is explicitly betting that trading a small early fidelity gap for chip-style manufacturability is the faster path to the 10,000+ qubit systems fault tolerance requires. Neither bet is resolved yet — this page tries to keep the demonstrated present and the roadmap future clearly separated.

Two ways to talk to an ion

The whole comparison in one picture. Left: ions wait in a storage ring, a pair is shuttled through a junction into a gate zone, and steered laser beams drive the gate. Right: the ions float above a chip and the electrodes underneath drive the gate, with no per-gate beam steering.



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

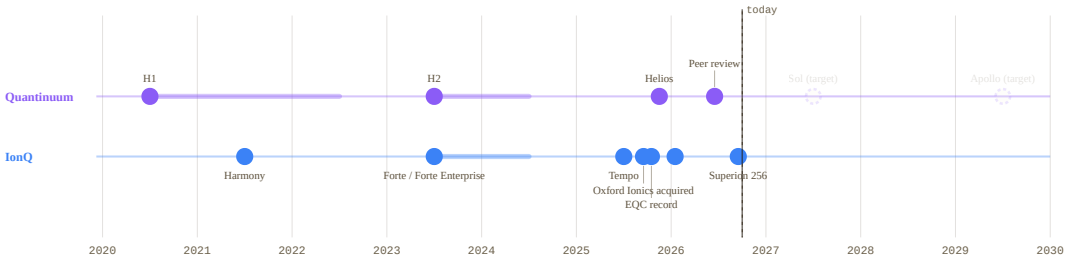
Schematic, not to scale. Left follows the page's description of Quantinuum Helios (rotatable storage ring, four-way junction, gate regions, laser gates); right follows IonQ's Electronic Qubit Control (gates driven by electrodes on the trap chip, removing most free-space beam-steering optics). Ion counts are illustrative.

33.3 Hardware Generations, Side by Side

Both companies iterate fast. Here is the generation-by-generation timeline, with what's demonstrated versus roadmap clearly marked.

Race replay

Both timelines below on one clock. Drag the playhead or press play; hollow markers are company roadmap targets, not results.



QUANTINUUM BY 2026

Peer review

Jun 2026 — Helios physical-layer results (Ransford, Bohnet et al., 200+ co-authors incl.

IonQ by 2026 **Superion 256** Sep 2026 — First EQC-based commercial platform launches: 256 qubits, server-rack footprint, sub-GPU-rack power draw.

Quantinuum H-Series

2020–2022 · H1

First commercial H-series systems (H1-1, H1-2); up to 20 qubits, linear QCCD trap, $^{171}Yb^{+}$ qubits.

2023–2024 · H2

56 fully-connected qubits (racetrack QCCD), >99.99% single-qubit and >99.9% two-qubit gate fidelity (company-quoted, consistent with arXiv:2305.03828); quantum volume 2^{25} = 33,554,432 reported September 2025.

Nov 2025 · Helios

Commercial launch: 98 qubits, $^{137}Ba^{+}$ qubits with $^{171}Yb^{+}$ sympathetic cooling, rotatable-ring QCCD with a four-way junction. Company claimed 94 entangled logical qubits at break-even and 48 fully error-corrected logical qubits at a 2:1 physical:logical ratio.

Jun 2026 · Peer review

Helios physical-layer results (Ransford, Bohnet et al., 200+ co-authors incl. Sandia) published in *Nature*: 99.9975% 1Q / 99.921% 2Q / 99.967% SPAM fidelity.

2027 (target) · Sol ROADMAP

Next-generation system on the path to fault tolerance; physical qubit counts expected to keep scaling incrementally from Helios.

2029 (target) · Apollo ROADMAP

Target: thousands of physical qubits, hundreds of logical qubits, universal fault-tolerant operation. Quantinuum's own messaging shifted this target from "by 2030" (2025 accelerated-roadmap release) to "by 2029" (later 2025 blog post) — both are company projections, not demonstrated results.

IonQ Product Line

2021 · Harmony

IonQ's first-generation cloud-accessible system; publicly traded on NYSE (IONQ) the same year via SPAC merger.

2023–2024 · Forte / Forte Enterprise

36 physical / algorithmic qubits (#AQ 36, 2024), 99.6% 2Q / 99.98% 1Q fidelity, AOD laser control, all-to-all connectivity. Enterprise variant built for on-premise/data-center deployment.

2025 · Tempo

100 physical qubits, many-to-many connectivity, 99.9% 2Q / 99.99% 1Q fidelity (company spec). IonQ reported an #AQ 64 benchmark result on Tempo in 2025, ahead of the system's broader commercial-availability target of late 2026 — worth separating "benchmark achieved" from "generally available."

Sep 2025 · Oxford Ionics acquired

~$1.075B deal closes, bringing Electronic Qubit Control (EQC) chip technology in-house.

Oct 2025 · EQC record (preprint)

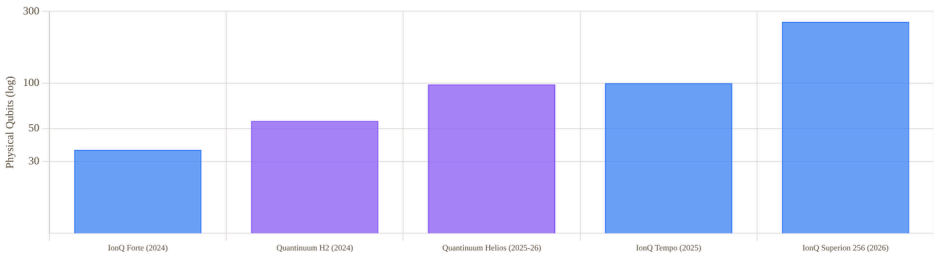
>99.99% two-qubit gate fidelity on a single ion pair without ground-state cooling — arXiv:2510.17286, not yet peer-reviewed, component-level not full-system.

Jan/May 2026 · SkyWater acquired

~$1.8B deal to acquire SkyWater Technology (semiconductor foundry) agreed January 2026, shareholder-approved May 2026, vertically integrating EQC chip manufacturing.

Sep 2026 · Superion 256

First EQC-based commercial platform launches: 256 qubits, server-rack footprint, sub-GPU-rack power draw. Orders open; customer deliveries targeted for 2027. Superion 10K (CMOS-integrated, aimed at IonQ's "Walking Cat" fault-tolerant architecture) is next in development ROADMAP.



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Demonstrated physical qubit counts for current, purchasable systems only (log scale). Roadmap targets (Sol, Apollo, Superion 10K) are shown separately in Section 07.

33.4 Gate Fidelity Deep Dive

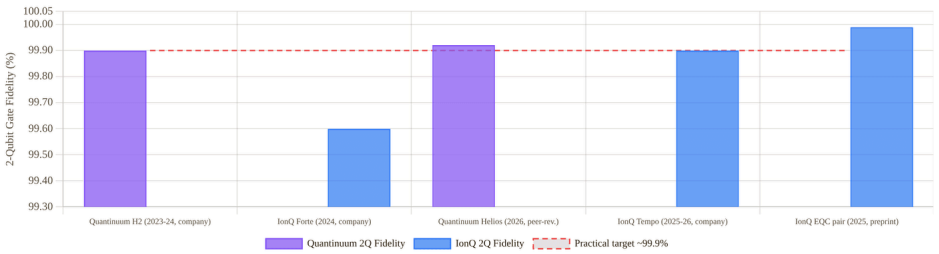
The single most important number in trapped-ion QC: how much error does one two-qubit gate introduce?

Two-qubit gate infidelity is the dominant error source in most trapped-ion circuits, and it sets the overhead needed for quantum error correction. Both companies report their numbers as average infidelity across many randomly-selected Clifford circuits (randomized benchmarking), which is the standard, defensible way to characterize a real device — but "average" and "worst-case" can differ, and a full-system number (averaged over many qubit pairs, over hours of drift) is a stronger claim than a single best-pair, best-moment number.

$$\mathcal{F}_{2Q} = 1 - \varepsilon_{2Q}, \qquad \varepsilon_{2Q}^{\text{Helios}} = 7.9(2) \times 10^{-4} \;\Rightarrow\; \mathcal{F}_{2Q}^{\text{Helios}} = 99.921\%$$

$$\varepsilon_{2Q}^{\text{IonQ EQC (pair)}} < 1 \times 10^{-4} \;\Rightarrow\; \mathcal{F}_{2Q}^{\text{IonQ EQC}} > 99.99\% \text{ (single ion pair, preprint)}$$

Number of sequential two-qubit gates a circuit can run before the accumulated error reaches order unity scales roughly as $N_{\text{gates}} \sim 1/\varepsilon_{2Q}$: about 1,270 gates at Helios's peer-reviewed rate, versus >10,000 gates at IonQ's preprint component-level rate — *if* that pair-level number holds up across a full multi-qubit system.



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Two-qubit gate fidelity by system and publication status. Dashed line marks the ~99.9% fidelity generally considered a practical threshold for useful near-term circuits.

Quantinuum, fidelity record

- H2 (2023–2024): >99.99% 1Q, >99.9% 2Q, all qubit pairs, company-quoted and consistent with peer-reviewed arXiv:2305.03828 PEER-REVIEWED
- Helios (2026): 99.9975% 1Q, 99.921% 2Q, 99.967% SPAM — all averaged across the full 98-qubit device, *Nature* 655, 81–86 PEER-REVIEWED
- Note: Quantinuum's own Helios launch blog rounds SPAM fidelity to "99.99%"; the peer-reviewed paper's SPAM infidelity of $3.3(5)\times10^{-4}$ corresponds to 99.967% — a good example of why this page prefers the peer-reviewed number when the two differ MARKETING ROUNDING

IonQ, fidelity record

- Forte / Forte Enterprise: 99.98% 1Q, 99.6% 2Q, full system, all-to-all COMPANY SPEC
- Tempo: 99.99% 1Q, 99.9% 2Q, full system, many-to-many COMPANY SPEC
- EQC record (Oct 2025): >99.99% 2Q fidelity on a single ion pair, without ground-state cooling, arXiv:2510.17286 PREPRINT, COMPONENT-LEVEL

Reading the fidelity race honestly: on *full, currently-purchasable systems*, Quantinuum's peer-reviewed 99.921% (Helios) modestly beats IonQ's company-quoted 99.9% (Tempo). On the *single best demonstrated number anywhere*, IonQ's preprint 99.99% two-ion EQC result is higher than anything Quantinuum has published — but it has not been shown across a multi-qubit processor, has not been peer-reviewed, and used a specific "no ground-state cooling" protocol whose behavior at scale is not yet public. Both statements are true at once; neither company has a peer-reviewed, full-system number above ~99.93%.

33.5 Logical Qubits & Fault Tolerance

Physical fidelity only matters because it sets how many physical qubits are needed per logical, error-corrected qubit.

$$p_L \sim A\left(\frac{p}{p_{\text{th}}}\right)^{\lfloor (d+1)/2 \rfloor}$$

Logical error rate p_L falls roughly exponentially with code distance d, provided the physical error rate p is below the code's threshold p_{th} (surface-code threshold is roughly 0.5–1%, well below both companies' physical 2Q error rates of order 10^{-3}–10^{-4}). This is why both companies are already in the "sub-threshold" regime and are now racing on *qubit overhead* — how many physical qubits it takes to reach a useful logical error rate — rather than racing to merely cross the threshold at all.

Quantinuum + Microsoft

- **Apr 2024:** 4 logical qubits on H2 (99.8% physical 2Q fidelity at the time), >14,000 circuits run without a single logical error reported by Microsoft; described as the largest reported gap between physical and logical error rates to date COMPANY BLOG
- **Sep 2024:** 12 logical qubits; circuit error rate 0.0011 (about 22× better than the 0.024 physical baseline); used two of the logical qubits for a real hybrid workflow estimating the ground-state energy of a catalytic P–N–N–P iron intermediate, combined with classical HPC (AutoCAS/AutoRXN) COMPANY BLOG, ARXIV PREPRINT REFERENCED
- **Nov 2025 (Helios launch):** company claims 94 fully-entangled logical qubits at break-even fidelity, 50 error-detected logical qubits above break-even, and 48 fully error-corrected logical qubits at a 2:1 physical-to-logical qubit ratio — the underlying physical-layer numbers are peer-reviewed (Section 04); the logical-qubit-count claims themselves are company-reported, not yet in a peer-reviewed logical-qubit paper COMPANY
- **Feb 2026 ("Skinny Logic," arXiv:2602.22211):** high-rate iceberg quantum error-detecting codes and two-level concatenated iceberg error-correcting codes run on the same 98-qubit Helios hardware — up to 94 error-detected logical qubits, and separately up to 48 error-corrected logical qubits, with GHZ-state fidelity ≈95% and a logical gate error rate around 10^{-4} PREPRINT, NOT YET PEER-REVIEWED
- **Sep 2026 ("Helix," arXiv:2609.03194):** a more compact code encoding a logical qubit into just 10 physical ions at code distance 6, reporting a logical Clifford-gate error of 2.8×10^{-4} (versus $\approx1.2\times10^{-3}$ for the physical baseline) and a logical memory error of 4.6×10^{-5} per logical qubit per cycle — both below the physical-layer baseline PREPRINT, NOT YET PEER-REVIEWED
- **Oct 2026 (color codes, arXiv:2610.24077):** triangular color codes at distances 3, 5, and 7 on Helios, using up to 37 physical qubits per logical qubit — logical memory error rate $5(+4/-2)\times10^{-6}$ with error suppression $\Lambda=12(+3/-2)$ per distance step, and $7(+3/-3)\times10^{-5}$ per logical two-qubit Clifford at distance 5 ($\Lambda=9$). Quantinuum's own Oct 19, 2026 press release describes this as a 10× improvement over its own Helix record and compares it to "100× lower than the best published result from other hardware types" (a company comparison, not independently benchmarked here) PREPRINT, NOT YET PEER-REVIEWED

IonQ qLDPC Break-Even

- **Jun 2026:** arXiv:2606.06455 reports a trapped-ion demonstration of logical-qubit *break-even* (logical memory lifetime exceeding the underlying physical qubit lifetime) using a bivariate-bicycle quantum low-density parity-check (qLDPC) code
- Configuration: 40 physical ion qubits encoding 4 logical qubits via the "BB5" code (a weight-5 generalized qLDPC matrix, "4 logical qubits in 18 physical containers"); logical memory lifetime 3.95 s vs 1.1 s for the physical qubits
- Reports logical error rates roughly 4× (Z-basis) and 9× (X-basis) lower than a comparable superconducting surface-code result reported elsewhere in the literature — *not* a direct comparison against Quantinuum's own logical-qubit numbers above
- Two independent secondary sources describing this preprint both report the physical qubits used were $^{133}Ba^{+}$ ions, not IonQ's commercial-product $^{171}Yb^{+}$/EQC line — worth flagging as likely a separate research demonstration rather than a Superion-family result PREPRINT, UNUSUAL ION SPECIES, VERIFY BEFORE REUSE

> **qLDPC vs surface codes:** qLDPC codes promise a much better logical-qubit-per-physical-qubit ratio than the surface code Google and others have used (potentially 10× or better at large scale), at the cost of requiring longer-range qubit connectivity to implement the code's check operators — a requirement that trapped-ion QCCD architectures, with their native all-to-all or many-to-many connectivity, are unusually well-suited to satisfy compared with fixed nearest-neighbor architectures. Both Quantinuum and IonQ are pursuing this connectivity advantage, though via different codes and different overhead claims; treat exact "N× better than superconducting" comparisons cautiously since code choice, decoder, and hardware error models all differ between the compared systems.

33.6 Full Head-to-Head Comparison

Every major dimension, side by side. Advantage marked in bold color where the evidence supports a clear edge.

Tug of war

Every row of the table below ends with an *Edge* call. Pick the categories you care about and watch the rope move.

Quantinuum 6

IonQ 6

Across these categories it's **roughly even** (6 vs 6, 10 even or different).

Quantinuum edge · 6

Even / different · 10

IonQ edge · 7

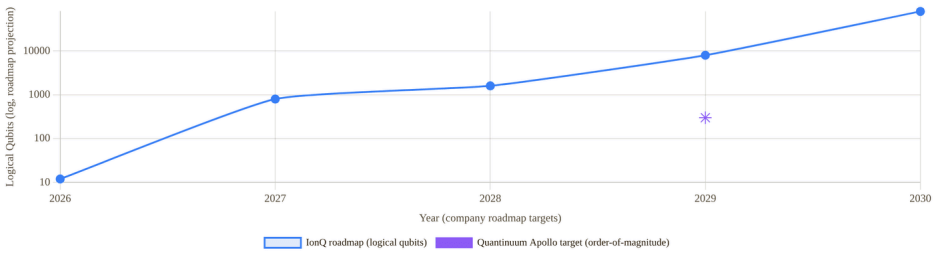
A tally of this page's own *Edge* column, not a score: rows are not equally important, which is exactly why you choose. Conditional calls ("if it scales", "pending review") count half and are hatched. Click any row to jump to it in the table.

CATEGORY	METRIC	QUANTINUUM	IONQ	EDGE
Company	Founded / origin	2021 merger of Honeywell Quantum Solutions + Cambridge Quantum	2015 spin-out of Duke University / University of Maryland	Different histories
Public listing	Nasdaq: QNT, IPO June 2026	NYSE: IONQ, public since 2021 (SPAC)	IonQ longer public track record	
Majority owner	Honeywell (~49% voting power post-IPO)	No single majority owner; widely held	Different structures	
Headquarters	Broomfield, Colorado / Cambridge, UK (dual HQ)	College Park, Maryland	Tie	
Architecture	Trap type	QCCD racetrack, rotatable storage ring + 4-way junction (Helios)	QCCD-style, chip-integrated electrodes (Superion)	Both QCCD family
Gate control	Laser-driven (AOD/EOM beam steering)	Electronic Qubit Control (EQC), chip-integrated, from Superion onward	IonQ for manufacturability; Quantinuum for maturity	
Current qubit species	$^{137}Ba^{+}$ (Helios), was $^{171}Yb^{+}$ (H1/H2)	$^{171}Yb^{+}$ (Forte/Tempo/Superion)	Different choices	
Connectivity	Full all-to-all, 98 qubits (Helios)	All-to-all (Forte) / many-to-many (Tempo)	Quantinuum slightly stronger claim at larger qubit count	
Fidelity	2Q fidelity, full system, peer-reviewed	99.921% (Helios, *Nature* 2026)	99.9% (Tempo, company spec only — not independently peer-reviewed)	Quantinuum
2Q fidelity, best demonstrated anywhere	99.921% (full 98-qubit system)	>99.99% (single ion pair, preprint)	IonQ if the pair-level result scales	
1Q gate fidelity	99.9975% (Helios, peer-reviewed)	99.99% (Tempo, company spec)	Quantinuum (narrowly)	
SPAM fidelity	99.967% (Helios, peer-reviewed; company blog rounds to 99.99%)	Not independently published for Tempo/Superion	Quantinuum (only side with a peer-reviewed number)	
Scale	Current physical qubits	98 (Helios)	256 (Superion 256)	IonQ on raw count
Manufacturing model	Lab-built, laser-aligned systems	Semiconductor-fab-built chips, targeting hundreds of units/year	IonQ (if realized)	
2030 roadmap physical qubit target	"Thousands" (Apollo, imprecise) ROADMAP	2,000,000 ROADMAP	Both unverified projections	

Error Correction	Logical qubits demonstrated	48 fully error-corrected (2:1 ratio, company-reported)	4 logical qubits, qLDPC break-even (preprint)	Quantinuum on count and maturity
Code approach	Color codes / QEC with Microsoft (details vary by demo)	Bivariate-bicycle qLDPC (BB5)	Different code families	
Business	Market cap / valuation	IPO Jun 2026 at $60/share, roughly $14–15B implied	~$14.7B market cap (Sep 2026)	Roughly comparable
Q2 2026 revenue	$8M (+279% YoY)	$80.1M (+287% YoY)	IonQ, roughly 10× larger	
FY2026 revenue guidance	$28–32M	$290M	IonQ	
Recent capital raised	$1.7B (June 2026 IPO)	~$2.9B combined for Oxford Ionics + SkyWater acquisitions APPROXIMATE	Both well-capitalized	
Cloud & Access	Cloud platforms	Microsoft Azure Quantum, Quantinuum's own Nexus platform, Oracle OCI (Helios)	AWS Braket, Microsoft Azure Quantum, Google Cloud Marketplace	IonQ broader multi-cloud reach
Software stack	TKET compiler/SDK, Guppy, Quantinuum Nexus	Qiskit, Cirq, and native SDK support via cloud partners	Both broadly compatible	

33.7 Roadmaps to Fault-Tolerant Quantum Computing

Both companies publish detailed, aggressive multi-year roadmaps. None of the future years below are demonstrated results.



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Company-published logical-qubit roadmap targets by year, log scale. IonQ publishes year-by-year figures through 2030; Quantinuum's Apollo target is a single 2029 milestone reported only as an order-of-magnitude ("hundreds" of logical qubits). Every point on this chart is a company projection, not a measurement.

YEAR	IONQ ROADMAP (PUBLISHED TARGETS)	QUANTINUUM ROADMAP
2025	64–100+ physical qubits, 99.9% fidelity	Helios launch (98 qubits) — achieved, not projected
2026	100–256+ physical, 99.99% fidelity, 12 logical qubits, logical error $<10^{-7}$	Sol development continues; Helios peer-reviewed (achieved)
2027	10,000 physical qubits, 800 logical qubits	Sol target system
2028	20,000 physical qubits, 1,600 logical qubits	Scaling toward Apollo
2029	200,000 physical qubits, 8,000 logical, error $<10^{-12}$, photonic interconnect introduced	Apollo target: thousands of physical qubits, hundreds of logical qubits, universal fault tolerance
2030	2,000,000 physical qubits, 80,000 logical qubits	No public target beyond Apollo as of this writing

How to read this table: IonQ's published curve implies roughly 10× qubit-count growth every one to two years sustained for half a decade — a pace with no historical precedent in any qubit modality, and the company itself frames it as "the fastest path to fault tolerance" rather than a guaranteed outcome. All entries in both columns for years after the current date are COMPANY ROADMAP claims; only the entries explicitly marked "achieved" reflect delivered hardware.

33.8 Business, Funding & Public Markets

Both companies are now public, giving an unusually transparent window into the economics of trapped-ion QC.

Quantinuum

- IPO priced June 3, 2026 on the Nasdaq, ticker **QNT**: 28M shares at $60, above a range already raised to $53–55, raising about $1.68B; trading began June 4. The Quantum Insider put the implied valuation at roughly $14–15B. Like every share-price figure, it has moved since. COMPANY FILINGS / PRESS
- Honeywell retains ~49.1% of voting power post-IPO; Cambridge Quantum-linked entities hold ~32.5%; founder Ilyas Khan is the largest individual shareholder, with a stake reported above $2B
- Q2 2026 (first earnings report since IPO, Aug 12 2026): revenue $8M, up 279% YoY from $2M; full-year 2026 guidance $28–32M
- GAAP net loss $597M in Q2 2026 (vs $57M in Q2 2025) — a jump that most likely reflects non-cash, IPO-related stock-compensation and transaction charges rather than a sudden change in operating cash burn; adjusted EBITDA loss was $68M, widened from $43M a year earlier
- Ended Q2 2026 with $2.1B in cash and short-term investments after the $1.7B June IPO raise COMPANY FILINGS

IonQ

- Public since October 2021 via SPAC merger with dMY Technology Group III, trading on the NYSE as **IONQ**
- Market capitalization ~$14.72B as of September 15, 2026 — down 7.4% from ~$15.90B at end of 2025, illustrating typical volatility for the sector rather than a steady climb
- Q2 2026 revenue: record $80.1M, up 287% year-over-year; full-year 2026 guidance raised to $290M
- Two major 2025–2026 acquisitions reshaped the balance sheet: Oxford Ionics (~$1.075B, closed Sept 2025) and SkyWater Technology (~$1.8B, agreed Jan 2026, approved May 2026) — both aimed at vertically integrating EQC chip design and fabrication rather than at near-term revenue
- Other recent acquisitions (quantum networking): ID Quantique and Qubitekk, expanding into quantum-safe networking and entanglement distribution alongside the core computing business COMPANY FILINGS
- Announced October 12, 2026: an all-stock agreement to acquire Zero Trust security firm Elisity for up to $1.3B (incl. up to $150M in earnouts), pairing IonQ's post-quantum cryptography/QKD hardware with Elisity's identity-based microsegmentation software; expected to close Q4 2026, subject to regulatory approval — not yet completed COMPANY FILINGS (8-K)

Reading the financials: IonQ's reported revenue is currently roughly 10× Quantinuum's on a quarterly basis, and its FY2026 guidance is about 10× larger too — but both companies remain deeply loss-making, and quantum-computing revenue recognition (hardware access credits, government contracts, long-term deals) varies enough between companies that quarter-to-quarter revenue comparisons should be read as directional, not as apples-to-apples SaaS-style metrics. Market capitalization for both currently sits in the same $14–15B range, despite very different revenue bases — a sign that public markets are pricing both largely on roadmap execution risk rather than current financials.

33.9 Companies & Ecosystem

Partners, customers, academic roots, and how each platform maps onto the broader quantum-computing job market.

Quantinuum Ecosystem

Honeywell ↗

Majority owner and original parent (Honeywell Quantum Solutions). Retains ~49% of voting power post-IPO. Long history in precision control systems and aerospace, which fed directly into Honeywell's original ion-trap control electronics.

Microsoft / Azure Quantum ↗

Deep logical-qubit collaboration: the April 2024 (4 logical qubits, 800× error reduction) and September 2024 (12 logical qubits, chemistry simulation) demonstrations were joint Microsoft–Quantinuum work. Azure Quantum also offers Quantinuum hardware access to cloud customers.

Sandia National Laboratories

Co-authors on the peer-reviewed Helios *Nature* paper (2026); a recurring government/national-lab collaborator across Quantinuum's fidelity benchmarking work.

JPMorgan Chase, Amgen, Mitsui & Co., NVIDIA

Named commercial and technology partners across financial random-number generation, drug discovery, and hybrid quantum-classical (CUDA-Q) workflows — illustrative of Quantinuum's enterprise-first go-to-market, distinct from IonQ's broader multi-cloud self-serve access model.

Cambridge Quantum heritage

The 2021 merger partner contributed the TKET compiler and quantum chemistry/NLP software lines that remain part of Quantinuum's software stack today.

IonQ Ecosystem

Duke University / University of Maryland

IonQ's founding academic roots: co-founders Christopher Monroe (Duke) and Jungsang Kim (Duke) built on decades of trapped-ion AMO physics. Strong ongoing recruiting pipeline from both universities' AMO and quantum-information groups.

Oxford Ionics ↗

UK spin-out from Oxford University's ion-trap group, acquired by IonQ in September 2025 (~$1.075B). Originated the Electronic Qubit Control (EQC) technology now central to the Superion product line.

SkyWater Technology

US (Minnesota) semiconductor foundry, agreed acquisition ~$1.8B (Jan 2026, shareholder-approved May 2026). Now fabricates IonQ's EQC ion-trap chips, cutting design-cycle time from nine months to two.

AWS, Microsoft Azure, Google Cloud

IonQ hardware is available through all three major cloud marketplaces — a broader multi-cloud footprint than Quantinuum's more enterprise-direct + Azure-centric model, and a real differentiator for researchers wanting quick, no-contract access.

ID Quantique & Qubitekk

Quantum-networking and entanglement-distribution acquisitions, positioning IonQ beyond standalone computing hardware into the broader "quantum internet" and quantum-safe communications space.

Elisity

Zero Trust security firm (IdentityGraph platform, 100+ enterprise/government deployments); all-stock acquisition announced Oct 12, 2026, up to $1.3B including earnouts, expected to close Q4 2026. Pairs IonQ's PQC/QKD hardware with Elisity's microsegmentation software — not yet completed as of this writing.

Career angle: both companies heavily recruit from AMO physics PhD programs — laser cooling, vacuum systems, precision spectroscopy, and microwave/RF control are directly transferable skills. Quantinuum's Helios generation shift to barium ions and IonQ's shift to chip-integrated electronic control both signal that semiconductor and photonics-integration experience is becoming as valuable as classic tabletop AMO experience for future hires at either company. See the Quantum Industry Map for a broader view of how trapped-ion roles compare to neutral-atom, superconducting, and photonic quantum computing careers.

33.10 Key Papers & References

Primary sources for every claim above: peer-reviewed papers, preprints, and company/press material, clearly separated.

Quantinuum, Landmark Papers & Releases

A 98-qubit trapped-ion quantum computer with all-to-all connectivity

Ransford, Bohnet et al. (200+ co-authors incl. Sandia National Laboratories), Nature 655, 81–86 (2026)

The Helios physical-layer paper: 98 $^{137}Ba^{+}$ qubits with $^{171}Yb^{+}$ sympathetic cooling, QCCD architecture with a rotatable ion-storage ring and four-way junction. Single-qubit infidelity $2.5(1)\times10^{-5}$, two-qubit infidelity $7.9(2)\times10^{-4}$, SPAM infidelity $3.3(5)\times10^{-4}$.

nature.com ↗

Demonstration of the trapped-ion quantum CCD computer architecture

Quantinuum/Honeywell team, arXiv:2305.03828 (2023)

Peer-reviewed-adjacent technical description underlying the H2 56-qubit, all-to-all connected system and its 99.9%-class two-qubit gate fidelity claims.

arXiv:2305.03828 ↗

Advancing science: Microsoft and Quantinuum demonstrate the most reliable logical qubits on record

Microsoft & Quantinuum, Official Microsoft Blog (Apr 3, 2024)

4 logical qubits, error rate reported as 800× better than physical qubits, >14,000 circuits without a logical error, on H2-generation hardware at 99.8% physical 2Q fidelity.

blogs.microsoft.com ↗

Microsoft and Quantinuum create 12 logical qubits and demonstrate a hybrid, end-to-end chemistry simulation

Microsoft Azure Quantum Blog (Sep 10, 2024)

12 logical qubits, circuit error rate 0.0011 (~22× improvement over physical baseline); two logical qubits used to estimate the ground-state energy of a catalytic iron intermediate in a hybrid quantum + classical HPC workflow.

azure.microsoft.com ↗

Introducing Helios: commercial launch announcement

Quantinuum, company blog & press release (Nov 5, 2025)

Commercial launch claims for Helios, including the 94 entangled / 50 error-detected / 48 fully error-corrected logical qubit figures, ahead of the June 2026 peer-reviewed physical-layer paper above.

quantinuum.com ↗

Computing with many encoded logical qubits beyond break-even

Dasu, DeCross, Guo, Lavasani, Behrends et al., arXiv:2602.22211 (Feb 25, 2026)

"Skinny Logic": high-rate iceberg error-detecting and two-level concatenated error-correcting codes on the 98-qubit Helios processor. Up to 94 error-detected logical qubits, and separately up to 48 error-corrected logical qubits; GHZ-state fidelity ≈95%. Preprint, not yet peer-reviewed.

arXiv:2602.22211 ↗

Experimental validation of a compact fault-tolerant architecture for trapped ions

Quantinuum, arXiv:2609.03194 (Sep 2, 2026)

"Helix": a logical qubit encoded into 10 physical ions at code distance 6 on Helios. Logical Clifford-gate error 2.8×10^{-4} (vs $\approx1.2\times10^{-3}$ physical) and logical memory error 4.6×10^{-5} per cycle — both below the physical-layer baseline. Preprint, not yet peer-reviewed.

arXiv:2609.03194 ↗

Quantum error correction well below the color code threshold

Lavasani, Reed, Dreiling et al. (23 authors), arXiv:2610.24077 (Oct 16, 2026)

Triangular color codes at distances 3/5/7 on Helios: logical memory error $5(+4/-2)\times10^{-6}$, $\Lambda=12(+3/-2)$ suppression per distance step; $7(+3/-3)\times10^{-5}$ per logical 2Q Clifford at distance 5 ($\Lambda=9$). Compares real-time neural-network decoders against algorithmic decoders. Preprint, not yet peer-reviewed; a company blog/press release (Oct 19, 2026) frames it as a 10× improvement over the Helix result above.

arXiv:2610.24077 ↗

Quantinuum Expands IPO as Valuation Climbs Above $14 Billion

The Quantum Insider (Jun 2, 2026)

IPO pricing, share count, Honeywell/Cambridge Quantum voting-power breakdown, and comparison to IonQ's market value at the time.

thequantuminsider.com ↗

IonQ, Landmark Papers & Releases

Trapped-ion two-qubit gates with >99.99% fidelity without ground-state cooling

Oxford Ionics / IonQ collaboration, arXiv:2510.17286 (2025)

The headline EQC fidelity record: two-qubit gate fidelity exceeding 99.99% on a single ion pair, achieved without ground-state cooling. Component-level result; full-system reproduction not yet public. Not yet peer-reviewed.

arXiv:2510.17286 ↗

Experimental demonstration of breakeven qLDPC and block codes on a trapped-ion architecture

IonQ, arXiv:2606.06455 (2026)

4 logical qubits from 40 physical ion qubits via a bivariate-bicycle "BB5" qLDPC code; logical memory lifetime 3.95 s vs 1.1 s physical (break-even). Secondary reporting describes the ions used as $^{133}Ba^{+}$, distinct from IonQ's commercial $^{171}Yb^{+}$ product line — flagged here as worth independent verification. Preprint, not peer-reviewed.

arXiv:2606.06455 ↗

IonQ Debuts Superion 256 Quantum Computing Platform

IonQ, company press release (Sep 8, 2026)

Launch details for the first EQC-based commercial platform: 256 qubits, server-rack footprint, SkyWater fabrication, 2027 customer deliveries, >300× projected cost-per-qubit reduction across the Superion roadmap.

ionq.com ↗

IonQ Quantum Systems — Compare Systems

IonQ, product specification page (accessed Sep 2026)

Official specification sheet for Forte, Forte Enterprise, and Tempo: qubit counts, connectivity, and 1Q/2Q fidelity figures used throughout this page's comparison table.

ionq.com ↗

IonQ Roadmap

IonQ, company roadmap page (accessed Sep 2026)

Year-by-year physical/logical qubit and fidelity targets through 2030, cited in Section 07. Entirely forward-looking; treat every figure as a company projection.

ionq.com ↗

IonQ Reports Record Second Quarter 2026 Financial Results

IonQ Investor Relations (Aug 2026)

Q2 2026 revenue $80.1M (+287% YoY), FY2026 guidance raised to $290M, and detail on the SkyWater and Oxford Ionics integration.

investors.ionq.com ↗

Independent Surveys & Reviews

Top Trapped-Ion Quantum Computing Companies in 2026

The Quantum Insider (Jun 12, 2026)

Independent survey comparing Quantinuum, IonQ, AQT, eleQtron, Universal Quantum, and QUDORA. Cross-checked against primary sources throughout this page rather than taken at face value.

thequantuminsider.com ↗

A new ion-based quantum computer makes error correction simpler

MIT Technology Review (Nov 5, 2025)

Independent technology-press analysis of the Helios launch and its error-correction claims, useful context alongside Quantinuum's own launch materials.

technologyreview.com ↗

CHAPTER 33 SUMMARY

Trapped ions currently set the high-water mark for peer-reviewed two-qubit gate fidelity, with modular scaling as the central challenge. We next compare the two titans of superconducting qubits.

CHAPTER 34 · CAREER STEP 06

34

Superconducting Processors: Google vs. IBM

🖥️ Google vs IBM · Deep dive

The superconducting-qubit showdown between the two longest-running programs and their two ideas of progress: Google's "beyond-classical" physics milestones versus IBM's "utility" claims and dated roadmap to fault tolerance. Each headline claim is shown next to the classical-simulation papers that challenged it.

> *Superconducting transmon processors operate at nanosecond gate speeds inside dilution refrigerators, with Google Quantum AI and IBM pursuing contrasting error-correction and modularity strategies.*

IN THIS CHAPTER

Prerequisites: Ch. 21, Ch. 31

- Contrast Google's 2D surface-code focus (Willow, below-threshold QEC) with IBM's move from heavy-hex lattices to the square-lattice Nighthawk and its qLDPC (gross-code) roadmap.
- Interpret headline quantum-advantage and utility claims alongside subsequent classical-simulation benchmarks.
- Identify the cryogenic wiring, coherence, and wafer-yield bottlenecks facing large-scale superconducting systems.

WEBSITE COMPANION Live interactive page ↗

34.1 Why These Two Companies?

The two research programs that between them have made both of the biggest public claims in superconducting quantum computing: "quantum advantage" and "quantum utility."

Superconducting transmon qubits are tiny LC circuits made from aluminum films on silicon or sapphire, cooled to mil-likelvin temperatures in a dilution refrigerator, and manipulated with microwave pulses. It is the oldest continuously-funded qubit modality at industrial scale, and it is dominated by two very differently structured organizations: **Google Quantum AI** (a research division inside Alphabet/Google, not separately publicly traded, that treats each generation of hardware as a physics-demonstration vehicle) and **IBM Quantum** (a business unit of a NYSE-listed technology company founded in 1911, that has published detailed multi-year commercial roadmaps and announced in June 2026 that it will invest more than $10B in quantum computing over five years). Both are chasing the same underlying physics — the surface code, or a related error-correcting code, layered on top of a scalable transmon lattice — but from opposite institutional starting points.

The core trade-off: Google has made the two headline-grabbing physics claims in the field: the original (and long-disputed) 2019 "quantum supremacy" result on the 53-qubit Sycamore chip, and a 2025 "verifiable quantum advantage" claim using the Quantum Echoes algorithm on the newer 105-qubit Willow chip, peer-reviewed and published in *Nature*. IBM, by contrast, has focused on a narrower, more conservative "utility" claim — that a 127-qubit device (Eagle) could produce numerically useful results on certain physics simulations *before* full error correction — a claim that was itself subsequently challenged by classical tensor-network simulations. IBM has also published by far the most detailed public fault-tolerance roadmap in the industry, naming every intermediate chip through 2029. **The tension this page is built around:** Google currently holds the stronger, more recent, peer-reviewed advantage claim; IBM has the more transparent, more skeptically stress-tested roadmap and a hardware philosophy (bivariate-bicycle qLDPC codes) that promises roughly 10× less physical-qubit overhead than the surface code Google has used to date. Neither company has yet demonstrated a fault-tolerant, commercially useful quantum computer.

13,000× **Google Quantum Echoes speedup vs. best classical algorithm (Nature 2025)**	105 **Willow qubits (two chip variants, Dec 2024)**	40× **Magic-state error reduction via cultivation (*Science* 2026)**
120 **IBM Nighthawk qubits (square lattice, Nov 2025)**	2029 **IBM's public target date for Starling, its first FTQC system**	$10B+ **IBM quantum investment over five years (announced Jun 2026)**

Calculator output at the page's default settings.

> **Source-confidence rule for this page:** PEER-REVIEWED numbers come directly from journal papers; PREPRINT numbers are posted but not yet peer-reviewed; COMPANY ROADMAP numbers are product specs, press releases, or future targets. Both companies' most famous advantage/utility claims were later challenged by independent classical-simulation papers — those challenges are presented alongside the original claims throughout this page rather than omitted, because that back-and-forth is itself the most instructive part of the story.

Who leads where, at a glance

Over the whole table the calls lean **IBM** (Google 3, IBM 10, 3 even or different), but the categories matter more than the total. Each card counts the *Edge* calls in one block of the full comparison table; click one to see its rows.

Conditional calls ("if it scales", "pending review") count half. Rows are not equally important: read the rows before quoting a winner.

34.2 Architecture & Qubit Control

Both use transmon superconducting qubits and microwave control — the difference is in lattice topology and how couplers are wired.

🔵 Google: Fixed-Frequency, Square/Diagonal Lattice

Google's transmons are arranged on a lattice with roughly grid-like nearest-neighbor connectivity (Willow's spec sheet reports an average connectivity of 3.47 — close to, but not exactly, a uniform 4-way square grid, since edge and corner qubits have fewer neighbors). Two-qubit gates use tunable couplers between fixed-frequency qubits, a design first introduced on Sycamore and refined through Willow.

- **Fixed-frequency transmons + tunable couplers:** reduces frequency-crowding-related errors relative to fully tunable-qubit designs, at the cost of needing careful frequency allocation across the chip during fabrication ESTABLISHED TECHNIQUE
- **Two chip variants at the Willow generation:** Google fabricates and tunes separate chips optimized for different demonstrations — one tuned for quantum error correction (better two-qubit fidelity, shorter T1), one tuned for random-circuit-sampling/advantage demonstrations (longer T1, even better two-qubit fidelity) — rather than a single chip serving both roles equally well COMPANY SPEC SHEET
- Packaging: flip-chip bump-bonded architecture with through-silicon vias for wiring, allowing control and readout lines to be routed from beneath the qubit plane rather than crowding the chip edges — a packaging approach Google has used since Sycamore and refined for Willow's denser lattice
- Google Quantum AI is a research division inside Alphabet, not a separately reported business unit — there is no public product line comparable to IBM's cloud-accessible System One/System Two hardware tiers

🟣 IBM: Heavy-Hex, Now Pivoting to Square

From Eagle (2021) through Heron R2 (2024), every mainline IBM chip used a **heavy-hexagonal lattice** — a topology chosen specifically because it is compatible with the surface code while minimizing frequency collisions between neighboring qubits, at the cost of below-4 average connectivity. With Nighthawk (Nov 2025), IBM switched to a denser **square lattice**

with tunable couplers.

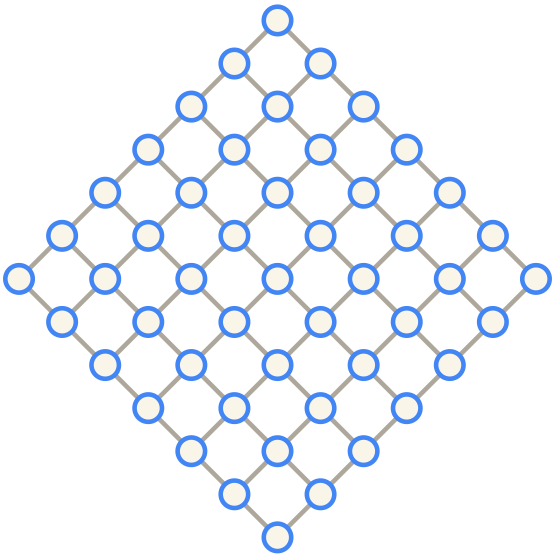
- **Heavy-hex generations (2021–2024):** Eagle (127q), Osprey (433q), Condor (1,121q), Heron/Heron R2 (133q/156q) — connectivity capped at 3 for most qubits, chosen to suppress two-qubit crosstalk and frequency collisions COMPANY
- **Nighthawk (2025–):** 120 qubits, square lattice, tunable couplers, each qubit connecting to up to 4 neighbors via 218 coupler pairs — IBM reports this is a 20% increase in inter-qubit connections over Heron, letting circuits use ~30% more two-qubit gate layers at comparable fidelity by cutting the number of SWAP gates needed to route non-adjacent qubits together
- Two-qubit gate fidelity above 99.9% is reported for over half of tested qubit pairs on Nighthawk-class hardware, with circuit-layer-operations-per-second (CLOPS) reaching 330,000, a 65% improvement over IBM's 2024 hardware COMPANY SPEC
- IBM explicitly frames the heavy-hex → square-lattice shift as a deliberate strategic pivot: the earlier 1,121-qubit Condor chip proved raw qubit count could scale, but IBM subsequently prioritized fidelity and connectivity per qubit (the Heron/Nighthawk lineage) over further count-scaling before returning to larger systems

> **Why this is the real story of this comparison:** both companies converged on transmon qubits with tunable couplers, but they made opposite bets on lattice density at different points in time — Google has favored denser near-square connectivity since Sycamore, while IBM spent 2021–2024 deliberately using a *lower*-connectivity heavy-hex lattice to buy fidelity headroom, only pivoting to a denser square lattice once its two-qubit fidelities were reliably above 99.9%. Nighthawk is, in a real sense, IBM catching up to a topology choice Google made years earlier — but arriving there with a more mature qLDPC error-correction strategy already public (Section 05).

Lattice lab: why connectivity matters

Both chips can only run a two-qubit gate between *coupled* neighbours. Anything farther apart must first be moved together with SWAP gates, and every SWAP adds error. Hover a qubit, click two to route a gate, or roll a random one on both lattices at once. Keyboard: Tab to a lattice, arrow keys to move, Enter to pick.

Square lattice Google Sycamore → Willow · IBM Nighthawk (2025–)

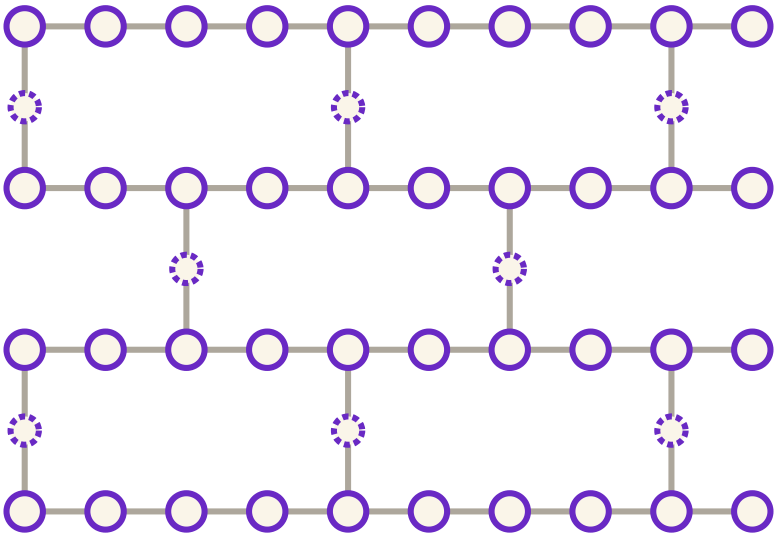


3.43avg. neighbours

4max neighbours

3.7avg. SWAPs per random gate

Heavy-hex latticeIBM Eagle → Heron R2 (2021–2024)



2.17avg. neighbours

3max neighbours

5.7avg. SWAPs per random gate

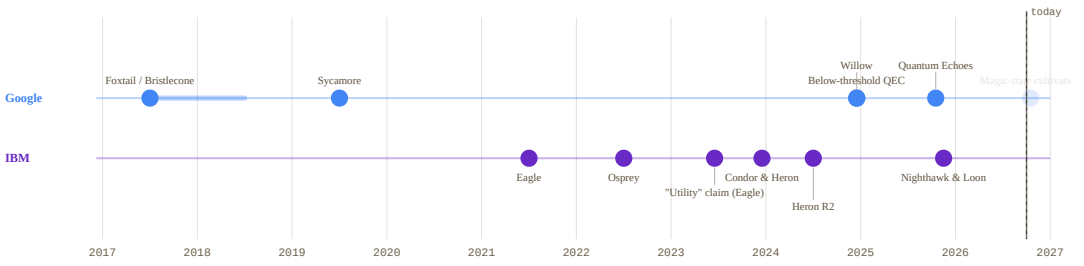
Schematic ~50-qubit patches, not the real chip maps. Willow has 105 qubits with an average connectivity of 3.47 (its spec sheet; edge qubits have fewer than 4); Nighthawk has 120 qubits with up to 4 neighbours; the heavy-hex Heron chips have 133/156. Real compilers route more cleverly than one SWAP chain per gate, but the ranking holds.

34.3 Hardware Generations, Side by Side

Very different scaling philosophies: Google iterates slowly with big physics claims; IBM iterates fast with named chips on a public cadence.

Race replay

Both timelines below on one clock. Drag the playhead or press play; hollow markers are company roadmap targets, not results.



Google by 2026 **Quantum Echoes** Oct 2025 — Peer-reviewed in Nature: a verifiable out-of-time-order-correlator (OTOC) algorithm run on Willow, reported 13,000× faster than the best known classical algorithm on a…

IBM BY 2026

Nighthawk & Loon

Nov 2025 — Nighthawk: 120-qubit square-lattice production chip (Section 02).

Google Quantum AI Chips

2017–2018 · Foxtail / Bristlecone

20-qubit and 72-qubit prototypes establishing the fixed-frequency, tunable-coupler transmon architecture Google has used ever since.

2019 · Sycamore

53 usable qubits (54 fabricated, 1 non-functional). Basis of the original "quantum supremacy" claim (Section 06) — later disputed by classical-simulation papers.

Dec 2024 · Below-threshold QEC

Published in *Nature*: first experimental demonstration that increasing surface-code distance (from 3 to 5 to 7) *decreases* logical error rate, the long-sought "below threshold" regime for quantum error correction.

Dec 2024 · Willow

105 qubits, average connectivity 3.47, fabricated as two tuned variants (QEC-optimized and RCS/advantage-optimized — see Section 04 for the differing specs). Random circuit sampling benchmark completed in under 5 minutes, a task Google estimated would take a leading classical supercomputer roughly 10^{25} years.

Oct 2025 · Quantum Echoes

Peer-reviewed in *Nature*: a verifiable out-of-time-order-correlator (OTOC) algorithm run on Willow, reported 13,000× faster than the best known classical algorithm on a leading supercomputer, cross-validated against real nuclear magnetic resonance (NMR) measurements on two molecules with UC Berkeley.

Oct 2026 · Magic-state cultivation

Peer-reviewed in *Science*: a "cultivation" protocol that prepares T-states (a resource consumed by non-Clifford logical gates) with roughly 40× lower error than preparing them without cultivation, reaching a state fidelity of at least 0.9999(1) on accepted (post-selected) attempts, about 8% of the total. Demonstrated on Willow, including a distance-5 surface-code "grafting" experiment using 54 of its 105 qubits.

IBM Quantum Chips

2021 · Eagle

127 qubits, first IBM chip to break the 100-qubit barrier and the first to use the heavy-hex lattice at scale.

2022 · Osprey

433 qubits, heavy-hex, primarily a packaging/scaling demonstration rather than a fidelity milestone.

Jun 2023 · "Utility" claim (Eagle)

127-qubit Eagle ran a 60-gate-layer kicked-Ising-model simulation with error mitigation, published on the cover of *Nature* — later challenged by classical tensor-network papers (Section 06).

Dec 2023 · Condor & Heron

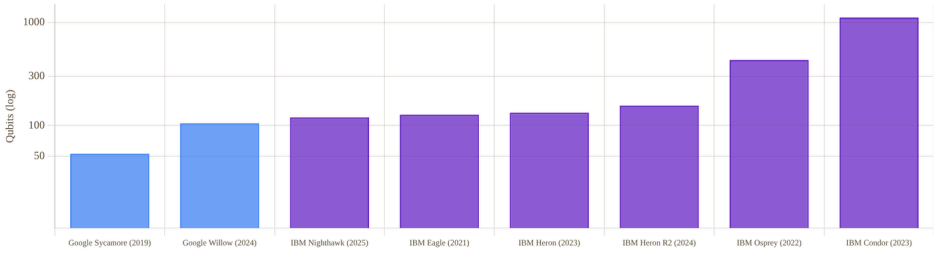
Condor: 1,121 qubits, heavy-hex — IBM's largest qubit-count chip to date and, so far, its last "more qubits" flagship. Heron (133 qubits): first chip to drop the fixed-coupler design in favor of tunable couplers, prioritizing fidelity over count.

2024 · Heron R2

156 qubits, refined tunable-coupler heavy-hex design; became IBM's primary cloud-accessible production chip through 2025.

Nov 2025 · Nighthawk & Loon

Nighthawk: 120-qubit square-lattice production chip (Section 02). Loon: a separate, non-production chip introducing long-range "c-couplers" needed for qLDPC codes — first rung of IBM's fault-tolerance roadmap (Section 05).



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Qubit count by named generation for each company (log scale). Note IBM's Condor (1,121 qubits, 2023) was a scale-only demonstration; IBM's subsequent chips (Heron, Nighthawk) deliberately used far fewer, higher-fidelity, better-connected qubits instead of continuing to scale count.

34.4 Gate Fidelity Deep Dive

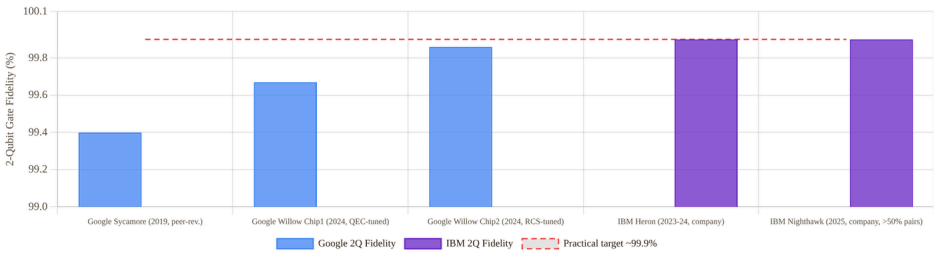
Willow shipped as two different chips with two different fidelity profiles — a nuance worth getting right before comparing to IBM.

Google's own Willow spec sheet (published alongside the December 2024 announcement) reports different calibration numbers for "Chip 1," tuned for the below-threshold quantum error correction demonstration, and "Chip 2," tuned for random circuit sampling / advantage demonstrations. Conflating the two — quoting Chip 2's better two-qubit fidelity as if it were the QEC chip's number, or vice versa — is a common but avoidable mistake, so both are reported separately below.

$$\varepsilon_{2Q}^{\text{Willow Chip1 (QEC)}} = 0.33\% \pm 0.18\% \;\Rightarrow\; \mathcal{F}_{2Q} \approx 99.67\%, \qquad T_1 = 68\,\mu s \pm 13\,\mu s$$

$$\varepsilon_{2Q}^{\text{Willow Chip2 (RCS)}} = 0.14\% \pm 0.052\% \;\Rightarrow\; \mathcal{F}_{2Q} \approx 99.86\%, \qquad T_1 = 98\,\mu s \pm 32\,\mu s$$

Both figures are randomized-benchmarking averages across the full 105-qubit device. Google's spec sheet attributes the T1/fidelity split between the two chips to "a tradeoff between optimizing qubit geometry for electromagnetic shielding and maximizing coherence" — i.e. the two chips are deliberately different devices, not the same hardware measured twice.



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Two-qubit gate fidelity by system. Dashed line marks ~99.9%, a commonly cited practical threshold for useful near-term circuits. IBM's ">99.9% for over half of tested pairs" is a distribution claim, not a single average — shown here as its reported representative figure.

Google, fidelity record

- Sycamore (2019): ~99.4% average 2Q fidelity on the RCS benchmark circuits used for the original supremacy claim PEER-REVIEWED, NATURE 574 (2019)
- Willow Chip 1 (QEC-tuned, 2024): 99.965% 1Q / 99.67% 2Q / 99.23% readout, averaged across the device COMPANY SPEC SHEET
- Willow Chip 2 (RCS-tuned, 2024): 99.964% 1Q / 99.86% 2Q / 99.33% readout COMPANY SPEC SHEET

IBM, fidelity record

- Eagle (2021–2023): fidelities in the low-to-mid 99% range for 2Q gates, sufficient for the error-mitigated (not error-corrected) 2023 utility experiment, but not separately publicized as a headline spec COMPANY
- Heron / Heron R2 (2023–2024): tunable couplers pushed typical 2Q fidelity into the "Heron-class" baseline IBM continues to cite, roughly 99.9% on well-behaved pairs COMPANY SPEC
- Nighthawk (2025): >99.9% 2Q fidelity reported for over 50% of tested qubit pairs on the new square lattice, alongside a 65% CLOPS improvement — IBM reports this as a distribution across pairs rather than a single device-wide average COMPANY SPEC

Reading the fidelity race honestly: Google's RCS-tuned Willow chip (99.86% 2Q, company spec) and IBM's Nighthawk (>99.9% for a majority of pairs, company spec) are in the same performance class, and neither company has published a directly comparable, identically-measured, full-device average that would let this page declare a clean winner. What differs more than raw fidelity is what each company chose to *build* on top of that fidelity: Google used its QEC-tuned chip to cross the below-threshold error correction milestone (Section 05), while IBM used its fidelity headroom to justify committing to a specific, named fault-tolerance roadmap (Section 05) years in advance.

34.5 Fault Tolerance & Error Correction

Google has demonstrated the below-threshold milestone; IBM has published the more detailed public roadmap and a fundamentally different, lower-overhead code family.

$$p_L(d) \approx A\left(\frac{p}{p_{\text{th}}}\right)^{\lfloor (d+1)/2 \rfloor}, \qquad \Lambda \equiv \frac{p_L(d)}{p_L(d+2)} \approx 2.14 \pm 0.02 \text{ (Google, Willow, 2024)}$$

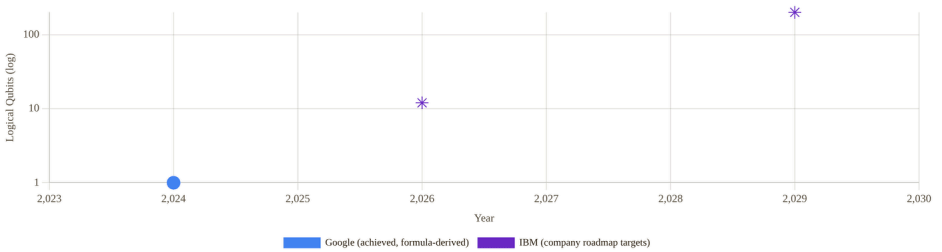
Google's Dec 2024 result measured the error-suppression factor Λ directly: each increase in surface-code distance by 2 (from $d = 3$ to $d = 5$ to $d = 7$) reduced the logical error rate by a consistent factor of about 2.14 — the first time this exponential suppression had been shown to continue below the code's error threshold, rather than saturating or reversing due to correlated/leakage errors as in earlier smaller-distance experiments.

Google: Below-Threshold Surface Code

- **Dec 2024 (Nature 638, 920–926; print Feb 2025):** distance-3, -5, and -7 surface codes on Willow's QEC-tuned chip; logical error rate roughly halved with each larger distance, and the best logical qubit outperformed the best constituent physical qubit's lifetime for the first time PEER-REVIEWED
- Uses the standard rotated surface code — well-understood theoretically but with a relatively high physical-to-logical qubit overhead (order 100+ physical qubits per logical qubit at useful error rates), compared with qLDPC alternatives
- No public Google roadmap names specific future chips or years the way IBM's does; Google's public messaging emphasizes physics milestones (below-threshold, verifiable advantage) over named hardware-generation dates

IBM: qLDPC Bivariate-Bicycle Codes

- **Bivariate-bicycle "gross code":** $[[144, 12, 12]]$ — 144 physical data qubits plus 144 syndrome qubits (288 total) encode 12 logical qubits; IBM reports roughly 10× fewer physical qubits than a surface code would need for comparable logical error rates, because the code's long-range checks pack more logical information per physical qubit COMPANY RESEARCH
- **Loon (2025):** first chip with the long-range "c-couplers" needed to physically realize qLDPC check operators — a proof-of-concept for the wiring, not yet a full logical-qubit demonstration
- **Kookaburra (2026, target):** first module combining qLDPC memory with an attached logical processing unit (LPU) ROADMAP
- **Cockatoo (2027, target):** entanglement between separate qLDPC modules via a "universal adapter" — the modular-scaling step IBM's whole roadmap depends on ROADMAP
- **Starling (2028–2029, target):** 2028 — magic-state injection across multiple modules; 2029 — full system target of 100 million quantum gates across 200 logical qubits, IBM's first claimed fault-tolerant machine ROADMAP



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Logical-qubit figures that are actually quantified in each company's own publications or roadmap statements — not every named milestone has a published qubit count, so only the ones that do are plotted. Google's point is a textbook-formula estimate (a distance-7 rotated surface code encodes 1 logical qubit in $2d^2-1 = 97$ physical qubits) applied to its 2024 below-threshold demonstration, not a number stated in the paper itself.

Demonstrated result vs. named roadmap: Google's below-threshold result is a completed, peer-reviewed physics demonstration on real hardware today. IBM's qLDPC/bivariate-bicycle overhead advantage is real and independently credible as *theory* (the code's distance and rate properties are published and reviewed), but the full modular architecture needed to use it at scale — Loon through Starling — is still almost entirely a forward-looking roadmap as of this writing, with only the first rung (Loon's c-couplers) built. Neither company has yet run a real, useful algorithm on a fault-tolerant logical processor.

34.6 Advantage & Utility Claims — and Their Rebuttals

Both companies' most famous headline results were later contested by classical-computing researchers. That back-and-forth is the most important context for reading either claim.

Google's Claims

- **2019 "Quantum Supremacy" (Sycamore):** a random-circuit-sampling task Google estimated would take the Summit supercomputer ~10,000 years, completed in 200 seconds. Published in *Nature* PEER-REVIEWED, but the classical-runtime estimate was disputed almost immediately — IBM itself published a rebuttal estimating 2.5 days with better classical algorithm/storage use, and subsequent papers (including a 2022 result using a large GPU cluster) narrowed or reversed the gap further for that specific circuit size.
- **2025 "Verifiable Quantum Advantage" (Quantum Echoes, Willow):** an out-of-time-order-correlator (OTOC) algorithm reported as 13,000× faster than the best known classical algorithm on a leading supercomputer, published in *Nature* PEER-REVIEWED. Google emphasizes this is the first *verifiable* advantage claim — the same algorithm was cross-validated on a second, independent Willow-class device and against real NMR spectroscopy measurements on two molecules (15 and 28 atoms) with UC Berkeley, rather than only checked against a classical simulation that becomes intractable at scale.
- As of this writing, no classical-simulation rebuttal of the October 2025 Quantum Echoes claim has been published — a meaningfully different situation from 2019, though the field's track record (below) suggests such attempts are likely and worth watching for.

IBM's Claim

- **2023 "Utility" (Eagle, kicked Ising model):** a 60-gate-layer, 127-qubit circuit, run with zero-noise-extrapolation error mitigation, reported as producing results "more accurate than leading classical approximation methods," published on the cover of *Nature* PEER-REVIEWED. IBM was explicit at the time that this was *not* a claim of outright classical-impossibility — only that the quantum device beat the classical approximation methods tested at that point.
- **Rebuttal:** within months, multiple independent groups showed classical tensor-network methods could match or beat the quantum device's accuracy on the same circuit. The most cited, Tindall, Fishman et al. (arXiv:2306.14887, later published in *PRX Quantum* 5, 010308, 2024), used belief-propagation tensor networks exploiting the heavy-hex lattice's tree-like correlation structure to produce results the authors report as "significantly more accurate and precise" than IBM's noise-mitigated quantum output — and argued their method scales to simulate the same physics in the thermodynamic (infinite-qubit) limit classically.
- IBM's own framing anticipated this: the paper stated the team "fully expect[ed] that the classical computing community will develop methods that verify the results" — which is exactly what happened. The net effect: the 2023 "utility" claim is best read as a genuine, peer-reviewed milestone in useful near-term computation, not as a surviving claim of classical intractability.

> **How to read both claims fairly:** "peer-reviewed" is not the same as "unchallenged" — both of these headline results passed peer review and were subsequently contested by other peer-reviewed or arXiv-posted classical-simulation work, which is normal, healthy scientific practice rather than evidence of fraud on either side. Google's 2025 claim differs from both companies' earlier ones mainly in that it was independently cross-checked against a real physical measurement (NMR) rather than solely against classical simulation — a stronger form of verification, though not immune in principle to a future, more efficient classical algorithm for the same task.

34.7 Full Head-to-Head Comparison

Every major dimension, side by side. Advantage marked in bold color where the evidence supports a clear edge.

Tug of war

Every row of the table below ends with an *Edge* call. Pick the categories you care about and watch the rope move.

Google 3

IBM 10

Across these categories the table leans **IBM** (10 vs 3, 3 even or different).

Google edge · 3

Even / different · 3

IBM edge · 10

A tally of this page's own *Edge* column, not a score: rows are not equally important, which is exactly why you choose. Conditional calls ("if it scales", "pending review") count half and are hatched. Click any row to jump to it in the table.

CATEGORY	METRIC	GOOGLE QUANTUM AI	IBM QUANTUM	EDGE
Organization	Program founded	~2012 (Google Quantum AI, initially the Quantum Artificial Intelligence Lab with NASA/USRA)	IBM's quantum computing research dates to the 1980s-90s theoretically; IBM Q commercial cloud access launched 2016	IBM longer institutional history
Structure	Research division within Alphabet/Google; no separate stock or public financials	Business unit of NYSE-listed IBM Corp (ticker: IBM)	IBM transparency	
Public roadmap detail	Physics-milestone driven; no named future chips or dated commercial targets published	Detailed, named-chip roadmap (Loon→Kookaburra→Cockatoo→Starling) through 2029	IBM roadmap clarity	
Investment disclosed	Not separately disclosed (folded into Alphabet's broader R&D)	$10B+ over five years, announced June 2026	IBM disclosed	
Hardware	Flagship chip / qubit count	Willow, 105 qubits (two tuned variants)	Nighthawk, 120 qubits	IBM qubit count
Lattice topology	Near-square, avg. connectivity 3.47, used consistently since Sycamore	Heavy-hex 2021–2024, pivoted to square lattice with Nighthawk (2025)	Tie converged	
Best reported 2Q fidelity	99.86% (Willow Chip 2, RCS-tuned, company spec)	>99.9% for >50% of pairs (Nighthawk, company spec)	Too close to call not directly comparable	
Headline Claims	Marquee physics result	Below-threshold QEC (2024) + verifiable quantum advantage (2025), both *Nature*, both currently unrebutted	"Utility" claim (2023), *Nature*, subsequently matched/beaten by classical tensor-network simulation	Google currently standing
Independent verification method	Cross-checked against real NMR spectroscopy data, not just classical simulation	Checked only against classical simulation, which subsequently caught up	Google stronger verification	
Error-correction code family	Standard surface code (higher physical-qubit overhead, mature theory)	Bivariate-bicycle qLDPC (≈10× lower overhead, theory published, hardware unproven at scale)	IBM overhead theory	
Fault Tolerance	Demonstrated milestone	Below-threshold logical error suppression, distance 3→5→7, achieved and peer-reviewed	Long-range c-couplers for qLDPC demonstrated (Loon); full logical-qubit system not yet shown	Google ahead, demonstrated
Target date, full FTQC system	No public date committed	Starling, 2029 (100M gates, 200 logical qubits) — company target	IBM only company with a date	

Facilities, Software & Access	Primary lab / campus	Quantum AI campus, Santa Barbara, CA — integrated hub combining a quantum data center, chip fabrication facility, and R&D under one roof	Distributed across multiple IBM Quantum System Two sites (including a dedicated Quantum Data Center) rather than one integrated campus	Tie different models
Primary open-source SDK	Cirq — used internally and by researchers, smaller ecosystem than Qiskit	Qiskit — widely regarded as the most broadly adopted open-source quantum SDK, with the largest third-party tooling and tutorial ecosystem	IBM ecosystem size	
External hardware access	Limited; primarily research collaborators rather than open self-serve cloud access	IBM Quantum Network: a large, decade-old (since 2016) consortium of enterprise, government, and academic partners with broad cloud access	IBM access breadth	
Standardized throughput metric published	Not published as a standing benchmark	CLOPS (circuit layer operations per second): 330,000 on Nighthawk, tracked release-over-release	IBM only one with a tracked metric	

34.8 Business, Investment & Roadmaps

One company reports quantum results inside a trillion-dollar conglomerate's research budget; the other reports them to public shareholders.

Google / Alphabet

- Google Quantum AI is organized as a research division within Alphabet (parent of Google), alongside Google Research and DeepMind — it has never been reported as, or spun out into, a separately-traded entity or disclosed business segment COMPANY STRUCTURE
- Alphabet does not break out Quantum AI revenue or spending in its quarterly SEC filings; quantum computing costs are absorbed within Alphabet's overall "Other Bets"/core-research R&D spending, which is not itemized by project
- No public target date for commercial quantum revenue has been stated by Google, in contrast to IBM's explicit 2028–2029 guidance (below) — consistent with Google's general public messaging emphasis on scientific milestones over commercialization timelines
- Alphabet's overall scale (quarterly revenue in the tens of billions of dollars) means Quantum AI's budget, whatever its size, is a research-and-development line item rather than a business the company is under investor pressure to monetize on any fixed near-term schedule INFERRED FROM PUBLIC FILINGS STRUCTURE

IBM

- IBM Quantum is a unit within IBM's broader Technology segment; as with Google, quantum-specific revenue is not yet broken out as its own reported line in IBM's quarterly filings COMPANY FILINGS
- June 2, 2026: IBM announced it will invest more than $10B in quantum computing over the next five years, funding its roadmap from current commercial systems through the targeted first fault-tolerant machines COMPANY PRESS RELEASE
- IBM's CEO stated in a July 2026 earnings call that the company expects quantum computing to have a "measurable impact" on IBM's earnings by 2028–2029, tied to the Starling launch target, and separately projected a trillion-dollar addressable market for quantum computing by the late 2030s — both are forward-looking company statements, not realized revenue COMPANY GUIDANCE
- IBM offers cloud-accessible commercial hardware today (IBM Quantum Network, System Two) with a large existing enterprise/government customer base built over a decade, a more mature commercialization posture than Google currently discloses publicly

Reading the business picture honestly: this is not a fair apples-to-apples financial comparison the way the Quantinuum-vs-IonQ page can offer, because neither Google nor IBM discloses standalone quantum-computing financials — Alphabet doesn't break out Quantum AI at all, and IBM's quantum unit sits inside a much larger technology segment. What can be compared is public communication style: IBM has chosen to publish detailed, dated, named-chip commercial commitments (and a specific revenue-impact year), while Google has chosen to publish detailed, dated *physics* milestones and left commercialization timing unstated. Both are legitimate strategies for a research program embedded inside a much larger company; neither is inherently more credible than the other on business grounds alone.

34.9 Companies & Ecosystem

Academic roots, partners, and how each program maps onto the AMO/condensed-matter physics job market.

Google Quantum AI Ecosystem

Alphabet Inc. ↗

Parent company (NASDAQ: GOOGL). Quantum AI is funded as part of Alphabet's broader research investment, alongside DeepMind and Google Research, without a separately reported budget.

UC Berkeley (Pines Magnetic Resonance Center)

Academic partner on the 2025 Quantum Echoes verification work, providing the independent nuclear magnetic resonance measurements used to cross-check the algorithm's output on real molecules.

Quantum AI campus, Santa Barbara

Google's integrated hardware hub (built substantially from the UC Santa Barbara group led by John Martinis, who ran the original Sycamore effort) houses chip fabrication, a quantum data center, and R&D under one roof — Google Quantum AI traces to around 2012, when the effort began as the Quantum Artificial Intelligence Lab in partnership with NASA and USRA.

Cirq (open source)

Google's open-source quantum SDK, used internally and by researchers. It has a smaller third-party tooling and tutorial ecosystem than IBM's Qiskit, which is more commonly cited as the field's default entry point for new quantum-software developers.

Google Cloud

Provides limited external/API access to Google's quantum processors for select research partners; far less broadly commercialized than IBM's decade-old public cloud-access model.

IBM Quantum Ecosystem

IBM Quantum Network ↗

A large, long-running (since 2016) consortium of enterprise, government, and academic partners with cloud access to IBM hardware — one of the broadest and oldest quantum-cloud-access programs in the industry.

US national labs & universities

IBM has published joint physics results with numerous academic and national-lab partners across its Eagle/Heron/Nighthawk generations, and its qLDPC code work builds on published academic error-correction theory (bivariate-bicycle codes).

Enterprise customers

IBM markets quantum access bundled with its broader enterprise software/consulting relationships (finance, materials science, logistics use cases), reflecting its business-unit structure versus Google's research-first posture.

Open-source: Qiskit

IBM's Qiskit SDK is one of the most widely used open-source quantum programming frameworks, a significant non-hardware contribution to the broader field's tooling and pedagogy.

Career angle: superconducting-qubit hardware roles draw heavily from condensed-matter and AMO physics, cryogenic engineering, and microwave/RF electronics — dilution-refrigerator experience and Josephson-junction fabrication are directly transferable skills at either company. Google's research-division structure tends to favor more academic-style hiring (PhD-heavy, publication-oriented); IBM's larger, longer-running commercial program offers a wider range of roles spanning hardware, software (Qiskit), and customer-facing quantum consulting. See the Quantum Industry Map for how superconducting roles compare to trapped-ion, neutral-atom, and photonic quantum computing careers, and Quantinuum vs IonQ for the equivalent comparison in the trapped-ion sector.

34.10 Key Papers & References

Primary sources for every claim above: peer-reviewed papers, preprints, and company/press material, clearly separated.

Google, Landmark Papers & Releases

Quantum supremacy using a programmable superconducting processor

Arute, Arya, Babbush et al., Nature 574, 505–510 (2019)

The original Sycamore "quantum supremacy" claim: 53-qubit random circuit sampling completed in ~200 seconds, estimated at ~10,000 years on a classical supercomputer. Subsequently disputed by IBM and later classical-simulation papers (Section 06).

nature.com ↗

Quantum error correction below the surface code threshold

Google Quantum AI & Collaborators, Nature 638, 920–926 (2025; published online Dec 2024)

The below-threshold demonstration on Willow's QEC-tuned chip: distance-3/5/7 surface codes, consistent ~2.14× logical error suppression per distance increase.

nature.com ↗

Meet Willow, our state-of-the-art quantum chip

Google Quantum AI, company blog (Dec 9, 2024)

Launch announcement and accompanying spec sheet distinguishing Willow's two tuned chip variants (QEC vs RCS), the primary source for the fidelity figures used in Section 04.

blog.google ↗

Our Quantum Echoes algorithm is a big step toward real-world applications for quantum computing

Google Quantum AI, company blog, with Nature publication (Oct 22, 2025)

The 2025 verifiable quantum advantage claim: 13,000× speedup on an OTOC algorithm, cross-validated against real NMR spectroscopy data with UC Berkeley on two molecules (15 and 28 atoms).

blog.google ↗

Magic state cultivation on a superconducting quantum processor

Rosenfeld, Gidney, Roberts et al., Google Quantum AI, *Science* (Oct 2026), DOI: 10.1126/science.aee7961 · theory: Gidney, Shutty & Jones, arXiv:2409.17595 (2024)

Magic-state cultivation: an imperfect T-state encoded in a small group of qubits, checked and post-selected, then "grafted" into a larger code (a distance-5 surface-code graft used 54 of Willow's 105 qubits). Reports ~40× lower T-state preparation error than uncultivated preparation and a state fidelity of at least 0.9999(1) on accepted attempts, keeping 8% of attempts (a lower bound, measured via kickback tomography). Peer-reviewed. What magic states are and why they matter: Learn Quantum → Magic States.

science.org ↗

Discover the Quantum AI campus

Google Quantum AI, company page (accessed Sep 2026)

Source for the Santa Barbara, CA campus description (integrated fabrication, quantum data center, and R&D hub) used in Sections 02 and 09.

quantumai.google ↗

IBM, Landmark Papers & Releases

Evidence for the utility of quantum computing before fault tolerance

Kim, Eddins, Anand et al., Nature 618, 500–505 (2023)

IBM's 2023 "utility" claim: a 127-qubit, 60-gate-layer kicked-Ising-model simulation on Eagle with error mitigation, reported cover story of the June 15, 2023 issue.

nature.com ↗

Efficient tensor network simulation of IBM's Eagle kicked Ising experiment

Tindall, Fishman, White, Stoudenmire, PRX Quantum 5, 010308 (2024); arXiv:2306.14887 (2023)

The primary rebuttal: belief-propagation tensor networks classically reproduce IBM's kicked-Ising results with higher accuracy, and the authors argue the method scales to the thermodynamic limit — directly challenging the classical-intractability implication of the 2023 utility claim.

arXiv:2306.14887 ↗

IBM Unveils Nighthawk and Loon: milestones toward quantum advantage and fault tolerance

IBM, company press release / blog (Nov 12, 2025)

Launch of the 120-qubit square-lattice Nighthawk processor and the Loon c-coupler proof-of-concept chip; source for the connectivity and CLOPS figures in Sections 02–04.

newsroom.ibm.com ↗

IBM lays out clear path to fault-tolerant quantum computing

IBM Quantum Computing Blog (Jun 10, 2025)

The full named-chip roadmap (Loon → Kookaburra → Cockatoo → Starling) and bivariate-bicycle qLDPC code details (the [[144, 12, 12]] gross code) used throughout Section 05.

ibm.com ↗

IBM Commits More Than $10 Billion to Quantum Computing

IBM Newsroom (Jun 2, 2026)

The investment commitment and framing used in Section 08, funding IBM's roadmap from current commercial systems through the targeted Starling fault-tolerant system.

newsroom.ibm.com ↗

Independent Surveys & Reviews

Questions and Concerns About Google's Quantum Supremacy Claim

Various authors, arXiv:2305.01064

A survey of the technical objections raised against the 2019 Sycamore supremacy claim, useful independent context alongside Google's own paper.

arxiv.org ↗

IBM CEO Expects Quantum Computing to Drive Revenue by 2020s, Trillion-Dollar Value by End of 2030s

The Quantum Insider (Jul 31, 2026)

Independent reporting on IBM's 2028–2029 revenue-impact guidance and long-term market-size projection, cross-referenced against IBM's own statements in Section 08.

thequantuminsider.com ↗

CHAPTER 34 SUMMARY

Superconducting qubits lead in gate speed and below-threshold surface-code demonstrations, while cryogenic wiring remains their defining scaling hurdle. Finally, we compare the two leading photonic approaches.

CHAPTER 35 · CAREER STEP 07

35

Photonic Architectures: PsiQuantum vs. Xanadu

💡 PsiQuantum vs Xanadu · Deep dive

The photonic showdown, and an asymmetric one: PsiQuantum has raised more money and published record component fidelities but no public full system, while Xanadu has run complete photonic machines in public, with peer-reviewed results, and is now listed. This page keeps that difference explicit instead of comparing unlike numbers.

> *Photonic quantum computers aim to leverage room-temperature silicon-photonics foundries and fiber interconnects, with PsiQuantum and Xanadu representing discrete-variable fusion and continuous-variable/GKP paradigms.*

IN THIS CHAPTER

Prerequisites: Ch. 28, Ch. 31

- Compare PsiQuantum's discrete-variable fusion-based architecture (Omega chipset) with Xanadu's continuous-variable / GKP-precursor networked racks (Borealis, Aurora).
- Distinguish component-level foundry demonstrations from full-system peer-reviewed benchmarks.
- Evaluate photon loss thresholds, SNSPD cryogenics, and optical switching as the primary photonic scaling bottlenecks.

WEBSITE COMPANION Live interactive page ↗

35.1 Why These Two Companies?

The two best-funded photonic quantum computing companies — and a genuinely different kind of comparison than this site's other head-to-heads.

Photonic quantum computing encodes information in single particles of light, which can travel through ordinary optical fiber and (unlike trapped ions or superconducting circuits) largely dodge the need for extreme cryogenic cooling of the qubits themselves. In exchange, photons are hard to store, hard to make interact with each other for two-qubit gates, and are lost from the system at a much higher rate than other modalities' error rates — loss, not just gate infidelity, is the central engineering problem. Two companies dominate the well-funded end of this approach: **PsiQuantum** (founded 2015, Palo Alto, still private, pursuing a fusion-based architecture built entirely on silicon-photonic chips fabricated at GlobalFoundries) and **Xanadu** (founded 2016, Toronto, publicly traded on Nasdaq and the TSX under **XNDU** since March 2026, pursuing continuous-variable squeezed-light and, more recently, fusion-based gate-model architectures).

The core asymmetry — and why this page is structured differently: PsiQuantum's public results are explicitly component- and subsystem-level: a February 2025 *Nature* paper reported record photonic-component fidelities (99.98% state preparation/measurement, 99.5% two-photon interference visibility) measured "conditional on photon detection" — i.e. only counting trials where the photons were not lost, which is standard practice for characterizing components but is not the same as demonstrating a working computer that tolerates loss end-to-end. PsiQuantum has not publicly operated a complete, multi-qubit quantum computer. Xanadu, by contrast, has operated and published peer-reviewed, full-system results twice: Borealis (2022, Gaussian boson sampling quantum computational advantage) and Aurora (2025, a smaller but "universal," gate-based photonic system). **Because of this asymmetry, this page does not try to force a symmetric fidelity-vs-fidelity comparison** the way the Quantinuum-vs-IonQ and Google-vs-IBM pages do — instead it keeps "what has each company actually run, end-to-end, and shown to the public" as the central organizing question throughout.

99.5%
PsiQuantum 2-photon interference (Nature 2025, component-level)

$7B
PsiQuantum valuation (Series E, Sep 2025, private)

216
Xanadu Borealis squeezed modes (Nature 2022, peer-reviewed advantage)

12
Xanadu Aurora qubits (2025, full universal system)

XNDU
Nasdaq & TSX ticker (public since Mar 27, 2026)

Calculator output at the page's default settings.

> **Source-confidence rule for this page:** PEER-REVIEWED numbers come directly from journal papers; COMPANY ROADMAP numbers are product specs, press releases, or future targets. Every PsiQuantum fidelity number on this page is explicitly flagged as **component-level, conditional on photon detection** where that is the case — this is the single most important distinction to preserve when reading photonic quantum computing claims, since a beautiful component fidelity says little on its own about a full system's usable error rate once real-world photon loss is included.

Who leads where, at a glance

Over the whole table the calls lean **Xanadu** (PsiQuantum 4, Xanadu 6, 4 even or different), but the categories matter more than the total. Each card counts the *Edge* calls in one block of the full comparison table; click one to see its rows.

Conditional calls ("if it scales", "pending review") count half. Rows are not equally important: read the rows before quoting a winner.

35.2 Architecture: Two Different Bets on Light

Both now use "fusion-based" measurement patterns to build entanglement, but they generate and route photons very differently.

🟣 PsiQuantum: Silicon-Photonic Fusion

PsiQuantum's architecture (Fusion-Based Quantum Computing, FBQC) builds large entangled cluster states out of small, pre-generated "resource states" of a few photons each, stitched together by "fusion" measurements — a strategy specifically designed to tolerate photon loss by using redundancy across many small resource states rather than requiring any single long-lived photon to survive the whole computation.

- **Omega chipset (Feb 2025):** silicon-photonic chips manufactured at GlobalFoundries' commercial fab in New York, incorporating a superconducting-nanowire single-photon detector layer and a barium-titanate electro-optic modulator layer (made by PsiQuantum in San Jose) bonded onto the silicon photonics COMPANY
- Photon sources: heralded single-photon sources via spontaneous four-wave mixing, integrated on-chip alongside Mach-Zehnder interferometers and ring resonators for routing and spectral filtering
- Chip-to-chip interconnect demonstrated over standard telecom fiber at distances up to 250 meters, with 99.72% interconnect fidelity — a key requirement since a million-qubit system will necessarily span many physically separate chips and racks
- Manufacturing claim: PsiQuantum reports characterizing millions of devices across thousands of wafers, roughly half a million measurements per month, framing the approach explicitly around semiconductor-style volume manufacturing rather than one-off lab assembly COMPANY

🟢 Xanadu: Squeezed Light → Fusion-Based Gates

Xanadu's earlier systems (X8, Borealis) used continuous-variable (CV) squeezed-light states and Gaussian boson sampling — a specialized, non-universal sampling task, not general-purpose computation. Aurora (2025) marks a shift toward a modular, networked, universal continuous-variable (CV) / GKP-precursor architecture built from room-temperature photonic racks.

- **Borealis (2022):** 216 squeezed light modes with three-dimensional (loop-based time-multiplexed) connectivity, generated by a pulsed optical parametric oscillator at 1,550 nm; performs Gaussian boson sampling, a fixed sampling task rather than universal computation PEER-REVIEWED
- **Aurora (2025):** 4 independent, interconnected server racks built from 35 photonic chips and roughly 13 km of fiber optics, operating at room temperature; described by Xanadu as its first "universal" photonic quantum computer, supporting gate operations, real-time error correction, and decoding rather than a single fixed sampling task — but with only 12 qubits deployed at launch, far fewer than Borealis's 216 modes NATURE 638, 912–919 (2025)
- Photonic chips for Aurora use commercially available fabrication techniques (details of the exact photon-source technology are less granularly public than PsiQuantum's Omega disclosures)
- Modular philosophy: Aurora's four-rack design is explicitly meant to scale by adding more networked racks — Xanadu states the same architecture could in principle extend to "thousands" of racks COMPANY

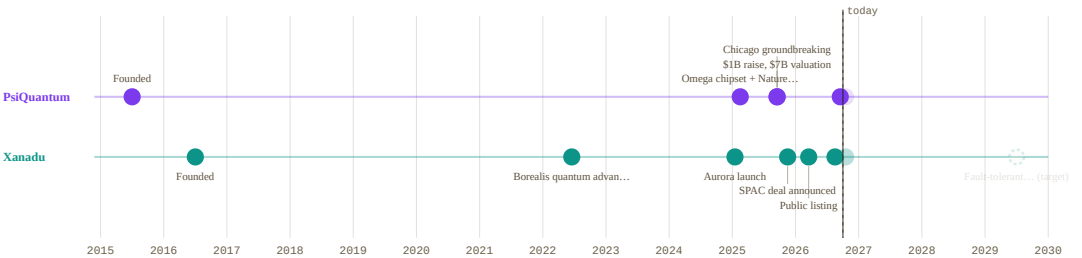
> **Why this is the real story of this comparison:** both companies have converged on fusion-based, loss-tolerant architectures — a sign the field has broadly settled on this as photonic quantum computing's best path to fault tolerance. The difference is scope of demonstration: PsiQuantum has shown excellent physics at the component and few-chip level but has not yet wired those components into a multi-qubit computer it has shown working end-to-end in public; Xanadu has repeatedly built and run smaller but genuinely complete systems, accepting a much lower qubit count in exchange for a real, working, peer-reviewed machine at each step.

35.3 Milestones, Side by Side

PsiQuantum's timeline is dominated by funding and infrastructure; Xanadu's by working systems.

Race replay

Both timelines below on one clock. Drag the playhead or press play; hollow markers are company roadmap targets, not results.



PSIQUANTUM BY 2026
$100M CHIPS R&D award finalized Sep 2026 — The U.S.

Xanadu by 2026 **CAD $195M Canadian government investment** Aug 28, 2026 — "Project OPTIMISM": federal funding for advanced domestic photonic-chip manufacturing facilities (integration, packaging, wafer-level testing, module assembly), part of…

PsiQuantum

2015 · Founded

Founded in Palo Alto by Jeremy O'Brien, Terry Rudolph, and others from the UK photonic-quantum-computing academic community.

Feb 2025 · Omega chipset + Nature paper

"A manufacturable platform for photonic quantum computing," Nature 641, 876–883: record component fidelities (99.98% SPAM, 99.5% two-photon interference, 99.22% two-qubit fusion, 99.72% chip-to-chip interconnect), all conditional on photon detection; component/subsystem-level only.

Sep 2025 · $1B raise, $7B valuation

Funding round led by BlackRock, Temasek, and Baillie Gifford.

Sep 2025 · Chicago groundbreaking

Broke ground on the Illinois Quantum and Microelectronics Park site, targeting a million-qubit, fault-tolerant system; phased plan begins with an intermediate-scale test system for DARPA's Quantum Benchmarking Initiative.

Sep 2026 · $100M CHIPS R&D award finalized

The U.S. Department of Commerce finalized a $100M CHIPS research-and-development award to PsiQuantum, following a May 2026 letter of intent.

Oct 8, 2026 · $150M conditional DoW loan

The U.S. Department of War's (the renamed Department of Defense) Office of Strategic Capital conditionally committed up to $150M to expand manufacturing/prototyping capacity (test equipment, assembly/integration infrastructure, cryogenic systems) at PsiQuantum's Milpitas, CA hub. Not guaranteed — final loan execution depends on PsiQuantum meeting financial, legal, technical, and regulatory conditions.

— · No public full-system demo AS OF THIS WRITING

PsiQuantum has not published a peer-reviewed or company demonstration of a complete, multi-qubit photonic quantum computer running an algorithm — its public record remains component/subsystem-level.

Xanadu

2016 · Founded

Founded in Toronto by Christian Weedbrook and collaborators, initially focused on continuous-variable photonic quantum computing and the open-source PennyLane software framework.

Jun 2022 · Borealis quantum advantage

"Quantum computational advantage with a programmable photonic processor," Nature 606, 75–81: 216-mode Gaussian boson sampling; one sample in 36 microseconds versus an estimated ~9,000 years on the Fugaku supercomputer for classical simulation. Made available for public cloud access after publication.

Jan 2025 · Aurora launch

First "universal" (gate-based) photonic quantum computer: 4 networked server racks, 35 photonic chips, ~13 km of fiber, 12 qubits, room-temperature operation; accompanying Nature methods paper on scaling and networking modular photonic systems.

Nov 2025 · SPAC deal announced

Agreed to go public via merger with Crane Harbor Acquisition Corp., a special-purpose acquisition company, at a $3B pre-money valuation, targeting roughly $500M in gross proceeds for expansion.

Mar 27, 2026 · Public listing

SPAC merger closed; began trading as XNDU on Nasdaq and the Toronto Stock Exchange, becoming the world's first publicly traded pure-play photonic quantum computing company.

Aug 28, 2026 · CAD $195M Canadian government investment

"Project OPTIMISM": federal funding for advanced domestic photonic-chip manufacturing facilities (integration, packaging, wafer-level testing, module assembly), part of up to CAD $390M in potential federal support announced earlier in March 2026.

Oct 6, 2026 · GlobalFoundries manufacturing partnership

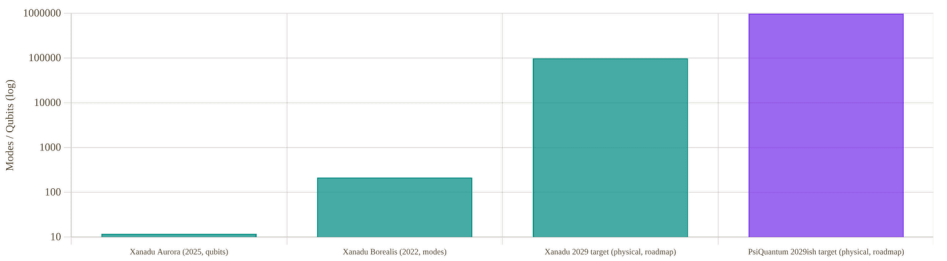
Multi-year deal to manufacture Xanadu's superconducting-nanowire single-photon detectors (SNSPDs) and ultra-low-loss silicon-nitride photonics at GlobalFoundries' 300mm fab in Malta, NY — the same commercial-fab-over-custom-cleanroom strategy PsiQuantum uses (Section 08), though at GF's Malta site rather than its New York City-area Omega fab.

Oct 19, 2026 · DARPA funding extension

Up to $10M provisionally approved under DARPA's Quantum Benchmarking Initiative — a one-year contract extension following Xanadu's Stage B work, funding roadmap-acceleration work with ongoing DARPA progress reporting.

2029 (target) ROADMAP

Company target: fault-tolerant photonic quantum computer with roughly 100,000 physical qubits encoding on the order of 1,000 logical qubits via error-correcting codes.



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Photonic modes/qubits in each company's largest publicly demonstrated system (log scale). PsiQuantum's largest disclosed numbers are chip-count and wafer-throughput figures, not a qubit count for a working computer — shown here as "not applicable" to avoid implying a false comparison.

35.4 Component Fidelity Deep Dive

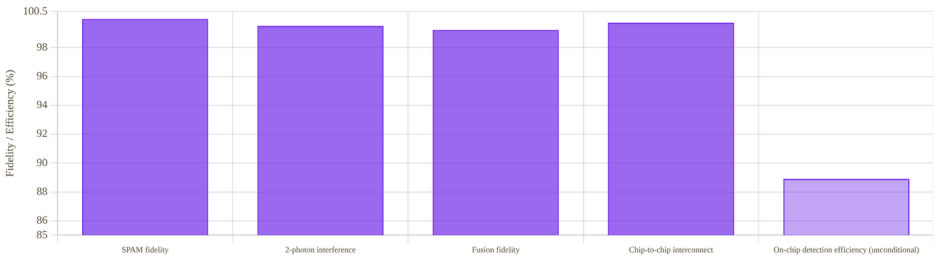
PsiQuantum's headline numbers are genuinely excellent — but they describe components, not a working computer's end-to-end error rate.

"Conditional on photon detection" is the load-bearing phrase in PsiQuantum's 2025 Nature paper: every fidelity number below was computed only from trials where the relevant photons were successfully detected, which excludes the dominant photonic error source — photon loss — from the reported figure. This is standard, legitimate practice for characterizing a component in isolation, but it means these numbers cannot be directly compared to a full-system fidelity the way Willow's or Nighthawk's two-qubit gate fidelities (which do include loss/error in their reported averages) can be on the Google-vs-IBM page.

$$\mathcal{F}_{\text{SPAM}}^{\text{PsiQuantum}} = 99.98\% \pm 0.01\%, \quad V_{2\gamma}^{\text{PsiQuantum}} = 99.50\% \pm 0.25\%, \quad \mathcal{F}_{\text{fusion}}^{\text{PsiQuantum}} = 99.22\% \pm 0.12\% \text{ (all conditional on detection)}$$

$$\eta_{\text{detect}}^{\text{PsiQuantum}} = 88.9\% \pm 3.5\%$$

That is the on-chip photon-detection efficiency: roughly 1 in 9 photons is lost before detection, unconditionally.
The detection efficiency η_{detect} is the number that actually captures loss, and it is the more informative figure for judging how close the platform is to a usable, unconditional error rate — an 88.9% detection efficiency implies a substantial fraction of photons are lost even in this component-level test, which the "conditional" fidelities above do not reflect.



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

PsiQuantum's four headline component fidelities (all conditional on photon detection) alongside its own loss-inclusive detection efficiency, for contrast. Xanadu has no directly equivalent component-fidelity figures published to this granularity — its results are reported as full-system outcomes instead, which is why no Xanadu bars appear here.

PsiQuantum, component record

- State preparation/measurement fidelity: 99.98% ± 0.01% PEER-REVIEWED, CONDITIONAL
- Two-photon interference visibility: 99.50% ± 0.25% PEER-REVIEWED, CONDITIONAL
- Two-qubit fusion fidelity: 99.22% ± 0.12% PEER-REVIEWED, CONDITIONAL
- Chip-to-chip interconnect fidelity: 99.72% ± 0.04%, over up to 250 m of telecom fiber PEER-REVIEWED, CONDITIONAL
- On-chip detection efficiency: 88.9% ± 3.5% — the loss-inclusive figure PEER-REVIEWED

Xanadu, system-level record

- Borealis: squeezing parameters and loss values sufficient to demonstrate a task classically intractable at the ~9,000-supercomputer-year scale — an end-to-end system result, inherently loss-inclusive, since the advantage claim depends on the full chain of squeezing, routing loss, and detection PEER-REVIEWED
- Aurora: individual component fidelities are less granularly publicized than PsiQuantum's Omega disclosures; Xanadu's public claim is architectural (universal gate-based operation, real-time decoding) rather than a headline fidelity number COMPANY + NATURE METHODS PAPER
- Because Xanadu's systems are operated as complete machines, their reported results are unconditional in the sense that matters most for usability — they reflect what the system actually delivers at its output, loss and all

Reading the fidelity picture honestly: PsiQuantum's conditional component fidelities are excellent by any standard and represent real, peer-reviewed progress in silicon-photonic manufacturing. But "conditional on detection" component numbers and "unconditional, full-system" numbers are not interchangeable, and this page treats them as answering different questions rather than ranking them on the same scale. The fairer framing: PsiQuantum has shown its parts can be excellent; Xanadu has shown a complete, if much smaller, machine can work end-to-end. Both are necessary; neither alone is sufficient for a fault-tolerant photonic quantum computer.

35.5 Fault Tolerance & Loss-Tolerant Codes

Photonic fault tolerance is primarily a loss-tolerance problem — different in character from the gate-fidelity-driven codes used by matter-qubit platforms.

$$p_{\text{loss}}^{\text{eff}} = 1 - \eta_{\text{total}} = 1 - \prod_i \eta_i \text{ (source, routing, coupling, detection losses compound multiplicatively)}$$

Because photon loss compounds across every optical component a photon passes through, fusion-based architectures are built to tolerate a nonzero, fixed loss rate per photon by using redundant resource states — the code "gives up" on any individual lost photon and reroutes around it, rather than requiring every photon to survive. This is a fundamentally different error model from the gate-infidelity-dominated codes used in trapped-ion or superconducting systems (Sections 05 of the other comparison pages on this site), which is why direct code-overhead comparisons across modalities should be treated cautiously.

PsiQuantum: Fusion-Based QC Theory

- PsiQuantum's FBQC architecture was developed specifically around loss tolerance from small resource states, described in earlier PsiQuantum-affiliated theoretical papers predating the 2025 hardware results
- A June 2025 PsiQuantum-affiliated study specifically mapped theoretical paths to greater loss tolerance in photonic fault tolerance, refining the resource-state and fusion-network design further COMPANY RESEARCH ANNOUNCEMENT
- The stated end goal — a million physical qubits at the Chicago and Brisbane sites — is explicitly a fault-tolerant target, but no public logical-qubit count, logical error rate, or intermediate fault-tolerance milestone has been demonstrated and published as of this writing ROADMAP
- DARPA's Quantum Benchmarking Initiative is cited as an external validation step for PsiQuantum's intermediate-scale test system ahead of the full million-qubit build — a third-party checkpoint worth watching for independently verified results

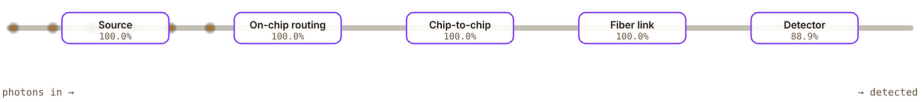
Xanadu: Modular, Networked Fault Tolerance

- Aurora's four-rack, fiber-networked design is explicitly built to demonstrate the modularity a fault-tolerant photonic computer will need — real-time error correction and decoding are cited as already integrated into Aurora's operation, though at a small (12-qubit) scale
- Xanadu's 2029 target (~100,000 physical qubits, ~1,000 logical qubits) implies an assumed physical-to-logical overhead around 100:1 — broadly consistent with, though not more favorable than, surface-code-class overheads reported by matter-qubit platforms ROADMAP
- As a newly public company, Xanadu's fault-tolerance roadmap is now subject to the same public-market scrutiny and quarterly reporting cadence as IBM's, a level of external accountability PsiQuantum (still private) does not yet face

> **Demonstrated result vs. roadmap, again:** neither company has published a photonic logical-qubit demonstration comparable to Quantinuum's, IonQ's, or Google's below-threshold result. PsiQuantum's advantage here is a more mature, purpose-built loss-tolerant code theory and a much larger capital base to build toward its million-qubit target; Xanadu's advantage is that its modular, networked architecture is not purely theoretical — Aurora is a working instance of the same networking idea at small scale, running today.

Loss budget: follow the photons

Every optical stage keeps a photon with some efficiency η, and the losses multiply. Set each stage and watch how many photons make it, and how often a whole multi-photon resource state survives intact.



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

88.9%of photons survive the chain

111lost per 1,000 photons

62.5%of 4-photon resource states arrive complete

—Monte Carlo tally below

$\eta_{total} = \Pi\ \eta_i$, as in the formula above; a k-photon resource state survives intact with η_{total}^{k}. The only company figure used is PsiQuantum's reported 88.9% on-chip detection efficiency (the first preset); no per-stage losses are attributed to either company. Fusion-based designs are built to tolerate losing some photons, which is exactly why this number matters.

35.6 Full Head-to-Head Comparison

Every major dimension, side by side. Several rows are intentionally marked "not comparable" rather than forced to a false tie.

Tug of war

Every row of the table below ends with an *Edge* call. Pick the categories you care about and watch the rope move.

PsiQuantum 4

Xanadu 6

Across these categories the table leans **Xanadu** (6 vs 4, 4 even or different).

PsiQuantum edge · 4

Even / different · 4

Xanadu edge · 6

A tally of this page's own *Edge* column, not a score: rows are not equally important, which is exactly why you choose. Conditional calls ("if it scales", "pending review") count half and are hatched. Click any row to jump to it in the table.

CATEGORY	METRIC	PSIQUANTUM	XANADU	EDGE
Company	Founded / origin	2015, Palo Alto; founders from the UK photonic-QC academic community	2016, Toronto; founders from Canadian/academic photonic-QC and open-source software backgrounds	Tie similar vintage
Public market status	Private; no public shares or standalone financial disclosure	Public since Mar 27, 2026 (Nasdaq & TSX: XNDU), first pure-play photonic QC public company	Xanadu public, first mover	
Most recent valuation	$7B (Series E, Sep 2025, private round)	$3B pre-money at SPAC announcement (Nov 2025); public market cap floats with XNDU share price thereafter	PsiQuantum higher valuation	
Demonstrated Hardware	Largest working full system shown publicly	None publicly demonstrated — component/subsystem-level only	Aurora: 12 qubits, universal/gate-based, 4 networked racks (2025)	Xanadu has a working system
Peer-reviewed system-level advantage claim	None published	Borealis, Nature 606 (2022): 216-mode Gaussian boson sampling quantum computational advantage	Xanadu only one with a system claim	
Best reported component fidelity	99.98% SPAM, 99.5% two-photon interference (conditional, Nature 2025)	Not separately publicized to the same granularity — system-level results emphasized instead	PsiQuantum component records	
Manufacturing	Fabrication approach	GlobalFoundries commercial semiconductor fab (New York); millions of devices characterized across thousands of wafers	Commercially available photonic-chip fabrication techniques; less granular public detail on fab partners/volume	PsiQuantum disclosed manufacturing scale

Physical infrastructure commitment	Two Quantum Compute Centers under construction (Brisbane, Australia; Chicago, USA), targeting 1,000,000 physical qubits	Modular rack-based systems (Aurora), targeting ~100,000 physical qubits / ~1,000 logical qubits by 2029	PsiQuantum larger stated target	
Fault Tolerance	Architecture family	Fusion-based quantum computing (FBQC), silicon-photonic, purpose-built for loss tolerance	Networked, modular fusion/gate-based (Aurora), evolved from earlier CV/Gaussian boson sampling systems	Tie converged approach
Public target date, full FTQC system	No single dated target published; phased Chicago/Brisbane build with DARPA-linked intermediate milestone	2029 (~100,000 physical / ~1,000 logical qubits) — company target	Xanadu has a dated target	
Government & Software	Sovereign / government investment	~$1B+ combined across Queensland/Australian federal co-investment (Brisbane site) and Illinois state partnership (Chicago site)	CAD $195M from the Government of Canada (Aug 2026, "Project OPTIMISM"), part of up to CAD $390M in potential federal funding for advanced photonic manufacturing facilities	Tie both sovereign-backed
Open-source software	No public SDK or open-source framework released	PennyLane — widely used open-source quantum machine learning / differentiable programming framework	Xanadu only one with public software	
Public cloud hardware access	None offered to date	Borealis was opened to public cloud access following its 2022 publication	Xanadu only one with public access	
Primary manufacturing/facility partner	GlobalFoundries (New York, commercial semiconductor fab) for chips; company-built Quantum Compute Centers for full systems	New Canada-funded advanced photonic manufacturing facilities (chip integration, packaging, wafer testing) to be built domestically	Tie different manufacturing strategies	

35.7 Business, Funding & Public Markets

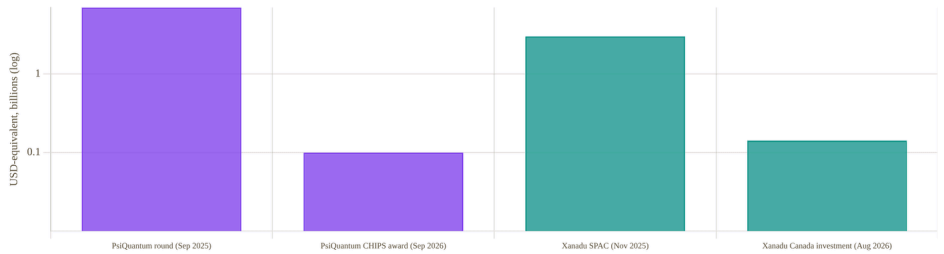
One company has raised more private capital than almost any startup in the sector; the other just became the industry's first pure-play public stock.

PsiQuantum

- Remains privately held as of this writing; latest priced round: a $1B Series E in September 2025 (led by BlackRock, Temasek and Baillie Gifford, with Nvidia's NVentures among new investors) at a $7B valuation. Some reports quote that round as "$1.5B at $10.5B": those are the same figures in Australian dollars, not a later round COMPANY FILINGS/PRESS
- Investors reported across funding history include Microsoft and Nvidia's venture arm, alongside the 2025 round's institutional backers
- Major capital commitments: Australian federal and Queensland state government support tied to the Brisbane Quantum Compute Center (reported in the ~$1B range for that site specifically); a Series E raise exceeding $1B tied to the Chicago site groundbreaking (Sep 2025)
- No public revenue figures have been disclosed — consistent with a private, pre-commercial-product company whose near-term "product" is the physical construction of its compute centers rather than sold hardware or cloud access INFERRED FROM ABSENCE OF DISCLOSURE
- October 8, 2026: the U.S. Department of War's Office of Strategic Capital conditionally approved up to $150M in loan financing for PsiQuantum's Milpitas, CA manufacturing/prototyping hub — a conditional commitment, not a disbursed grant, separate from the September 2026 $100M CHIPS Act award GOVERNMENT PRESS RELEASE

Xanadu

- Went public via SPAC merger with Crane Harbor Acquisition Corp.; deal announced November 3, 2025 at a $3B pre-money valuation, targeting roughly $500M in gross proceeds; merger closed and XNDU began trading on Nasdaq and the TSX on March 27, 2026 COMPANY/SEC FILINGS
- As a newly public foreign private issuer, Xanadu now files periodic reports (e.g., Form 6-K) disclosing quarterly results, including a widening reported net loss alongside growing quantum-sector funding activity broadly, per its FY2025 year-end filing SEC FILING
- Public listing gives Xanadu a transparency requirement PsiQuantum does not share — quarterly financials, material-event disclosures, and analyst coverage all now apply, similar to IonQ's and Quantinuum's positions among trapped-ion companies
- Xanadu's PennyLane open-source software framework, widely used in academic and industry quantum-machine-learning research, is a significant non-hardware asset and community touchpoint distinct from anything PsiQuantum has published publicly
- August 28, 2026: the Government of Canada announced a CAD $195M investment in Xanadu under "Project OPTIMISM," funding advanced domestic manufacturing facilities for fault-tolerant photonic quantum computers (photonic-chip integration, packaging, wafer-level testing, and quantum module assembly) — part of up to CAD $390M in potential federal funding announced in March 2026 GOVERNMENT/COMPANY PRESS RELEASE
- October 6, 2026: multi-year manufacturing partnership with GlobalFoundries (300mm Malta, NY fab) to industrialize Xanadu's SNSPDs and silicon-nitride photonics, with future work packages under discussion for higher-performance quantum modules COMPANY PRESS RELEASE
- October 20, 2026: DARPA provisionally approved up to $10M under a one-year Quantum Benchmarking Initiative extension, following Xanadu's Stage B work COMPANY/DARPA PRESS RELEASE



Snapshot of an interactive figure at its default settings. On the website it responds to the controls.

Disclosed private-capital and government-backed funding milestones by announcement date (log scale, USD-equivalent at time of announcement; CAD figures converted at an approximate 0.73 USD/CAD rate for comparability only — treat as order-of-magnitude, not a precise conversion). Valuations and investment amounts are point-in-time announcements, not cumulative totals.

> **Reading the business picture honestly:** PsiQuantum has raised more total private capital and carries a higher valuation than Xanadu's public-market starting point, reflecting investor confidence in its component results and its aggressive Chicago/Brisbane infrastructure build-out. But private valuation is not the same kind of signal as a public listing — Xanadu is now subject to ongoing public disclosure and market pricing that will make its progress (and setbacks) far more visible, in something closer to real time, than PsiQuantum's periodic funding-round announcements allow.

35.8 Companies & Ecosystem

Government infrastructure partnerships loom unusually large for both companies, more so than in trapped-ion or superconducting QC.

PsiQuantum Ecosystem

GlobalFoundries

Commercial semiconductor foundry (New York) manufacturing PsiQuantum's Omega silicon-photonic chipset — a rare example of a quantum computing hardware company relying on an established, high-volume commercial fab rather than its own custom cleanroom.

Government of Australia / Queensland

Public co-investment in the Brisbane Quantum Compute Center, part of a roughly $1B state/federal commitment — one of the largest government-backed quantum computing infrastructure deals globally.

Illinois Quantum and Microelectronics Park / State of Illinois

Host site and public partner for the Chicago Quantum Compute Center, described by PsiQuantum as America's largest quantum computing infrastructure project.

DARPA

Quantum Benchmarking Initiative provides an independent evaluation pathway for PsiQuantum's intermediate-scale test system ahead of full-scale construction.

Microsoft, Nvidia (venture arm)

Named investors across PsiQuantum's funding history, reflecting continued strategic technology-sector interest in the company's silicon-photonic approach.

U.S. Department of War, Office of Strategic Capital

Conditionally committed up to $150M (Oct 2026) toward expanding PsiQuantum's Milpitas, CA manufacturing/prototyping hub — separate from, and in addition to, the Department of Commerce's CHIPS Act award and DARPA's Quantum Benchmarking Initiative support. Not yet a completed, disbursed loan.

Xanadu Ecosystem

Crane Harbor Acquisition Corp.

The SPAC vehicle Xanadu merged with to go public (closed March 2026), now dissolved into the combined public entity trading as XNDU.

Government of Canada

CAD $195M investment announced August 2026 ("Project OPTIMISM"), part of up to CAD $390M in potential federal funding, for advanced domestic photonic quantum-computing manufacturing facilities — a sovereign-backing parallel to PsiQuantum's Australian and Illinois state partnerships.

PennyLane (open source)

Xanadu's widely-adopted open-source quantum machine learning and differentiable-programming framework, used well beyond Xanadu's own hardware and a major source of academic mindshare and hiring pipeline.

GlobalFoundries

Multi-year manufacturing partnership announced Oct 2026 (300mm Malta, NY fab) for Xanadu's SNSPDs and silicon-nitride photonics — the same commercial-fab strategy PsiQuantum uses (left column), now shared by both companies in this comparison.

Cloud access partners

Borealis was made available for public cloud access following its 2022 publication, an unusually open posture for a quantum-advantage-class result and one PsiQuantum has not matched with any comparable public system access.

Academic roots

Xanadu's continuous-variable photonics approach traces to a broader Canadian and international academic community in linear/Gaussian quantum optics, informing both its hardware and its PennyLane software philosophy.

Career angle: photonic quantum computing draws heavily on integrated photonics, nonlinear and quantum optics, and (increasingly, at PsiQuantum in particular) semiconductor fabrication and packaging engineering — skills that overlap only partially with the laser-cooling/trapping AMO toolkit central to trapped-ion and neutral-atom platforms. Xanadu's open-source PennyLane project is also a notable route into the field for software- and algorithms-focused physicists. See the Quantum Industry Map for how photonic roles compare to trapped-ion, neutral-atom, and superconducting quantum computing careers, and Quantinuum vs IonQ or Google vs IBM for the equivalent comparisons in those sectors.

35.9 Key Papers & References

Primary sources for every claim above: peer-reviewed papers, preprints, and company/press material, clearly separated.

PsiQuantum, Landmark Papers & Releases

A manufacturable platform for photonic quantum computing

PsiQuantum Team, Nature 641, 876–883 (2025)

The component-level results underlying Sections 02 and 04: SPAM, two-photon interference, fusion, and chip-to-chip interconnect fidelities, all conditional on photon detection; 88.9% on-chip detection efficiency.

nature.com ↗

PsiQuantum Announces Omega, a Manufacturable Chipset for Photonic Quantum Computing

PsiQuantum, company press release (Feb 26, 2025)

Omega chipset details: GlobalFoundries fabrication, superconducting-nanowire detectors, barium-titanate modulators, and the 250-meter chip-to-chip interconnect demonstration.

psiquantum.com ↗

PsiQuantum Breaks Ground on America's Largest Quantum Computing Project in Chicago

PsiQuantum, company press release (Sep 30, 2025)

Chicago Quantum Compute Center groundbreaking, million-qubit target, phased build plan, and DARPA Quantum Benchmarking Initiative link, cited in Sections 03 and 05.

psiquantum.com ↗

PsiQuantum company & funding profile

Sacra, company research profile (accessed Sep 2026)

Funding-round history and valuations used in Sections 03 and 07: $7B Series E (Sep 2025) and investor list. Caution: its text also mentions a "May 2026, $10.5B" round that we could not confirm; it matches the Sep 2025 round in Australian dollars.

sacra.com ↗

Xanadu, Landmark Papers & Releases

Quantum computational advantage with a programmable photonic processor

Madsen, Laudenbach, Askarani et al., Nature 606, 75–81 (2022)

The Borealis result: 216-mode Gaussian boson sampling, 36-microsecond sampling time versus an estimated ~9,000 classical-supercomputer-years, subsequently opened to public cloud access.

nature.com ↗

Scaling and networking a modular photonic quantum computer

Xanadu Team, Nature 638, 912–919 (2025)

The technical methods paper underlying Aurora's networked, modular rack architecture and loss-tolerance approach, cited throughout Sections 02–05.

nature.com ↗

Xanadu Introduces Aurora: World's First Scalable, Networked and Modular Quantum Computer

Xanadu, company press release (Jan 22, 2025)

Launch announcement for Aurora: 4 racks, 35 chips, 13 km fiber, 12 qubits, room-temperature operation, and the "universal" framing distinguishing it from Borealis.

xanadu.ai ↗

Xanadu Quantum SPAC Deal Approved, Targets $302 Million Raise

The Quantum Insider (Mar 19, 2026)

Coverage of the Crane Harbor SPAC merger's shareholder approval ahead of the March 27, 2026 public listing, cited in Section 07.

thequantuminsider.com ↗

Xanadu Announces Fourth Quarter and Full Year 2025 Results

Xanadu Investor Relations (Apr 9, 2026)

First public financial disclosures following the SPAC close, referenced in Section 07's business discussion.

investors.xanadu.ai ↗

Government of Canada Invests CAD $195 Million in Xanadu to Build the Quantum Supply Chain

Xanadu / Government of Canada, joint press release (Aug 28, 2026)

The "Project OPTIMISM" federal investment for advanced domestic photonic manufacturing facilities, cited in Sections 03, 07, and the comparison table's "Government & Software" row.

xanadu.ai ↗

Independent Surveys & Reviews

Global Photonic Quantum Computing Market Report 2026–2036

GlobeNewswire / industry market report (Aug 17, 2026)

Independent market survey evaluating linear-optics and continuous-variable strategies across PsiQuantum, Xanadu, Quandela, and ORCA — useful cross-check against both companies' own framing of their competitive position.

globenewswire.com ↗

PsiQuantum

Wikipedia (accessed Sep 2026)

General-purpose independent overview cross-referenced for founding details, funding timeline, and site locations.

en.wikipedia.org ↗

CHAPTER 35 SUMMARY

Photonic platforms trade deterministic matter-qubit interactions for native telecom networking and semiconductor manufacturability. Together, Chapters 1–35 span the complete journey from a single laser-cooled atom to the modern quantum computing landscape.

PART V

Teach with this Book

The chapters, laid out as a one-semester course for senior undergraduates and new graduate students.

V.1 How the course works

Fourteen weeks, two lectures a week. Lecture 1 sets up the physics; Lecture 2 turns it into the engineering decision a lab actually makes. Each week uses one to three chapters as its reading and its demonstrations, and every chapter ends with problems of three kinds: USE THE TOOL on the interactive page, BY HAND from the formulas, and DESIGN, a decision to make and defend. Prerequisites are undergraduate quantum mechanics and electromagnetism and enough Python to read a short script.

A suggested weighting is 40% weekly problems, 20% a midterm after week 7, 30% a final project and 10% paper summaries and participation. Project ideas: design a lab end to end with Part I; reproduce a published figure with a chapter and then in your own code; write and test a calculator in Python; rebuild the error budget of a published two-qubit gate (Chapter 25); run a full optimisation workflow on a simulated stage (Chapter 10); or teach one week's topic yourself.

V.2 The fourteen weeks

WEEK	TOPIC AND LECTURES	CHAPTERS
01	**Atoms, light and the two-level system** Two-level atom, Rabi oscillations, the Bloch sphere, detuning and the generalised Rabi frequency. Choosing an atom: linewidths, recoil, hyperfine structure, bosons and fermions, what each species is good for.	Ch. 21, Ch. 1
02	**Lasers and polarisation** Which lasers an experiment needs and why: cooling, repump, imaging, trap and Rydberg light; AOM and fibre schemes. Jones vectors, wave plates, Stokes parameters; circular light and the quantisation axis.	Ch. 2, Ch. 5
03	**Locking lasers to atoms and cavities** Saturated absorption, error signals, PID servos and what limits them. Offset locks with beat notes, Pound–Drever–Hall locking to a cavity, ultrastable references.	Ch. 4
04	**Vacuum and the lab bench** Why ultra-high vacuum: background collisions, pumping chains, conductance, bake-out. Everyday optics and RF: fibre coupling, AOMs, dBm budgets, Zeeman shifts, Clebsch–Gordan coefficients.	Ch. 3, Ch. 12, Ch. 20
05	**Laser cooling and the MOT** Radiation pressure, optical molasses, the Doppler limit; sub-Doppler and gray-molasses cooling. The magneto-optical trap: capture velocity, detuning and gradient trade-offs, species differences.	Ch. 13, Ch. 6
06	**Imaging clouds and measuring temperature** Absorption imaging: Beer–Lambert, optical depth, saturation and atom-number calibration. Time of flight: ballistic expansion, fitting a temperature, the recoil and Doppler scales.	Ch. 14, Ch. 15
07	**Controlling the experiment** Hardware timing, TTL/DAC/ADC/DDS/AWG signals, real device delays, pulse shapes and their spectra. ARTIQ, labscript and NI systems; writing a shot in Python; git for a shared experiment. Midterm review.	Ch. 7
08	**Optical tweezers and wavefronts** Dipole potentials, trap depth and frequencies, photon scattering from trap light, array geometry. Aberrations and Zernike modes, Strehl ratio, SLM holograms and the Gerchberg–Saxton algorithm.	Ch. 8, Ch. 9
09	**Seeing and measuring single atoms** Fluorescence imaging of one atom: scattering rate, collection efficiency, camera noise, detection fidelity. Release and recapture: measuring a single atom's temperature from a survival curve.	Ch. 16, Ch. 17
10	**Ground-state cooling and the full cycle** Lamb–Dicke regime, resolved and unresolved sideband cooling, Raman and gray-molasses schemes. A whole experimental cycle: loading, imaging, rearrangement and the exponential wall of large arrays.	Ch. 18, Ch. 19
11	**Coherence and dynamical decoupling** T_1, T_2^* and T_2; dephasing versus dissipation; how each is measured. Spin echo, CPMG and XY sequences; filter functions and noise spectroscopy.	Ch. 22, Ch. 23

12	**Rydberg atoms and two-qubit gates** Rydberg scaling laws, van der Waals interactions, the blockade radius and array geometry. CZ gates and their error budget: decay, Doppler dephasing, laser noise, blockade leakage, SPAM.	Ch. 24, Ch. 25
13	**Benchmarking, processors and automation** Randomized and interleaved benchmarking; zoned architectures and atom transport. Running a lab as a system: feedback and feedforward, calibration loops, optimisers, machine learning and AI agents.	Ch. 26, Ch. 27, Ch. 10
14	**Light–matter interfaces, networks and the field** Cavity QED: coupling rates, cooperativity, Purcell enhancement; atom–photon entanglement and remote links. Where the field is going: reading the literature, platforms and companies. Final project presentations.	Ch. 11, Ch. 28, Ch. 29

> The live course page, amotoolkit.com/pages/teach.html, has the in-class demonstrations for each week, printable problem sets and the solutions request for instructors.

A · Python notebooks

Each notebook reproduces one chapter's calculator with the page's own constants and formulas, adds a scan, fit or simulation the page does not show, and ends with a cell that checks its numbers against the live page. They need only numpy, scipy and matplotlib. The folder, with a README, is at github.com/saumitraphatak/amo-career/tree/main/notebooks.

CHAPTER	CALCULATOR	NOTEBOOK
Ch. 14	Absorption Imaging Lab	Colab · absorption-imaging.ipynb
Ch. 10	Automating the Experiment	Colab · automation.ipynb
Ch. 11	Cavity QED Coupling Lab	Colab · cavity-qed.ipynb
Ch. 7	Experiment Control Systems	Colab · control-systems.ipynb
Ch. 13	Laser Cooling Simulator	Colab · cooling-simulator.ipynb
Ch. 23	Dynamical Decoupling	Colab · dd-playground.ipynb
Ch. 25	Rydberg Gate Error Budget	Colab · fidelity-budget.ipynb
Ch. 16	Single-atom Imaging	Colab · imaging-calculator.ipynb
Ch. 20	Quick Lab Console	Colab · lab-calculators.ipynb
Ch. 18	Single-atom Cooling	Colab · laser-cooling.ipynb
Ch. 4	Laser Locking Guide	Colab · laser-locking.ipynb
Ch. 6	MOT Designer	Colab · mot-designer.ipynb
Ch. 26	Randomized Benchmarking	Colab · rb-explorer.ipynb
Ch. 17	Single-atom Temperature	Colab · release-recapture.ipynb
Ch. 28	Remote Entanglement	Colab · remote-entanglement.ipynb
Ch. 24	Rydberg Blockade Lab	Colab · rydberg-calculator.ipynb
Ch. 15	MOT Temperature	Colab · tof-calculator.ipynb
Ch. 8	Tweezer Array Design Lab	Colab · tweezer-designer.ipynb
Ch. 3	Vacuum Systems Guide	Colab · vacuum-systems.ipynb

B · Constants used in this book

CODATA values (2018 adjustment, with the 2019 SI's exact constants marked). Individual pages sometimes round these to four or five figures, which the notebooks reproduce exactly; species data (wavelengths, linewidths, saturation intensities) are given in the chapters that use them, with their sources.

QUANTITY	SYMBOL	VALUE
Speed of light	c	299 792 458 m/s (exact)
Planck constant	h	$6.626\,070\,15 \times 10^{-34}$ J s (exact)
Reduced Planck constant	$\hbar$	$1.054\,571\,817 \times 10^{-34}$ J s
Boltzmann constant	k_B	$1.380\,649 \times 10^{-23}$ J/K (exact)
Elementary charge	e	$1.602\,176\,634 \times 10^{-19}$ C (exact)
Atomic mass unit	u	$1.660\,539\,066\,60 \times 10^{-27}$ kg
Bohr magneton	μ_B	$9.274\,010\,078 \times 10^{-24}$ J/T ($\mu_B/h \approx 1.399\,6$ MHz/G)
Bohr radius	a_o	$5.291\,772\,109 \times 10^{-11}$ m
Vacuum permittivity	ε_o	$8.854\,187\,813 \times 10^{-12}$ F/m
Hartree energy	E_h	$4.359\,744\,722 \times 10^{-18}$ J
Standard gravity	g	9.806 65 m/s²

C · Citing this book and the work behind it

If this book or the website helps your research, please cite the relevant chapter or page. For results, derivations or material from the author's doctoral research, please also cite the thesis or the relevant paper.

This book

S. S. Phatak, *Building with Atoms: a training kit for experimental AMO and neutral-atom researchers*, edition 1, November 2026. AMO Toolkit, amotoolkit.com/book/.

Doctoral thesis

Phatak, S. S. (2026). *Cooling Lithium and Cesium Single Atoms in Optical Tweezers* (Doctoral dissertation, Purdue University Graduate School). hammer.purdue.edu

Paper (sideband cooling limits, used in Chapters 18 and 13)

S. S. Phatak *et al.*, *Phys. Rev. A* **110**, 043116 (2024). doi:10.1103/PhysRevA.110.043116

D · Using this material responsibly

- **Models, not measurements.** Most calculators use simple, clearly stated models (two-level atoms, Gaussian beams, linear error budgets). Each chapter says what its model leaves out. Use them for scale and for design trade-offs, then measure your own system.
- **Sources.** Numbers quoted from papers are as their authors report them; preprints are tagged and may change. Follow the references to the original before relying on a value.
- **Dated facts.** Company, funding and roadmap facts in Part IV change quickly. Each carries the date it was checked; check again before quoting.
- **Safety.** Lasers, high-current coils, RF amplifiers and vacuum systems can injure people and damage equipment. Nothing here replaces your institution's safety training, interlocks and procedures.
- **This is a snapshot.** The website is corrected and extended continuously; the book is regenerated from it. The edition date is on the cover. Report errors with the "Found a problem?" panel on any page.