Fault-tolerant operation and materials science with neutral atom qubits (Nature npj QI, 2025)

Portability dossier - app

agQSL portability pipeline

2026-06-07

Field	Value
Slug	`2026-nature-fault-tolerant-operation-materials-science-neutral`
Source	journal
Link	https://www.nature.com/articles/s41534-025-01095-w
Category	app
Triaged	2026-06-07 by port__until__julien__parallel
Bootstrapped	2026-06-07

Paper: fetch failed, see `paper.url`

Contents

1 Source

The published title (npj QI vol. 11, art. 193, 2025) is *Fault-tolerant operation and materials science with neutral atom logical qubits*. The open-access preprint is **arXiv:2412.07670**. Authorship (Radnaev, Bedalov, Carnahan, Chong, et al.) and the companion hardware paper place this on **Atom Computing**'s universal neutral-atom processor: ^{171}Yb nuclear-spin qubits with individual optical addressing and non-destructive mid-circuit readout. The work realises the Gottesman-2016 fault-tolerance benchmarking protocol on $[[4, 2, 2]]$ (C4) error-detecting logical qubits, reporting logical-over-physical error reduction of $\approx 12\times$ (Bell state) and $\approx 15\times$ (random circuits), plus a prototype Anderson Impurity Model ground-state solver for materials science (up to $\approx 6\times$, run non-fault-tolerantly).

1.1 Domain classification

Primary: **materials**. The application demonstration is an Anderson Impurity Model (AIM) ground-state solver, the canonical impurity problem at the heart of dynamical mean-field theory (DMFT) for strongly correlated electron materials. Secondary scope is **other (fault-tolerance / QEC hardware demonstration)**: the bulk of the paper is a logical-qubit benchmarking result, and the AIM solver is positioned as a prototype application riding on top of it rather than the main scientific claim.

1.2 Expert persona for Julien

A condensed-matter and quantum-simulation physicist who knows the Anderson Impurity Model and its role in DMFT for correlated materials, and who can read a small AIM ground-state circuit as a variational or state-preparation primitive. Critically, this reviewer also needs calibrated intuition for neutral-atom logical-qubit machinery: $[[4, 2, 2]]$ C4 detection codes, transversal gates, post-selection versus genuine fault-tolerant error correction, mid-circuit non-destructive readout and qubit reuse, and how a logical-versus-physical error ratio propagates into application-level accuracy. The hardware fingerprint Julien must extract (atom counts, single- and two-qubit gate fidelities, shot budgets, logical cycle depth) will sit largely in the supplementary material, so comfort mining SI is essential.

1.3 Related prior work (brief bibliography)

- arXiv:2408.08288, *A universal neutral-atom quantum computer with individual optical addressing and non-destructive readout* (the Atom Computing ^{171}Yb platform paper; same vendor and hardware).
- arXiv:2411.11822, *Logical computation demonstrated with a neutral atom quantum processor* (companion logical-qubit demonstration from the same group).
- arXiv:2312.03982, Bluvstein et al., *Logical quantum processor based on reconfigurable atom arrays* (QuEra/Harvard; same primitive, logical qubits on a different neutral-atom vendor, the key cross-vendor portability comparison).
- arXiv:1910.04735, Rungger et al., *Dynamical mean field theory algorithm and experiment on quantum computers* (same AIM/DMFT application on superconducting hardware; application-level portability baseline).

- arXiv:1610.03507, Gottesman, *Quantum fault tolerance in small experiments* (the benchmarking protocol this paper claims to realise).

In-corpus sibling: `dossiers/2026-arxiv-2605.21276-pasqal-logical-kernel-de-solver` is a Pasqal logical neutral-atom dossier (logical qubits on a different neutral-atom vendor, application-level residuals), and `2026-arxiv-2605.04025-fermi-hubbard-speedup` covers the correlated-electron application side. Both are useful framing for Brillant's later portability pass.

1.4 Difficulty estimate

high. The drivers are paper length (30 authors, a substantial hardware paper with heavy SI dependence) and exotic primitives: the hardware fingerprint is wrapped in QEC machinery ($[[4, 2, 2]]$ logical encoding, transversal operations, the post-selection versus fault-tolerant distinction, mid-circuit readout) rather than a bare gate sequence, and the atom counts, fidelities and circuit depths Julien needs will mostly live in the supplement. The application slice itself (a prototype AIM solver) is small.

No paywall blocks the work, but note the action item: the folder's `paper.url` points only at the auth-walled Nature landing page, and the bootstrap fetch failed. **Recommend the triager attach the open arXiv:2412.07670 PDF and its supplementary material to this folder before Julien runs**, so the fingerprint extraction works from full text rather than the abstract alone.

2 Extraction

Worked from the open-access arXiv preprint **arXiv:2412.07670v1** (full text, TeX source bundle, and figures), the version of record for the published *npj Quantum Information* article (vol. 11, art. 193, 2025; DOI `s41534-025-01095-w`). The Nature landing page is authentication-walled and the dossier bootstrap fetch failed, so the PDF, TeX source, and figure rasters were downloaded fresh from arXiv into this folder.

2.1 What the paper does (one paragraph)

The authors run logical qubits on a neutral-atom quantum computer and show that the encoded (logical) circuits make fewer errors than the bare (physical) circuits on the same machine. Two logical qubits are encoded in the $[[4,2,2]]$ error-detecting code (the C_4 code) using Cesium atoms held in optical tweezers. They report the first complete run of the Gottesman 2016 logical-qubit benchmarking protocol, finding a 15x reduction in output-distribution error for random circuits, plus a Bell-state result ($\approx$ 12x error reduction) and a prototype materials-science application: preparing the ground state of the single-impurity Anderson model (the impurity problem at the heart of dynamical mean-field theory) on the logical qubits, with up to $\approx$ 6x lower energy error than the unencoded version. The materials application is run non-fault-tolerantly and uses post-selection, not active error correction.

2.2 Quantum hardware used

- **Vendor / machine:** Infleqtion (formerly ColdQuanta), "Sqale" neutral-atom processor. Caution: Romain's sourcing note attributes this to Atom Computing with ^{171}Yb qubits; that is incorrect. The paper is Infleqtion's and uses ^{133}Cs atoms (Sec. I; Fig. 1b; companion hardware paper `radnaev2024universal` = arXiv:2408.08288, also Infleqtion). See Caveats.
- **Qubit count:** 6 physical qubits used, encoding 2 logical qubits via $[[4,2,2]]$ (Fig. 1a, p. 2). The code uses 4 data qubits; fault-tolerant $|00\rangle_L$ preparation adds a 5th (flag) qubit in a ring; the encoded Anderson-model circuit uses all 6 (4 data + 2 ancilla flags). Unencoded baselines run on 2 physical qubits.
- **Connectivity:** Triangular / hexagonal lattice at 6 μm atom spacing (Fig. 1b, p. 2). Two-qubit coupling is via the Rydberg interaction between atoms placed within range; the lattice geometry is chosen to give a 5-qubit ring (required for fault-tolerant $|00\rangle_L$ prep) plus the extra connectivity the Anderson-model circuit needs (Sec. II, p. 2). Not all-to-all; the coupling graph is set by atom placement, which is reconfigurable via the AOD tweezer array (App. B).
- **Gate set:** Native neutral-atom gateset (Sec. III.A): $GR_{\theta,\phi}$ global rotation gates that rotate every atom by θ about an in-plane axis ϕ; single-site R_z rotations (e.g. Z, S, $S^\dagger$) via individual optical addressing; and CZ entangling gates between connected atoms. Circuits compiled with Superstaq integrated with NVIDIA CUDA-Q.
- **Notable features leveraged:** Global rotation gate (one pulse drives all atoms, so the compiler inserts GR/R_z sandwiches to leave idling and flag qubits undisturbed, Sec. III.A); individual optical addressing for single-site R_z; non-destructive state-selective readout (NDSSR) followed by an atom-occupancy (loss) readout, enabling atom-loss post-selection (App. B); reconfigurable atom

placement to realise the 5-qubit ring. Deliberately not used: mid-circuit measurement and reset (all measurements are terminal), and the $XXXX$ stabiliser is never measured (Sec. I, p. 2).

2.3 Computational primitive

Two primitives, intertwined:

1. **QEC benchmarking / logical Clifford circuits + post-selection** (the bulk of the paper): transversal logical gates on the $[[4, 2, 2]]$ detection code, with terminal measurement and parity post-selection (Sec. III, IV).
2. **VQE / Hamiltonian Variational Ansatz state preparation** for the Anderson model (Sec. V). Important nuance: the variational angles were optimised **classically** in simulation (BFGS via SciPy/Qulacs, or CUDA-Q); the hardware then ran fixed, pre-optimised circuits. There is no hardware-in-the-loop variational feedback, so on the device this reduces to shallow fixed state preparation plus measurement.

2.4 Resource fingerprint

Metric	Value	Source (page / equation / SI)
Qubits (logical)	2 ($k = 2$ from the $[[4, 2, 2]]$ C_4 code)	Sec. III, Eq. (1); Fig. 1a (p. 2)
Qubits (physical)	6 used; $[[4, 2, 2]]$ = 4 data + 1 flag (5-qubit ring, Gottesman) or + 2 flags (6 qubits, Anderson-model encoded circuit); unencoded baseline 2	Fig. 1a (p. 2); Sec. II (p. 2); Sec. IV (p. 4); Fig. 6 (encoded AIM)
Circuit depth	Gottesman protocol: up to $T = 8$ gate layers (each circuit is $t + 2$ logical steps); fault-tolerant $\|00\rangle_L$ prep = 5 CZs (encoded) vs 0 (unencoded); Anderson ansatz = 2 variational angles. No single aggregate depth reported	Sec. IV (pp. 3 to 4); Sec. III (p. 3)

Metric	Value	Source (page / equation / SI)
2Q-gate count (total)	Not reported as a total. $\|00\rangle_L$ prep uses 5 CZs; logical $X/Z/HH/CX$ use virtual (relabelling) SWAPs, so no physical 2Q gates after state preparation	Sec. III (pp. 3 to 4)
Measurement shots	Gottesman: ≈ 1,050/circuit baseline + 7,000 each on 9 overlapping indices; > 450,000 shots total across 147 circuits. Anderson model: 1,000 shots/circuit (Trial 1). Bell tomography: 81 measurement settings × 2,000 shots = 162,000 (pre-loss-selection)	Sec. IV (p. 4); Sec. V (p. 5); Fig. 7 / Sec. VI (p. 6)
Classical loop iter.	Variational optimisation done classically, “a few dozen iterations at most” to converge; zero hardware-in-the-loop iterations (hardware ran fixed angles)	Sec. V (p. 5)
Wall-clock runtime	Not reported	Not reported

Supporting hardware fidelities (read from Fig. 1b, p. 2, and App. A noise model):

Quantity	Value	Source
CZ fidelity (per connection, atom-loss post-selected)	99.35% to 99.81%; array-median ≈ 99.48%	Fig. 1b (p. 2)
R_z fidelity (array-median)	99.81%	Fig. 1b (p. 2)

Quantity	Value	Source
GR fidelity (array-median)	99.96% (Fig. 1b); 99.97% by randomized benchmarking	Fig. 1b (p. 2); App. A
SPAM fidelity (array-median)	97.3%	Fig. 1b (p. 2)
Atom loss per site / measurement depumping	$\approx$ 1% state-average atom-loss per site; $\approx$ 3% bright-to-dark depumping	App. B
Prior single-pair CZ (1×2 grid, ref. machine)	99.35(4)%	App. B (`radnaev2024universal`)

2.5 Assumptions and results

- **Claims.** First complete realisation of the Gottesman 2016 benchmarking protocol on two logical qubits, with every logical circuit at least matching the physical circuits: 15x mean TVD reduction for $|00\rangle_L$ random circuits (10.5% → 0.7%), 6.7x for $|0+\rangle_L$, 7.7x for the Bell initialisation (Sec. IV). A logical Bell state with up to 12.4(23)x error reduction (physical fidelity 91.7(2)%; logical 96.5(3)% after $ZZZZ$ post-selection and tracing the unmeasured $XXXX$ parity; 99.3(1)% projected if both stabilisers were available) (Sec. VI). A prototype Anderson-model ground-state solver with up to $\approx$ 6x (5.7x to 6x) lower relative energy error (13% → 2.3%), run non-fault-tolerantly, the first time a strongly correlated materials-science subroutine is improved by an error-detection code (Sec. V).
- **Error bars / noise.** TVD 68% confidence intervals from Dirichlet / Markov-chain Monte Carlo sampling (10,000 draws, App. C); Bell 95% CIs from a Metropolis-Hastings tomographic reconstruction; Anderson-model averages reported as geometric mean and geometric standard deviation. A device-realistic circuit-level noise model (5-level per-atom system capturing leakage and atom loss, simulated with CUDA-Q density-matrix simulation) reproduces the experiment; the GR overrotation (34.5 mrad in the fitted model) is identified as the parameter that sets a pseudothreshold (App. A).
- **Comparison to classical baseline.** The headline comparison is logical-versus-physical on the same device, not quantum-versus-classical. The Anderson model used here is a two-qubit problem whose exact ground-state energies are computed classically (the parameter sweep $U \in \{1, 5, 9\}$, $V \in \{-9, -1, 7\}$ eV); the paper claims no quantum speedup and frames the application as a step towards DMFT on future fault-tolerant hardware (Sec. V).

2.6 Portability flags

- **Global rotation (GR) gate has no per-qubit equivalent.** A single GR pulse rotates every atom at once. Gate-based per-qubit architectures (superconducting, trapped-ion digital) have no native all-atom rotation; the GR/R_z “sandwich” compilation that protects idling and flag qubits is specific to this constraint and would be replaced by independent single-qubit rotations elsewhere.

- **Rydberg-mediated CZ with placement-defined connectivity.** The coupling graph is set by where atoms are positioned (reconfigurable via AOD tweezers), not by a fixed coupling map. The 5-qubit ring needed for fault-tolerant $|00\rangle_L$ prep comes for free from the lattice geometry; on fixed-connectivity hardware (heavy-hex, linear ion chain) the ring-plus-flag topology requires SWAP routing.
- **Atom-loss post-selection is a neutral-atom-specific channel.** An occupancy readout discards shots where atoms were lost; reported fidelities are post-loss-selection. Superconducting and trapped-ion platforms have no atom-loss channel, so this post-selection step does not transfer.
- **Individual optical addressing required** for single-site R_z and for embedding the encoded Anderson-model circuit into the lattice.
- **Terminal measurement only (portability plus).** The protocol deliberately avoids mid-circuit measurement and reset, so it does not require MCM. The trade-off is that it cannot do repeated syndrome extraction; scaling beyond detection to true correction would need MCM + reset.
- **Only the $ZZZZ$ stabiliser is enforced.** Error suppression is via Z-basis parity post-selection; $XXXX$ is never measured. Any portability comparison should note that the demonstrated suppression covers a single stabiliser, and the doubly-stabilised 99.3% Bell fidelity is a projection, not a measured result.
- **Detection-plus-post-selection, not correction.** Value depends on the dominant hardware error being detectable; throughput drops with the discarded-shot fraction. This is platform-agnostic in principle but its benefit is tied to each platform's error profile.
- **Application angles pre-optimised classically.** The Anderson-model demonstration runs fixed circuits; portability of the "application" is really portability of a shallow, fixed state-preparation plus measurement, not of a quantum variational optimiser.
- **Compilation toolchain.** Superstaq + NVIDIA CUDA-Q. Not hardware-locking, but the GR-aware compilation pass is specific to the global-gate platform.

2.7 Caveats

- **Vendor / species correction (material).** The dossier sourcing note states "Atom Computing's universal neutral-atom processor: ^{171}Yb nuclear-spin qubits". This is wrong for this paper. The authors are all Infleqtion (plus NVIDIA and university collaborators), the substrate is ^{133}Cs in optical tweezers, and the machine is Infleqtion's "Sqale" system (Sec. I; Fig. 1b; App. B; and the `vendor-app/edges.json` entry `e045`, which already records this as the first Infleqtion entry). The Atom Computing / Yb confusion most likely arose from the contemporaneous prior-work entry `reichardt2024logical` (arXiv:2411.11822, "Nov 2024, Neutral Atom (Atom Comp.)" in Tab. I), which is a different group's logical-qubit demonstration cited here, not the hardware of this paper. The companion hardware paper `radnaev2024universal` (arXiv:2408.08288) is also Infleqtion. Brillant's portability pass should treat this as Infleqtion / Cesium, and the in-corpus Atom Computing and Pasqal logical-qubit siblings as cross-vendor comparisons, not as this paper's machine.
- **Source provenance.** Extraction is from arXiv:2412.07670v1 (10 Dec 2024). The published *npj QI* version (2025) may differ in minor wording or rounding, but the preprint is the open-access record and carries the full appendices (noise model, experimental

details) that would otherwise sit in supplementary material. No separate SI file was needed or unreachable.

- **Figure-derived numbers.** The per-connection CZ values and the array-median R_z / GR / SPAM fidelities were read from the Figure 1b raster (legible: CZ 99.35% to 99.81%; R_z 99.81%; GR 99.96%; SPAM 97.3%).
- **Not reported.** Wall-clock runtime, total two-qubit gate counts, and per-circuit depths as aggregate numbers are not given. Where the paper does not state a value, the fingerprint table says so rather than estimating.

3 Portability matrix