

Certified quantum simulation of an Ising quench via error detection on the $[[6,4,2]]$ iceberg code

Portability dossier - app

agQSL portability pipeline

2026-05-31

Field	Value
Slug	2026-draft-iceberg-642-ising-quench
Source	local draft (inbox)
Link	local draft, no public link
Category	app
Triaged	2026-05-31 by manual-pipeline-test
Bootstrapped	2026-05-31

Paper: local draft at `dossiers/2026-draft-iceberg-642-ising-quench/paper.pdf` (no arXiv/DOI link).

Contents

1	Source	3
1.1	Domain classification	3
1.2	Expert persona for Julien	3
1.3	Related prior work (brief bibliography)	3
1.4	Difficulty estimate	4

2	Extraction	5
2.1	What the paper does (one paragraph)	5
2.2	Quantum hardware used	5
2.3	Computational primitive	5
2.4	Resource fingerprint	6
2.5	Assumptions and results	7
2.6	Portability flags	8
2.7	Caveats	8
3	Portability matrix	10
3.1	Summary	10
3.2	IonQ Forte – runs as-is (<i>anchor</i>)	11
3.3	Quantinuum H2 – runs with translation	12
3.4	AQT LYNX – runs with translation	12
3.5	IBM Heron r2 – runs with translation	13
3.6	Atom Computing AC1000 – runs with translation (after compiler retargeting)	14
3.7	QuEra Gemini – runs with translation (after compiler retargeting)	15
3.8	Pasqal – architectural rewrite	15
3.9	PsiQuantum – infeasible today	16
3.10	Xanadu – infeasible today	17
3.11	Quantum Brilliance – infeasible today	17
3.12	Residual caveats	19

1 Source

1.1 Domain classification

Primary: materials (condensed-matter many-body dynamics). The workload is a quantum quench of the transverse-field Ising chain $H = -J \sum_i Z_i Z_{i+1} - g \sum_i X_i$ on four spins, with the average magnetisation $\langle M_z \rangle$ tracked across the ferromagnetic to paramagnetic crossover at $g/J \approx 1$. It sits in the same family as the IBM kicked-Ising utility edge (e008) and the Quantinuum H2 Floquet-Ising demonstration already in `edges.json`.

Secondary axis is methodological rather than a separate application domain: quantum error detection (the $[[6, 4, 2]]$ iceberg code) plus circuit accreditation. That machinery is the paper’s real novelty, but it is cross-cutting infrastructure, not an application domain from the taxonomy. Classify the edge by its physics payload (materials) and let Julien fingerprint the certification overhead separately.

1.2 Expert persona for Julien

A quantum-information scientist with hands-on near-term trapped-ion experience, comfortable on both sides of the encoding boundary. They should know native Mølmer-Sørensen two-qubit gates and IonQ Forte error budgets ($p_{1q} \approx 2 \times 10^{-4}$, $p_{2q} \approx 3 \times 10^{-3}$), Trotterised TFIM dynamics, and small quantum error-detection codes (iceberg / $[[4, 2, 2]]$ family) together with accreditation and randomised-benchmarking methodology. The key skill here is separating the physics-bearing data circuit (the four-qubit quench, bare or $[[6, 4, 2]]$ -encoded) from the certification scaffolding (Pauli-twirled trap circuits, Ferracin physical vs Mills logical accreditation) so the hardware fingerprint reflects the simulation circuit and not the trap overhead. Fluency in reading two-qubit gate-count and shot-budget tables is essential.

1.3 Related prior work (brief bibliography)

- C. N. Self, M. Benedetti, D. Amaro, *Protecting Expressive Circuits with a Quantum Error Detection Code*, Nat. Phys. (2024). arXiv:2211.06703. Defines the iceberg $[[2m, 2m - 2, 2]]$ code family that the paper’s $[[6, 4, 2]]$ instance belongs to, including the trapped-ion native two-qubit logical rotations the draft exploits.
- S. Ferracin, T. Kapourniotis, A. Datta, *Accrediting outputs of noisy intermediate-scale quantum computing devices*. arXiv:1811.09709. The bare-circuit accreditation bound (the paper’s ref [3]); its single-qubit non-Clifford assumption is the constraint driving the CNOT-VC vs native-MS compilation tradeoff.
- J. Mills, A. Sireesh, D. Leichtle, J. Roffe, E. Kashefi, *Logical accreditation: a framework for efficient certification of fault-tolerant computations*. arXiv:2508.05523. The encoded-circuit accreditation used on the $[[6, 4, 2]]$ -encoded runs (the paper’s ref [4]).
- D. Gottesman, *Quantum fault tolerance in small experiments*. arXiv:1610.03507. Source of the Gottesman short-circuit criterion the draft uses as its $n_s = 1$ sanity check.

- **A. Vairogs et al., *Iceberg Beyond the Tip: Co-Compilation of a Quantum Error Detection Code and a Quantum Algorithm*.** arXiv:[2504.21172](#). Same iceberg code co-compiled with an algorithm on trapped ions; a direct portability comparison for the native-MS encoding strategy.
- **Performance of Quantum Approximate Optimization with Quantum Error Detection**, Commun. Phys. (2025). arXiv:[2409.12104](#). Same code family, different application (optimisation) and a postselection-vs-bare comparison on trapped-ion hardware, useful as the cross-application portability reference.

1.4 Difficulty estimate

Medium. The paper is fully readable: a complete 10-page draft (plus appendix Tables E1/E2) sits in the folder, so there is no paywall, no missing SI, and no fetch problem. The hardware fingerprint is unusually well tabulated. Forte noise parameters, qubit counts (4 bare, 8 to 10 encoded), per-cell two-qubit gate counts, trap and data shot budgets, and the \$6,726 package total are all laid out explicitly in Tables E1 and E2. What pushes this above **low** is conceptual density rather than missing data: Julien must hold three distinct compilations (CNOT-VC, MS-VC, native MS-Trotter) and two accreditation protocols in view at once, and must not conflate the certification scaffolding with the simulation circuit when extracting the fingerprint. No exotic primitives beyond the iceberg code and Mølmer-Sørensen gates, both of which are well documented in the bibliography above.

2 Extraction

2.1 What the paper does (one paragraph)

The paper simulates a quantum quench of a four-site transverse-field Ising chain and asks whether wrapping the circuit in a small quantum error-detecting code buys you anything on a realistic trapped-ion device. It encodes the four logical qubits into the $[[6, 4, 2]]$ “iceberg” code, runs the Trotterised dynamics, and keeps only the shots whose stabiliser syndrome is trivial (postselection). The central claim is that, with the two-qubit gates compiled to native Mølmer-Sørensen rotations (which roughly halves the encoding overhead), this syndrome postselection beats the unencoded “bare” circuit at the ferromagnetic-to-paramagnetic crossover under simulated IonQ Forte noise. Separately, it attaches a trap-based accreditation certificate (Ferracin for bare circuits, Mills logical for encoded) that bounds how far the noisy output sits from the ideal answer without classically computing that answer. All numbers are simulator-projected; a “hardware-ready” circuit package for IonQ Forte is specified but not run.

2.2 Quantum hardware used

- **Vendor / machine:** IonQ Forte (simulated; noise parameters and pricing taken from Forte). No hardware execution is reported.
- **Qubit count:** 4 physical for the bare circuit; 8 physical for the $[[6, 4, 2]]$ -encoded circuit per the submission manifest (Table E2). Note discrepancy below.
- **Connectivity:** All-to-all. IonQ Forte is a static linear $^{171}\text{Yb}^+$ chain with all-to-all coupling through shared collective motional modes (platform property; cross-checked against `vendor-notes/trapped_ion_qubits.tex`, not stated explicitly in the paper).
- **Gate set:** Native Mølmer-Sørensen two-qubit rotations $ZZ(\theta)$ and $XX(\theta)$, plus single-qubit R_X , R_Z rotations. A CNOT-based compilation (CNOT-VC) is the alternative baseline the paper compiles against.
- **Notable features leveraged:** Syndrome-based postselection (error detection); native MS $ZZ(\theta)/XX(\theta)$ rotations; Pauli-twirled accreditation traps. A two-round variant uses one mid-circuit stabiliser round, but the paper states Forte has **no native mid-circuit measurement (MCM)** and so realises it by **deferred measurement** (Fig. 1 caption).

2.3 Computational primitive

Trotterised Hamiltonian simulation. The workload is a second-order (symmetric) Suzuki-Trotter evolution of the transverse-field Ising Hamiltonian, quenched from $|0000\rangle$. A variationally-compiled (VC) circuit family (fixed $R_{ZZ}(\pi/2)$ brickwork plus parametric single-qubit rotations) is used as an alternative compilation that approximates the same $U(g, t)$; it is still simulation of the quench, not a separate VQE/QAOA optimisation primitive. The error-detection and accreditation machinery is cross-cutting certification scaffolding, not the simulation primitive.

The simulated Hamiltonian (paper Eq. 1):

$$H = -J \sum_{i=1}^2 Z_i Z_{i+1} - g \sum_{i=0}^3 X_i$$

with observable the average magnetisation $\langle M_z \rangle = \frac{1}{4} \sum_i \langle Z_i \rangle$.

2.4 Resource fingerprint

Values are for the simulation (data) circuit, separated from the trap / accreditation scaffolding. Where a quantity varies across the cell grid, the range is given with the spanning cells. “Shots” has several distinct meanings in this paper (analysis simulation vs accreditation traps vs the proposed hardware allocation); see the notes under the table.

Metric	Value	Source (page / equation / SI)
Qubits (logical)	4	Abstract; §II (“four qubits”)
Qubits (physical)	4 bare / 8 encoded (manifest); abstract & §V-VI say “10 physical qubits”	Table E2 “qubits” column (4, 8); abstract, §V, §VI
Circuit depth	Not reported as a depth; parameterised by Trotter steps $n_s \in \{1, 2, 3, 4, 5\}$ and VC ansatz depth $D \in \{2, 3, 4, 6\}$	§II–III; Figs. 1, 3; Table E2
2Q-gate count (total)	Data circuit, N_{2q}^{data} : bare MS-Trotter $6 \rightarrow 30$ ($n_s=1 \rightarrow 5$); encoded MS-Trotter $31 \rightarrow 71$; bare CNOT-VC $3 \rightarrow 9$; encoded CNOT-VC $75 \rightarrow 183$	Table E2 (cells 1–20); §II.A scaling $10n_s+21$ (native MS enc.), $6n_s+30$ (CNOT enc.)
Measurement shots	Analysis sims: 32,768 shots $\times$ 20 seeds. Accreditation validation: 100 traps $\times$ 512 shots. Proposed hardware: $K=20$ traps $\times$ 500 + 1,000 data shots per cell (11,000/cell)	§III (“32 768 shots each...20...seeds”); §III.A (100 traps $\times$ 512); Table E2

Metric	Value	Source (page / equation / SI)
Classical loop iter.	Not reported (VC angles “optimised per (g, J, t) ”; no iteration count or optimiser given)	§III.A
Wall-clock runtime	Not reported (all simulated; cost quoted in dollars instead: \$6,726 for the 20-cell package, \$336.30/cell)	Table E2; §VII

Notes on the table:

- **Qubit-count discrepancy.** The abstract, §V and §VI repeatedly state “ten physical qubits”, but Table E2’s manifest lists 8 qubits for every encoded cell, consistent with §II.A (six-qubit code block: two parity c_0, c_5 + four data c_1-c_4) plus two ancilla qubits for the stabiliser round. The “10 physical qubits” figure is not broken down anywhere; I report 8 from the manifest and flag the inconsistency rather than reconcile it. The dossier difficulty note already anticipates “8 to 10 encoded”.
- **Multiple shot/budget figures.** Beyond the table row, §III.C describes an alternative \$10,000 Braket allocation ($S_{\text{phys}} = 1.5 \times 10^4$ data shots, $M = 150$ traps $\times$ 200 shots, $\approx$ \$9,200, 5 g -points), Fig. E6 uses 30 traps $\times$ 1,000 shots, and §V/§VI mention a smaller “312 circuits, $\approx$ \$1,700” package. These are different proposed allocations, not a single executed run.
- **2Q-gate overhead.** Native-MS encoding overhead over bare is $\approx 2.5-5\times$ at $n_s = 3-5$ (§II.A); CNOT compilation of the encoded circuit is far heavier (75–183 two-qubit gates, $\approx 25\times$ the bare CNOT-VC count, §V).

2.5 Assumptions and results

- **Claim.** Under simulated IonQ Forte noise ($p_{2q} = 3 \times 10^{-3}$), $[[6, 4, 2]]$ syndrome postselection with native-MS compilation produces a tighter certified envelope on $\langle M_z \rangle$ than the bare circuit at the Ising crossover, but only above a pseudothreshold of $N_{\text{bare}} \sim 5-6$ two-qubit gates; below it (and for the CNOT-VC compilation), bare + Ferracin wins. Trap-based accreditation certifies the output to total-variation distance $\gamma = 2\hat{f}_{\text{trap}} < 0.25$ without computing the ideal answer, and this bound is shown empirically to survive the extension from single-qubit (Ferracin) to native two-qubit (Mills logical) non-Clifford rotations. The QED advantage is strongly observable-dependent: the abstract reports nearest-neighbour correlators $\langle Z_i Z_{i+1} \rangle$ have margins an order of magnitude larger than the magnetisation.
- **Error bars / noise assumptions.** Gate-type-dependent depolarising channels at $p_{1q} = 2 \times 10^{-4}$, $p_{2q} = 3 \times 10^{-3}$, $p_{\text{ro}} = 3 \times 10^{-3}$ (§III). Statevector / density-matrix simulation. Main results averaged over 20 seeds; envelopes use 95% confidence half-widths ($2\sigma_{95\%}$); the bare-vs-encoded magnetisation separation is reported significant at $p < 0.001$ (abstract). The QED margin itself is

small, $\Delta \sim 0.002$, because Forte sits near the pseudothreshold; the leading-order weight-2 estimate gives $p^* \approx 0.014$ at $n_s = 3$ vs Forte’s $p_{2q} = 0.003$, i.e. within a factor ≈ 5 (§V).

- **Classical baseline.** At four logical qubits the ideal quench is classically computable, and the exact $\langle M_z \rangle$ from a statevector simulation of $U(g, t)$ is used throughout as ground truth (Figs. 2, 4, E1–E3; Table E1 “ $\langle M_z \rangle_{\text{exact}}$ ” column). The point of the accreditation certificate is precisely that it does **not** require this baseline; the 4-qubit exact answer is available here only because the system is small enough to validate the certificate against.

2.6 Portability flags

Features the paper leans on that will matter for Brillant’s portability analysis:

- **Native MS $XX(\theta)/ZZ(\theta)$ two-qubit rotations are essential to the headline.** The “QED beats bare” result holds only for the native-MS compilation; under CNOT decomposition the same encoded circuit carries $\approx 2.5\text{--}25\times$ more two-qubit gates and falls below the pseudothreshold (bare wins). Any gate set lacking a native parametric two-qubit rotation must decompose these and likely loses the advantage.
- **All-to-all connectivity is assumed.** The iceberg stabilisers $S_Z = Z^{\otimes 6}$, $S_X = X^{\otimes 6}$ are weight-6 operators; the ancilla must couple to all six code qubits, and the native-MS encoder/logical rotations couple arbitrary pairs. On limited connectivity (e.g. heavy-hex superconducting) these require substantial SWAP overhead not modelled here.
- **No native mid-circuit measurement on the target.** The two-round QED variant needs an MCM stabiliser round, but Forte lacks native MCM, so the paper uses deferred measurement (extra qubits/depth). Platforms with native MCM (Quantinuum H-series, IBM) could realise the $\approx 10\%$ two-round gain more cheaply; conversely the deferred-measurement cost is Forte-specific.
- **Accreditation assumption depends on the gate set.** Ferracin’s bound assumes single-qubit non-Clifford gates (satisfied by the bare/VC circuits whose non-Clifford content is single-qubit R_X/R_Z); native-MS compilation makes the non-Clifford gates two-qubit, requiring the stronger gate-independence assumption underlying Mills logical accreditation. Porting the certificate across gate sets changes which accreditation theorem applies.
- **Operates deliberately near the pseudothreshold.** The benefit exists only in a narrow error-rate window: at much higher p_{2q} encoding does not help, and at much lower p_{2q} the bare circuit is already good enough. A target with materially different two-qubit error rate would land on a different side of the $N_{\text{bare}} \sim 5\text{--}6$ crossover.

2.7 Caveats

- **No hardware data.** This is entirely a simulation / simulator-projected study. Every $\langle M_z \rangle$, every γ , and every survival fraction comes from statevector or density-matrix simulation under Forte-rate depolarising noise. The paper itself states “On hardware, absolute γ values will likely be larger; the direction of each inequality is what the claim rests on (§VII), and labels the headline” Simulator-projected”. The “hardware submission” / “hardware-ready circuit package” is a specified but unexecuted manifest. Treat the fingerprint as a *proposed* resource budget, not a measured one.

- **Access.** Full access, no paywall: a complete ~10-page draft plus appendix Tables E1/E2 and Figs. E1–E6 are present in the folder. No sections were unreachable. There is no arXiv source bundle (local draft, no public link), so equations were transcribed from the PDF per the Math-notation rules rather than copied from TeX source.
- **Unresolved internal inconsistency.** The 8-qubit manifest figure (Table E2) versus the “10 physical qubits” statements (abstract, §V, §VI) is not reconciled in the text; flagged above, not invented away.

3 Portability matrix

Anchor: IonQ Forte (36-qubit static linear $^{171}\text{Yb}^+$ chain, all-to-all through shared motional modes). The source draft executes **nothing on hardware**: it specifies a “hardware-ready” circuit package for Forte (Table E2 manifest: 4 physical qubits bare, 8 encoded) and projects every result under a Forte-rate depolarising noise model ($p_{1q} = 2 \times 10^{-4}$, $p_{2q} = 3 \times 10^{-3}$, $p_{ro} = 3 \times 10^{-3}$; §III). The headline, that $[[6, 4, 2]]$ syndrome postselection with native-MS compilation produces a tighter certified envelope on $\langle M_z \rangle$ than the bare circuit at the Ising crossover, is therefore a **simulator-projected** claim (§VII: “On hardware, absolute γ values will likely be larger”), reported with 95% confidence half-widths over 20 seeds and a $p < 0.001$ bare-vs-encoded separation, but with a small magnetisation margin ($\Delta \approx 0.002$, an order of magnitude larger on nearest-neighbour $\langle Z_i Z_{i+1} \rangle$ correlators). Two claims must stay separate per the hardware-demo-vs-resource-analysis rule: (a) the executed object is a *proposed* \$6,726 / 20-cell package, never run; and (b) the QED-beats-bare inequality is a simulated projection, not a measurement. Verdicts below concern whether the encoded, certified workload (not merely the 4-qubit quench, which is classically trivial and runs anywhere) ports to each vendor, and how.

Julien’s fingerprint: four logical qubits encoded into the $[[6, 4, 2]]$ iceberg code (8 physical per the Table E2 manifest; abstract/§V/§VI say “10 physical”, unreconciled, flagged not invented away); second-order Suzuki-Trotter TFIM quench $H = -J \sum_{i=1}^2 Z_i Z_{i+1} - g \sum_{i=0}^3 X_i$ from $|0000\rangle$, $n_s \in \{1..5\}$ steps. Four structurally binding elements drive the matrix: (i) **native Mølmer-Sørensen $ZZ(\theta)/XX(\theta)$ two-qubit rotations are essential to the headline** (under CNOT decomposition the encoded circuit carries ≈ 2.5 to $25\times$ more two-qubit gates, 75 to 183 vs the native-MS 31 to 71, and falls below the pseudothreshold so bare wins); (ii) **all-to-all connectivity** (weight-6 stabilisers $S_Z = Z^{\otimes 6}$, $S_X = X^{\otimes 6}$ plus an all-pairs native-MS encoder); (iii) **no native mid-circuit measurement on the target** (the two-round variant is realised by deferred measurement on Forte; native-MCM platforms could take the $\approx 10\%$ two-round gain more cheaply); (iv) the workload **operates deliberately near the pseudothreshold** ($p^* \approx 0.014$ at $n_s = 3$ vs Forte’s 0.003, within a factor ≈ 5), so a target’s two-qubit error rate, not just its gate-set compatibility, decides which side of the $N_{\text{bare}} \sim 5$ to 6 crossover it lands on. A fifth thread runs through every row: the **accreditation theorem** depends on the gate set (Ferracin assumes single-qubit non-Clifford content; native-MS makes the non-Clifford gate two-qubit, invoking the stronger Mills logical bound). See `extraction.md` for full detail.

3.1 Summary

Modality	Vendor	Verdict	2026 fit
Trapped ion (static chain) (anchor)	IonQ Forte	runs as-is	yes (native package; simulator-projected, unrun)
Trapped ion (QCCD)	Quantinuum H2	runs with translation	partial: native fit strong; QED window narrows at H2’s lower p_{2q}
Trapped ion (static chain)	AQT LYNX	runs with translation	partial: public device specs thin; p_{2q} unplaced

Modality	Vendor	Verdict	2026 fit
Superconducting (heavy-hex)	IBM Heron r2	runs with translation	partial: data circuit yes; QED advantage not expected to survive
Neutral atom (digital)	Atom Computing AC1000	runs with translation (after compiler retargeting)	partial (pending native lowering)
Neutral atom (digital)	QuEra Gemini	runs with translation (after compiler retargeting)	partial: pending native lowering; p_{2q} above the QED window
Neutral atom (analog Rydberg)	Pasqal	architectural rewrite	no
Photonic (FBQC / MBQC)	PsiQuantum	infeasible today	no
Photonic (GBS / CV prototype)	Xanadu	infeasible today	no
NV / diamond	Quantum Brilliance	infeasible today	no

3.2 IonQ Forte – runs as-is (*anchor*)

Modality: Trapped ion (static chain).

Forte is the anchor because its documented native instruction set is exactly what the paper compiles to: a single parametric two-qubit entangler $ZZ(\theta) = \exp(-i\pi\theta Z \otimes Z)$ (the Mølmer-Sørensen-mediated rotation; $XX(\theta)$ is the same gate in a rotated single-qubit frame), plus GPi/GPi2/virtual-Z single-qubit rotations, on a 36-qubit all-to-all chain. The encoded circuit’s 8 (or 10) physical qubits and its all-pairs native-MS encoder and weight-6 stabiliser couplings sit well inside 36 all-to-all qubits, so there is no SWAP tax. Forte exposes **no native mid-circuit measurement**, which is why the two-round QED variant is realised by deferred measurement (Fig. 1 caption) at extra qubit/depth cost. Because the non-Clifford content is the two-qubit MS rotation, the encoded runs invoke **Mills logical** accreditation, not Ferracin. Forte’s $p_{2q} = 3 \times 10^{-3}$ is what places the workload near the pseudothreshold; the QED advantage is a property of this error rate, and the whole study is a simulator projection of it, not an executed run.

Translation: (anchor row, no translation needed).

Evidence:

- `dossiers/2026-draft-iceberg-642-ising-quench/extraction.md` (Julien’s fingerprint, Table E2 manifest, §III noise model, §VII simulator-projected caveat); local draft `paper.pdf`, no public link.
- [docs.ionq.com – getting started with native gates](https://docs.ionq.com/getting-started-with-native-gates) – Forte native set = GPi/GPi2/virtual-Z plus parametric $ZZ(\theta) = \exp(-i\pi\theta Z \otimes Z)$; MS gate listed as Aria-class only; no mid-circuit measurement documented for Forte, retrieved 2026-05-31.
- ionq.com/quantum-systems/forte – 36 qubits, any-to-any (all-to-all) entanglement, retrieved 2026-05-31.
- `vendor-notes/trapped_ion_qubits.tex` §6 – IonQ $^{171}\text{Yb}^+$ hyperfine, static linear chain, all-to-all via shared modes.

3.3 Quantinuum H2 – runs with translation

Modality: Trapped ion (QCCD).

H2 is the strongest second platform on gate-set grounds and the only target that clears all four binding fingerprint elements natively. Its native two-qubit set exposes a fixed ZZMax, a parameterised ZZPhase/RZZ, and a general SU(4) entangler $\text{TK2}(\alpha, \beta, \gamma) = \exp(-\frac{i}{2}\pi\alpha \text{XX}) \exp(-\frac{i}{2}\pi\beta \text{YY}) \exp(-\frac{i}{2}\pi\gamma \text{ZZ})$ (QASM Rxxyyzz), so the native-MS $ZZ(\theta)/XX(\theta)$ rotations the headline depends on lower one-to-one onto ZZPhase/TK2 with no decomposition tax. All-to-all QCCD connectivity carries the weight-6 stabilisers and the all-pairs encoder with no SWAP overhead, and unlike Forte H2 has **native mid-circuit measurement with qubit reuse**, so the two-round stabiliser variant runs natively rather than by deferred measurement (the $\approx 10\%$ two-round gain is available here). The non-Clifford content stays two-qubit, so Mills logical accreditation applies as on the anchor.

Translation: (i) retarget the native-MS data circuit from IonQ $ZZ(\theta)$ to H2 ZZPhase/Rxxyyzz (identity-shaped, one parameterised entangler per logical rotation); (ii) realise the stabiliser round with native MCM + reuse instead of deferred measurement, removing the Forte-specific ancilla/depth overhead; (iii) re-place the workload on the pseudothreshold curve. This last step is the catch and the reason the 2026 fit is **partial**, not **yes**: H2’s two-qubit infidelity is materially **below** Forte’s $p_{2q} = 3 \times 10^{-3}$, and the QED benefit exists only in a window around $p^* \approx 0.014$; at a lower error rate the bare circuit is closer to “already good enough”, so the *expected* QED margin on H2 is **smaller**, not larger, than on Forte. Preconditions for promoting this row off **partial** (costing-as-precondition, not footnotes): compiled encoded two-qubit-gate count and depth on the ZZPhase/TK2 basis; H2’s current p_{2q} versus the n_s -dependent pseudothreshold; and the shot budget for a target $\langle M_z \rangle$ envelope half-width at H2’s QCCD shot-rate, which is materially lower than a static-chain device. The “H2 is the better QED platform” intuition is a **hypothesis pending an actual compile-and-run**, and the pseudothreshold argument is a concrete reason it could fail.

Evidence:

- docs.quantinuum.com – H2 parameterised-angle 2-qubit gates – ZZMax, ZZPhase/RZZ, general SU(4) TK2 = Rxxyyzz(α, β, γ) native set, retrieved 2026-05-31.
- Quantinuum H2 Product Data Sheet v4.00 (2025-09-17), docs.quantinuum.com/systems/data_sheets – $N \geq 56/56$ qubits, all-to-all QCCD, native MCM with qubit reuse, retrieved 2026-04-21 (data-sheet page returned HTTP 403 on a 2026-05-31 re-fetch; figures carried from the precedent dossier’s verified read, see [dossiers/2026-arxiv-2604.16164-nonlinear-spectroscopy/](#)).
- [vendor-notes/trapped_ion_qubits.tex](#) §6 – Quantinuum QCCD, mid-circuit measurement.

3.4 AQT LYNX – runs with translation

Modality: Trapped ion (static chain).

AQT shares the substrate that makes the headline work: $^{40}\text{Ca}^+$ optical qubits in a static linear chain with a Mølmer-Sørensen-type entangler and all-to-all coupling within the chain, so the native parametric two-qubit rotation the paper depends on is substrate-native

rather than something to decompose. The current LYNX series reports a record Quantum Volume of 32 768 ($= 2^{15}$, i.e. of order 15 “QV qubits”), enough headroom for the 8-to-10 physical-qubit encoded circuit. The obstacle is not architecture but disclosure: AQT’s public English-language device specifications (exact addressable qubit count, two-qubit fidelity, and crucially mid-circuit-measurement support) are thinner than IonQ’s or Quantinuum’s, and cloud access is announced as “coming soon” rather than generally available.

Translation: (i) retarget the native-MS data circuit to AQT’s MS entangler (same modality, expected identity-shaped lowering, but the named gate API is not publicly pinned); (ii) the two-round variant’s stabiliser round needs either native MCM (undocumented publicly, so assume deferred measurement as on Forte) or the deferred-measurement fallback; (iii) Mills logical accreditation applies as the non-Clifford content is the two-qubit MS rotation. Preconditions for promoting off **partial**: a published two-qubit infidelity to place LYNX on the $N_{\text{bare}} \sim 5$ to 6 pseudothreshold curve (without it the QED-vs-bare verdict is undecidable for AQT), a confirmed addressable qubit count ≥ 8 for the encoded cell, an MCM statement for the two-round gain, and generally available cloud access. Treat the substrate-level MS and connectivity claims as provisional per the vendor note’s own uncertainty flag.

Evidence:

- aqt.eu/quantum-computing and aqt.eu/lynx – LYNX series, 19-inch rack-mounted trapped-ion system, record Quantum Volume 32 768, cloud readiness announced as forthcoming, retrieved 2026-05-31.
- `vendor-notes/trapped_ion_qubits.tex` §6 – AQT $^{40}\text{Ca}^+$ optical qubit, static linear chain, MS-type entangler; note’s own caveat that AQT device parameters (chain length, fidelities, gate times) are publicly thinner than IonQ/Quantinuum and “should be treated as provisional”.

3.5 IBM Heron r2 – runs with translation

Modality: Superconducting (heavy-hex).

The bare 4-qubit quench compiles to any Heron device trivially; the *certified, encoded* workload is where IBM struggles, on two sourced grounds. First, **connectivity**: Heron r2 is a 156-qubit heavy-hexagonal lattice with nearest-neighbour coupling only, so the weight-6 iceberg stabilisers and the all-pairs native-MS encoder require substantial SWAP insertion not modelled in the paper, inflating the encoded two-qubit-gate count past the pseudothreshold in the same way the CNOT-VC compilation does (and the paper already shows that compilation loses to bare). Second, **gate set**: IBM does expose a native parametric two-qubit rotation, the fractional gate $R_{ZZ}(\theta)$, but only for $0 < \theta \leq \pi/2$, and IBM’s documentation states fractional gates are **incompatible with dynamic circuits**. The two ingredients the headline leans on, a native parametric two-qubit rotation *and* a mid-circuit stabiliser round, therefore cannot coexist on Heron: use fractional R_{ZZ} and you lose native MCM (back to deferred measurement); use dynamic-circuit MCM and you lose the fractional entangler (back toward CX decomposition and its $\approx 25\times$ two-qubit overhead). $R_X(\theta)$ for the single-qubit content is unrestricted, so the bare/VC circuits whose non-Clifford content is single-qubit would use Ferracin accreditation; a fractional- R_{ZZ} encoded circuit would move to Mills logical.

Translation: the data circuit ports, so the verdict is **runs with translation**, but the headline QED advantage is **not expected to survive** and the 2026 fit is scoped to **partial** accordingly. Concretely: (i) route the weight-6 stabilisers and the encoder over heavy-hex with SWAPs and cost the resulting encoded two-qubit count, (ii) choose between the fractional- R_{ZZ} (no native MCM) and CX-plus-dynamic-circuit (no fractional entangler) branches, neither of which reproduces the Forte native-MS-plus-deferred-MCM configuration the headline rests on, (iii) place IBM’s p_{2q} (comparable to Forte at a few $\times 10^{-3}$) against the SWAP-inflated N_{bare} , which lands on the wrong side of the crossover. These are preconditions, not caveats: until the SWAP-inflated encoded two-qubit count and the chosen branch are costed, no positive QED verdict on IBM is defensible.

Evidence:

- [quantum.cloud.ibm.com – processor types](https://quantum.cloud.ibm.com/processor/types/heron/r2) – Heron r2 = 156 qubits, heavy-hexagonal lattice (nearest-neighbour), retrieved 2026-05-31.
- [quantum.cloud.ibm.com – fractional gates](https://quantum.cloud.ibm.com/processor/types/heron/fractional-gates) – native $R_{ZZ}(\theta)$ for $0 < \theta \leq \pi/2$ and $R_X(\theta)$ for any θ on Heron; fractional gates listed as incompatible with dynamic circuits; experimental feature, retrieved 2026-05-31.
- [vendor-notes/fractional_gates.tex](#) – native $R_{ZZ}(\theta)$ pulse replaces the CX--RZ--CX decomposition (one entangler vs two), Heron-class basis extension.

3.6 Atom Computing AC1000 – runs with translation (after compiler retargeting)

Modality: Neutral atom (digital).

The active primitive is a routine gate-model encoded Trotter circuit with a stabiliser round, no analogue evolution, so there is no paradigm change to execute on a digital neutral-atom machine; the gap is a missing *public* native lowering, which is a compiler-retargeting obstacle, not an architectural one. AC1000 advertises 1,200+ physical qubits, all-to-all connectivity, 1Q > 99.9%, 2Q > 99.6%, **native mid-circuit measurement with reset and reuse**, real-time conditional branching, and an OpenQASM/QIR API on ^{171}Yb nuclear-spin qubits. Capacity, connectivity, and the MCM/control-flow surface are all generous: the native-MCM and branching support means the two-round QED variant could run natively rather than by deferred measurement, the one fingerprint element where AC1000 is strictly better placed than the anchor.

Translation: the binding constraint is that the public product page names no Rydberg-mediated two-qubit generator and documents no native lowering of the parametric $ZZ(\theta)/XX(\theta)$ MS-equivalent rotation the headline needs; neutral-atom two-qubit gates are natively Rydberg CZ-type, so the “native-MS halves the encoding overhead” mechanism is a trapped-ion property that does not transfer for free. The port is therefore (i) submit the OpenQASM/QIR encoded circuit, obtain Atom’s native compiled circuit, (ii) check the compiled two-qubit count and depth against the > 99.6% 2Q-fidelity envelope, (iii) place the resulting p_{2q} ($< 4 \times 10^{-3}$, comparable to Forte) against the n_s -dependent pseudothreshold, and (iv) re-derive the shot budget against the achievable shot rate. Until the native lowering of the parametric two-qubit rotation is published or emitted, the QED-vs-bare margin cannot be predicted, so the 2026 fit is **partial**

(pending native lowering). If Atom discloses a parametric two-qubit rotation or a documented SU(4) compiler path, this row moves toward runs with translation without the retargeting qualifier.

Evidence:

- atom-computing.com/ac1000 – 1,200+ physical qubits, all-to-all connectivity, 1Q > 99.9%, 2Q > 99.6%, mid-circuit measurement with qubit reset and reuse, real-time conditional branching, OpenQASM/QIR API, ^{171}Yb nuclear-spin qubits, retrieved 2026-05-31.
- Modality framing cross-checked against dossiers/2026-arxiv-2604.16164-nonlinear-spectroscopy/ (digital-neutral-atom compiler-retargeting precedent).

3.7 QuEra Gemini – runs with translation (after compiler retargeting)

Modality: Neutral atom (digital).

Gemini is a 260-qubit digital gate-mode ^{87}Rb system with all-to-all connectivity through Dynamic Qubit Array shuttling, a QASM circuit-level interface, and a published 99.2% global two-qubit fidelity at 1 shot/s. As with AC1000 the workload is a routine gate-model circuit, so the obstacle is compiler retargeting, not paradigm change: capacity (260 qubits for an 8-to-10 qubit encoded cell) and connectivity are comfortable, but the public page names no native two-qubit generator and documents no lowering of the parametric MS-equivalent $ZZ(\theta)/XX(\theta)$ rotation. The page does not document native mid-circuit measurement either, so assume the two-round variant uses deferred measurement as on Forte.

Translation: (i) lower the encoded native-MS data circuit to QuEra’s documented native gate family and cost the compiled two-qubit count and depth on the actual Gemini compiler; (ii) realise the stabiliser round by deferred measurement absent a documented MCM path; (iii) place Gemini’s p_{2q} on the pseudothreshold curve, and here the number is decisive: a 99.2% global 2Q fidelity is $p_{2q} \approx 8 \times 10^{-3}$, **above** Forte’s 3×10^{-3} and on the “much higher p_{2q} , encoding does not help” side of the QED window, so even a clean native lowering would likely not reproduce the QED-beats-bare result; (iv) at 1 shot/s the 11,000-shots-per-cell hardware allocation is hours of beam time per cell before queueing. The 2026 fit is therefore **partial: pending native lowering; p_{2q} above the QED window**. Preconditions to promote it: the published native lowering, and a device p_{2q} inside the pseudothreshold window.

Evidence:

- quera.com/gemini – 260 physical ^{87}Rb qubits, digital gate-mode, all-to-all via Dynamic Qubit Array shuttling, 99.2% global 2Q fidelity, 1 shot/s, QASM circuit-level interface, retrieved 2026-05-31.

3.8 Pasqal – architectural rewrite

Modality: Neutral atom (analog Rydberg).

Pasqal’s production mode is analogue: programmed Rydberg evolution under a global laser drive with a native Ising-type generator ($\sim C_6/R_{ij}^6$ pairwise interaction). The 4-qubit TFIM quench itself is close to the native Rydberg Ising Hamiltonian and could be re-expressed as analogue evolution, but that is a different computational paradigm from the paper’s gate-model encoded circuit, and crucially the paper’s **actual novelty, the $[[6, 4, 2]]$ error-detection and the accreditation certificate, has no analogue form:** syndrome extraction on weight-6 stabilisers and the native-MS logical rotations are gate-model primitives, and stabiliser postselection requires per-shot syndrome readout that the global-drive analogue mode does not provide.

Translation: a port is a research project, not a transpilation. One would re-express the bare quench against the native Rydberg Ising generator (losing the transverse-field structure unless XY/digital-analogue mode is invoked, which is emulator/backend-dependent in current Pasqal material rather than a production QPU promise), and there is no analogue route to the iceberg stabiliser postselection or to the native-MS two-qubit non-Clifford rotations that the Mills logical accreditation assumes. None of the four binding fingerprint elements (native-MS rotations, weight-6 stabiliser extraction, MCM, pseudothreshold-windowed QED) survive the move to analogue. Pasqal’s digital-analogue mode, once it exposes arbitrary-angle individually-addressable two-qubit rotations and mid-circuit syndrome readout on a production QPU, would move this toward **runs with translation**; today it does not.

Evidence:

- docs.pasqal.com – QoolQit Rydberg model – analogue Ising-mode Hamiltonian with C_6/R_{ij}^6 interaction; no gate-model stabiliser-extraction or parametric two-qubit-rotation mode, retrieved 2026-04-21 (via [dossiers/2026-arxiv-2604.16164-nonlinear-spectroscopy/](#)).
- [dossiers/2026-arxiv-2604.16164-nonlinear-spectroscopy/](#) – analogue-Rydberg architectural-rewrite precedent for a gate-model Trotter workload.

3.9 PsiQuantum – infeasible today

Modality: Photonic (FBQC / MBQC).

PsiQuantum’s commercial roadmap is a fault-tolerant photonic machine built on fusion-based measurement-based computation with telecom-band single photons. There is no 2026 programmable gate-model QPU exposing user-defined real-time encoded Trotter circuits on 8-to-10 physical qubits with native parametric two-qubit rotations, so none of the fingerprint elements have a 2026 target. The workload needs a near-term NISQ-style substrate with cheap shot loops and a documented entangling-gate path; PsiQuantum is explicitly not building that. Horizon: once the utility-scale fault-tolerant systems come online (targeted late-decade), the same application re-expressed under an FTQC cost model, with logical-qubit overhead replacing postselection shot loops, is a different port and a different research question; post-2030 if at all on this workload’s terms.

Translation: infeasible today; uncertain horizon, needs follow-up only once a programmable gate-model photonic QPU exists. The gap is sourced below.

Evidence:

- psiquantum.com/technology – telecom-band single photons, fusion-based FTQC roadmap, component-level fidelities (not a programmable 8-to-10-qubit gate-model QPU), retrieved 2026-04-21 (via [dossiers/2026-arxiv-2604.16164-nonlinear-spectroscopy/](#)).

3.10 Xanadu – infeasible today

Modality: Photonic (GBS / CV prototype).

Per the multi-device rule, Xanadu’s two public substrates are split, and neither runs this workload in 2026:

- **Borealis / X-series (Gaussian Boson Sampler).** A photonic sampler, not a programmable gate-model QPU; it cannot express an encoded TFIM Trotter circuit with stabiliser postselection at all (wrong machine). No access path to the fingerprint primitives. Status: available but architecturally disqualified.
- **Aurora (continuous-variable networked prototype).** A 12-mode CV research prototype (Madsen et al., *Nature* **638**, 912-919, 2025: four modular racks, 35 chips, 84 squeezers, 36 photon-number-resolving detectors), not a user-programmable QPU on cloud in 2026, with the near-term focus stated as optical-loss reduction before scaling. Even if opened, a discrete-variable to continuous-variable encoding layer (e.g. GKP) would have to precede any port of the gate-model iceberg circuit. Status: prototype, not user-accessible.

Translation: infeasible today on both devices; the GBS branch is a permanent modality mismatch for a gate-model encoded circuit, the Aurora branch tracks substrate availability and a missing DV-to-CV encoding layer. Horizon: once Xanadu exposes a programmable DV substrate the Aurora row would move to **architectural rewrite**, not to a routine port.

Evidence:

- xanadu.ai – [Aurora announcement \(2025-01-22\)](#) – four-rack modular CV prototype, optical-loss-reduction focus, retrieved 2026-04-21 (via [dossiers/2026-arxiv-2604.16164-nonlinear-spectroscopy/](#)).
- Madsen et al., *Nature* **638**, 912-919 (2025), “Scaling and networking a modular photonic quantum computer”; Borealis/X-series GBS distinction.

3.11 Quantum Brilliance – infeasible today

Modality: NV / diamond.

Quantum Brilliance’s public material describes room-temperature NV-diamond technology, NV arrays, and chip integration, but does not commit a 2026 device with the 8-to-10 addressable gate-model qubits, the all-to-all native-MS- equivalent two-qubit rotations, the

weight-6 stabiliser-extraction route, or the two-qubit fidelity envelope this encoded workload needs. No public material documents an NV gate-model system with those control features today, so the fingerprint has no 2026 target on this substrate.

Translation: infeasible today; uncertain horizon, needs follow-up. Plausible only after a publicly specified multi-qubit NV module with a documented entangling path and the qubit count for an encoded cell; no public roadmap commits to one. The gap is sourced below.

Evidence:

- quantumbrilliance.com/quantum-computing and quantumbrilliance.com/quantum-brilliance-hardware – descriptive room-temperature NV pages without a committed 2026 gate-model qubit count or entangling-gate specification, retrieved 2026-04-21 (via dossiers/2026-arxiv-2604.16164-nonlinear-spectroscopy/).

Commentary (\$ \$360 words). The workload is two objects with very different portability. The 4-qubit TFIM quench is classically trivial and runs on any gate-model machine; the paper’s real payload is the *certified*, $[[6, 4, 2]]$ -*encoded* circuit, and that is where portability narrows. The decisive axis is not gate-set compatibility but **pseudothreshold-windowing**: the QED-beats-bare result exists only in a band around $p^* \approx 0.014$, so each target’s two-qubit error rate, not just whether it can express the circuit, decides the verdict. This produces an anti-hype inversion worth flagging to the curator. The gate-set-cleanest second platform, **Quantinuum H2**, lowers the native-MS data circuit one-to-one onto **ZZPhase/Rxxyyzz**, carries the weight-6 stabilisers on all-to-all QCCD with no SWAP tax, and adds native MCM for the two-round variant, yet its *lower* two-qubit infidelity likely yields a *smaller* QED margin than Forte, because the bare circuit is already closer to good enough. A better device does not imply a bigger error-detection win here. Conversely **QuEra Gemini**’s 99.2% 2Q fidelity ($p_{2q} \approx 8 \times 10^{-3}$) sits above the window, so even a clean native lowering would likely not reproduce the advantage.

Among the strongest second platforms: **Quantinuum H2** on capability (native parametric entangler, all-to-all, native MCM, Mills-logical-compatible two-qubit non-Clifford content), with the caveat that its headline benefit is a hypothesis pending an actual compile-and-run and may be smaller than on the anchor. The structural gaps: **IBM Heron r2** cannot hold a native parametric two-qubit rotation and a mid-circuit stabiliser round at once (fractional gates are documented as incompatible with dynamic circuits), and its heavy-hex connectivity imposes a SWAP tax on the weight-6 stabilisers that the paper does not model, so the QED advantage is not expected to survive; **Atom Computing AC1000** and **QuEra Gemini** are capacity-ample but await a public native lowering of the parametric two-qubit rotation, a trapped-ion-specific mechanism that does not transfer free to Rydberg CZ. The analogue **Pasqal** row loses the entire certification payload (no analogue syndrome extraction); the photonic (**PsiQuantum FTQC**, **Xanadu GBS-plus-Aurora-prototype**) and **Quantum Brilliance NV** rows are not 2026 ports.

Two disciplines bind every positive non-anchor row, and the curator should read them as gating, not advisory. First, all numbers are **simulator-projected**: no circuit was executed on any hardware, so each verdict concerns expressibility and the projected pseudothreshold placement, never a measured run. Second, the **compiled encoded two-qubit count, depth, and the shot budget against a chosen ε** are preconditions for promoting any row off **partial**; the paper supplies these only for the Forte native-MS package. Modalities outside the gate-model frame (D-Wave annealing, Fujitsu/Toshiba quantum-inspired, HRL silicon-spin) are deliberately omitted as rows: none

can express an encoded gate-model Trotter circuit with stabiliser postselection, so they are infeasible by the same reasoning as the photonic GBS branch and are noted here rather than given separate rows.

3.12 Residual caveats

- The **IonQ Forte** anchor remains a simulator-projected, unrun package; the matrix separates the proposed \$6,726 / 20-cell resource package from any measured hardware result, but readers should not treat the QED advantage as experimentally observed.
- The **IonQ Forte** anchor carries the paper’s unreconciled 8-vs-10 physical-qubit discrepancy for the encoded cell.
- The **Quantinuum H2** row is a capability-clean translation, but its QED benefit is framed as a hypothesis because H2’s lower p_{2q} can shrink the bare-vs-encoded margin.
- The **AQT LYNX** row rests on substrate and Quantum Volume evidence while public qubit-count, two-qubit-infidelity, MCM, and cloud-access details remain thin.
- The **IBM Heron r2** row leaves the QED claim at partial because heavy-hex routing and the fractional-gate versus dynamic-circuit branch have not been costed for the encoded circuit.
- The **Atom Computing AC1000** row depends on an unpublished native lowering for the MS-equivalent parametric two-qubit rotations.
- The **QuEra Gemini** row depends on an unpublished native lowering, and its published 99.2% global two-qubit fidelity and 1 shot/s rate make the QED margin and shot budget unfavourable.
- The **Pasqal** row treats the bare TFIM as analogue-adjacent, but the iceberg stabiliser postselection and Mills-style gate-model certificate have no documented analogue production-QPU route.
- The **PsiQuantum** row is infeasible on the matrix’s NISQ terms because the public roadmap is utility-scale FBQC rather than a 2026 programmable gate-model QPU.
- The **Xanadu** row correctly separates Borealis/X-series from Aurora; Borealis remains a sampler mismatch, while Aurora remains a prototype whose discrete-variable encoding and user-programmable access path are not available in the cited public material.
- The **Quantum Brilliance** row lacks a public 2026 multi-qubit gate-model specification with the qubit count, entangling operations, all-to-all route, and error rates required by the encoded workload.
- Across positive non-anchor rows, the compiled encoded two-qubit count, depth, and shot budget remain preconditions rather than measured portability results.