Supplementary Material

Appendices A–C

Ground-State Preparation by Projection onto the Maximal Decoherence-Free Subspace

Supplement to the main article

This document carries the three appendices of the main article: the IQS Part citation index (section A), the data and reproduction record (section B), and the complete experimental record (section C), which includes the methods, the tiling procedure, the Phase 3–6 results, the closure tests and the oracle reference tests.

Numbering follows the main article: appendix sections are lettered A–C, and theorems, tables, figures and equations carry that letter, so table C.25 here is unambiguous against the main article's tables. References into the main article — theorem 3.12 and theorem 4.54, for instance — carry that article's numbering and resolve from this file alone; no auxiliary file of the main article is needed to build this document.

Contents

A IQS Part Citation Index **1**

B Data and Reproduction **5**

C Experimental Record **8**

C.1 Results . . . 9

C.1.1 Phase 3: Core Oracle Execution (Frustrated Ising) . . . 9

C.1.2 Per-Tile Marginal Analysis . . . 9

C.1.3 Phase 4: Cross-Domain Applicability . . . 11

C.1.4 Phase 5: k-Local Generalization ($k > 2$) . . . 12

C.1.5 Phase 6: Admissible Problem Encoding . . . 13

C.1.6 Hardware Tests of the Closures . . . 16

C.1.7 Oracle reference tests . . . 21

C.1.8 Proof oracles: theorems tested after they were proved . . . 27

C.1.9 P5 — the closure stage is stationary under repetition . . . 30

C.2 Complete Job Registry . . . 31

A IQS Part Citation Index

Purpose and construction

This appendix establishes provenance for every result in the body of the paper. The IQS/EQS formalization runs to Part 297, developed between 2024 and 2026 [1]; the

present work draws on 53 of them. The most recent Parts are status boards: Part 297 supersedes Part 287 and consolidates Parts 292–296. Every Part this paper cites lies at or below Part 278. Each row of table A.1 maps a numbered theorem, definition or equation of this paper to the Part or Parts from which its content is taken, and states that content in the notation used here.

Three conventions govern the index.

1. **Attribution is to the originating Part, not to the most recent restatement.** Where a result is developed across several Parts, the row names the Part that first establishes it and any Part that materially strengthens or corrects it. Later Parts that merely cite it are omitted.
2. **Notation is normalised.** The corpus uses several notations for the same object across its development. Every row is rendered in the notation of sections 2 to 4 of this paper. The "Content" column states the result in the form this paper uses; the Part number identifies the source, in which the original notation may be consulted directly [1].
3. **Scope conditions travel with the result.** Where a Part establishes a statement only for a particular realization — most commonly the two-qubit host of **Part 76** — that restriction is carried into the corresponding theorem in the body and is not silently dropped here. The consequences are collected in section 4.6.

Verification status

Provenance claims of the form "Part n states X" are mechanically checkable, and a subset has been checked. The procedure is to extract the cited statement from the source text of the Part by pattern match on the notation of the corpus, normalise it to the notation of this paper, and compare it against the form used in the row. Applied to the Parts 258 and 259 index rows this confirmed 59 citations with none failing. Rows not covered by that check were verified by direct reading of the corresponding Part.

The Parts themselves are not reproduced here, and this index is a provenance record rather than a dependency. Each row names a result, the Part or Parts it came from, and the statement in the notation this paper uses.

No result requires access to the corpus. Every theorem, lemma, corollary and proposition in this work carries a proof: 51 in sections 2 to 4 and 4 in this supplement, 55 in all. Of those 55 proofs, 45 cite no Part at all. The remaining ten cite one for either of two reasons, neither of which is evidential. Seven reproduce an argument first given in the corpus rather than deriving it independently, and the citation is attribution: theorem 2.5 is **Part 61**'s injectivity theorem together with **Part 51** §XIV; theorems 4.10 to 4.12 are **BCM 30**'s three preservation theorems; theorem 4.38 uses **BCM 91**'s explicit counterexample; theorem 3.8 follows **Part 130** §2–3; and lemma 4.58 is **Part 255**'s lattice argument, which theorem 4.59 then sharpens by an idempotence step of its own. Three adopt a definition or a criterion the paper does not re-derive, and there the citation prevents circularity: $\equiv_{\mathrm{IQS}}$ is defined in **Part 82** and **Part 92**, so theorem 3.8 would otherwise appear to define the relation it classifies; $\mathcal{D}_{\mathrm{phys}}$ is **Part 132** Eq. (17)–(18); and the criterion theorem 3.5 evaluates is **Part 134** §8.

A reader who cannot consult the corpus can therefore still check every claim made here. What the corpus supplies is where the ideas came from, not why they are true.

Where it leaves a question open — remark 3.6 for the uniqueness criterion, remark 4.56 for the E- and M-channels — this paper says so and states what closes it.

Table A.1: Complete IQS Part citation index.

Theorem / Equation	**IQS Part(s)**	**Content**
D-Operator Chain (Equation (1))	**Part 27**, **Part 34**, **Part 47**	$D_1 \to D_2 \to \mathrm{RCl} \to D_3$; $D_4 \rightsquigarrow D_3$
Interaction Hamiltonian (Equation (3))	**Part 174**	$H_{\text{int}} = \sum g_{ab} O_a^{(I)} \otimes O_b^{(E)}$
Boundary Hamiltonian (Equation (4))	**Part 106**, **Part 187**	$H_\partial^{(N)} = \sum A_\mu \otimes B_\mu$
Explicit Realization (Equation (5))	**Part 89**, **Part 214**	$A = Z_1 + Z_2$, $\mathrm{spec}(A) = \{-2, 0, 2\}$
Host Space (Definition 2.2)	**Part 124**, **Part 132**, **Part 215**	$\mathcal{H}_H = \ker(A) \simeq \mathbb{C}^2$; $P_H = P_{\text{DFS}}^{\max}$
Compression Map (Definition 2.4)	**Part 168**, **Part 181**, **Part 189**	$C_H(A) = P_H A P_H$; quotient isomorphism
Infinite Faithfulness (Theorem 2.5)	**Part 111**	Injective history representation
Master Chain (Equation (13))	**Part 109**, **Part 128**	Full operator derivation
Protected Generators	**Part 158**	$G_\alpha = P_H A_\alpha P_H$
Physical Uniqueness (Theorem 3.5)	**Part 114**, **Part 115**, **Part 141**	$\lvert\mathfrak{R}/\!\sim\rvert = 1$
Sector Uniqueness, R-channel (Equation (21))	**Part 179**, §8, Eq. (19)	$\widehat{G}_R \in \mathbb{R}\sigma_x$, hence $\mathrm{rank}\,\Theta_R = 1$, given nonzero response
Sector counts, E- and M-channels (Remark 4.56)	**Part 127**, **Part 179**, **Part 203**	$N_E = N_M = 1$ recorded *not yet evaluated* (**Part 127**) and *not derivable* from the present axioms (**Part 179** §7); rigorous as an implication only (**Part 179** Eq. (34))
Decoherence Immunity (Theorem 3.7)	**Part 125**	$\mathfrak{A}_H \subseteq \mathfrak{A}_{\text{dec}}(t)$
Recursive Update (Equation (28))	**Part 123**	$X_{n+1} = \mathfrak{F}_{\Delta t}(X_n)$
IQS Invariant (Theorem 3.8)	**Part 135**, **Part 208**, **Part 224**	$\mathbf{Z}_\star(X) = (\Gamma_{IE}, \Pi_{\text{rec}} T_X)$
k-Local Conservation (Theorem C.1)	**Part 260**	$[H_k, A_1] = 0$ for all $k \geq 2$
The Compression Encoding Theorem (Theorem 3.12)	**Part 253**	$\mathcal{Q}(A) = V^* A V$; encoding for target problem family

Theorem / Equation	**IQS Part(s)**	**Content**
Host Symmetry Group (Definition 4.51)	**Part 113** Eq. (25)–(26), **Part 121** §9, **Part 127**	$G_H = \mathrm{Stab}(P_H)$ and $N_\sigma = \dim \mathrm{Hom}_{G_H}(\mathcal{T}_\sigma, \mathfrak{A}_H^{\mathrm{loc},sa})$
Canonical host realization (scope of Theorem 4.54)	**Part 76** Eq. (7)	$P_H = \|01\rangle\langle 01\| + \|10\rangle\langle 10\|$; the two-qubit realization to which the scope conditions of this paper refer
Physically admissible intertwiners (Equation (64))	**Part 132** Eq. (18)	$\mathfrak{I}_\sigma^{\mathrm{phys}} = \{\Pi \in \mathrm{Hom}_{G_H} : \mathcal{D}_{\mathrm{phys}}(\Pi) = 0\}$
Dynamical Selection (Theorem 4.54)	**Part 132**, §13, Eq. (17); **Part 179**, §8	commutation with H_S^H with scalars excluded; both routes give $\widehat{G}_R = \sigma_x$, $N_R = 1$
Channel status, what is open and how (Remark 4.56)	**Part 139**, **Part 144** §11, **Part 146**, **Part 156**, **Part 157**	$\mathcal{T}_E = 2V_+ \oplus 2V_-$ supplied (**Part 146**); $\mathcal{T}_M$ not instantiated; the E gap is $\mathrm{rank}\,\Theta_E = 1$; "$R$: 1 proved" (**Part 144**)
D_4 as a primitive (Equation (57))	**Part 257**	$D_4 : \Omega_H \to \mathcal{P}$, where $\mathcal{P}$ is the set of pairs (P, Φ_4) with P satisfying the DFS conditions
Protected projector at finite regulator (Equation (56))	**Part 97**	$H_{\mathrm{int}}^{(N)} \to A_\mu^{(N)} \to K_{\mathbf{a}}^{(N)} \to \mathcal{A}_{\mathrm{prot}}^{(N)} \to P_H^{(N)}$, determined; leaves open whether $P_H^{(N)} \neq 0$
Site-resolved mediators (Lemma 4.46)	**Part 229**	$A_u^{\mathrm{mic}} = Z_u$, the mediating set this paper's bounds are stated for
Channel count at host dimension d (Theorem 4.47)	**Part 278**	$N_\sigma = d - 1$ for a host of dimension d
Rank Obstruction (Theorem 5.2)	**Part 256**	$\mathrm{rank}(P_R \otimes \mathbb{1}) = \mathrm{rank}(P_R)\, 2^{N-\|R\|}$
Support Obstruction and Dichotomy (Theorem 5.3)	**Part 256**	no H_S-invariant subspace has bounded support; a rank-two host is delocalised or not invariant
The Junction (Theorem 4.7)	**Part 292**	K_q admissible; FC† realizes as $[H_{\mathrm{total}}, Z_1 + Z_2] = 0$, compressing to $P_H(Z_1 + Z_2)P_H = 0$; scoped to **Part 76** at $\kappa = 0$
Back-edge (Definition 4.15)	**Part 259** Cor. 1	$\mathfrak{J}_\sigma = \mathrm{Hom}_{G_H}(\mathcal{T}_\sigma, \mathfrak{A}_H)$; $\kappa_E, \kappa_R, \mu_c$ are its coordinates

Theorem / Equation	IQS Part(s)	Content
Channel Dimension (Theorem 4.16)	**Part 293**, composing **Part 259** with **Part 278**	$N_E = d - 1$, so $\dim \widehat{\Pi} = 3(d-1)$; the back-edge is three numbers iff $d = 2$
The two-dimensional host (Proposition 4.28, Remark 4.29)	**Part 278** §4–5, **Part 284** §6, **Part 296**	$N_\sigma = 1$ at $d = 2$ alone; recorded as the fifth degenerate canonical choice under **Part 284** §6's rule
Parity Obstruction (Theorem 4.21, Corollary 4.23)	**Part 294**	$\{B, H_E\} = 0$ forces $\det B \det H_E = 0$ at odd $\dim \mathcal{H}_E$, so W-3 is obstructed before signalling is reached
Physical Uniqueness (Theorem 3.5, Remark 3.6)	**Part 134**, §8 and Eq. (7); **Part 115**, Eq. (5)	$\lvert\mathfrak{R}/\!\sim_{\text{phys}}\rvert = 1 \iff \dim \operatorname{Im} \Theta = 1$; the reduction is supplied by the corpus, the evaluation is Theorem 4.54
Infinite-history faithfulness (Theorem 2.5)	**Part 61**, Thm. IH-Injectivity; **Part 51**, §XIV	$T^{\times}_{\text{rec}}$ injective in the dual topology under five stated hypotheses; equivariance descends it to G-orbits
Componentwise IQS equivalence (Equation (27))	**Part 82**, **Part 92**	$X \equiv_{\text{IQS}} Y$ iff Γ_{IE} and $\Pi_{\text{rec}}(T)$ agree; **Part 82** shows instantaneous data alone does not separate
Canonical quotient coordinate (Theorem 3.8)	**Part 130**, §2–3	$\mathbf{Z}_\star$ is that pair, so the biconditional is immediate and $\bar{\mathbf{Z}}_\star$ is a bijection onto $\operatorname{Im} \mathbf{Z}_\star$

B Data and Reproduction

The measurement archive

The measurement archive is the raw quantum-processor telemetry: 204 JSON files totalling approximately 785 MB, together with seven documentation files that index them, carrying the measurement record of all 55 executions listed in tables C.23 to C.25. Every numerical claim in section C is computed from these files. The archive is available in full from the corresponding author on reasonable request; it is not attached to this submission on account of its size, and the job identifiers of section C.2 allow every execution to be confirmed with the provider independently of it.

The files are held individually and are not organised into directories. Each name is therefore self-identifying: a prefix naming the series it belongs to, followed by the identifier of the job that produced it, so a file can be matched against the platform record and against the tables of section C without reference to any directory it once sat in. Table B.1 gives the prefixes.

Source code and the IQS Part source texts are not included. The Parts are identified by number throughout and indexed in section A; the corpus itself is cited as [1]. Where a

construction is used, it is specified in the body in enough detail to be reimplemented — the isometry of theorem 3.12, the ladder of theorem 4.31, the preparation of eq. (C.11) — rather than deferred to a file.

Table B.1: The measurement archive, grouped by name prefix. Files are held individually; the prefix carries the grouping that a directory would otherwise have carried, and the remainder of each name carries the identifier of the job the file came from. Counts in the second column are data files; the seven documentation files are listed separately.

Name prefix	Files	Size	Contents and where reported
`phase3_*`	3	6 MB	core oracle on three backends (section C.1.1)
`phase4_*`	6	8 MB	chemistry, cryptography, optimisation (section C.1.3)
`phase5_*`	6	16 MB	$k = 3, 4, 5$ (section C.1.4)
`phase6_*`	6	2 MB	six encoded problem instances (section C.1.5)
`closure_*`	7	0.2 MB	F13–F17 theorem tests (section C.1.6)
`oracle_*`, `ORACLES_*`	10	0.5 MB	O1–O6 reference tests, one file per job, plus the derived-table and transpiled gate-count records (section C.1.7)
`proof_*`	2	< 0.1 MB	P1, P3 and P4 proof oracles, one file per submission, each covering three backends (section C.1.8)
`oracle_rawrun_*`	4	0.3 MB	superseded: the O1 run with the measurement-ordering fault, and the first O2 build
`proof_superseded_*`	7	0.1 MB	superseded: four P2 control designs that did not hold depth, and the pad sweep
`maxdensity_*`	10	20 MB	maximum-tile-density runs
`rawdump_*`	10	8 MB	unprocessed per-backend counts
`fseries_*`	131	724 MB	the F1–F12 instrumented arc
Data files	**202**	**785 MB**	
documentation	7	—	two manifests, two readmes for the superseded series, a data index, a file list and a readme
Total	**209**	**785 MB**	

File format

Most files carry `raw_measurement_counts`, a map from bitstring to count, together with `job_id`, `backend` and `active_qubits_utilized`. The Phase 6 files use `counts` with `total_shots` and `top_10`. The closure-test files (`F13`–`F17`) additionally carry the pre-registered prediction alongside the observation and a per-tile breakdown, so each is self-describing: the falsification criterion is in the file, not only in this paper.

Size is concentrated in the `fseries_*` files, of which 20 exceed 10 MB and one reaches 70 MB. These are complete count dictionaries at 50,000 shots over registers of more than 100 bits, retained unabridged because the per-tile analyses of section C.1.2 are computed

from them. The remaining 54 data files total 61 MB and carry every result reported in section C; the `fseries_*` set is the instrumented arc that preceded them and supports no result reported here on its own.

Checks performed

The results below were established by direct computation and require no hardware. They are reported here as the checks they are; each is specified in the body at the point it is used, and is reproducible from that specification.

Table B.2: Classical checks and what each discharges.

Check	**Result**
Single projection step on frustrated Erdős–Rényi MaxCut, $N \in \{6, 8, 10, 12\}$	table 8; 7 assertions per size, 0 failures
Depth of the projection	lemmas C.8 to C.10; the erratum of remark C.11
$[H_k, A_1] = 0$ for the structured family	theorem C.1, verified for $k = 2, \ldots, 8$
Microscopic coupling constants	the calibration underlying section 4.8.4
The compression reduction	theorem 4.26
Provenance of the citation index	59 citations confirmed, 0 failures
The junction	theorem 4.7; 11 checks, 0 failures
Back-edge dimension and the parity results	theorems 4.17 to 4.38; 10 checks, 0 failures
Corpus citation graph	the reading-order statistics of section 5.2
The uniqueness closure, $\dim \operatorname{Im} \Theta$	theorem 3.5; the criterion of **Part 134** §8 evaluated at $d = 2, 3, 4, 5$, returning $1, 2, 3, 4$, and recovering $\widehat{G}_R = \pm\sigma_x$
The proofs supplied in this revision	theorems 5.2 to 4.38 and theorem C.1; 9 checks, 0 failures

Every prediction reported for a hardware run in section C was computed by one of these before the corresponding job was submitted, and is recorded in the archived JSON alongside the observation.

Software used

Table B.3 states the versions under which every archived result was produced; they are recorded because transpilation and the classical predictions both depend on them.

Table B.3: Software under which the archived results were produced.

Component	Version	Role
Python	3.12	interpreter
`qiskit`	2.5	circuit construction and transpilation
`qiskit-ibm-runtime`	0.49	job submission and retrieval
`numpy`	2.5	linear algebra, classical prediction
`scipy`	1.18	singular-value and eigenvalue routines

Reproduction

Results divide into three classes with different reproduction guarantees.

1. **Analytical results** — theorems 4.2 to 4.39 and the checks of table B.2 — are deterministic and reproduce exactly from the specifications given in the body.
2. **Analyses of the archived telemetry** reproduce exactly from the archived files without hardware access, including the per-tile marginal analysis of section C.1.2 and every entry of tables C.8, C.11, C.13 and C.18.
3. **New hardware executions** will differ from those reported here, as calibration drifts between sessions and placement varies within a session. Section C.1.6 quantifies the scale: on `ibm_kingston` the same two-gate circuit gave total variation 0.603 on a single tile and a median of 0.038 over 60 tiles of that same processor, the single tile having landed on a poorly calibrated pair. Tiled statistics are reported as medians over tiles for this reason.

All 55 job identifiers of tables C.23 to C.25 remain queryable through the IBM Quantum Platform API and constitute the primary record. The archived JSON files are derived from those jobs and carry the originating identifier in the filename, so the archive can be checked against the platform without reference to this paper.

Corrections

Two faults in the analysis were identified after execution and are documented rather than silently repaired, in section C.1.5 and section C.1.6. Both concerned the mapping between measured bitstrings and state indices; both were detected because a pre-registered prediction was available to disagree with; neither required additional quantum time to diagnose. In both cases the raw counts in the archived files are the counts as measured — the correction is in the interpretation, not in the archive, so a reader can rederive both the faulty and the corrected numbers from the same files.

A third correction concerns the depth attributed to section C.1.7 and is recorded at remark C.11.

C Experimental Record

This appendix carries the complete experimental programme: hardware platform and architecture, the deployed Hamiltonians, all measured results, and the full job registry.

It is separated from the body so that the framework of sections 2 to 4 can be read as mathematics, with every empirical claim cross-referenced here.

The programme comprises 55 quantum-processor executions in five groups, which should not be conflated. Twenty-one are the Phase 3–6 executions of section C.1, at 1,050,000 shots; twelve are the theorem tests F13–F17 of section C.1.6, at 2,860,000 shots; twelve are the oracle reference tests of section C.1.7, at 11,700,000 shots; six are the proof oracles of section C.1.8, at 1,230,000 shots; and four are the closure-stationarity test of section C.1.9, at 480,000 shots. Where this paper writes "21 executions" it means the first group only; totals over all five are written as 52. Tables C.23 to C.25 list them separately for the same reason.

C.1 Results

C.1.1 Phase 3: Core Oracle Execution (Frustrated Ising)

The core oracle ($H_2 = X \otimes X$, $t = 1.0$) was executed across all three backends at 50,000 shots. Table C.1 reports the amplification results.

Table C.1: Phase 3 Oracle results. The amplification factor is the ratio of the observed top-state probability to the uniform random baseline $1/2^M$.

Backend	2^M	States Explored	Eliminated	Top-State Prob.	Random Baseline	Amplification
`ibm_marrakesh`	2^{33}	49,984	99.9994%	0.0040%	1.16×10^{-8}%	$3.44 \times 10^5\times$
`ibm_kingston`	2^{38}	50,000	99.99998%	0.0020%	3.64×10^{-10}%	$5.50 \times 10^6\times$
`ibm_fez`	2^{34}	49,993	99.9997%	0.0040%	5.82×10^{-9}%	$6.87 \times 10^5\times$

Key observation. The amplification factor *increases* with tile count M. On `ibm_kingston` ($M = 38$), the oracle achieves $5.5 \times 10^6\times$ amplification over the random thermal floor. This ratio is reported for continuity with the tiling literature, but it should not be read as evidence of non-uniformity on its own: at 50,000 shots over 2^{33}–2^{38} states it is dominated by the shot budget rather than by the prepared state, as section C.1.2 shows. The non-uniformity claim is established instead in section C.1.2, by a per-tile statistic that is independent of M and that rejects the uniform null at 158σ.

C.1.2 Per-Tile Marginal Analysis

Why a second statistic is needed The amplification factor of table C.1 is the ratio of the observed top-state probability to the uniform baseline 2^{-M}. That statistic has a structural limitation which should be stated plainly: at $M = 33$–38 and 50,000 shots, the expected number of coincident outcomes under a *uniform* distribution is $n^2/2^{M+1}$, namely 0.146, 0.073 and 0.005 for $M = 33, 34, 38$. The maximum count is therefore 1 or 2 with overwhelming probability whatever the underlying state, and the observed distinct-outcome counts (49,991, 49,997, 49,999 of 50,000) are what uniform sampling produces. Evaluating the same ratio on simulated uniform draws returns 1.88×10^5, 3.61×10^5 and 5.50×10^6 at $M = 33, 34, 38$ — the same order as the measured values.

The amplification factor is thus dominated by the shot budget and M rather than by the prepared state, and it cannot by itself establish non-uniformity. We therefore report a second statistic which is free of that limitation and which the existing telemetry already supports at no additional QPU cost.

The observable The tiled circuit measures exactly one qubit per tile — the host member of each disjoint pair — into its own classical bit, so the M-bit register is a vector of M *independent single-qubit outcomes*, not a joint sample from a 2^M-dimensional space. The natural estimator is therefore the per-tile marginal

$$\hat{p}_j \;=\; \frac{1}{n}\sum_{s=1}^{n} b_j^{(s)}, \qquad \text{s.e.}(\hat{p}_j) = \sqrt{\hat{p}_j(1-\hat{p}_j)/n}, \tag{C.1}$$

with $n = 50{,}000$, giving a standard error of about 0.2% per tile *independently of* M. Where the amplification factor has no power at this shot budget, eq. (C.1) has a great deal.

For the Phase 3 circuit the prediction is exact. With $U_2 = e^{-iH_2 t}$, $H_2 = X \otimes X$ and $t = 1$ acting on $|00\rangle$,

$$e^{-itX\otimes X}\,|00\rangle = \cos t\;|00\rangle \;-\; i\sin t\;|11\rangle\,, \tag{C.2}$$

so measuring the host qubit gives

$$\boxed{P(b_j = 1) = \sin^2 t = 0.7081} \tag{C.3}$$

against a uniform null of 1/2. Equation (C.3) is a parameter-free prediction: no quantity in it is fitted.

Result Table C.2 reports the per-tile marginals against eq. (C.3).

Backend	M	**mean** $\hat{p}$	**sd**	**tiles with** $\lvert z\rvert > 5$	**visibility** v
`ibm_fez`	34	0.6952	0.0142	34/34	0.938 ± 0.012
`ibm_kingston`	38	0.6978	0.0145	38/38	0.950 ± 0.012
`ibm_marrakesh`	33	0.6953	0.0073	33/33	0.939 ± 0.006
Pooled	105	**0.6961**	0.0127	**105/105**	$\mathbf{0.943 \pm 0.006}$

Table C.2: Per-tile marginals for the Phase 3 oracle. z is measured against the uniform null $p = 1/2$. Visibility is defined by $\hat{p} = \frac{1}{2} + v(\sin^2 t - \frac{1}{2})$.

Three statements follow, in increasing order of strength.

The distribution is not uniform. Every one of the 105 tile marginals exceeds the uniform null by more than 5σ individually, and the pooled mean $\hat{p} = 0.6961 \pm 0.0012$ rejects $p = 1/2$ at 158σ. This is the claim the amplification factor was intended to support, now established by a statistic that a uniform distribution cannot reproduce.

The deviation is in the predicted direction and of the predicted size. The measured value is 98.3% of $\sin^2 t = 0.7081$. Writing the deficit as a contrast loss, $\hat{p} = \frac{1}{2} + v(\sin^2 t - \frac{1}{2})$, gives a pooled visibility $v = 0.943 \pm 0.006$ — that is, 94.3% of the predicted contrast survives, for a 5.7% deficit. With Heron median two-qubit error 6–8×10^{-3} (table 5) and single-qubit readout error of order 1–2%, a few-percent contrast loss is the expected magnitude. No parameter was fitted to obtain it.

The residual is structured, not noise. The three backends order as $v_{\text{kingston}} > v_{\text{marrakesh}} \approx v_{\text{fez}}$, and `ibm_kingston` has both the lowest median CNOT error and the longest T_2 of the three. The per-tile spread is also informative: `ibm_marrakesh` shows half the dispersion of the other two (sd = 0.0073 against 0.0142 and 0.0145), indicating that tile quality is more uniform across that device.

Phase 6 encodings under the same statistic The encoded instances of section C.1.5 are a different regime, and the marginal statistic separates them cleanly. Equation (C.3) does not apply there — each instance carries its own H_{custom} — but the marginals remain well defined.

Instance	M	**distinct outcomes**	**top prob.**	**mean $\hat{p}$**
Factoring-15	29	617	68.45%	0.0129
MaxCut-Triangle	29	607	68.07%	0.0131
Portfolio-2Asset	29	629	68.34%	0.0129
TSP-2City	29	629	65.55%	0.0142
Frustrated Ising	29	6,880	20.43%	0.0558
H_2 Molecule	29	48,089	0.03%	0.2351

Table C.3: Per-tile marginals for the six encoded instances on `ibm_fez`. Unlike Phase 3, four of the six concentrate on a single outcome at high probability.

Four of the six instances place 65–68% of all shots on one bitstring and explore only ~ 620 of 2^{29} states. That is a concentration the amplification statistic of section C.1.2 *would* have detected, and which is unambiguous here: a top-state probability of 0.68 at 50,000 shots is not a shot-budget artefact by any measure. The frustrated Ising instance sits between the two regimes, and the H_2 instance behaves like Phase 3, with 48,089 distinct outcomes and marginals near 0.235.

We note without interpretation that the four concentrated instances share a common marginal value to three decimal places (0.0129–0.0142), which warrants checking that their encoded Hamiltonians are in fact distinct at the transpiled-pulse level. Section 5.5 records this as an open item; the raw telemetry needed to settle it is included in the supplementary archive.

Summary The per-tile marginal is the appropriate statistic for a tiled architecture that measures one qubit per tile. It is independent of M, costs no additional QPU time, and on the existing Phase 3 telemetry it rejects uniformity at 158σ while landing within 2% of a parameter-free prediction. We recommend it in preference to the amplification factor wherever the number of shots is small compared with 2^M.

C.1.3 Phase 4: Cross-Domain Applicability

Three domain-specific Hamiltonians were injected into the identical oracle architecture. All executions completed successfully at 50,000 shots (Table C.4).

Table C.4: Cross-domain oracle executions on live IBM hardware.

Domain	**Backend**	**Job ID**	**Tiles**	**Status**
H_2 Chemistry	`ibm_marrakesh`	`da8ljd9q...kj50`	33	**DONE**
Cryptographic Factoring	`ibm_fez`	`da8ljmce...q860`	34	**DONE**
Portfolio Optimization	`ibm_kingston`	`da8lkt6r...im0`	38	**DONE**

The results demonstrate that the oracle architecture is *domain-agnostic*: the same $\mathcal{O}(1)$-depth tiled circuit resolves ground states across molecular simulation, number theory, and combinatorial optimization.

C.1.4 Phase 5: k-Local Generalization ($k > 2$)

The IQS formalization (**Part 260**) proves that the conservation law $[H_{\text{total}}, A_1] = 0$ holds for *arbitrary* N-qubit systems, not only $N = 2$. This establishes the construction for the structured k-local family considered below, for arbitrary locality k. We verified this computationally and experimentally.

For each locality $k \in \{3, 4, 5\}$, we construct the k-local Hamiltonian:

$$H_k = \underbrace{Z \otimes Z \otimes \cdots \otimes Z}_{k} + \frac{1}{2} \sum_{i=1}^{k-1} (X_i X_{i+1} + Y_i Y_{i+1}) \tag{C.4}$$

where the first term is a k-body plaquette and the remaining terms are the Part 76 exchange primitives.

Theorem C.1 (Conservation for the structured k-local family, **Part 260**). *For every $k \geq 2$, let H_k be a Hamiltonian of the form given in Eq. (C.4). Then*

$$[H_k, A_1] = 0, \qquad A_1 = \sum_{x=1}^{k} Z_x. \tag{C.5}$$

Consequently, the Host projector P_H is well-defined for this structured family.

Table C.20 tests the conservation on hardware against a matched-depth arm whose bonds do not commute with A_1: the two separate by 0.20–0.43 at $t = 0.8$, in the predicted ratio.

Proof. Write $H_k = J_Z \Pi_k + C_k$ as in eq. (49). The plaquette $\Pi_k = Z^{\otimes k}$ is diagonal in the computational basis, as is $A_1 = \sum_x Z_x$, so $[\Pi_k, A_1] = 0$. For the chain it suffices to check one bond, since terms on disjoint sites commute with the rest of A_1 trivially. On sites $i, i+1$,

$$X_i X_{i+1} + Y_i Y_{i+1} = 2\big(\sigma_i^+ \sigma_{i+1}^- + \sigma_i^- \sigma_{i+1}^+\big), \tag{C.6}$$

which moves one excitation from i to $i+1$ or back and so leaves $Z_i + Z_{i+1}$ unchanged; explicitly $[\sigma_i^+ \sigma_{i+1}^-, Z_i + Z_{i+1}] = 0$ because $[\sigma^\pm, Z] = \mp 2\sigma^\pm$ and the two contributions cancel. Summing over bonds gives $[C_k, A_1] = 0$, hence $[H_k, A_1] = 0$ for every $k \geq 2$. Since A_1 is then a conserved quantity, its spectral projectors commute with H_k and P_H is well defined. Verified numerically at $k = 2, \ldots, 6$ to machine precision. □

We verified $\|[H_k, A_1]\| = 0$ numerically for $k = 3, 4, 5$ and deployed the oracles on the heavy-hex coupling graph using disjoint k-chains (Table C.5).

Table C.5: k-local oracle deployments on 156-qubit IBM hardware. The conservation law $[H_k, A_1] = 0$ is verified for all k.

k	Backend	Job ID	Chains	Qubits	Space	Status
3	`ibm_marrakesh`	`da8m87mr...n4tg`	51	153	2^{51}	**DONE**
3	`ibm_kingston`	`da8mbjhq...la4g`	49	147	2^{49}	**DONE**
3	`ibm_fez`	`da8mboce...qv7g`	49	147	2^{49}	**DONE**
4	`ibm_marrakesh`	`da8mbv4j...trdg`	32	128	2^{32}	**DONE**
4	`ibm_kingston`	`da8mc5se...qvlg`	32	128	2^{32}	**DONE**
5	`ibm_marrakesh`	`da8mccke...qvv0`	29	145	2^{29}	**DONE**

Key result. The oracle architecture extends to this specific structured family of k-local Hamiltonians without modification to the theoretical framework. The conservation law is a structural property of the $A_1 = \sum Z_x$ operator, which commutes with any product of Z operators by construction, rather than an unstructured arbitrary k-local mapping.

C.1.5 Phase 6: Admissible Problem Encoding

This phase deploys the Compression Encoding Theorem, theorem 3.12, together with the admissibility conditions of definition 3.9 and the host–ground-space correspondence of lemma 3.11; all three are stated and proved in section 3.7. What follows is their instantiation on six representative encoded problem instances spanning combinatorial optimization, number theory, chemistry, finance, logistics and frustrated magnetism, deployed on `ibm_fez` with 29 tiles each (table C.6).

Table C.6: Admissible encoding: 6 representative encoded problem instances deployed on `ibm_fez`.

Problem	Type	$E_{\min}$	Tiles	Depth	Job ID
MaxCut-Triangle	Combinatorial	-1.000	29	10	`da8mrh6r...novg`
Factoring-15	Number Theory	-2.000	29	9	`da8mrhmr...np0g`
H_2 Molecule	Chemistry	-1.500	29	11	`da8mri6r...np10`
Portfolio-2Asset	Finance	-1.200	29	9	`da8mrike...rf8g`
TSP-2City	Logistics	-1.118	29	15	`da8mrj1q...lqjg`
Frustrated Ising	Condensed Matter	-3.000	29	14	`da8mrjse...rfag`

Key result. The Compression Encoding Theorem (Theorem 3.12) maps problem Hamiltonians to the oracle algebraically. While the classical synthesis of the resulting unitary remains computationally dependent on the problem size, the quantum execution of the projection $C_H(A) = P_H A P_H$ is achieved in $\mathcal{O}(1)$ depth on the hardware.

What the six deployments establish, and what they do not The six executions of table C.6 were deployed by passing

$$U_x \;=\; \exp(-iH_x t), \qquad t = 1, \tag{C.7}$$

as the `custom_unitary` argument of the tiled circuit. Theorem 3.12 is a statement about a different operator, the projector $P_H(x) = V_x V_x^*$ of eq. (36). The distinction is material and we record it here rather than leaving it implicit.

Three instances execute a trivial circuit. MaxCut-Triangle ($Z \otimes Z$), Factoring-15 ($2Z \otimes Z - Z \otimes \mathbb{1} - \mathbb{1} \otimes Z$) and Portfolio-2Asset ($0.7\, Z \otimes Z + 0.3\, Z \otimes \mathbb{1} - 0.2\, \mathbb{1} \otimes Z$) are diagonal in the computational basis. For any such H_x,

$$\exp(-iH_x t)\,|0\cdots0\rangle \;=\; e^{-i(H_x)_{00}t}\;|0\cdots0\rangle\,, \tag{C.8}$$

so eq. (C.7) acts as the identity up to a global phase. The measured all-zeros outcome is therefore the untouched input state, and the 65–68% concentration of table C.3 is the survival probability of an identity operation under hardware noise — a per-tile infidelity of about 1.4% across 29 tiles, consistent with the readout error of table 5 — rather than evidence of projection.

The two instances carrying genuine off-diagonal weight, H_2 ($Z \otimes Z + 0.5\, X \otimes X$) and Frustrated Ising ($X \otimes X + Y \otimes Y + Z \otimes Z$), are precisely the two that do *not* concentrate, with 48,089 and 6,880 distinct outcomes.

The prepared state does not minimise the problem Hamiltonian. Evaluating $\langle H_x \rangle$ on the state eq. (C.7) actually prepares:

Instance	$E_{\min}$	$\langle H_x \rangle$ **of output**	
MaxCut-Triangle	−1.000	+1.000	maximum eigenvalue
Portfolio-2Asset	−1.200	+0.800	maximum eigenvalue
Frustrated Ising	−3.000	+1.000	maximum eigenvalue
Factoring-15	−2.000	0.000	intermediate
TSP-2City	−1.118	+1.000	intermediate
H_2 Molecule	−1.500	+1.000	intermediate

Table C.7: Energy of the state prepared by eq. (C.7), against the classical minimum. No instance attains $E_{\min}$; three return the highest-energy configuration.

The $E_{\min}$ column of table C.6 is the correct classical ground-state energy of each encoded Hamiltonian. It is a property of H_x, computed classically, and not a measurement of the corresponding run.

Consequently. Lemma 3.11 and theorem 3.12 are unaffected as mathematics: their proofs are algebraic and are given in section C.1.5. What must be withdrawn is the inference from tables C.3 and C.6 to those theorems. The six executions establish that the tiled architecture accepts an arbitrary two-qubit unitary and runs it at $\mathcal{O}(1)$ depth across 29 tiles on live hardware; they do not establish ground-state preparation, because the operator deployed was not the projector the theorems concern.

The corrected deployment. $P_H(x)$ is a projector and cannot be submitted directly. Two realizations are admissible. The first is the measurement-conditioned channel of theorem 3.3, whose depth is $\mathcal{O}(1)$ and whose expected repetition count scales as $1/p_{\text{succ}}(x)$. The second, which requires no post-selection, is a preparation unitary: any U_x whose first column lies in Ground(H_x) satisfies

$$\boxed{U_x \left|0 \cdots 0\right\rangle \in \text{Ground}(H_x)} \tag{C.9}$$

exactly and deterministically. Since U_x is again a two-qubit unitary, it transpiles to the same `CZ/ZSX` block: the circuit depth, the tiling and the shot budget of section 6 are unchanged, and the entire problem-dependence sits in the classical synthesis of V_x — which is what theorem 3.12 already bounds by $T_{\text{classical}}(x) \leq r(|x|)$.

We verified eq. (C.9) for all six instances: each attains $E_{\min}$ exactly, with ground-space degeneracies $g = 2, 2, 1, 2, 1, 1$ in the order of table C.7. The implementation is the construction of eq. (C.9) itself, which replaces the encoding call used in eq. (C.7) without changing anything else about the deployment. Re-running the six instances through it is the experiment that would test theorem 3.12, at the same hardware cost as the runs reported here.

Hardware verification of the corrected deployment The construction of eq. (C.9) was executed on `ibm_fez`, job `da98vhce74ec73ajglgg`, at 91.07 quantum-seconds: six circuits, $M = 36$ disjoint tiles, 50,000 shots each.

Design. Two changes distinguish this run from the Phase 6 deployments. First, the operator submitted is U_x of eq. (C.9), not $\exp(-iH_x t)$. Second, *both* qubits of every tile are measured into a $2M$-bit register, so the observable is the full per-tile distribution over $\{00, 01, 10, 11\}$ rather than a single marginal. The second change matters for a reason established the hard way: an earlier execution (job `da96dj6rbfbs73chaifg`) measured only the host qubit, and a qubit-ordering discrepancy between the `numpy` Kronecker convention and Qiskit's little-endian register inverted the prediction. Four instances returned $1 - p_{\text{pred}}$ and the two whose prediction was $1/2$ were unaffected, which is the signature that identified the fault. A four-outcome distribution admits no such silent inversion.

All predicted distributions were computed classically and printed before submission. They are pre-registration in the strict sense: the run can confirm or refute them, and nothing in the analysis was chosen after seeing the data.

Result. Table C.8 gives, for each of the six instances, the pre-registered distribution, the measured distribution, and the distance from each to the two competing hypotheses.

Instance	**predicted**	**observed**	$\mathrm{TV}_{\mathrm{thy}}$	$\mathrm{TV}_{\mathrm{null}}$	**tiles**
MaxCut-Triangle	[.014 .957 .014 .014]	[.021 .959 .010 .011]	0.026	0.979	36/36
Factoring-15	[.014 .957 .014 .014]	[.021 .959 .010 .010]	0.026	0.979	36/36
Portfolio-2Asset	[.014 .014 .957 .014]	[.024 .024 .942 .010]	0.026	0.976	36/36
TSP-2City	[.014 .907 .014 .064]	[.013 .905 .001 .081]	0.028	0.986	36/36
H_2 Molecule	[.014 .486 .486 .014]	[.022 .485 .474 .019]	0.035	0.959	36/36
Frustrated Ising	[.014 .486 .486 .014]	[.027 .478 .482 .013]	0.035	0.973	36/36
Pooled			**0.029**	**0.975**	**216/216**

Table C.8: Per-tile outcome distributions for the corrected deployment, ordered $[00, 01, 10, 11]$ with the host qubit least significant. $\mathrm{TV}_{\mathrm{thy}}$ and $\mathrm{TV}_{\mathrm{null}}$ are total-variation distances to the pre-registered prediction and to the distribution the original $\exp(-iH_x t)$ deployment would produce. The last column counts tiles closer to theory than to null.

Every one of the 216 tile-instance pairs lies closer to the prediction than to the null, with mean total variation 0.029 against 0.975 — a separation of more than thirty-fold. The three instances whose original deployment was the identity now concentrate on the correct ground state at weight 0.94–0.96, and H_2 and Frustrated Ising reproduce their degenerate $\frac{1}{2}, \frac{1}{2}$ structure to within 0.012.

The circuits are distinct. The transpiler flagged coincident operation profiles across several circuits, which section 5.5 records as a hazard. The measured outcomes settle it: Portfolio-2Asset peaks on $|10\rangle$ at weight 0.942 while MaxCut-Triangle and Factoring-15 peak on $|01\rangle$ at 0.959, and Frustrated Ising peaks opposite H_2. Equal gate counts with different rotation angles produce different physics; the circuits are not the same object.

Contrast. Inverting the depolarizing model instance by instance recovers a visibility of 0.931 on average, against 0.943 measured for single-qubit marginals in section C.1.2. Reading both qubits of a tile costs roughly one further percent of contrast, which is the expected direction and magnitude for an additional readout per tile.

Standing. Theorem 3.12 now has an experimental test that it passes. The claim supported is precise: for six encoded instances spanning combinatorial optimization, number theory, chemistry, finance, logistics and frustrated magnetism, the compression construction prepares the ground space of the encoded Hamiltonian on 36 disjoint tiles of a 156-qubit processor at $\mathcal{O}(1)$ circuit depth, matching pre-registered distributions to a mean total variation of 0.029. The classical synthesis of V_x remains bounded as in theorem 3.12; what the hardware establishes is the quantum execution step.

C.1.6 Hardware Tests of the Closures

Three of the results of section 4.7 admit a hardware test. They were run on `ibm_kingston` (job `da99d0urbfbs73chdce0`) and `ibm_fez` (job `da99ickjbipc73fek51g`), ten circuits each at 10,000 shots, with all predictions computed and printed before submission. The outcome is mixed and is reported as such: one leg confirms, one is inconclusive for an identifiable reason, and one exposed an error in the circuit rather than in the theorem.

Removing the k-body cost Each locality $k \in \{3, 4, 5\}$ was built twice — once through the plaquette ladder of eq. (50), once by handing the same $2^k \times 2^k$ unitary to generic synthesis — and both were run.

k	**structured 2q**	**generic 2q**	**ratio**	$\mathrm{TV}_{\mathrm{struct}}$	$\mathrm{TV}_{\mathrm{generic}}$
3	11	37	3×	0.231	0.120
4	18	206	11×	0.248	0.465
5	31	909	29×	0.249	**0.877**

Table C.9: Transpiled two-qubit gate counts (identical on both processors) and total-variation distance to the pre-registered distribution, `ibm_fez`.

The gate-count claim of theorem 4.31 is confirmed directly, and the consequence is not marginal. At $k = 5$ the generic circuit carries 909 two-qubit gates and returns TV = 0.877 — statistically indistinguishable from the uniform distribution. The structured circuit carries 31 and returns 0.249. The structured build is what makes the experiment possible at all above $k = 3$.

The structured residual is a Trotter error, and it is ours. The TV ≈ 0.23–0.25 floor in the structured column does not come from the hardware. The circuit applied $e^{-iC_k t}$ as a product of two-site blocks, which remark 4.32 notes is not an identity. Evaluating that split classically gives TV = 0.220, 0.210, 0.210 at $k = 3, 4, 5$ — accounting for essentially the whole observed residual. The measurement is therefore consistent with theorem 4.31 as corrected, and inconsistent with the withdrawn $\mathcal{O}(k)$ claim.

Sector conservation Three inputs were prepared at $A_1 = +2, 0, -2$ and evolved under H_S.

sector	**predicted**	**observed** (`ibm_fez`)	TV
$A_1 = +2$	$\|00\rangle$ at 1.000	$\|00\rangle$ at 0.924	0.076
$A_1 = 0$	$\|10\rangle$ 0.827, $\|01\rangle$ 0.173	$\|10\rangle$ 0.767, $\|01\rangle$ 0.153	**0.081**
$A_1 = -2$	$\|11\rangle$ at 1.000	$\|11\rangle$ at 0.918	0.082

Table C.10: Sector conservation on `ibm_fez`. No population leaks between A_1 sectors.

All three reproduce to within 0.082, with the non-trivial $A_1 = 0$ split recovered at 0.767/0.153 against a predicted 0.827/0.173 and no leakage into the $A_1 = \pm 2$ states. Theorem 4.36 is confirmed: the sector is fixed by the conserved value on the input, and the context is not chosen.

`ibm_kingston` returned TV = 0.500, 0.326, 0.238 on the same three circuits. These are two gates deep, so the degradation is not a depth effect; it points to a poor qubit pair on the path selected by the placement routine on that device. We report it rather than averaging it away, and note that a calibration-aware placement would be the appropriate control.

The even-dimensional environment: inconclusive The $\dim \mathcal{H}_E = 4$ tile of theorem 4.34 was built by generic synthesis on four qubits: depth 314, 98 two-qubit gates. It returned TV = 0.815 and 0.813 on the two processors.

Given table C.9 — where a 909-gate circuit returns 0.877 and is uniform — a 98-gate, depth-314 circuit returning 0.81 is a depth failure, not a measurement of the underlying claim. Theorem 4.34 is a statement about the existence of an antiunitary satisfying eq. (51), which was verified constructively and numerically in section 4.7.2; this run neither supports nor refutes it. A structured build for the four-qubit tile, on the pattern of section C.1.6, is the test that would settle it.

Full-device tiled verification Section C.1.6 ran one tile per circuit and produced an anomaly. On the three two-gate Leg C circuits, `ibm_kingston` returned total variations of 0.603, 0.378 and 0.272, where `ibm_fez` reproduced the same three at 0.075, 0.075 and 0.080. Depth cannot account for that. A single tile, however, cannot separate a bad qubit pair from a failed prediction — it yields one sample.

The experiment was therefore repeated with every circuit tiled across the full device at 50,000 shots, reporting the median over tiles together with the spread. These are rows 28 and 29 of table C.24, on `ibm_kingston` at 135.4 quantum-seconds and on `ibm_fez`.

leg	circuit	tiles	kingston	fez	good K	good F
A	`k3_structured`	50	0.113	0.133	36/50	27/50
A	`k3_generic`	50	0.177	0.224	15/50	2/50
A	`k4_structured`	28	0.167	0.230	9/28	3/28
A	`k4_generic`	28	0.538	0.662	0/28	0/28
A	`k5_structured`	25	0.285	0.411	1/25	0/25
A	`k5_generic`	25	0.881	0.884	0/25	0/25
B	`envdim4_structured`	28	**0.094**	**0.126**	25/28	18/28
C	`sector_A1p2`	60	**0.039**	**0.058**	56/60	54/60
C	`sector_A1z`	60	**0.038**	**0.054**	57/60	58/60
C	`sector_A1m2`	60	**0.051**	**0.069**	55/60	52/60

Table C.11: Median total variation over tiles, 50,000 shots. "Good" counts tiles with TV < 0.15. Between 112 and 150 qubits are engaged per circuit.

The anomaly was a placement artefact. Leg C on `ibm_kingston` moves from $0.603, 0.378, 0.272$ single-tile to $0.039, 0.038, 0.051$ tiled, with 55–57 of 60 tiles below 0.15 and maxima of $0.61, 0.36, 0.51$. The single-tile run had landed on one of the few bad pairs. Nothing about the device or the prediction was at fault, and the tiled statistic is robust to it by construction.

Sector conservation. Theorem 4.36 is confirmed on both processors in all three sectors, median TV between 0.038 and 0.069, with the non-trivial $A_1 = 0$ split reproduced and no population leaking to $A_1 = \pm 2$. This is the strongest of the three results.

Even-dimensional environment. Theorem 4.34 is confirmed: $\dim \mathcal{H}_E = 4$ gives median TV of 0.094 and 0.126 with 25/28 and 18/28 tiles good, from a nine-gate structured build. The generic build of the same unitary needed 89–98 gates and returned 0.80 in section C.1.6.

Removing the k-body cost. The comparative claim of theorem 4.31 holds at every locality on both devices, and the margin widens with k. At $k = 5$ the structured circuit

uses 1,350 two-qubit gates against 23,628 for generic synthesis — a factor of 17.5 — and the generic build has *zero* tiles below 0.15 on either processor, sitting at 0.88, which is the uniform distribution.

Remark C.2 (What is not established)**.** The structured $k = 5$ circuit is itself degraded: median 0.285 and 0.411, with 1 and 0 good tiles. Theorem 4.31 is a statement about the k-body *cost*, and as such it is confirmed — structured dominates generic everywhere and is the only build that survives above $k = 3$. It is not a claim that $k = 5$ is executable at useful fidelity on this hardware, and table C.11 does not support one. The residual is a two-local Trotter and depth problem, which remark 4.32 places outside the theorem.

Remark C.3 (A step count chosen by an untested cost model)**.** The run tabulated in table C.9 applies the chain as a single product of two-site blocks, at 11, 18 and 31 two-qubit gates for $k = 3, 4, 5$. A second run (jobs `da99orke74ec73ajhdu0` and `da99p91qtnsc73d1bqq0`) used a second-order step count of 2, chosen by a cost model balancing Trotter error against an assumed linear hardware term, which raised the counts to 23, 42 and 64.

The outcome was mixed rather than uniformly worse, and the pattern is the informative part:

k	one step	two steps	
3	0.231	**0.100**	finer step helps
4	0.248	**0.144**	finer step helps
5	**0.249**	0.326	finer step hurts

Below $k = 5$ the Trotter term dominates and refining it pays; at $k = 5$ the doubled gate count costs more than the refinement buys. The cost model predicted a total of 0.089 at $k = 5$ and the measurement gave 0.326, so it under-predicted the hardware term by roughly a factor of four. Section C.1.6 therefore reverts to a single step, and $k = 5$ improves to 0.285 despite the coarser approximation. The episode is recorded because the model looked principled and was not tested before being acted on.

The mined theorems on hardware. Theorems 4.41, 4.43 and 4.47 were tested on `ibm_kingston` (job `da9almhqtnsc73d1cmm0`) and `ibm_fez` (job `da9bsvsjbipc73femhq0`): fifteen circuits each at 20,000 shots, tiled across the device, 5,296 two-qubit gates per backend.

Sector totality and parity. One representative was prepared per A_1 sector at $N = 3$ and $N = 4$ and the retention measured — the fraction of shots satisfying $\sum_u Z_u = \mathbf{z}$, median over tiles. Table C.12 gives every value.

Table C.12: Median sector retention, F16. Every sector admitted by theorem 4.41 was populated and retained on both processors. The $N = 3$ register carries no $\mathbf{z} = 0$ row because no three-qubit state can satisfy $\sum_u Z_u = 0$; that absence is theorem 4.43 observed directly, not a gap in the measurement.

N	**Sector z**	**Tiles**	`ibm_kingston`	`ibm_fez`
3	+3	50	0.9092	0.8559
3	+1	50	0.9135	0.8704
3	−1	50	0.9035	0.8652
3	−3	50	0.8846	0.8310
4	+4	28	0.8428	0.7499
4	+2	28	0.8652	0.7795
4	0	28	0.8626	0.7703
4	−2	28	0.8543	0.7664
4	−4	28	0.8401	0.7341

All $N + 1$ sectors are present at each N, as theorem 4.41 requires, and the parity structure of theorem 4.43 is visible in the row labels themselves: the admissible eigenvalues step by two and share the parity of N.

The mediator-count bound. The six half-filling states of $N = 4$ were prepared and every qubit measured, giving the 6×4 response matrix of theorem 4.47. Table C.13 lists its singular values.

	σ_1	σ_2	σ_3	σ_4
predicted	1.4635	1.0779	0.7948	**0.0000**
`ibm_kingston`	1.4102	1.0902	0.9689	**0.0293**
`ibm_fez`	1.3224	1.0208	0.9334	**0.0176**

Table C.13: Singular values of the mediator response matrix. The vanishing fourth value is the content of theorem 4.47. The predicted row is computed with the exchange chain taken as $\sum_i (X_i X_{i+1} + Y_i Y_{i+1})$, without the factor $\frac{1}{2}$ that eq. (49) carries; under eq. (49) as written the same construction gives $(2.2542, 1.0797, 0.2496, 0.0000)$ at the same t. The choice rescales the first three values and changes neither the rank nor σ_4, which is why remark C.4 applies to both, but the row as printed is the un-halved convention and is reproducible only under it.

The decisive ratio is σ_4/σ_1: predicted ≈ 0, measured 0.0208 and 0.0133. Independent mediators would give ≈ 0.5, so the observation sits a factor of 24 to 38 below the null. The rank of the response matrix is $3 = N - 1$, not $5 = d - 1$, on both processors. Row sums, which lemma 4.46 requires to vanish, measured 0.009–0.079 in absolute value.

Theorem 4.47 is therefore confirmed, and with it corollary 4.49: the back-edge dimension is set by the number of physical mediators, not by the dimension of the protected sector.

Remark C.4 (Why this test is robust)**.** The rank deficiency follows from $\sum_u P_{\mathbf{z}} Z_u P_{\mathbf{z}} = \mathbf{z} P_{\mathbf{z}}$, which holds for *any* unitary commuting with A_1. Every XY-bond Trotter step does, so σ_4 remains at machine zero for step counts 1, 2 and 3 — verified classically before submission. Trotter error and gate noise therefore degrade the entries of the matrix

without restoring the missing rank, which is why the measured σ_4/σ_1 is small rather than merely smaller than the null. The quantity being measured is protected by the same symmetry the theorem is about.

An error in the analysis, and how it was caught The first pass of this analysis mapped measured bitstrings to state indices as `int(bs[::-1][:k][::-1], 2)` rather than `int(bs[::-1], 2)`. Re-analysing the same counts with the correct ordering moved table C.10's $A_1 = 0$ row from TV = 0.674 to 0.081 and the $k = 5$ structured circuit from 0.964 to 0.249. No additional quantum time was required to find it.

This is the second ordering fault in this arc; the first is recorded at section C.1.5. Both were caught because a pre-registered prediction was available to disagree with. Neither would have been visible from the data alone.

C.1.7 Oracle reference tests

The results above were obtained by scripts written alongside the theorems they test. To close the gap between what the paper states and what was executed, the constructions were re-implemented directly from the numbered equations — one function per object, named for it — and deployed as six oracles on all three processors at 50,000 shots with full-device tiling. Every prediction was computed classically and fixed before submission. Table C.14 lists the six and their outcomes.

Table C.14: The six oracles, what each deploys, and its outcome. Tile counts are full-device: 78×2, 52×3, 39×4, 26×6 and 13×12 all occupy 156 qubits.

	Object deployed	**Circ.**	**Outcome**
O1	theorem 3.12, eq. (C.9) — U_x on the six encoded instances	6	confirmed; TV = 0.030–0.055 on all three
O2	theorem 4.31 — structured gadget against generic synthesis	6	confirmed; structured beats generic 9/9
O3	eq. (C.11) by lemma C.9, against the ladder	8	confirmed; depth one beats the ladder 12/12
O4	theorem 4.47 — the half-filling response matrix	6	confirmed; $\sigma_4/\sigma_1 = 0.057$–$0.073$
O5	lemma 4.46 — the sector sum rule	5	confirmed at $d = 2$; inconclusive at $d = 3$, see remark C.6
O6	theorem 4.54, corollary 4.55 — the conserved direction	15	confirmed; $V = 0.0006$–0.0012 against a null of 1

O3 measures the depth erratum. Remark C.11 argued that the $\mathcal{O}(N)$ of section C.1.7 belonged to the ladder rather than to the construction. O3 puts the two routes on the same instances, the same tiles and the same shot budget, and the argument becomes a measurement. Table C.15 gives every comparison.

Table C.15: Fraction of shots returning an exact maximum cut, median over tiles. "depth 1" is lemma C.9 — one layer of X gates, no two-qubit gates; "ladder" is the $(N-1)$-CNOT route of section C.1.7. The null is $g/2^N$.

N	**null**	`ibm_kingston`		`ibm_fez`		`ibm_marrakesh`	
		depth 1	ladder	depth 1	ladder	depth 1	ladder
6	3.13×10^{-2}	0.9288	0.8838	0.8953	0.8354	0.8991	0.8400
8	8.59×10^{-2}	0.9345	0.8896	0.9007	0.8657	0.9204	0.8663
10	1.95×10^{-3}	0.8485	0.7502	0.8396	0.7388	0.8014	0.6574
12	4.88×10^{-4}	0.8307	0.6931	0.8402	0.6732	0.8022	0.6160

Depth one wins in all twelve comparisons, and the margin grows with N on every processor — on `ibm_marrakesh` from +0.059 at $N = 6$ to +0.186 at $N = 12$, on `ibm_fez` from +0.060 to +0.167, and on `ibm_kingston` from +0.045 to +0.138 — which is the signature of the ladder's accumulating two-qubit error and not of anything else the two circuits differ in. Lemma C.8 says a Z-basis measurement cannot separate their outputs; on hardware it separates them by exactly the ladder's cost.

O2 required a rebuild, and the first build is retained. As first written, O2 emitted $e^{-iC_k t}$ as a generic k-qubit unitary in *both* branches. The two therefore agreed with each other — which is theorem 4.31 being an identity — but carried the same $\mathcal{O}(4^k)$ cost, so no advantage could appear. The rebuild removes every k-qubit unitary from the structured branch: the exact plaquette gadget of eq. (50) plus the chain as two-local bond blocks under second-order Trotter, which is what table 2 means by one block per bond. Table C.16 reports the rebuilt run against the branch-wise predictions.

Table C.16: Total variation to each branch's own exact prediction, median over tiles. The structured branch carries Trotter error and the generic branch does not; the floor column is that error as a total variation between the Trotterized and exact predictions, computed classically. The cost column is the two-qubit synthesis cost of table 2 — $5(k-1)$ for the structured route against 4^k for generic synthesis — which is what theorem 4.31 predicts and what the archives record. The transpiled column is that circuit rewritten into the native basis and routed onto the coupling map, per tile, at the optimization level used for submission and with the routing seed pinned; see remark C.5.

k	**branch**	`kingston`	`fez`	`marrakesh`	**Trotter floor**	**cost**	**transpiled**
3	structured	**0.1106**	**0.1336**	**0.1621**	0.0095	10	15
3	generic	0.1687	0.2478	0.2698	—	64	37
4	structured	**0.1433**	**0.2371**	**0.2063**	0.0101	15	30
4	generic	0.5765	0.7425	0.6629	—	256	206
5	structured	**0.2378**	**0.3259**	**0.3467**	0.0101	20	53
5	generic	0.8943	0.8974	0.8947	—	1024	957

The structured branch is closer to its prediction in all nine comparisons. At $k = 4$ it places 17 of 28 tiles below TV = 0.15 on `ibm_kingston` where the generic branch places none; at $k = 5$ the generic branch sits at the noise floor on every processor while the structured one still carries signal. The trade is the one remark 4.32 describes: at $k = 5$

the measured 0.2378 sits against a Trotter floor of 0.0101, so approximately 96% of the residual is hardware error and 4% is the approximation.

Remark C.5 (The transpiled advantage is real but smaller than the synthesis count)**.** The two cost columns of table C.16 measure different things and neither substitutes for the other. The synthesis cost is what theorem 4.31 predicts; the transpiled count is what a processor is actually asked to execute. Transpilation moves both branches toward each other: it inflates the structured route by routing and basis translation ($10 \to 15$, $15 \to 30$, $20 \to 53$) and deflates the generic one by optimising the synthesised block ($64 \to 37$, $256 \to 206$, $1024 \to 957$). The ratio between the branches is therefore $2.5\times$, $6.9\times$ and $18.1\times$ at $k = 3, 4, 5$, against the $6.4\times$, $17.1\times$ and $51.2\times$ of table 2. The advantage survives and still grows with k, but a reader should take the transpiled ratio, not the synthesis ratio, as the one a device sees.

All three processors return identical counts, being Heron r2 parts with the same coupling topology and basis. The figures are reproducible only with the routing seed fixed — at this optimization level the pass is stochastic, and an unpinned rerun moves $k = 5$ generic between 23,871 and 23,925 gates over the whole tiled circuit — and they track the target in force when they are generated. `oracle_gatecounts.py` regenerates them and records the seed; the values are archived in `GATE_COUNTS.json`.

O5 and O6: the sum rule, and the test it cannot perform. O5 prepares each host basis state, evolves, and reads the response matrix of $\langle Z_u \rangle$. The sector sum rule of lemma 4.46 is confirmed directly at $d = 2$: row sums measured $-0.008, 0.019$ on `ibm_marrakesh`, $0.009, 0.005$ on `ibm_fez` and $-0.004, 0.019$ on `ibm_kingston`, against a predicted 0. At $d = 3$ the same sums read 0.63–0.80 against a predicted 1, a 20–37% deficit: the three-qubit circuit does not hold its sector well enough for the measurement to bear weight.

Remark C.6 (Why O5 does not test the channel count)**.** It is tempting to read the rank of the O5 response matrix as a measurement of N_σ. It is not. Lemma 4.46 forces every row to sum to $\mathbf{z}$, so the row space lies in a $(k-1)$-dimensional subspace as a matter of algebra, and rank $\leq d-1$ follows for *any* matrix whatever. Centring the rows drives the last singular value to exactly zero by construction. O5 therefore measures the sum rule — which table C.13 already measures — and says nothing independent about theorem 4.54. O6 was built to remove this confound.

O6 tests the conserved direction directly. Corollary 4.55 says exactly one direction of $\mathfrak{A}_H$ survives the dynamical selection, and names it σ_x. O6 prepares a host state with weight on both x and z, evolves it, and performs full host tomography at five times, reading $\langle \sigma \rangle = (P(01) - P(10))/P(\text{host})$ so that sector leakage divides out instead of biasing the estimate. The statistic is

$$V \;=\; \frac{\mathrm{Var}_t \langle \sigma_x \rangle}{\frac{1}{2}\Big(\mathrm{Var}_t \langle \sigma_y \rangle + \mathrm{Var}_t \langle \sigma_z \rangle\Big)}, \tag{C.10}$$

which theorem 4.54 predicts to vanish and which would be of order unity if no direction were singled out.

Table C.17: Host tomography. $\langle\sigma_x\rangle$ is predicted constant at -0.7071; $\langle\sigma_y\rangle$ and $\langle\sigma_z\rangle$ rotate. The `ibm_kingston` columns are two independent executions in separate sessions. The $\langle\sigma_y\rangle, \langle\sigma_z\rangle$ columns are the `ibm_kingston` run-1 values; the other two runs agree with them to within 0.03 at every time.

	`ibm_kingston`		`ibm_marrakesh`		
t	run 1	run 2		$\langle\sigma_y\rangle$	$\langle\sigma_z\rangle$
0.0	−0.6772	−0.6801	−0.6773	0.0177	0.6915
0.4	−0.6540	−0.6623	−0.6409	0.4861	0.4837
0.8	−0.6712	−0.6604	−0.6771	0.6768	−0.0188
1.2	−0.6782	−0.6744	−0.6758	0.4472	−0.5174
1.6	−0.6834	−0.6868	−0.6780	−0.0413	−0.6727
Var_t	0.00010	0.00010	0.00021	0.0786	0.2865
V	**0.0006**	**0.0006**	**0.0012**	null ≈ 1	

The y and z columns are the control and they track the prediction closely — 0.6768 against 0.7068 at $t = 0.8$, -0.0413 against -0.0413 at $t = 1.6$ — so the dynamics is being followed accurately and the flatness of $\langle\sigma_x\rangle$ is a conservation rather than a dead channel. Host retention was 0.94–0.96 throughout, so the renormalisation does very little work. The uniform 5% amplitude contraction, -0.677 against a predicted -0.7071 and 0.6915 against 0.7071 at $t = 0$, is a global depolarising visibility factor acting on all three axes alike, not a directional effect.

Remark C.7 (Session-to-session reproducibility)**.** Section C.1.6 records that the same two-gate circuit gave median total variation 0.603 on one `ibm_kingston` tile and 0.038 over 60 tiles of the same processor, which is the scale of drift this architecture is subject to. The two `ibm_kingston` columns of table C.17 are separate sessions with independent calibration and independent tile placement, and they agree on V to four decimal places. Whatever else drifts, this statistic does not.

A third ordering fault, and how it was caught. The first O1 execution applied the constructed unitaries with the qubit list unreversed, Qiskit taking the first qubit as least significant where the algebra of section 3 takes it as most significant. Four of the six instances returned TV ≈ 0.99 and the two whose predictions are symmetric under bit reversal were unaffected at TV ≈ 0.04. That pattern identifies the fault uniquely. Rescoring against the bit-reversed prediction recovered all four with no additional quantum time (0.9949 → 0.0396, 0.9956 → 0.0370, 0.9715 → 0.0722, 0.9275 → 0.0752), and the subsequent executions of the corrected circuit returned 0.030–0.055 directly. O2 through O6 were unaffected: O2 because H_k is symmetric under site reversal, O3 and O6 because they are built from per-site gates only, and O4 and O5 because singular values are invariant under the permutation the fault induces.

This is the third fault of its kind in this work and the third caught the same way. None was visible in the data alone; each was visible only against a prediction fixed beforehand.

The oracle projection at $N > 2$ Theorem 3.12 was tested at the two-qubit tile in section C.1.5. The classical check of section 6 carries the same construction to $N = 6, 8, 10, 12$; this run puts those instances on hardware, on `ibm_kingston` (job

`da9fo6ke74ec73ajncb0`) and `ibm_marrakesh` (job `da9fqt6rbfbs73chjrb0`), six circuits each at 20,000 shots, 497 two-qubit gates per backend.

What is prepared, and why it is exact. For the $N = 6, 10, 12$ instances the ground space is two-dimensional and closed under complement, $\text{Ground}(H_x) = \text{span}\{|b\rangle, |\bar{b}\rangle\}$, so the projector's output on the uniform superposition is

$$P_H(x)\ |+\rangle^{\otimes N} \ \propto\ \tfrac{1}{\sqrt{2}}\big(|b\rangle + |\bar{b}\rangle\big), \tag{C.11}$$

which is a GHZ-type state. It was realized here with one Hadamard, a sequential $(N-1)$-CNOT ladder, and an X on every site where b is 1. The prepared state *is* the projector's output — no approximation, no post-selection. The ladder is an implementation choice, and a poor one; see remark C.11. The $N = 8$ instance is 22-fold degenerate; preparing its full uniform superposition would need a general eight-qubit state preparation at $\mathcal{O}(2^N)$ entangling gates, so three of its eleven complementary pairs were prepared instead. Each lies inside $\text{Ground}(H_x)$, so those three rows test that what is prepared is optimal, not that the whole manifold is reached, and are marked accordingly. Table C.18 reports both processors.

Table C.18: Oracle projection on hardware. "p_{opt}" is the median over tiles of the fraction of shots whose bitstring is an exact maximum cut; the null is the same quantity for a uniformly random assignment, $g/2^N$. "cut" is the median cut value achieved, against the optimum and against the random-assignment expectation $|E|/2$. Rows marked † prepare a two-dimensional subspace of a larger ground space.

circuit	N	**tiles**	p_{opt} **king.**	p_{opt} **marr.**	**null**	**cut king.**	**max/rand**
N6	6	17	0.8753	0.8287	3.13×10^{-2}	7.64	8 / 4.5
N8†	8	12	0.9000	0.8649	8.59×10^{-2}	9.73	10 / 7.0
N8†	8	12	0.8796	0.8110	8.59×10^{-2}	9.73	10 / 7.0
N8†	8	12	0.8502	0.7873	8.59×10^{-2}	9.64	10 / 7.0
N10	10	8	0.7642	0.4893	1.95×10^{-3}	14.22	15 / 9.5
N12	12	8	0.7160	0.4258	4.88×10^{-4}	21.59	23 / 14.5

Result. At $N = 12$, 71.6% of shots on `ibm_kingston` land on an exact maximum cut of a 29-edge frustrated graph, against a uniform null of 4.88×10^{-4} — a factor of 1,470. The median cut is 21.59 of a possible 23 where random assignment gives 14.5, closing 83.5% of the optimality gap. The corresponding factors are 391 at $N = 10$ and 28 at $N = 6$; the ratio grows with N because the null falls as 2^{-N} while the prepared state does not.

The degradation with N is orderly — 0.875, 0.900–0.850, 0.764, 0.716 — and tracks the length of the CNOT ladder rather than anything structural. That is a statement about the ladder, and remark C.11 shows the ladder was avoidable. Total variation inside each prepared pair stays between 0.012 and 0.038, so the two complementary optima are returned in balance as eq. (C.11) requires.

The two processors differ, and the difference is not small. `ibm_marrakesh` matches `ibm_kingston` to within 0.05 at $N \leq 8$ but falls to 0.489 and 0.426 at $N = 10$ and 12, closing 74.0% and 61.2% of the gap against kingston's 85.9% and 83.5%. Both

remain orders of magnitude above the null — $\times 251$ and $\times 872$ — so the qualitative result is reproduced on both, but the quantitative margin is a property of the device and not of the construction. This is the same lesson as section C.1.6, at a different scale: a single-processor number would have misrepresented the spread.

An erratum on depth. An earlier version of this subsection inferred from the ladder that "the circuit depth here is $\mathcal{O}(N)$, not the $\mathcal{O}(1)$ of the tiled two-qubit oracle". That inference was wrong. It contradicted theorem 3.3, which gives the coherent depth of the projection as $\mathcal{O}(1)$ with the problem's difficulty carried by $1/p_{\text{succ}}$, and it contradicted section 5.2.5, which establishes that the Lieb-Robinson bound of [2] constrains finite-depth *unitary* circuits and that eq. (70) is not one. The $\mathcal{O}(N)$ was a property of the ladder that was written, not of the oracle projection. Three lemmas locate it exactly.

Lemma C.8 (Z-basis sufficiency)**.** *Let H_x be diagonal in the computational basis, and let O be any observable that is a function of the measured bitstring — the cut value is one. Then the statistics of O depend on the prepared state ρ only through* $\operatorname{diag}\rho$*. In particular the coherent pair $\frac{1}{\sqrt{2}}\big(|b\rangle + |\bar{b}\rangle\big)$ and the classical mixture $\frac{1}{2}\big(|b\rangle\langle b| + |\bar{b}\rangle\langle\bar{b}|\big)$ are indistinguishable by that measurement, exactly and not merely to within noise.*

Proof. A computational-basis measurement returns outcome a with probability $\langle a|\,\rho\,|a\rangle$, a functional of the diagonal alone. The two states have the same diagonal — weight $\frac{1}{2}$ on each of b and $\bar{b}$, zero elsewhere — and differ only in the off-diagonal entry $\rho_{b\bar{b}}$, which equals $\frac{1}{2}$ for the pure state and 0 for the mixture. No diagonal functional reads it. □

Lemma C.9 (Depth-one realization)**.** *The mixture of lemma C.8 is prepared by a single layer of X gates on a uniformly chosen element of $\{b, \bar{b}\}$: circuit depth 1, no two-qubit gates, independent of N.*

Proof. Sample $r \in \{0,1\}$ uniformly by a classical coin and let $c = b$ if $r = 0$ and $c = \bar{b}$ otherwise. Apply X to exactly those qubits u with $c_u = 1$, starting from $|0\rangle^{\otimes N}$. The gates act on disjoint qubits, so they form a single layer: depth 1, no two-qubit gate, and the layer is the same size whatever N is. The resulting ensemble is $\frac{1}{2}|b\rangle\langle b| + \frac{1}{2}|\bar{b}\rangle\langle\bar{b}|$, which is the mixture of lemma C.8. □

Together these say that the observable of table C.18 is reproducible at depth one. The $(N-1)$-CNOT ladder was therefore not merely suboptimal for this measurement; it was unnecessary.

Lemma C.10 (The ladder is not forced even where coherence is wanted)**.** *Where the coherent superposition is itself the object, a sequential ladder is still not required. $\frac{1}{\sqrt{2}}\big(|b\rangle + |\bar{b}\rangle\big)$ is produced* exactly *by a doubling tree of $\lceil \log_2 N \rceil$ CNOT layers. Under nearest-neighbour unitary gates alone $\Omega(N)$ is a genuine lower bound, by the light-cone argument of [2]; under measurement and feed-forward the depth is $\mathcal{O}(1)$, the route [3, 4] identify and which theorem 3.3 and remark 5.4 already place this construction in.*

Proof. Apply H to qubit 0, giving $\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$ on that qubit and $|0\rangle$ elsewhere. Then apply $\lceil \log_2 N \rceil$ layers, the s-th consisting of the CNOTs $(j, j+s)$ for $0 \le j < s$ with $j + s < N$, each layer acting on disjoint pairs. After layer s the value of qubit 0 has been copied to qubits $0, \ldots, \min(2s, N) - 1$, so the count doubles per layer and all N qubits carry it after $\lceil \log_2 N \rceil$ layers, leaving $\frac{1}{\sqrt{2}}(|0\rangle^{\otimes N} + |1\rangle^{\otimes N})$. A final layer of X gates

on the sites where b is 1 maps this to $\frac{1}{\sqrt{2}}(|b\rangle + |\bar{b}\rangle)$ exactly. The lower bound is the light-cone argument: a nearest-neighbour unitary circuit of depth D correlates sites at distance at most $\mathcal{O}(D)$, so producing correlation across N sites needs $D = \Omega(N)$ [2]; the tree evades it by using non-adjacent CNOTs, and measurement with feed-forward evades it altogether. □

All three were checked by direct computation: the Z-basis total variation between the pair and the mixture is 1.1×10^{-16} at every size while the X-basis total variation is 0.5, confirming the two states are genuinely distinct and separated only by a measurement this experiment does not perform; and the doubling tree reproduces the ladder's state to 4.4×10^{-16} in 3, 3, 4, 4 layers against the ladder's 5, 7, 9, 11.

Remark C.11 (What the run does and does not establish about depth)**.** Table C.18 is evidence for theorem 3.12 above $N = 2$ and for nothing about depth. The depth figures in it are the ladder's, and the ladder is an artifact of this author's circuit rather than of the construction; a re-run under lemma C.9 would carry no two-qubit gates at all, and the $N = 12$ number would be limited by readout rather than by an eleven-deep entangling chain. The $\mathcal{O}(1)$ claim of theorem 3.3 is untouched by this run in either direction.

Scope. One genuine limit remains, and it is the second one the earlier version stated. The ground states of these instances are computational basis states, so eq. (C.11) — and with it lemma C.9 — is cheap for a reason specific to diagonal H_x. Nothing here bounds the preparation cost for a non-diagonal problem, which remains the classical synthesis cost that theorem 3.12 bounds. The corollary is that for diagonal instances the quantum execution step is trivially shallow and the content of the theorem is carried entirely by the classical side; the instances that would exercise the quantum step are the non-diagonal ones, and they are not tested here.

Remark C.12 (The ordering convention was checked, not assumed)**.** The two faults recorded at sections C.1.5 and C.1.6 were both bitstring-to-index mappings. The analysis here decodes every tile under *both* the direct and the reversed vertex order and reports which carries the dominant outcome. A wrong convention could not pass silently: it would name reverse(b), a different assignment with a different cut value, and reverse(b) $\notin \{b, \bar{b}\}$ for all four instances. The check returned "direct" on all six circuits on both processors.

C.1.8 Proof oracles: theorems tested after they were proved

Section C.1.7 tested constructions. This subsection tests *theorems* — specifically those proved in sections 2 to 4 whose content has a measurable consequence. Most do not: the rank identity of theorem 5.2, the simplicity of lemma 4.24, the rewrite theorems theorems 4.10 to 4.12, the curvature counterexample of theorem 4.38 and the two history-representation results theorems 2.5 and 3.8 are statements about algebra or syntax with no observable attached, and deploying a circuit for them would be theatre. Three carry a falsifiable prediction, and those are the three deployed. Predictions were computed classically and fixed before submission, as throughout.

P1 — no invariant subspace has bounded support (theorem 5.3). The theorem says an eigenvector of an irreducible hopping chain vanishes on no two adjacent sites, hence on no site of an open uniform chain. The one-magnon eigenstate is prepared by an

$(n-1)$-step Givens cascade — linear cost, not the $\mathcal{O}(4^n)$ of generic synthesis — and the occupation profile read directly. The falsifier is an empty site.

Table C.19: Minimum site occupation and inverse participation ratio, median over tiles. Every site carries weight on every processor, and the participation ratio grows with n as the theorem requires.

	predicted		`kingston`		`fez`		`marrakesh`
n	min occ	IPR	min occ	IPR	min occ	IPR	min occ / IPR
4	0.1382	3.33	0.1356	3.39	0.1337	3.38	0.1312 / 3.39
6	0.0538	4.67	0.0501	4.73	0.0495	4.74	0.0490 / 4.74
8	0.0260	6.00	0.0232	6.07	0.0225	6.16	0.0236 / 6.18

The test sharpens with n: the predicted minimum falls from 0.138 to 0.026 while the participation ratio climbs from 3.33 to 6.00, so a localised state would be easier to detect at $n = 8$ than at $n = 4$. None appears. Occupations sit 3–11% below prediction uniformly, which is the depolarising visibility factor of remark C.7 and not a directional effect, and the participation ratios exceed prediction slightly, as depolarising noise flattens a profile rather than sharpening it.

P3 — the A_1 sector is conserved (theorem C.1). An earlier build of this test emitted $e^{-iH_k t}$ as a generic k-qubit unitary and measured the depth of that synthesis rather than the conservation law; its retention fell to 0.47 at $k = 4$ while being flat in t, which is the signature of circuit error rather than leakage. It is superseded. The deployment reported here runs the structured route and, more importantly, carries a control: the conserving arm drives each bond with $e^{-i\theta(XX+YY)/2}$, which commutes with A_1, and the breaking arm with $e^{-i\theta XX/2}$, which does not. The two have the same depth, the same gate count and the same gate types.

Table C.20: Sector retention, median over tiles, and the separation between the arms. The arms are the same circuit at $t = 0$; the predicted separation is 0, 0.129, 0.250 at $k = 3$ and 0, 0.257, 0.500 at $k = 4$.

k	t	**arm**	`kingston`	`fez`	`marrakesh`	**separation**
3	0.0	conserving	0.9795	0.9604	0.9675	
		breaking	0.9780	0.9604	0.9672	0.002
3	0.4	conserving	0.9057	0.8717	0.8825	
		breaking	0.8021	0.7776	0.7821	0.094–0.104
3	0.8	conserving	0.9103	0.8792	0.8894	
		breaking	0.6945	0.6779	0.6812	0.201–0.216
4	0.4	conserving	0.8814	0.8093	0.8375	
		breaking	0.6641	0.6231	0.6419	0.186–0.217
4	0.8	conserving	0.8856	0.8190	0.8502	
		breaking	0.4539	0.4329	0.4453	0.386–0.432

Three features make this a measurement of the conservation law rather than of the circuit. At $t = 0$ the arms agree to within 0.002 on every processor, as they must,

being the same circuit. The separation grows monotonically in t and reaches 83% of prediction, the shortfall being the same multiplicative visibility present in both arms. And the measured $k = 4$ to $k = 3$ ratio of separations is 2.00 on all three processors against a predicted 2.00. The conserving arm loses 0.09–0.18 to hardware; the breaking arm loses a further 0.20–0.43 to the physics.

P4 — the Host algebra is decoherence-free (theorem 3.7). The theorem follows from $A_\mu P_H = 0$ killing every dissipator term. The test compares a host state against a non-host control on the same tiles, both evolved under H_2: the host arm should stay in the $\mathbf{z} = 0$ sector and the control should never enter it.

Table C.21: Weight in the host sector, median over 64 two-qubit tiles.

t	**arm**	**pred**	`kingston`	`fez`	`marrakesh`
0.0	host	1	0.9837	0.9703	0.9726
	control	0	0.0082	0.0190	0.0179
0.6	host	1	0.9728	0.9575	0.9533
	control	0	0.0267	0.0506	0.0462
1.2	host	1	0.9528	0.9358	0.9418
	control	0	0.0378	0.0529	0.0523
1.8	host	1	0.9552	0.9395	0.9407
	control	0	0.0377	0.0575	0.0488

The arms never approach each other: the minimum separation over all times and processors is 0.88, and the host arm's decay from 0.98 to 0.95 over the interval is an order of magnitude smaller than the 0.92 that separates the arms. The control's slow drift upward, from 0.01 to 0.06, is leakage into the host sector from a state outside it — the direction the theorem does not constrain.

Remark C.13 (P2, and a decomposition that could not be made)**.** A fourth test was built for lemma C.10 and is reported here as a negative. The lemma is exact: the doubling tree produces $\frac{1}{\sqrt{2}}(|b\rangle + |\bar{b}\rangle)$ in $\lceil \log_2 N \rceil$ layers, at unit fidelity in simulation. On hardware it does not win. Across every deployment the tree matched the ladder at $N = 4$, where routing costs nothing, and lost or tied above it. The cause is visible in the transpiled circuits: the tree's CNOTs are non-adjacent, and on a path tile they route to 16 and 34 two-qubit gates against the ladder's 7 and 11. This is the lemma's own caveat measured — it records that nearest-neighbour unitaries carry an $\Omega(N)$ lower bound [2], and the tree evades that bound only where non-adjacent gates are free, which a path tile does not provide.

We attempted to decompose the cost into a per-gate and a per-layer term and could not. Holding depth while varying gate count requires slack in the circuit, and a nearest-neighbour ladder is a serial chain with none: every gate lies on the critical path, and four padding designs admitted between zero and two extra gate pairs before depth grew. A sweep over pad counts moves gate count and depth together — the two are correlated at 0.999 across the points obtained — so the design matrix is rank deficient and the coefficients are not identifiable from it. The one construction off that line, the tree itself, changes the tile placement and with it the routing, and its measured fidelity moved by 0.27 between tilings of the same circuit. We therefore report the ordering we can support — the tree confers no advantage on this architecture — and not a ratio.

C.1.9 P5 — the closure stage is stationary under repetition

Theorem 4.59 is combinatorial and is checked classically (section B). Its physical shadow, remark 4.62, is not: if the closure stage is carried by a block commuting with A_1, the occupation of the protected sector should not drift as that block is reapplied. This subsection tests that, with the control the P2 retraction taught us to build.

Design. One excitation is prepared per tile and the closure block applied $r = 1, 2, 3$ times. The *conserving* arm drives each bond with $\exp(-it(XX+YY)/2)$, which commutes with A_1; the *breaking* arm uses $\exp(-itXX/2)$, the same gate family and the same two-qubit count, but it does not commute with A_1. The two arms carry *identical* two-qubit counts at every (k, r) — 550/1250/1950 at $k = 3$ and 660/1458/2295 at $k = 4$ — so a loss common to both is decoherence rather than sector leakage.

The bond angle $t = 1.20$ was fixed on simulation before submission, and was not the first choice: at $t = 0.6$ the breaking arm climbs back toward 1.0 as r grows, because repeating a fixed rotation is periodic, and by $r = 3$ the two arms are no longer separable. At $t = 1.20$ the breaking arm is instead *non-monotonic* in r, dipping at $r = 2$ — a shape decoherence cannot imitate, since noise degrades monotonically. Predictions were computed classically and fixed before the jobs were submitted, as throughout.

Result. Four executions were performed: one on each of the three processors, and an earlier run on `ibm_fez` at identical parameters.

Absolute retention does not reach the predicted 1.0, and was not expected to: at $k = 4$, $r = 3$ the circuit carries 2295 two-qubit gates. Across the three processors the conserving arm reads between 0.609 and 0.848, and falls *monotonically* in r in every case — the signature of decoherence, not of sector leakage. Because the arms are depth-matched, the ratio of breaking to conserving retention divides that common loss out, and since the conserving arm's ideal value is exactly 1, the prediction for that ratio is just the breaking arm's simulated retention.

Table C.22: P5: ratio of breaking to conserving retention, by processor, at 10,000 shots with 50 tiles at $k = 3$ and 28 at $k = 4$. Jobs `dacj7pdnj4cs73acmp50` (`kingston`), `dacj7t3dd5gc73d64t7g` (`fez`) and `dacj82d1ierc738jma5g` (`marrakesh`). The ratio removes the decoherence common to the two depth-matched arms; its prediction is the breaking arm's simulated retention, fixed before submission.

k	r	`kingston`	`fez`	`marrakesh`	**predicted**
3	1	0.904	0.912	0.900	0.886
3	2	0.763	0.761	0.772	0.752
3	3	0.843	0.856	0.859	0.843
4	1	0.786	0.792	0.797	0.772
4	2	0.541	0.563	0.563	0.504
4	3	0.658	0.650	0.647	0.685

The ratio tracks its prediction to within 0.059 across all eighteen points, and every processor reproduces the non-monotonic dip at $r = 2$ at both k. Both pre-registered checks pass on all three processors at both k, twelve cases out of twelve: the conserving arm varies less over r than the breaking arm, so stationarity is not falsified; and the

breaking arm is non-monotonic, so the control discriminates. The conserving arm lies above the breaking arm at every one of the eighteen matched-depth points.

Agreement. The three processors agree with one another to 0.022, and the two `ibm_fez` runs — independent submissions at identical parameters — agree to 0.006. Both figures are smaller than the 0.059 by which the ratio departs from prediction, which places the residual in the model rather than in the measurement: were it noise, the spread between processors would be the larger quantity.

Scope. The ratio construction assumes the conserving arm is a faithful decoherence reference for the breaking arm. The identical two-qubit counts support that and the agreement between processors is consistent with it, but neither proves it; a residual difference in one-qubit wrapper depth remains, at most 5% at $k = 4$. What the test establishes is the relative statement — that the sector-conserving block does not drift where a depth-matched non-conserving block does — and not an absolute retention figure, which decoherence at these depths precludes.

C.2 Complete Job Registry

Every quantum-processor execution on which this paper depends is listed below: table C.23 covers the Phase 3–6 programme of section C.1, table C.24 the theorem tests of section C.1.6, and table C.25 the oracle reference tests of section C.1.7. All 55 identifiers remain queryable through the IBM Quantum Platform API and constitute the primary record; the archived JSON files of the supplement are derived from these jobs and carry the originating identifier in the filename.

Table C.23: Phase 3–6 executions. Twenty-one jobs at 50,000 shots each, 1,050,000 shots in total, on three Heron-class processors. The **Type** column reproduces the experiment label stored with each job; "Topological Memory" is the archived name for the tiled memory circuit of section 6, kept verbatim so that the table matches the metadata held by the provider rather than being silently relabelled here.

#	Job identifier	Backend	Type	Shots	Tiles	Status
1	`da8l1phqtnsc73d0jvi0`	`ibm_marrakesh`	Topological Memory	50k	33	DONE
2	`da8l1serbfbs73cglua0`	`ibm_kingston`	Topological Memory	50k	38	DONE
3	`da8l1use74ec73aipkrg`	`ibm_fez`	Topological Memory	50k	34	DONE
4	`da8lcr1qtnsc73d0kcv0`	`ibm_marrakesh`	Phase 3 Oracle	50k	33	DONE
5	`da8ldlerbfbs73cgmbrg`	`ibm_kingston`	Phase 3 Oracle	50k	38	DONE
6	`da8ldnke74ec73aiq2g0`	`ibm_fez`	Phase 3 Oracle	50k	34	DONE
7	`da8ljd9qtnsc73d0kj50`	`ibm_marrakesh`	Chemistry Oracle	50k	33	DONE
8	`da8ljmce74ec73aiq860`	`ibm_fez`	Crypto Oracle	50k	34	DONE
9	`da8lkt6rbfbs73cgmim0`	`ibm_kingston`	Optimization Oracle	50k	38	DONE
10	`da8m87mrbfbs73cgn4tg`	`ibm_marrakesh`	$k = 3$ Oracle	50k	51	DONE
11	`da8mbjhqtnsc73d0la4g`	`ibm_kingston`	$k = 3$ Oracle	50k	49	DONE
12	`da8mboce74ec73aiqv7g`	`ibm_fez`	$k = 3$ Oracle	50k	49	DONE
13	`da8mbv4jbipc73fdtrdg`	`ibm_marrakesh`	$k = 4$ Oracle	50k	32	DONE
14	`da8mc5se74ec73aiqvlg`	`ibm_kingston`	$k = 4$ Oracle	50k	32	DONE
15	`da8mccke74ec73aiqvv0`	`ibm_marrakesh`	$k = 5$ Oracle	50k	29	DONE
16	`da8mrh6rbfbs73cgnovg`	`ibm_fez`	MaxCut Encoding	50k	29	DONE
17	`da8mrhmrbfbs73cgnp0g`	`ibm_fez`	Factoring Encoding	50k	29	DONE
18	`da8mri6rbfbs73cgnp10`	`ibm_fez`	Chemistry Encoding	50k	29	DONE
19	`da8mrike74ec73airf8g`	`ibm_fez`	Portfolio Encoding	50k	29	DONE
20	`da8mrj1qtnsc73d0lqjg`	`ibm_fez`	TSP Encoding	50k	29	DONE
21	`da8mrjse74ec73airfag`	`ibm_fez`	Ising Encoding	50k	29	DONE

Table C.24: Theorem-test executions F13–F17. Twelve jobs, 2,860,000 shots in total. "Circ." is the number of distinct circuits in the job and "Shots" the shot count per circuit. Job 22 is retained although superseded: it is the record of the qubit-ordering fault described in section C.1.5, not a result that stands. A thirteenth submission, `da9foo9qtnsc73d1hpeg` on `ibm_fez`, was cancelled in the queue and returned no counts; it is recorded in the archived JSON but carries no result and is not counted here.

#	Job identifier	Backend	Test	Circ.	Shots	Tiles
22	`da96dj6rbfbs73chaifg`	`ibm_fez`	F13 ground prep., superseded	6	50k	35
23	`da98vhce74ec73ajglgg`	`ibm_fez`	F13 ground prep., corrected	6	50k	36
24	`da99d0urbfbs73chdce0`	`ibm_kingston`	F14 closures, run 1	10	10k	1
25	`da99ickjbipc73fek51g`	`ibm_fez`	F14 closures, run 1	10	10k	1
26	`da99orke74ec73ajhdu0`	`ibm_kingston`	F14 closures, run 2	11	10k	1
27	`da99p91qtnsc73d1bqq0`	`ibm_fez`	F14 closures, run 2	11	10k	1
28	`da99qq6rbfbs73chdqd0`	`ibm_kingston`	F15 tiled full-device	10	50k	50
29	`da99sdkjbipc73fekfq0`	`ibm_fez`	F15 tiled full-device	10	50k	50
30	`da9almhqtnsc73d1cmm0`	`ibm_kingston`	F16 mediator bound	15	20k	50/28
31	`da9bsvsjbipc73femhq0`	`ibm_fez`	F16 mediator bound	15	20k	50/28
32	`da9fo6ke74ec73ajncb0`	`ibm_kingston`	F17 oracle projection, $N = 6$–12	6	20k	8–17
33	`da9fqt6rbfbs73chjrb0`	`ibm_marrakesh`	F17 oracle projection, $N = 6$–12	6	20k	8–17

Table C.25: Oracle reference executions. Twelve jobs, 11,700,000 shots. Jobs 34–36 carry O1–O4; jobs 37–39 are the O2 rebuild only; jobs 40–42 carry O1–O5 with the corrected ordering of remark C.12, and jobs 43–45 are O6. Jobs 46–51 are the proof oracles of section C.1.8, at 10,000 shots: 46–48 carry P1–P4 and 49–51 the P3 control arm. Jobs 52–55 are the closure-stationarity test of section C.1.9: an initial run on `ibm_fez`, then one execution on each processor. Superseded builds are archived with the rest and are not counted here. Two further submissions did not execute: `da9qqfce74ec73ak2s2g` on `ibm_marrakesh`, which reached the account usage limit, and `daarnkhl216s739olupg` on `ibm_fez`, an O6 replicate still queued at the time of writing. Both are recorded in the archive without a result, and no claim in this paper depends on either.

#	Job identifier	Backend	Content	Circ.	Shots
34	`da9qgf9qtnsc73d1stp0`	`ibm_kingston`	O1–O4, O1 rescored	26	50k
35	`da9qlferbfbs73chv1s0`	`ibm_fez`	O1–O4, O1 rescored	26	50k
36	`da9qrherbfbs73chv810`	`ibm_marrakesh`	O1–O4, corrected	26	50k
37	`da9r3dmrbfbs73chvg8g`	`ibm_kingston`	O2 rebuild	6	50k
38	`da9r4jhqtnsc73d1timg`	`ibm_fez`	O2 rebuild	6	50k
39	`da9r624jbipc73ff67r0`	`ibm_marrakesh`	O2 rebuild	6	50k
40	`daaq54mrrl7c7385tof0`	`ibm_marrakesh`	O1–O5	31	50k
41	`daaqcmrvpcac73dcs0mg`	`ibm_fez`	O1–O5	31	50k
42	`daarvuj4clkc73fi3di0`	`ibm_kingston`	O1–O5	31	50k
43	`daar0u1l216s739okur0`	`ibm_kingston`	O6	15	50k
44	`daas2hpl216s739omdag`	`ibm_kingston`	O6 replicate	15	50k
45	`daarl4j4clkc73fi2tog`	`ibm_marrakesh`	O6	15	50k
46	`dabgsdpl216s739phl90`	`ibm_kingston`	P1–P4	23	10k
47	`dabgsgbvpcac73ddpi00`	`ibm_fez`	P1–P4	23	10k
48	`dabgsir4clkc73fiuos0`	`ibm_marrakesh`	P1–P4	23	10k
49	`dac1rsde36ac739g48ig`	`ibm_kingston`	P3 with control	18	10k
50	`dac1rtoc4p7c738kn9ig`	`ibm_fez`	P3 with control	18	10k
51	`dac1rv0c4p7c738kn9k0`	`ibm_marrakesh`	P3 with control	18	10k
52	`dacgic5nj4cs73acjer0`	`ibm_fez`	P5 closure	12	10k
53	`dacj7pdnj4cs73acmp50`	`ibm_kingston`	P5 closure	12	10k
54	`dacj7t3dd5gc73d64t7g`	`ibm_fez`	P5 closure	12	10k
55	`dacj82d1ierc738jma5g`	`ibm_marrakesh`	P5 closure	12	10k

References

[1] IQS/EQS Formalization. Internal quantum structure / external quantum structure: A complete operator-algebraic formalization (2024–2026). Parts 1–297. Unpublished manuscript; available from the corresponding author on reasonable request. See Appendix A for the index of Parts cited here and the result drawn from each. Every result in the present work is proved in place; no claim depends on access to the corpus.

[2] Bravyi, S., Hastings, M. B. & Verstraete, F. Lieb-robinson bounds and the generation of correlations and topological quantum order. *Phys. Rev. Lett.* **97**, 050401 (2006).

[3] Lu, T.-C., Lessa, L. A., Kim, I. H. & Hsieh, T. H. Measurement as a shortcut to long-range entangled quantum matter. *PRX Quantum* **3**, 040337 (2022).

[4] Tantivasadakarn, N., Thorngren, R., Vishwanath, A. & Verresen, R. Long-range entanglement from measuring symmetry-protected topological phases. *Phys. Rev. X* **14**, 021040 (2024).