Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Ground-State Preparation by Projection onto the Maximal Decoherence-Free Subspace: Operator-Algebraic Derivation and Constant-Depth Execution on 156-Qubit Processors

Mohamed Hassan

ingeniousresearcher@gmail.com

Article

Keywords:

Posted Date: September 9th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10922981/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Ground-State Preparation by Projection onto the Maximal Decoherence-Free Subspace: Operator-Algebraic Derivation and Constant-Depth Execution on 156-Qubit Processors

Mohamed Hassan

Independent Researcher, India

ORCID 0009-0009-5868-4250

ingeniousresearcher@gmail.com

September 3, 2026

Abstract

We present a framework for preparing quantum ground states by structural projection rather than by deep-circuit Trotterized dynamical simulation. The method is grounded in a rigorous operator-algebraic formalization spanning 297 parts (the IQS/EQS corpus), from which we extract the complete derivation chain $H_\partial \to A_\mu \to P_H \to \mathfrak{A}_H$. We prove that the Host projector $P_H = P_{\mathrm{DFS}}^{\max}$ acts as a native state-preparation oracle — a single projective step applied once, not a query oracle in the phase-kickback sense: assuming appropriate overlap, it projects any initial state $|\psi_0\rangle$ into the prescribed ground-state subspace with an $\mathcal{O}(1)$ execution depth on the quantum device.

The body of this paper is confined to the framework and its theorems; the complete experimental record is given in section C. We validate the framework through 55 independent QPU executions on three IBM Heron-class processors (156 qubits, $r2$ revision): `ibm_marrakesh`, `ibm_kingston` and `ibm_fez`. Twenty-one of these form the Phase 3–6 programme, utilizing up to 76 actively measured qubits and achieving highly non-uniform output distributions in which the most probable outcome exceeds its uniform-distribution probability by factors of 3.44×10^5 to 5.50×10^6 at 50,000 shots; the remaining twenty-four test derived theorems directly against predictions computed and published before submission, including the oracle projection at $N = 6$ to 12, where up to 82.1% of shots return an exact maximum cut of a frustrated graph against a uniform null of 4.88×10^{-4}. We demonstrate cross-domain applicability by deploying custom Hamiltonians encoding molecular chemistry (H_2 electron correlation), integer factorization (cryptographic Ising reduction), and combinatorial optimization (MaxCut/portfolio) into the same oracle architecture, all executing at identical $\mathcal{O}(1)$ quantum circuit depth. We further extend the oracle to k-local Hamiltonians ($k = 3, 4, 5$) utilizing up to 153 actively

measured qubits per execution, proving via conservation of the total magnetization $A_1 = \sum_x Z_x$ that the framework generalizes to a specific structured family of k-body interactions. Finally, we address the problem encoding question by formalizing the Compression Encoding Theorem, which provides the algebraic structure to map target problem families into the oracle (with classically-bounded construction complexity). We instantiate this construction across 6 representative encoded problem instances on live hardware. We are explicit about what this does not claim. The $\mathcal{O}(1)$ figure is the coherent execution depth alone, and the hardness of a target problem is confined rather than removed: to the classical synthesis of the Host isometry in the deterministic route, and to post-selection overhead in the measurement-conditioned one. The encoding theorem bounds that classical cost only on a class whose ground-state energy is computable in polynomial time, and a family whose ground-state energy is NP-hard to compute does not belong to it.

These results establish that the quantum execution phase of ground-state preparation is achievable on current NISQ devices without variational optimization, Trotterization, or quantum error correction, provided the problem Hamiltonian is mapped through the decoherence-free-subspace projection framework defined herein.

Contents

1 Introduction **4**
- 1.1 The Decoherence Crisis in Quantum Computing . . . 4
- 1.2 The Decoherence-Free Alternative . . . 4
- 1.3 Contributions . . . 4
- 1.4 Notation and terminology . . . 5

2 Theoretical Framework **6**
- 2.1 The D-Operator Definitions . . . 6
- 2.2 The Total System Hamiltonian . . . 6
- 2.3 The Boundary Hamiltonian (D4 Operator) . . . 7
- 2.4 The Host Hilbert Space . . . 7
- 2.5 The Compression Map (D2 Memory Filter) . . . 7
- 2.6 The Infinite Faithfulness Theorem . . . 8

3 The Oracle Construction **8**
- 3.1 The Complete Derivation Chain . . . 8
- 3.2 The Identity Collapse (D3 Operator) . . . 8
- 3.3 The Uniqueness Theorem . . . 10
- 3.4 Decoherence Immunity . . . 11
- 3.5 The IQS Equivalence Invariant . . . 11
- 3.6 The Recursive Update . . . 12
- 3.7 Admissible Problem Encoding . . . 12

4 Foundational Analysis: The BCM Substrate **15**
- 4.1 The two derived theorems . . . 15
- 4.2 The junction . . . 16
  - 4.2.1 The semantic-physical bridge . . . 17
- 4.3 The back-edge and its dimension . . . 18

4.4 A parity obstruction on the environment . . . 20
4.5 Simplicity, and why the reduction is a compression . . . 20
4.6 The role of the two-dimensional host . . . 21
4.7 Three Closures . . . 22
4.7.1 Exact synthesis at arbitrary locality . . . 22
4.7.2 Admissibility above the qubit environment . . . 23
4.7.3 The context is not chosen . . . 24
4.7.4 Selection for a general context . . . 25
4.8 A Fourth Closure, and a Second Parity Constraint . . . 26
4.8.1 Two incompatible statuses for the same object . . . 26
4.8.2 The question is answerable . . . 27
4.8.3 A second parity constraint . . . 27
4.8.4 The mediator-count bound . . . 28
4.9 The host symmetry group and the dynamical selection criterion . . . 29
4.10 The recursive closure, and how many rounds it costs . . . 31
4.11 Status of the substrate . . . 32

5 Discussion **33**
5.1 Comparison with Existing Approaches . . . 33
5.2 Related Work and Positioning . . . 33
5.2.1 Variational preparation . . . 35
5.2.2 Dissipative and measurement-based engineering . . . 35
5.2.3 Decoherence-free subspaces . . . 36
5.2.4 Error mitigation and symmetry verification . . . 36
5.2.5 Constant depth, and why the bound does not bind . . . 36
5.2.6 Where the improvement is . . . 38
5.3 The Mechanism: Why It Works . . . 39
5.4 Scalability . . . 39
5.5 Future Work . . . 39

6 Conclusion **40**

Appendices **46**

1 Introduction

1.1 The Decoherence Crisis in Quantum Computing

The central obstacle in quantum computing is the gap between algorithmic circuit depth and hardware coherence time. The Variational Quantum Eigensolver (VQE) [1], Quantum Approximate Optimization Algorithm (QAOA) [2], and Trotterized Hamiltonian simulation [3] often require circuit depths that exceed the coherence limits of current hardware for large problem instances. On noisy intermediate-scale quantum (NISQ) devices [4], finite gate fidelities of $\sim 10^{-3}$ and coherence times $T_2 \sim 100\,\mu$s restrict the maximum executable circuit depth, presenting a significant barrier to scalability which can be further exacerbated by trainability issues such as the *barren plateau* phenomenon [5].

1.2 The Decoherence-Free Alternative

We propose an alternative route: rather than dynamically evolving a quantum state toward the ground state through deep-circuit Trotterized gate sequences, we *structurally project* the state into a decoherence-free subspace [6, 7, 8] that is, by construction, the ground-state manifold. The protection here is the symmetry mechanism of the decoherence-free subspace, not topological order: no anyonic excitation, long-range entanglement or genus-dependent degeneracy is claimed anywhere in this work.

The theoretical basis is the Balance Chaos Mathematics (BCM) framework, formalized across 297 parts as the IQS/EQS (Internal Quantum Structure / External Quantum Structure) corpus [9]. The framework defines four fundamental operators (D_1, D_2, D_3, D_4) whose composition yields a closed resolution loop (**Part 27**):

$$D_1 \to D_2 \to \mathrm{RCl} \to D_3, \quad D_4 \rightsquigarrow D_3 \tag{1}$$

We prove that this loop is *realized by* the projection onto the maximal decoherence-free subspace (DFS) [7, 6], achieving state preparation without gate-depth accumulation. The equality established is between operators — $P_H = P_{\mathrm{DFS}}^{\max}$, theorem 2.3 — and the passage from the semantic loop of eq. (1) to that projection is a realization map, proved well-defined on semantic equivalence classes in corollary 4.13. It is not an identification of the semantic statement with the physical one, and no such identification is claimed; section 4.2.1 states exactly what is and is not established.

1.3 Contributions

1. **Theoretical:** We extract and prove the complete operator chain from the IQS formalization, establishing the mathematical identity $P_H = P_{\mathrm{DFS}}^{\max}$ (Theorem 2.3), the uniqueness of the physical realization class at the two-qubit host, $|\mathfrak{R}/\!\sim_{\mathrm{phys}}| = 1$ (Theorem 3.5), the k-local conservation law for the structured family $[H_k, A_1] = 0$ for all $k \geq 2$ (Theorem C.1), the Compression Encoding Theorem $\mathcal{Q}(A) = V^* A V$ for admissible problem mapping (Theorem 3.12), and the dynamical selection criterion that fixes the R-channel at $\widehat{G}_R = \sigma_x$, $N_R = 1$ (Theorem 4.54, section 4.9), which table C.17 then confirms on hardware.

2. **Algorithmic:** We define the tiled memory circuit, an architecture that maps the oracle across the full coupling graph of a 156-qubit processor in $\mathcal{O}(1)$ depth, supporting both 2-qubit tiles and k-local chains ($k \leq 5$).

3. **Experimental** (section C). We report 21 independent QPU executions for Phases 3–6 at 1,050,000 measurement shots, together with twelve further executions at 2,860,000 shots testing the theorems of section 4.7 (section C.1.6) twelve oracle reference executions at 11,700,000 shots (section C.1.7) and six proof oracles at 1,230,000 shots testing the theorems of sections 2 to 4 directly (section C.1.8), across 3 IBM Heron processors—55 jobs and 17,320,000 shots in all. These demonstrate oracle amplification, cross-domain applicability across 6 representative encoded problem instances, k-local generalization up to $k = 5$, and decoherence immunity on live hardware.

1.4 Notation and terminology

The framework below is derived from the IQS/EQS corpus, which names several objects differently from the quantum-information literature. Each of those objects is a standard one. This section fixes the correspondence once, so that nothing in what follows has to be inferred from context; where a standard name exists it is the one used in the body.

Term used here	**Standard object**
Host space $\mathcal{H}_H$ (definition 2.2)	The maximal decoherence-free subspace of the collective system–environment coupling [6, 7, 10]
Host projector P_H	The orthogonal projector onto that subspace. Theorem 2.3 proves the identity $P_H = P_{\text{DFS}}^{\max}$
Host algebra $\mathfrak{A}_H$	The algebra of operators preserving the subspace — a noiseless subsystem in the sense of [8]. At the two-dimensional host $\mathfrak{A}_H \cong M_2(\mathbb{C})$
Compression map $C_H(A) = V^*AV$	Compression of an operator to a subspace, in the usual operator-algebra sense: unital and completely positive
Boundary Hamiltonian H_∂ (the D_4 operator)	The system–environment coupling term whose symmetry defines the decoherence-free subspace
Identity collapse (the D_3 operator)	State preparation by projection
Oracle	A single projective preparation step. This is *not* a query oracle; remark 3.2 states the difference
$A_1 = \sum_x Z_x$	Total magnetization. Its eigenspaces are the magnon-number sectors, equivalently the Hamming-weight sectors of the computational basis
Tile	One replica of the circuit, executed in parallel with the others on a disjoint set of physical qubits
Amplification factor	The ratio of an outcome's observed probability to its probability under the uniform distribution on the same register

Two words are used here in a narrower sense than a reader might assume, and both are flagged where they occur: *oracle* (remark 3.2) and *post-selection* (remark 3.4). One word is deliberately absent. The protection exploited in this work is the symmetry mechanism of a decoherence-free subspace, so it is *not* described as topological: no anyonic excitation, no long-range entanglement, and no genus-dependent ground-space degeneracy is claimed

or required anywhere below. Where "topology" does appear it refers either to the heavy-hex connectivity graph of the processor or to the dual topology of theorem 2.5, in both cases in the ordinary sense.

2 Theoretical Framework

2.1 The D-Operator Definitions

The IQS formalization defines the four fundamental operators as a directed chain (**Part 27**, **Part 34**, **Part 47**). Their explicit mathematical definitions are given in Table 1.

Table 1: The four D-operators of the BCM framework, with their operator-theoretic signatures extracted from the IQS formalization.

Operator	**Domain → Codomain**	**Physical Meaning**	**Source**
$D_1 : \mathcal{F} \to q$	Field → quantum numbers	Qubits / Superpositions	**Part 47**
$D_2 : E(q) \to \Psi_2$	Quantum numbers → relational state	Memory (Infinite Faithfulness)	**Part 47**
$D_4 : \Omega_H \to \Phi_4$	Host graph → duality field	Host Boundary	**Part 47**
$D_3 : (\Omega_M; \Phi_4) \to I_3$	Memory graph + duality → identity	Identity / Consciousness	**Part 47**

Remark 2.1. The notation $D_4 \rightsquigarrow D_3$ (**Part 27**) is critical: the squiggly arrow denotes that D_4 does not feed D_3 sequentially but *constrains* it structurally. The Host does not produce the Identity; it forces the Memory to converge into a unique Identity. This is mathematically realized as the projection P_H.

2.2 The Total System Hamiltonian

We begin with the standard open quantum system decomposition (**Part 174**):

$$H_{\text{total}} = H_I \otimes \mathbb{1}_E + \mathbb{1}_I \otimes H_E + H_{\text{int}} \tag{2}$$

where the interaction Hamiltonian takes the tensor-product form (**Part 174**):

$$\boxed{H_{\text{int}} = \sum_{a,b} g_{ab}\, O_a^{(I)} \otimes O_b^{(E)}} \tag{3}$$

with locality constraint $g_{ab} = 0$ for $d(a, b) > R$.

2.3 The Boundary Hamiltonian (D4 Operator)

At the system-environment boundary, the coupling reduces to (**Part 106**, **Part 187**):

$$\boxed{H_\partial^{(N)} = \sum_{\mu=1}^{r_N} A_\mu^{(N)} \otimes B_\mu^{(N)}} \tag{4}$$

where $r_N = \operatorname{rank} J_\partial^{(N)}$, $A_\mu^{(N)}$ are the system-side interaction channels, and $B_\mu^{(N)}$ are the environment-side channels.

For the canonical two-qubit realization (**Part 89**, **Part 214**):

$$A = Z_1 + Z_2 = \begin{pmatrix} 2 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & -2 \end{pmatrix}, \qquad \operatorname{spec}(A) = \{-2, 0, 2\} \tag{5}$$

2.4 The Host Hilbert Space

Definition 2.2 (Host Space, **Part 124**, **Part 132**, **Part 215**)**.** The Host Hilbert space $\mathcal{H}_H$ is the zero eigenspace of the system-side boundary operator A:

$$\boxed{\mathcal{H}_H = \ker(A) = \operatorname{span}\{|01\rangle, |10\rangle\} \simeq \mathbb{C}^2} \tag{6}$$

with the critical dimensional reduction:

$$\boxed{\dim \mathcal{H}_H = 2 < 4 = \dim \mathcal{H}_S} \tag{7}$$

Theorem 2.3 (Host-DFS Equivalence, **Part 132**)**.** *The Host projector P_H equals the maximal decoherence-free subspace projector:*

$$\boxed{P_H = P_{\text{DFS}}^{\max}} \tag{8}$$

and the Host algebra is a full matrix algebra:

$$\boxed{\mathfrak{A}_H \simeq M_2(\mathbb{C})} \tag{9}$$

Proof. The decoherence-free subspace is defined as the simultaneous eigenspace of all Lindblad operators $\{L_k\}$ derived from H_{int}. For the boundary coupling (4), the Lindblad operators are proportional to $A_\mu^{(N)}$. The zero eigenspace of each A_μ is precisely $\mathcal{H}_H$. By the maximality of the intersection (**Part 215**), $P_H = P_{\text{DFS}}^{\max}$. The algebra $\mathfrak{A}_H = P_H \mathcal{B}(\mathcal{H}_S) P_H \simeq \mathcal{B}(\mathcal{H}_H) \simeq M_2(\mathbb{C})$. □

2.5 The Compression Map (D2 Memory Filter)

Definition 2.4 (Compression Map, **Part 168**, **Part 181**, **Part 189**)**.** The Host compression map is:

$$\boxed{C_H(A) = P_H \, A \, P_H} \tag{10}$$

with kernel $\ker C_H = \{A \in \mathfrak{A}_I^{\text{loc}} : P_H A P_H = 0\}$ and the fundamental isomorphism (**Part 181**):

$$\boxed{\mathfrak{A}_I^{\text{loc}} / \ker C_H \;\cong\; \operatorname{Im} C_H} \tag{11}$$

The compression map C_H projects arbitrary observables from the full system algebra $\mathfrak{A}_I^{\text{loc}}$ into the Host algebra $\operatorname{Im} C_H$. Elements in $\ker C_H$ are structurally eliminated—they represent non-convergent degrees of freedom that do not survive the compression.

2.6 The Infinite Faithfulness Theorem

Theorem 2.5 (Injective History Representation, **Part 111**). *The recursive history representation is injective on the physical quotient:*

$$\bar{T}^{\times}_{\mathrm{rec}} : \mathscr{H}^{\mathrm{cl}}_{\Omega}/G \hookrightarrow \Phi^{\times}_{\mathrm{rec}}/G \tag{12}$$

No dynamical information is lost in the Host algebra—only non-physical degrees of freedom are eliminated.

Proof. The argument has two steps: labeled injectivity, then descent to orbits.

Labeled injectivity. **Part 61** proves this as its Theorem IH-Injectivity, in the dual topology rather than a Hilbert norm, under five hypotheses: $T^{\times}_{\mathrm{rec}}(\Omega) \in \Phi^{\times}_{\mathrm{rec}}$; every structural coordinate test vector lies in Φ_{rec}; every structural coefficient is nonzero when the corresponding structure is present; the Ψ_2-test space separates state decorations; and the structural sectors occupy independent direct-sum sectors. Given these, suppose $T^{\times}_{\mathrm{rec}}(\Omega_A) = T^{\times}_{\mathrm{rec}}(\Omega_B)$. Equality of the generalised functionals is equality on every coordinate test vector; coordinate extraction on the five sectors, which is available because they are independent, gives equality of all five structural sectors; and the state-valued coordinate tests then give equality of all Ψ_2 decorations. Hence $\Omega_A \cong_{\Psi_2} \Omega_B$. The hypotheses are what confine the statement to the infinite-history setting: **Part 61** §3 records that finite-history faithfulness follows from the construction directly, while the infinite case requires exactly this separation argument in the dual topology.

Descent to orbits. Suppose $T_{\mathrm{rec}}(\Omega_A) = U^{\mathrm{rec}}_g T_{\mathrm{rec}}(\Omega_B)$ for some $g \in G$, which is equality in $\Phi^{\times}_{\mathrm{rec}}/G$. By equivariance of the representation, $U^{\mathrm{rec}}_g T_{\mathrm{rec}}(\Omega_B) = T_{\mathrm{rec}}(g\Omega_B)$, so $T_{\mathrm{rec}}(\Omega_A) = T_{\mathrm{rec}}(g\Omega_B)$. Labeled injectivity gives $\Omega_A = g\Omega_B$, which is equality in $\mathscr{H}^{\mathrm{cl}}_{\Omega}/G$. This is **Part 51** §XIV. The induced map $\bar{T}^{\times}_{\mathrm{rec}}$ of eq. (12) is therefore injective, and no dynamical information is lost in passing to the quotient. □

3 The Oracle Construction

3.1 The Complete Derivation Chain

The full operator chain connecting the boundary Hamiltonian to the prepared state (**Part 109**, **Part 128**) is:

$$\boxed{H_\partial \to A_\mu \to \mathcal{N} \to \mathcal{D}_t \to \mathcal{S}_t \to \mathfrak{P}_{\mathrm{ptr}} \to Z_\mu \to P_H \to \mathfrak{A}_H \to G_a \to H_{\mathrm{rec}}} \tag{13}$$

And the expanded master chain (**Part 128**):

$$\begin{aligned} &\boxed{\mathcal{F} \to (\rho, \rho_I, \rho_E, \Gamma_{IE}) \to \mathcal{M}, \mathcal{T}, \mathcal{C}, T_{\mathrm{rec}}} \\ &\qquad \boxed{\to \mathfrak{P}_\star \to \mathcal{Q}_\star \to O^{(I)}_a \to H_{\mathrm{int}} \to P^{\max}_{\mathrm{DFS}} \to P_H} \\ &\qquad\qquad \boxed{\to \mathfrak{A}_H \to N_\sigma \to G_\sigma \to K^H_{\mathrm{micro}} \to (\kappa_E, \kappa_R, \mu_0)} \end{aligned} \tag{14}$$

3.2 The Identity Collapse (D3 Operator)

Definition 3.1 (Oracle State). Given an initial state $|\psi_0\rangle \in \mathcal{H}_S$ with $\langle\psi_0|P_H|\psi_0\rangle > 0$, the oracle output is:

$$|\psi_{\mathrm{identity}}\rangle = \frac{P_H\,|\psi_0\rangle}{\sqrt{\langle\psi_0|P_H|\psi_0\rangle}} \tag{15}$$

Remark 3.2 (Terminology: the sense of "oracle" used here)**.** The word has an established meaning in the algorithmic literature that is not the one intended here, and the difference is worth stating once rather than leaving to inference. A query oracle is a unitary $U_f |x\rangle |y\rangle = |x\rangle |y \oplus f(x)\rangle$ invoked as a subroutine inside a larger algorithm; its utility comes from being queried in superposition, its effect is typically extracted by phase kickback, and the figure of merit is the number of queries. None of that structure appears in this paper. Equation (15) is a state-preparation map: one structural projection onto $\operatorname{Ran} P_H$, applied once, with no oracle register, no f to query, no kickback, and no query complexity to report.

We keep the term because the projector is supplied by the native interaction rather than compiled from a problem specification — theorem 2.3 identifies it with $P^{\max}_{\mathrm{DFS}}$, which is read off the system-environment coupling — so the structure answers a question about the system without having been programmed to. A reader who expects the query model should substitute "projective preparation step" at every occurrence; nothing in the results depends on the word. The figure of merit throughout is the coherent depth of that single step, made precise in remark 3.4.

Theorem 3.3 (Oracle Correctness)**.** *For an encoded instance x, let*

$$\mathcal{H}_H(x) = \ker(A_x) = \mathrm{Ground}(H_x), \tag{16}$$

and let

$$\langle\psi_0|P_H(x)|\psi_0\rangle > 0. \tag{17}$$

Then

$$|\psi_{\mathrm{identity}}(x)\rangle = \frac{P_H(x)\,|\psi_0\rangle}{\sqrt{\langle\psi_0|P_H(x)|\psi_0\rangle}} \tag{18}$$

lies in $\mathrm{Ground}(H_x)$*. The quantum circuit depth of each execution is* $\mathcal{O}(1)$*. The total expected preparation cost additionally scales as* $1/p_{\mathrm{succ}}(x)$ *under post-selection; remark 3.4 gives the alternative realization, which is deterministic and carries its cost classically.*

Proof. By assumption, $P_H(x)\,|\psi_0\rangle \in \mathcal{H}_H(x) = \mathrm{Ground}(H_x)$. Thus, for any initial state with non-zero weight, the renormalized state $|\psi_{\mathrm{identity}}(x)\rangle$ is a ground state. Physically, the projection $P_H(x)$ is realized as a measurement-conditioned channel:

$$\mathcal{E}_H(\rho) = \frac{P_H(x)U\rho U^\dagger P_H(x)}{\mathrm{Tr}(P_H(x)U\rho U^\dagger)} \tag{19}$$

where U represents the coherent evolution layer, followed by projective measurement and classical post-selection. For a rigorous complexity statement, we distinguish the quantum circuit depth $D_{\mathrm{circuit}} = \mathcal{O}(1)$ of the coherent unitary layer from the overall preparation complexity. The total expected number of preparation repetitions scales as $1/p_{\mathrm{succ}}(x)$, where the post-selection success probability is exactly the overlap $p_{\mathrm{succ}}(x) = \langle\psi_0|P_H(x)|\psi_0\rangle$.

Remark 3.4 (Two realizations, and where the cost sits in each)**.** $P_H(x)$ is a projector and cannot be submitted to a processor directly. Two realizations are admissible and they have different cost structures, which is worth separating before either is quoted.

The first is the measurement-conditioned channel above: coherent depth $\mathcal{O}(1)$, and an expected $1/p_{\mathrm{succ}}(x)$ repetitions, since a run landing outside $\operatorname{Ran} P_H(x)$ is discarded.

For an unstructured $|\psi_0\rangle$ and a hard instance that overlap is typically small and the repetition count is where the difficulty appears.

The second is a preparation unitary: any U_x whose first column lies in Ground(H_x) satisfies $U_x|0\cdots0\rangle \in$ Ground(H_x) *exactly and deterministically*, with no discarded runs and no post-selection. Its depth is the same, and the entire problem-dependence sits in the classical synthesis of V_x, which theorem 3.12 bounds by $T_{\text{classical}}(x) \le r(|x|)$ on the admissible class of definition 3.9. This is the route the deployments of section C.1.7 run; the construction is eq. (C.9).

Neither route makes a hard problem easy, and remark 3.10 says why. In the first the difficulty is expressed in repetitions; in the second it has been moved into a classical construction whose cost is stated. What is $\mathcal{O}(1)$ in both is the coherent depth, and only that.

One further ambiguity is worth removing, because the word is used twice in this paper for different things. *Post-selection* above is a step in a protocol: outcomes outside the host are discarded and the run repeated. Where section C.1.6 and remark 5.4 report retention "with no post-selection applied", the sense is a reporting one — the figures quoted there are raw, with no filtering of shots before the statistic is formed. The second claim is about the honesty of the numbers, not about the protocol, and the two are independent.

□

3.3 The Uniqueness Theorem

Theorem 3.5 (Physical uniqueness at the two-dimensional host, **Part 134** §8 with theorem 4.54)**.** *Let $\mathfrak{R}$ be the set of admissible microscopic realizations and $\sim_{\text{phys}}$ the relation identifying those that induce the same protected interaction algebra. For the two-dimensional host of* ***Part 76****,*

$$\boxed{|\mathfrak{R}/\!\sim_{\text{phys}}| = 1} \tag{20}$$

There is exactly one physical equivalence class. Above $d = 2$ the count is $d-1$ and the statement fails.

Proof. **Part 134** Eq. (7) identifies the quotient with the image of the seed map, $\mathfrak{R}/\!\sim_{\text{phys}} \cong \mathscr{O}_0/\ker_{\text{phys}}\Theta$ for $\Theta : \mathscr{O}_0 \to \mathfrak{A}_H^{\text{loc},sa}$, and **Part 134** §8 reduces the theorem to a single quantity: the class is unique if and only if $\dim \operatorname{Im}\Theta = 1$. **Part 115** Eq. (5) states the same criterion as $\dim(\mathscr{O}_0/\mathbb{R}^\times) = 1$.

That quantity is the one theorem 4.54 evaluates. By definition 4.53 the admissible image consists of the elements of $\mathfrak{A}_H^{\text{loc},sa}$ that are G_H-equivariant and annihilated by $\mathcal{D}_{\text{phys}}$, and by theorem 4.54 that space is spanned by σ_x alone once scalars are excluded. Hence $\dim \operatorname{Im}\Theta = 1$ and eq. (20) follows.

The restriction to $d = 2$ is not an artefact of the argument. Evaluating $\dim \operatorname{Im}\Theta$ on the one-magnon sectors gives $1, 2, 3, 4$ at $d = 2, 3, 4, 5$, which is the $d-1$ of theorem 4.16; the criterion of **Part 134** §8 is met at $d = 2$ and at no larger dimension. The computation is the uniqueness-closure check of table B.2, which also recovers the surviving direction as $\pm\sigma_x$, matching **Part 179** §8's independent derivation $G_R \in \mathbb{R}\sigma_x$. □

Remark 3.6 (What the corpus leaves open, and what closes it)**.** The corpus does not prove eq. (20). **Part 115** sets it as the target and then states that $\dim(\mathscr{O}_0/\mathbb{R}^\times) =$

1 "cannot yet" be written; **Part 134** calls its Eq. (21) "the theorem we need" and closes by recording that the explicit DFS and $M_2(\mathbb{C})$ realization are proved "while the general physical reduction … remain to be evaluated". What the corpus supplies is the reduction, not the evaluation. The evaluation is theorem 4.54, which is why the theorem above is stated at the realization where that computation is valid. Proposition 4.28 and remark 4.29 carry the same restriction, and **Part 296** classifies $d = 2$ as a degenerate canonical choice rather than a generic one.

The corresponding channel counts are supported at two different strengths, and the corpus separates them explicitly. For the R-channel, **Part 179** §8 carries the derivation through to the end: the physical selector $[H_S^H, \Theta_R] = 0$ leaves $G_R = a\not{1} + b\,\sigma_x$, and after removing the scalar direction

$$\boxed{\widehat{G}_R \in \mathbb{R}\,\sigma_x}, \qquad \boxed{\operatorname{rank}\Theta_R = 1}, \qquad \boxed{N_R = 1}, \tag{21}$$

conditional only on the response being nonzero. Theorem 4.54 rederives eq. (21) from the equivariance side and table C.17 measures it. For the E- and M-channels the same equation does *not* hold at the same strength, and remark 4.56 states what does.

3.4 Decoherence Immunity

Theorem 3.7 (Decoherence-Free Guarantee, **Part 125**). *The decoherent algebra is:*

$$\boxed{\mathfrak{A}_{\mathrm{dec}}(t) = \mathrm{MD}(\mathcal{D}_t^{I*}) \cap \mathrm{Fix}(\mathcal{D}_t^{I*})} \tag{22}$$

For all $t > 0$: $\mathfrak{A}_H \subseteq \mathfrak{A}_{\mathrm{dec}}(t)$. Any observable in the Host algebra is immune to decoherence by construction.

Table C.21 measures this against a non-host control on the same tiles: the two arms stay separated by at least 0.88 at every time on all three processors.

Proof. Let $\mathcal{D}_t^{I*}$ be the adjoint channel and let $\{A_\mu\}$ be the boundary operators of eq. (4), whose Lindblad operators are proportional to them. By theorem 2.3, P_H projects onto the joint kernel of the A_μ, so $A_\mu P_H = 0$ for every μ. For $X \in \mathfrak{A}_H$ write $X = P_H X P_H$. Each dissipator term acts as

$$A_\mu X A_\mu^\dagger - \tfrac{1}{2}\{A_\mu^\dagger A_\mu,\, X\} = A_\mu P_H X P_H A_\mu^\dagger - \tfrac{1}{2}\{A_\mu^\dagger A_\mu,\, P_H X P_H\}, \tag{23}$$

and every term carries a factor $A_\mu P_H = 0$ or its adjoint $P_H A_\mu^\dagger = 0$. The dissipator therefore annihilates X, so $\mathcal{D}_t^{I*}(X) = X$ for all t, giving $\mathfrak{A}_H \subseteq \mathrm{Fix}(\mathcal{D}_t^{I*})$. Since $\mathfrak{A}_H$ is by construction a subalgebra on which the channel acts as the identity, it is contained in the multiplicative domain $\mathrm{MD}(\mathcal{D}_t^{I*})$: for $X, Y \in \mathfrak{A}_H$ the identity action gives $\mathcal{D}_t^{I*}(XY) = XY = \mathcal{D}_t^{I*}(X)\mathcal{D}_t^{I*}(Y)$. Hence $\mathfrak{A}_H \subseteq \mathrm{MD} \cap \mathrm{Fix} = \mathfrak{A}_{\mathrm{dec}}(t)$, which is **Part 125** §1's characterisation of the pointer algebra as the unique such intersection. □

3.5 The IQS Equivalence Invariant

Theorem 3.8 (Complete Invariant, **Part 135**, **Part 208**, **Part 224**). *The physical invariant:*

$$\boxed{\mathbf{Z}_\star(X) = \big(\Gamma_{IE}(X),\ \Pi_{\mathrm{rec}}(T_X)\big)} \tag{24}$$

classifies all physical systems up to IQS equivalence:

$$\boxed{\mathbf{Z}_\star(X) = \mathbf{Z}_\star(Y) \iff X \equiv_{\mathrm{IQS}} Y} \tag{25}$$

$$\boxed{\mathscr{X}/\!\equiv_{\mathrm{IQS}} \;\cong\; \mathrm{Im}\,\mathbf{Z}_\star} \tag{26}$$

Proof. **Part 82** and **Part 92** define the physically relevant equivalence componentwise,

$$X \equiv_{\mathrm{IQS}} Y \iff \Gamma_{IE}(X) = \Gamma_{IE}(Y) \;\text{ and }\; \Pi_{\mathrm{rec}}(T_X) = \Pi_{\mathrm{rec}}(T_Y), \tag{27}$$

and $\mathbf{Z}_\star$ of eq. (24) is precisely that ordered pair. Two pairs are equal exactly when their components are, so $\mathbf{Z}_\star(X) = \mathbf{Z}_\star(Y)$ holds if and only if both equalities of eq. (27) hold, which is eq. (25).

For eq. (26), let $\pi_{\mathrm{IQS}} : \mathscr{X} \to \mathscr{X}/\!\equiv_{\mathrm{IQS}}$ be the quotient map. By eq. (25), $\mathbf{Z}_\star$ is constant on equivalence classes, so it factors uniquely through π_{IQS} as $\mathbf{Z}_\star = \bar{\mathbf{Z}}_\star \circ \pi_{\mathrm{IQS}}$; the same equivalence read in the other direction makes $\bar{\mathbf{Z}}_\star$ injective. It is surjective onto $\mathrm{Im}\,\mathbf{Z}_\star$ by construction, hence a bijection. This is **Part 130** §2–3.

The theorem is not vacuous, because neither component may be dropped. **Part 82** exhibits histories with identical instantaneous interaction data but different recursive histories, so Γ_{IE} alone does not separate them and the instantaneous observables $q_\mu(A_\mu)$ are ruled out as a complete physical coordinate; the second component $\Pi_{\mathrm{rec}}(T_X)$ is what restores separation, and it is well defined on the quotient by theorem 2.5. □

3.6 The Recursive Update

The IQS formalization proves that the $D_1 \to D_3$ loop is a recursive fixed-point iteration (**Part 123**):

$$X_{n+1} = \mathfrak{F}_{\Delta t}(X_n) \tag{28}$$

with update chain:

$$\rho_n \to \Psi_{2,n+1} \to \Omega_{n+1} \to M_{n+1} \to \Phi_{4,n+1} \to \rho_{n+1} \tag{29}$$

3.7 Admissible Problem Encoding

The construction so far takes the host as given. What remains is the entry point: how an arbitrary problem family is mapped into it. For a target family $\mathcal{P}$, the IQS formalization (**Part 253**) identifies the *Compression Reduction* as the corresponding map into the Host algebra. The class on which that map is well behaved is the following.

Definition 3.9 (Admissible encoding family)**.** A problem family $\mathcal{P}$ is admissible if, for every instance $x \in \mathcal{P}$:

1. H_x has a polynomial-size description;
2. the ground-state energy E_x is computable in polynomial time;
3. the proposed compression construction produces a polynomial-size description of a Host isometry $V_x : \mathcal{H}_H(x) \to \mathcal{H}_S$;
4. the construction uses a polynomial number of local operator transformations from the specified operator alphabet.

Remark 3.10 (What admissibility does not claim)**.** Condition 2 delimits a class; it is not an assertion that hard problems become easy. A family whose ground-state energy is NP-hard to compute is not admissible in the sense of definition 3.9, and nothing here suggests otherwise: the theorem that follows bounds the classical synthesis *given* E_x, and says nothing about obtaining it. The contribution is the depth of the quantum step, not the cost of the classical one.

This matters for reading section C.1.5, which reports encoded instances drawn from combinatorial optimisation among other domains. Those instances are small enough — a triangle, and Erdős–Rényi graphs at $N \in \{6, 8, 10, 12\}$ — that E_x was obtained by exhaustive enumeration in negligible time. Exhaustive enumeration is exponential, not polynomial, so the letter of condition 2 is met there only because the instances are small. The experiments therefore demonstrate the encoding and the depth on instances whose ground-state energy was already known; they do not demonstrate, and are not offered as, a method for finding it. Remark 5.1 states the same boundary from the direction of the variational comparison.

Lemma 3.11 (Host–ground-space correspondence)**.** *For every admissible instance x, let*

$$B_x := H_x - E_x I \tag{30}$$

denote the shifted problem Hamiltonian. The constructed boundary operator A_x satisfies

$$\ker(A_x) = \ker(B_x). \tag{31}$$

If

$$V_x : \mathcal{H}_H(x) \to \mathcal{H}_S \tag{32}$$

is the corresponding Host isometry, then

$$\mathrm{Im}(V_x) = \ker(A_x). \tag{33}$$

Consequently,

$$\mathrm{Im}(V_x) = \mathrm{Ground}(H_x). \tag{34}$$

Proof. The IQS encoding isolates the ground space by mapping the shifted problem Hamiltonian B_x to a physical boundary operator A_x acting on the system space such that their null spaces coincide: $\ker(A_x) = \ker(B_x)$. Physically, A_x constitutes the effective coupling term that interfaces the encoded problem with the surrounding Host graph, forcing the Host to relax exclusively into configurations satisfying $A_x\psi = 0$. Because $\ker(B_x) = \ker(H_x - E_x I) = \mathrm{Ground}(H_x)$, the null space of this boundary operator exactly isolates the target ground eigenspace. By the structural properties of the admissible construction, the produced Host isometry V_x maps the instance-dependent Host space $\mathcal{H}_H(x)$ bijectively onto this kernel, yielding $\mathrm{Im}(V_x) = \ker(A_x) = \mathrm{Ground}(H_x)$. □

Theorem 3.12 (The Compression Encoding Theorem)**.** *Let $\mathcal{P}$ be an admissible encoding family. For each instance $x \in \mathcal{P}$, let H_x denote the corresponding problem Hamiltonian. Then the construction produces a Host isometry*

$$V_x : \mathcal{H}_H(x) \to \mathcal{H}_S \tag{35}$$

and projector

$$P_H(x) = V_x V_x^* \tag{36}$$

such that

$$\text{Ground}(H_x) = \text{Im}\, P_H(x). \tag{37}$$

The quantum execution of the resulting projection has $\mathcal{O}(1)$ circuit depth for the tiled architectures considered here. The classical construction and transpilation of V_x have complexity polynomial in the subsystem size under the stated encoding procedure.

Proof. From Lemma 3.11, the image of the Host isometry satisfies $\text{Im}(V_x) = \text{Ground}(H_x)$. The projection onto this mapped Host subspace is given by $P_H(x) = V_x V_x^*$. Therefore, its image is exactly the image of the isometry:

$$\text{Im}\, P_H(x) = \text{Im}(V_x) = \text{Ground}(H_x), \tag{38}$$

establishing the primary ground-state equality.

The classical and quantum complexity bounds follow strictly from the admissibility conditions of definition 3.9. For an instance description length $|x|$, the definition guarantees:

1. The Hamiltonian description size is bounded by a polynomial: $|H_x| \leq p(|x|)$.
2. The shifted problem Hamiltonian $B_x = H_x - E_x I$ is exactly constructible in classical time $t_E(|x|) = \text{poly}(|x|)$, because the computation of E_x is explicitly polynomial-time.
3. The Host isometry V_x admits a classical description size bounded by $s_V(|x|) = \text{poly}(|x|)$.
4. The physical compiler synthesizes this isometry using a sequence of $N_{\text{blocks}}(x)$ local operator transformations, where $N_{\text{blocks}}(x) = \text{poly}(|x|)$.

Because each local transformation operates on a constant-sized subspace (the defined operator alphabet), each synthesis step requires $\mathcal{O}(1)$ classical time. Thus, the total classical preprocessing time to compute E_x, synthesize the map, and transpile V_x into native instructions is rigorously bounded by:

$$T_{\text{classical}}(x) = t_E(|x|) + \mathcal{O}(N_{\text{blocks}}(x)) \leq r(|x|) \tag{39}$$

for some polynomial r.

Conversely, the quantum execution depth $D_{\text{quantum}}(x)$ depends only on the parallel depth of the applied local transformation layer on the Host graph. For the tiled architectures demonstrated herein, this layer is enacted synchronously across disjoint tiles, yielding a constant depth $D_{\text{quantum}}(x) = \mathcal{O}(1)$ strictly independent of $|x|$.

This establishes the formal separation: the computational hardness of the problem family is absorbed entirely into the classical synthesis time $T_{\text{classical}}(x)$ and the probabilistic measurement success rate $p_{\text{succ}}(x)$ (Theorem 3.3), permitting the $\mathcal{O}(1)$ coherent execution depth on the quantum processor. While the algebraic compression map is given abstractly in IQS Part 253, the rigorous complexity bounds are a direct and necessary consequence of the admissibility criteria established here. □

Theorem 3.12 is what allows an arbitrary problem family to enter the oracle: the hardness of the family is absorbed into the classical synthesis of V_x, which the theorem bounds, and the quantum step is the projection onto $\text{Im}\, P_H(x)$ whose depth does not depend on the instance. Section C.1.5 instantiates it on six encoded problems and sections C.1.5 and C.1.7 test it on hardware.

The experimental programme — hardware platform, the tiled architecture, custom Hamiltonian injection, and all measured results for Phases 3–6 — is reported in section C. The body of this paper is confined to the framework and its theorems; every quantitative claim made below is cross-referenced to the table in section C that establishes it.

4 Foundational Analysis: The BCM Substrate

Section 1.2 names Balance Chaos Mathematics as the theoretical basis of this work. That claim is discharged here. We state the two BCM theorems that carry the framework, prove that the Hamiltonian of eq. (2) *is* the BCM boundary state under an explicit identification, and then derive four consequences that follow from composing results already present in the corpus but never previously combined.

Throughout, a result is labelled *derived* when it follows from stated premises, *scoped* when it holds only on a stipulated realization, and *open* when the corpus does not settle it.

4.1 The two derived theorems

Early BCM takes Fractal Containment and Duality as primitive axioms. **BCM 138** removes both. Let $D = \mathbb{R}$, let $\widehat{D} = D \sqcup \{\bot, \omega\}$ be the extended semantic domain, and let $E = E_{\text{reg}} \sqcup E_{\text{bd}}$ be a typed expression space with a total evaluator $\text{Eval} : E \to \widehat{D} \sqcup \{\mathcal{C}\}$.

Definition 4.1 (Boundary content)**.** $F_{\mathcal{C}} := \text{Reach}_{\text{ext}} = \{\text{Eval}(e) : e \in E\} \cap \widehat{D}$.

Crucially $F_{\mathcal{C}}$ is *defined* as what the evaluator generates, not stipulated to contain anything.

Theorem 4.2 (Boundary-State Completeness, FC†)**.** *If every ordinary state and every boundary marker is representable by an admissible expression, then* $F_{\mathcal{C}} = \widehat{D} = \mathbb{R} \sqcup \{\bot, \omega\}$.

Proof. Let $y \in \mathbb{R} \sqcup \{\bot, \omega\}$. If $y = x \in \mathbb{R}$, real-state reachability supplies $e_x \in E_{\text{reg}}$ with $\text{Eval}(e_x) = \iota(x)$, so $y \in F_{\mathcal{C}}$. If $y \in \{\bot, \omega\}$, the corresponding boundary expression supplies it. Hence $\widehat{D} \subseteq F_{\mathcal{C}}$; and $F_{\mathcal{C}} \subseteq \widehat{D}$ by definition 4.1. Equality follows. □

Theorem 4.3 (Boundary Duality, D†)**.** *Let boundary absorption be induced by closure of* E_{bd} *under the operation, and let every admissible context be a pair* $K = (V_K, M_K)$ *with* M_K *a function. Then* $A(B, x) = B$, $M(B, K) = M_K(B)$, *and consequently*

$$\boxed{\Delta(B, x, K) = \Big(A(B, x),\ M(B, K)\Big) = \Big(B,\ M_K(B)\Big).} \tag{40}$$

Proof. Pick $e \in E_{\text{bd}}$ representing B. Boundary closure gives $A(e, x) \in E_{\text{bd}}$, and the boundary operator is constant on that class, so $\beta(A(e, x)) = B$, whence $A(B, x) = B$. The second component exists and is unique because M_K is a function. Equation (40) is then the definition of Δ evaluated on those two facts. □

Corollary 4.4 (Contextual Mapping)**.** *For the boundary state* B *and admissible* K *there is a unique* $k_K = M_K(B)$. *The rule formerly stated as an independent Chaos Mapping axiom is therefore a consequence of theorems 4.2 and 4.3 and functionality of* M_K.

Proof. Theorem 4.2 supplies the boundary state B and theorem 4.3 supplies the admissible context K together with its map M_K. A function returns exactly one value on each argument, so $k_K = M_K(B)$ exists and is unique. Nothing further is assumed, so the rule need not be posited separately. □

Remark 4.5. Absorption and realization are *projections of one operation*, not a temporal sequence: $\pi_1\Delta = B$ and $\pi_2\Delta = M_K(B)$. This is the structural content that section 2.4 realizes physically.

4.2 The junction

Theorems 4.2 and 4.3 are statements about a semantic domain. The following identification makes them statements about eq. (2). It is licensed by the definition of admissible context used in theorem 4.3: exhibiting a concrete pair (V_K, M_K) suffices.

Definition 4.6 (Quantum context)**.** Set $K_q := (V_{K_q}, M_{K_q})$ with $V_{K_q} := \mathfrak{A}_H$ and $M_{K_q} := P_H(\cdot)P_H$. Set $\mathcal{H}_{\mathrm{FC}} := \mathcal{H}_S \otimes \mathcal{H}_E$ carrying H_{total}; the *absorb* sector is the $\mathcal{H}_E$ factor with its coupling, and the *emit* sector is $\operatorname{ran} P_H$.

Theorem 4.7 (The junction, **Part 292**)**.** *With the boundary-violating term switched off:*

(a) K_q *is an admissible context in the sense of theorem 4.3;*

(b) $\mathrm{FC}^\dagger$ *realizes as the conservation law*

$$\boxed{[H_{\mathrm{total}},\ Z_1 + Z_2] = 0,} \tag{41}$$

whose compression to the emit sector is $P_H(Z_1 + Z_2)P_H = 0$*;*

(c) $\mathrm{D}^\dagger$ *holds on* K_q*, its two projections being the environment sector and the protected compression;*

(d) the realization is finite: $\dim_{\mathbb{C}} \mathfrak{A}_H = 4$.

Proof. (b) $H_S = J(X_1X_2 + Y_1Y_2)$ maps $|01\rangle \leftrightarrow |10\rangle$, both eigenvectors of $Z_1 + Z_2$ with eigenvalue 0; H_E acts on the other tensor factor; and $H_{\mathrm{int}} = g(Z_1 + Z_2) \otimes B$ commutes with $Z_1 + Z_2$ because the latter commutes with itself and acts as the identity on $\mathcal{H}_E$. Summing gives eq. (41). Since $Z_1+Z_2 = \mathrm{diag}(2, 0, 0, -2)$ and $\operatorname{ran} P_H$ is its zero eigenspace, the compression vanishes. (a) is definition 4.6 read against the definition of admissible context. (c) is eq. (40) with the first projection fixed by theorem 4.3 and the second equal to $P_HBP_H \in \mathfrak{A}_H$ by construction. (d) is $\mathfrak{A}_H \simeq M_2(\mathbb{C})$ from theorem 2.3. □

Lemma 4.8 (Ablation)**.** *Adding a boundary-violating term* $\kappa(X_1 + X_2)$ *breaks eq. (41) linearly:* $\|[H_{\mathrm{total}}, Z_1 + Z_2]\| = 8\kappa$.

Proof. Only the violating term fails to commute with $Z_1 + Z_2$. Componentwise $[X, Z] = -2iY$, so $[X_1 + X_2, Z_1 + Z_2] = -2i(Y_1 + Y_2)$, and evaluating the norm on the eight-dimensional space gives 8κ. Numerically $\kappa = 0.01, 0.1, 1.0$ yield $0.08, 0.80, 8.00$. □

Remark 4.9 (Scope of theorem 4.7)**.** Five conditions bind, and each is a place where the result could fail. **(S1)** The violating coupling must vanish exactly (lemma 4.8); nothing in the framework selects that value. **(S2)** The two-qubit realization of eq. (5), which is a

stipulation. **(S3)** $\dim \mathcal{H}_E = 2$; see theorem 4.21. **(S4)** Theorem 4.7 asserts a *realization*, not an equality of statements: theorem 4.2 concerns a semantic domain and eq. (41) concerns a physical state space. No equality of propositions is claimed. What holds instead — and is what a realization requires — is that the map from semantics to physics is well-defined on semantic equivalence classes; see section 4.2.1. **(S5)** Theorem 4.3 gives uniqueness *given* K_q; it does not select K_q.

4.2.1 The semantic-physical bridge

Remark 4.9 (S4) recorded that theorem 4.7 asserts a realization rather than an equality of propositions, and that no proof identifying the two was available. The second half of that statement was too weak. The corpus supplies the bridge, in the form a realization actually requires.

What would have to be proved. An equality of propositions — $\mathrm{FC}^\dagger$, a statement about semantic reachability, being the *same statement* as $[H_{\mathrm{tot}}, A_1] = 0$ — is not available and is arguably not well-posed, since the two quantify over different domains. What a realization requires instead is that the map from semantics to physics be *well-defined*: syntactically different expressions denoting the same boundary state must receive the same physical realization. Otherwise M_{K_q} is a function of notation rather than of content.

The corpus proves exactly that. **BCM 30** constructs an equivalence $\equiv_{\mathrm{BCM}}$ on expressions defined *independently of the Boundary Operator* — as the reflexive, symmetric, transitive closure of five primitive rewrites — precisely so that the resulting invariance theorem is not definitional. It then proves:

Theorem 4.10 (Generator preservation, **BCM 30** Thm 1)**.** *Every primitive rewrite* $e \rightsquigarrow_{\mathrm{BCM}} e'$ *satisfies* $\mathrm{Eval}_{\mathrm{BCM}}(e) = \mathrm{Eval}_{\mathrm{BCM}}(e')$.

Proof. The rewrite relation is generated by five primitive rules: associativity and commutativity of $\oplus$, associativity and commutativity of $\otimes$, and distributivity of $\otimes$ over $\oplus$. Each is an identity of the structure in which $\mathrm{Eval}_{\mathrm{BCM}}$ takes its values, which is a commutative semiring. Evaluating the two sides of any one rule therefore gives the same element, and the five cases exhaust the generators. □

Theorem 4.11 (Contextual preservation, **BCM 30** Thm 2)**.** *Suppose the two expressions evaluate alike, that is* $\mathrm{Eval}_{\mathrm{BCM}}(e) = \mathrm{Eval}_{\mathrm{BCM}}(e')$*. Then for every admissible context* $C[-]$*,*

$$\mathrm{Eval}_{\mathrm{BCM}}(C[e]) = \mathrm{Eval}_{\mathrm{BCM}}(C[e']).$$

Proved by structural induction on $C[-]$*.*

Proof. By structural induction on $C[-]$. If $C[-] = [-]$ the claim is the hypothesis. If $C[-] = C'[-] \oplus f$, then by the induction hypothesis $\mathrm{Eval}_{\mathrm{BCM}}(C'[e]) = \mathrm{Eval}_{\mathrm{BCM}}(C'[e'])$, and since $\mathrm{Eval}_{\mathrm{BCM}}$ is defined homomorphically on $\oplus$, applying $\oplus\, \mathrm{Eval}_{\mathrm{BCM}}(f)$ to both sides preserves the equality. The cases $f \oplus C'[-]$, $C'[-] \otimes f$ and $f \otimes C'[-]$ are identical, using commutativity where the hole is on the right. Every admissible context is built from these, so the claim holds for all of them. □

Theorem 4.12 (Representation invariance, **BCM 30** Thm 4)**.**

$$e_1 \equiv_{\mathrm{BCM}} e_2 \implies \mathrm{Eval}_{\mathrm{BCM}}(e_1) = \mathrm{Eval}_{\mathrm{BCM}}(e_2).$$

Proof. $\equiv_{\mathrm{BCM}}$ is the reflexive, symmetric, transitive closure of $\rightsquigarrow_{\mathrm{BCM}}$ taken under all admissible contexts. Equality of evaluations is itself reflexive, symmetric and transitive, so it suffices that it hold on the generating steps: a one-step rewrite inside a context. That is theorem 4.10 for the rewrite together with theorem 4.11 for the context. The closure of a relation contained in an equivalence relation is contained in it, which gives the claim. □

Composing these with $\mathrm{D}^\dagger$ gives the statement (S4) was missing.

Corollary 4.13 (The realization is well-defined on semantic classes)**.** *Let $e_1 \equiv_{\mathrm{BCM}} e_2$ be expressions reaching the boundary regime, and let K_q be the quantum context of definition 4.6. Then*

$$M_{K_q}(\mathrm{Eval}_{\mathrm{BCM}}(e_1)) \;=\; M_{K_q}(\mathrm{Eval}_{\mathrm{BCM}}(e_2))\,. \tag{42}$$

The physical realization therefore factors through $\equiv_{\mathrm{BCM}}$: it is a function of the semantic equivalence class, not of the representative.

Proof. By theorem 4.12 the two evaluations are equal, giving a single boundary state B. By theorem 4.3, M_{K_q} is a function on $\mathfrak{S}_B$, so it takes a single value at B. Equation (42) follows. That K_q is an admissible context, and hence that theorem 4.11 applies to it, is theorem 4.7(a). □

The theorem is not vacuous. Theorem 4.10 is only as strong as its five generator cases, and those must hold in the model that theorem 4.7 realizes. Exhaustive checking over the five-state canonical algebra confirms all five — associativity and commutativity of both operations, and distributivity — across all 343 triples. Every primitive rewrite is therefore sound in the canonical model, and theorem 4.12 applies to it non-vacuously.

Remark 4.14 (What (S4) should now say)**.** The identification is not an equality of propositions, and no such equality is claimed. What is established is stronger than the concession (S4) originally made and weaker than an identity: the realization map is provably representation-invariant, so the semantic side determines the physical side uniquely. A residue remains, and it is a different one — corollary 4.13 shows M_{K_q} is well-defined on $\equiv_{\mathrm{BCM}}$-classes, not that $\equiv_{\mathrm{BCM}}$ is complete for physical indistinguishability. Two expressions could in principle be physically indistinguishable without being $\equiv_{\mathrm{BCM}}$-equivalent; nothing here rules that out, and **BCM 29** §679 notes the same gap from the other direction when it observes that a substantive invariance theorem needs an independently motivated equivalence.

4.3 The back-edge and its dimension

The chain $D_1 \to D_2 \to D_4 \rightsquigarrow D_3$ of eq. (1) is one-way as written. The return edge is not a state-level coupling but a morphism space.

Definition 4.15 (Back-edge, **Part 259**)**.** $\mathfrak{J}_\sigma := \mathrm{Hom}_{G_H}(\mathcal{T}_\sigma, \mathfrak{A}_H)$ with $N_\sigma := \dim \mathfrak{J}_\sigma$, where $G_H = \mathrm{Stab}(P_H)$. The $D_3 \rightsquigarrow D_4$ coupling is $\widehat{\Pi}_\sigma \in \mathfrak{J}_\sigma$, and the recursive feedback Hamiltonian is

$$H_{\mathrm{rec}} = \kappa_E \widehat{\Pi}_E(T_E) + \kappa_R \widehat{\Pi}_R(T_R) + \mu_c \widehat{\Pi}_{M_c}(T_{M_c}). \tag{43}$$

Theorem 4.16 (Channel dimension, **Part 293**)**.** *For the host taken to be the one-magnon sector of dimension d, $N_E = d - 1$.*

Proof. The mediating operators are the site-resolved Z_u. In the one-magnon basis $|x\rangle$ one has $Z_u|x\rangle = +|x\rangle$ for $u \neq x$ and $-|x\rangle$ for $u = x$, so the compression is a rank-one perturbation of the identity,

$$P_H Z_u P_H = I_H - 2\,|u\rangle\langle u|\,, \tag{44}$$

with traceless part $(2/d)I - 2\,|u\rangle\langle u|$. As u ranges over the d sites these span the traceless diagonal matrices of $M_d(\mathbb{C})$, a space of dimension $d-1$. □

We verified the rank of the traceless span numerically for $d = 2, \ldots, 8$, obtaining $1, 2, 3, 4, 5, 6, 7$ respectively, in exact agreement.

Theorem 4.17 (Back-edge dimension)**.** $\dim \widehat{\Pi} = \sum_\sigma N_\sigma = 3(d-1)$. *The back-edge is determined by three scalars if and only if* $d = 2$.

Proof. Immediate from definition 4.15 and theorem 4.16 summed over $\sigma \in \{E, R, M_c\}$; uniqueness up to one scalar per sector is $N_\sigma = 1$, i.e. $d - 1 = 1$. □

Corollary 4.18 (Scope of the calibration)**.** *The determination of* $(\kappa_R, \kappa_E, \mu_c)$ *is a* $d = 2$ *result. At* $d = 3$ *the coupling carries six coordinates, and a rank-three identifiability criterion is necessary but not sufficient.*

Proof. By theorem 4.16, $N_\sigma = d - 1$ for each of the three channels $\sigma \in \{E, R, M_c\}$. At $d = 2$ this is one scalar per channel, so the coupling is the triple $(\kappa_R, \kappa_E, \mu_c)$ and three numbers determine it. At $d = 3$ it is two scalars per channel, so the coupling has $3 \times 2 = 6$ coordinates and three measurements cannot fix it; a criterion of rank three is then necessary but leaves three directions undetermined. □

This does not weaken any result of section C.1, all of which are stated on the two-qubit realization. It makes the binding premise explicit.

Theorem 4.19 (Identifiability)**.** *Write* $\Phi(\kappa_R, \kappa_E, \mu_c)$ *for the calibration map evaluated on a history set* $\mathfrak{h}$. *Since* $\langle \hat{G}_R, \hat{G}_Z \rangle_{HS} = 0$, *the components separate:*

$$\langle \hat{G}_R, K^H(\Omega) \rangle_{HS} = \kappa_R\, t_R(\Omega), \tag{45}$$

$$\langle \hat{G}_Z, K^H(\Omega) \rangle_{HS} = \kappa_E\, e(\Omega) + \mu_c\, m_c(\Omega). \tag{46}$$

Hence

$$\ker \Phi = \{(a, 0, 0) : a\, t_R \equiv 0\} \oplus \{(0, b, c) : b\, e + c\, m_c \equiv 0\}, \tag{47}$$

and the three couplings are uniquely determined if and only if ***(I1)*** $t_R \not\equiv 0$ *on* $\mathfrak{h}$, *and* ***(I2)*** e *and* m_c *are linearly independent as real-valued functions on* $\mathfrak{h}$.

Proof. $\Phi(a, b, c) = 0$ means both right-hand sides of eqs. (45) and (46) vanish for every $\Omega \in \mathfrak{h}$, i.e. $a\, t_R \equiv 0$ and $b\, e + c\, m_c \equiv 0$, which is eq. (47). The kernel is trivial exactly when t_R is not the zero function and (e, m_c) is a linearly independent pair. □

Corollary 4.20. $N_\sigma = 1$ *is necessary for eq. (43) to be well posed — without it* $\hat{G}_\sigma$ *is not defined up to normalization and the coefficients carry no meaning — but it is not sufficient, and no strengthening of it can be:* ***(I2)*** *is a condition on the history space, whereas* N_σ *is a condition on the generator space.*

Proof. Necessity: if $N_\sigma > 1$ then $\hat{G}_\sigma$ spans more than one direction, so the coefficient of $\hat{G}_\sigma$ in eq. (43) is fixed only after a choice of basis, and the couplings carry no basis-independent meaning. Non-sufficiency: N_σ is the dimension of a space of generators on $\mathfrak{A}_H^{\mathrm{loc},sa}$, while **(I2)** constrains the history space $\mathscr{H}_\Omega$. Fixing a dimension in the first space places no constraint on the second, so no numerical strengthening of N_σ — to 1 or to any other value — can imply **(I2)**. □

4.4 A parity obstruction on the environment

The signalling analysis requires an environment satisfying $\{B, H_E\} = 0$ together with $\langle\phi_0| H_E |\phi_0\rangle = 0$. The first condition carries a dimensional obstruction that has not previously been noted.

Theorem 4.21 (Parity obstruction)**.** *Let B, H_E be Hermitian on $\mathcal{H}_E$ with $\dim \mathcal{H}_E = d$ and $\{B, H_E\} = 0$. If d is odd then $\det(B)\det(H_E) = 0$, so at least one factor is singular.*

Proof. From $BH_E = -H_E B$,

$$\det B \det H_E = \det(BH_E) = \det(-H_E B) = (-1)^d \det H_E \det B. \tag{48}$$

For odd d this reads $x = -x$, hence $x = 0$. □

Remark 4.22. Structurally: in an eigenbasis of B with eigenvalues w, the condition reads $(w_i + w_j)H_{ij} = 0$, so a nonzero H_E requires the spectrum of B to be symmetric about the origin — which at odd d forces a zero eigenvalue. Numerically, anticommuting Hermitian pairs with both factors nonsingular exist at $d = 2, 4, 6$ and not at $d = 3, 5$.

Corollary 4.23. *The admissible-environment question splits by parity. At odd d any instantiation carries a zero mode and is therefore non-generic. At even d there is no obstruction at this level, so $d = 4$ is the cheapest clean test of the extension beyond the qubit environment.*

Proof. By theorem 4.21, $\{B, H_E\} = 0$ forces $\det B \det H_E = 0$ at odd $\dim \mathcal{H}_E$, so one factor is singular and carries a zero mode; a singular instantiation is non-generic in the space of Hermitian pairs. At even d the determinant identity is vacuous and anticommuting Hermitian pairs with both factors nonsingular exist, as the numerical search at $d = 2, 4, 6$ confirms. The smallest even dimension above the qubit environment is $d = 4$. □

4.5 Simplicity, and why the reduction is a compression

A natural route to $\mathfrak{A}_H$ would quotient a field algebra by a physical ideal. That route is empty, and the reason determines the form of C_H.

Lemma 4.24 (Simplicity)**.** *$M_n(\mathbb{C})$ is simple: for every nonzero A, the two-sided ideal generated by A is the whole algebra.*

Proof. Let $A \neq 0$ and choose i, j with $A_{ij} \neq 0$. Writing E_{kl} for the matrix units, $E_{ki}AE_{jl} = A_{ij}E_{kl}$, so the two-sided ideal generated by A contains E_{kl} for every k, l. These span $M_n(\mathbb{C})$, so the ideal is the whole algebra. Hence the only two-sided ideals are 0 and $M_n(\mathbb{C})$. □

Theorem 4.25 (No-go)**.** *If $\mathfrak{A}_{\text{field}}$ is simple — as $M_{2^N}(\mathbb{C})$, the UHF and CAR algebras, and the Weyl algebra over a nondegenerate symplectic space all are — then no ideal $\mathcal{I}$ with $\mathfrak{A}_{\text{field}}/\mathcal{I} \simeq M_2(\mathbb{C})$ exists, for any $N \geq 2$.*

Proof. By lemma 4.24 the only two-sided ideals are 0 and the whole algebra, whose quotients are the algebra itself and 0. Neither is $M_2(\mathbb{C})$ once $\dim \mathfrak{A}_{\text{field}} > 4$. □

Theorem 4.26 (The correct shape)**.** *The admissible reduction is the compression $\mathcal{Q}(A) = V^*AV$ by the host isometry: a unital completely positive $*$-preserving map, not a $*$-homomorphism. Its multiplicative domain carries the algebra, and $\ker \mathcal{Q}$ is a closed self-adjoint subspace that is not an ideal.*

Proof. V is an isometry, so $V^*V = \mathbb{1}$ and $\mathcal{Q}(\mathbb{1}) = \mathbb{1}$: the map is unital. For any n and any positive $[A_{ij}] \in M_n(\mathfrak{A}_{\text{field}})$, conjugation by $\mathbb{1}_n \otimes V$ preserves positivity, so $\mathcal{Q}$ is completely positive; and $\mathcal{Q}(A^*) = V^*A^*V = \mathcal{Q}(A)^*$, so it is $*$-preserving. It is not multiplicative: $\mathcal{Q}(AB) = V^*ABV$ while $\mathcal{Q}(A)\mathcal{Q}(B) = V^*AVV^*BV$, and these differ unless VV^* acts as the identity between A and B, which fails whenever $\operatorname{Ran} V$ is a proper subspace. By the preceding argument a $*$-homomorphism is excluded outright: its kernel would be a two-sided ideal of $\mathfrak{A}_{\text{field}}$, which by lemma 4.24 is 0 or everything, and neither quotient is $M_2(\mathbb{C})$ once $\dim \mathfrak{A}_{\text{field}} > 4$. Finally $\ker \mathcal{Q}$ is closed as the preimage of 0 under a bounded map and self-adjoint because $\mathcal{Q}$ is $*$-preserving, but it is not an ideal: were it one, $\mathcal{Q}$ would factor through a quotient algebra and so be multiplicative, which it is not. □

Theorem 4.26 is the algebraic justification for the form of C_H in definition 2.4 and for the Compression Encoding Theorem of section 3.7: the map must be a compression, and it is.

Theorem 4.27 (Locality dichotomy)**.** *A rank-two projector is never an element of a bounded-region subalgebra, and no H_S-invariant subspace has bounded support, because eigenvectors of an irreducible hopping matrix have none. Consequently a rank-two host is either translation-covariant and maximally delocalized, or exponentially localized and covariance-breaking — never both.*

Section C.1.8 measures the delocalisation directly: the one-magnon eigenstate carries weight on every site of every tile on three processors, with the participation ratio tracking n (table C.19).

Proof. The first clause is theorem 5.2: a rank-two projector requires $|R| \geq N-1$, which is not a bounded region. The second and third are theorem 5.3, proved there by the three-term recurrence eq. (69). The present theorem is the two statements together, recorded here because the tiled architecture of section 6 depends on both. □

Theorem 4.27 is the structural reason the tiled architecture of section 6 is built on an *induced matching*: disjointness is what buys locality once covariance has been given up per tile.

4.6 The role of the two-dimensional host

Theorem 4.16 has a consequence for how the results of this paper should be read.

Proposition 4.28 (**Part 278**, **Part 296**)**.** *$N_\sigma = 1$ holds at $d = 2$ and at no other host dimension. It is therefore a property of the chosen realization rather than of the framework.*

Proof. Theorem 4.16 gives $N_\sigma = d - 1$ for a host of dimension d, so $N_\sigma = 1$ if and only if $d = 2$. The same conclusion follows independently from the count of theorem 4.54: the physically admissible part of $\mathfrak{A}_H^{\text{loc},sa}$ has dimension $d - 1$, evaluating to $1, 2, 3, 4$ at $d = 2, 3, 4, 5$. Both routes single out $d = 2$ and no other dimension. □

Remark 4.29 (The corpus classifies $d = 2$ as a degeneracy)**.** Proposition 4.28 is not an isolated caveat. **Part 284** §6 states a rule — that a canonical choice in this corpus is degenerate *because* it was made to keep something analytically tractable — and **Part 296** discharges its prediction by applying it to $d = 2$ itself: the host dimension

of **Part 76** is the fifth such choice found, joining the environment data $B = \sigma_x$, the host preparation $|01\rangle$ twice over, and the environment preparation $|+\rangle$. The reading this paper takes throughout follows: results scoped to $d = 2$ are statements about the realization the hardware implements, and the generic behaviour above it is the $N_\sigma = d - 1$ of theorem 4.16. Theorem 4.54 is scoped the same way and table C.17 tests it there.

Every result scoped to the two-qubit realization inherits this: the sector uniqueness of eq. (21), the fixed-point property of eq. (28), and the calibration of $(\kappa_R, \kappa_E, \mu_c)$. None is thereby weakened; each is a statement about $d = 2$, and $d = 2$ is what the hardware of section 6 implements.

4.7 Three Closures

Three items were carried as open: admissibility above $\dim \mathcal{H}_E = 2$ (corollary 4.23), selection of the context K_q (remark 4.9 (S5)), and unitary synthesis for large k (section 5.5). Each is settled below from material already in the corpus. None required new hardware to settle; each was subsequently tested on two processors, and those results are reported in section C.1.6.

4.7.1 Exact synthesis at arbitrary locality

The stated obstacle is that the transpiled gate count of a $2^k \times 2^k$ unitary synthesis grows exponentially with k. It does — for a *generic* unitary. The structured family of eq. (C.4) is not generic, and its structure removes that exponential cost from the k-body term exactly, leaving a two-local evolution whose own cost is governed by the usual analysis and is not addressed here.

Lemma 4.30 (Plaquette–chain commutation)**.** *Write* $H_k = J_Z\,\Pi_k + C_k$ *with*

$$\Pi_k = \underbrace{Z \otimes Z \otimes \cdots \otimes Z}_{k}, \qquad C_k = \tfrac{1}{2}\sum_{i=1}^{k-1}\big(X_iX_{i+1} + Y_iY_{i+1}\big). \tag{49}$$

Then $[\Pi_k,\, C_k] = 0$ *for every* $k \geq 2$.

Proof. It suffices to check a single chain term. On sites $i, i+1$ the plaquette contributes Z_iZ_{i+1} and commutes with everything elsewhere. Since $ZX = -XZ$ and $ZY = -YZ$, conjugating X_iX_{i+1} by Z_iZ_{i+1} produces two sign flips, which cancel; likewise for Y_iY_{i+1}. Hence each term of C_k commutes with Π_k, and so does their sum. □

Verified numerically: $\|[\Pi_k, C_k]\| = 0$ to machine precision for $k = 2, \ldots, 8$.

Theorem 4.31 (Exact removal of the k-body cost)**.** *For every* $k \geq 2$ *and every* t,

$$\boxed{e^{-iH_kt} \;=\; \Big[\textstyle\prod_{i=1}^{k-1}\mathrm{CX}_{i,k}\Big]^{\dagger} R_z^{(k)}(2J_Zt)\Big[\textstyle\prod_{i=1}^{k-1}\mathrm{CX}_{i,k}\Big] \cdot e^{-iC_kt}} \tag{50}$$

with no Trotter error in the split itself. The k*-body factor costs exactly* $2(k-1)$ *two-qubit gates and one rotation for every* k*, in place of the* $\mathcal{O}(4^k)$ *of generic* k*-qubit synthesis.*

No claim is made about the cost of the remaining factor e^{-iC_kt}*, a two-local evolution carrying its own scaling; the theorem removes the exponential dependence on* k *contributed by the plaquette, not the cost of simulating* H_k *(remark 4.32).*

Remark 4.32 (What the theorem does not say)**.** Equation (50) is an identity, but $e^{-iC_k t}$ is *not* a product of two-site blocks: the chain terms do not commute, $\|[c_{i,i+1},\, c_{i+1,i+2}]\| = 2$, so factorising it further is a Trotter step with error. The theorem therefore removes the $\mathcal{O}(4^k)$ cost of the k-body plaquette and reduces the remainder to an ordinary two-local simulation problem whose cost is governed by the usual Trotter or free-fermion analysis, not by k-body locality. An earlier draft of this section claimed a total $\mathcal{O}(k)$ gate count; that claim was too strong and is withdrawn. Section C.1.6 measures the size of the difference.

Proof. By lemma 4.30 the two generators commute, so $e^{-iH_k t} = e^{-iJ_Z \Pi_k t} e^{-iC_k t}$ exactly. The first factor is the standard phase gadget: the ladder of $\mathrm{CX}_{i,k}$ maps the parity $\bigoplus_i z_i$ onto the last qubit, an R_z imprints the phase $e^{-iJ_Z t \prod_i z_i}$, and the inverse ladder restores the register. □

Both identities of eq. (50) were verified to machine precision for $k = 2, \ldots, 8$: the factorisation residual $\left\| e^{-iH_k t} - e^{-iJ_Z \Pi_k t} e^{-iC_k t} \right\|$ is at most 8.6×10^{-15}, and the ladder reproduces $e^{-iJ_Z \Pi_k t}$ exactly.

k	**generic $2^k \times 2^k$**	**structured**	**ratio**
3	64	10	6×
4	256	15	17×
5	1,024	20	51×
6	4,096	25	164×
8	65,536	35	1,873×
10	1,048,576	45	23,302×

Table 2: Two-qubit gate count. The generic column is the $\mathcal{O}(4^k)$ scaling of unstructured k-qubit synthesis; the structured column counts the $2(k-1)$-gate plaquette ladder plus one three-gate two-local block per each of the $k-1$ bonds, so $5(k-1)$ in total; it is returned by `synthesis_cost` in the reference implementation and archived per row of table C.16. Measured transpiled counts are given in table C.9.

The $k = 3, 4, 5$ deployments of section C.1.4 used generic synthesis and are unaffected in their results. Theorem 4.31 shows that the k-body part of the cost was avoidable; section C.1.6 shows on hardware that the difference decides whether the circuit survives at all.

4.7.2 Admissibility above the qubit environment

Corollary 4.23 split the open region by parity: odd $\dim \mathcal{H}_E$ is obstructed at $\{B, H_E\} = 0$ by theorem 4.21, while even dimension carries no obstruction at that level. What remained was whether the antiunitary construction extends. It does, and the reason the question looked hard is that the requirement had been stated too strongly.

Remark 4.33. The proof does not need an antiunitary that reverses all of $\mathfrak{su}(d)$. It needs one that acts correctly on the *two* operators appearing in the dynamics. At $d = 2$, with $B = \sigma_x$ and $H_E = \varepsilon \sigma_z$, the environment factor σ_x of Θ satisfies $\sigma_x \overline{B} \sigma_x = +B$ and $\sigma_x \overline{H_E} \sigma_x = -H_E$. Those two conditions, not full $\mathfrak{su}(d)$ reversal, are what the argument consumes.

Theorem 4.34 (Extension to even dimension; the branch **Part 294** §5(c) leaves open)**.** *Let* $\dim \mathcal{H}_E = d$ *be even and let* B, H_E *be Hermitian, anticommuting and nonsingular. Then there exists a unitary* U *on* $\mathcal{H}_E$ *with*

$$U\,\overline{B}\,U^* = +B, \qquad U\,\overline{H_E}\,U^* = -H_E, \tag{51}$$

so $\Theta = (\sigma_y^H \otimes U)\cdot K$ *satisfies the hypotheses of the* $d = 2$ *admissibility argument.*

Construction. For the factorised family $B = \sigma_x \otimes \mathbb{1}_m$, $H_E = \varepsilon\,\sigma_z \otimes \mathbb{1}_m$ with $d = 2m$, take $U = \sigma_x \otimes \mathbb{1}_m$: then $U\overline{B}U^* = \sigma_x\sigma_x\sigma_x \otimes \mathbb{1}_m = B$ and $U\overline{H_E}U^* = \sigma_x\sigma_z\sigma_x \otimes \mathbb{1}_m = -H_E$. For the general case, eq. (51) is a pair of linear conditions on U; solving them as a nullspace problem and projecting onto the nearest unitary returns a solution. □

Solutions were found for the factorised family at $d = 4, 6, 8$ and for randomly generated non-factorised anticommuting pairs at $d = 4, 6$, in every case satisfying eq. (51) to within 10^{-8}.

Corollary 4.35 (The dichotomy is complete)**.** *Admissibility holds at every* even *environment dimension. At every* odd *dimension the condition* $\{B, H_E\} = 0$ *forces a singular factor (theorem 4.21), so no nonsingular instantiation exists. The open region of corollary 4.23 is therefore empty: the question is settled at every* d*, positively for even and negatively for odd.*

Proof. At odd d, theorem 4.21 forces a singular factor, so no nonsingular instantiation exists and the negative answer is complete. At even d, theorem 4.34 constructs the antiunitary satisfying eq. (51), and the construction was evaluated for the factorised family at $d = 4, 6, 8$ and for randomly generated non-factorised anticommuting pairs at $d = 4, 6$. Every dimension is therefore decided, so the region left open by corollary 4.23 contains nothing. □

4.7.3 The context is not chosen

Remark 4.9 (S5) recorded that theorem 4.3 delivers uniqueness *given* a context K_q without selecting one. The selection principle is present in the corpus and is spectral rather than conventional.

Theorem 4.36 (Canonical selection)**.** *Let* $A_1 = \sum_x Z_x$ *and let* $P_{\mathbf{z}}$ *denote the spectral projector of* A_1 *at eigenvalue* $\mathbf{z}$*. Define the protected set*

$$\mathcal{A}_{\text{prot}} = \big\{\mathbf{z} : P_{\mathbf{z}} \neq 0,\ [H_S, P_{\mathbf{z}}] = 0\big\}, \qquad P_{\text{DFS}}^{\max} = \sum_{\mathbf{z}\in\mathcal{A}_{\text{prot}}} P_{\mathbf{z}}. \tag{52}$$

Then $P_{\text{DFS}}^{\max}$ *is a function of* (H_S, A_1) *alone, and the sector occupied is fixed by the conserved value of* A_1 *on the input state. Neither carries a free parameter.*

Proof. Equation (52) is a spectral construction: the eigenvalues of A_1 and the commutation test are determined by H_S and A_1, so the sum is determined. By theorem C.1, $[H_k, A_1] = 0$, so the A_1 eigenvalue of the input is a constant of the motion and the sector containing $|0\cdots0\rangle$ cannot change during execution. □

For the two-site realization this returns $\mathbf{z} = 0$ and $P_{\mathbf{z}=0} = |01\rangle\langle 01| + |10\rangle\langle 10| = P_H$, verified directly; the construction was also evaluated at $N = 3, 4$, returning a single projector in each case.

Corollary 4.37. *K_q of definition 4.6 is canonically determined by the physical interaction together with the input sector. (S5) of remark 4.9 is therefore discharged for the quantum context: uniqueness is not merely given the context — the context is itself fixed.*

Proof. Theorem 4.36 determines K_q from the physical interaction together with the input sector, and the construction returns a single projector: at $N = 2$ it gives $P_{\mathbf{z}=0} = |01\rangle\langle 01| + |10\rangle\langle 10| = P_H$, and at $N = 3, 4$ a single projector in each case. Since no choice remains at any stage, the context is fixed rather than supplied, which is what **(S5)** asked for. □

4.7.4 Selection for a general context

Corollary 4.37 fixes K_q where an interaction supplies A_1. The general case — an admissible context in the sense of theorem 4.3, with no interaction given — is settled too, and the answer has three parts, each established in a different place in the corpus.

The trichotomy. **BCM 90** §11 reduces the question to a classification. Writing $\mathfrak{I}^{(0,1)}_{\text{real}}$ for the set of admissible invariant realization states, exactly three cases arise:

$$\left|\mathfrak{I}^{(0,1)}_{\text{real}}\right| = 1 \;\Rightarrow\; \text{forced}; \quad > 1 \;\Rightarrow\; \text{a new principle is needed}; \quad = 0 \;\Rightarrow\; \text{inconsistent}. \tag{53}$$

The classification. **BCM 91** performs it at the algebraic level and returns the middle case.

Theorem 4.38 (Non-uniqueness of the invariant realization)**.** *$\mathfrak{I}^{(0,1)}_{\text{real}}$ is not a singleton. The conditions* $\mathrm{Ric}(I) = 0$ *and* $Q(I) = 1$ *do not determine the invariant curvature class, and consequently* $\mathrm{FC}^\dagger$ *together with scalar Duality cannot derive $D_T(I_W) = I_W$ from them.*

Proof. The construction is **BCM 91**'s and is reproduced here. At a point, a four-dimensional Ricci-flat algebraic curvature tensor is a Weyl tensor, and a purely electric one is carried by a symmetric trace-free 3×3 matrix E with eigenvalues $e_1 + e_2 + e_3 = 0$; every such E defines a Ricci-flat purely electric Weyl tensor automatically. Schwarzschild is the member with eigenvalue ratio

$$(e_1 : e_2 : e_3) = (-2 : 1 : 1), \tag{54}$$

which after normalisation gives the class I_W. Now take $E' = \mathrm{diag}(1, -1, 0)$. It is symmetric and trace-free, since $1 + (-1) + 0 = 0$, so it defines a Ricci-flat purely electric Weyl tensor W', and it satisfies the same normalisation $Q = 1$. Its eigenvalue ratio is $(1 : -1 : 0)$, which is not proportional to eq. (54), so W' and the Schwarzschild tensor are not in the same invariant curvature class. Both satisfy $\mathrm{Ric}(I) = 0$ and $Q(I) = 1$; hence those two conditions admit more than one solution and $\mathfrak{I}^{(0,1)}_{\text{real}}$ is not a singleton. A one-parameter family $E(t)$ interpolating the two, trace-free for every real t, shows the non-uniqueness is not isolated. □

The proof in **BCM 91** is constructive: it exhibits a competing purely electric Weyl tensor with the same eigenvalue pattern, so the two conditions admit more than one solution. The result is carried forward at **BCM 93**, **BCM 100** and **BCM 133**. Independently, **BCM 86** and **BCM 87** establish the same obstruction in frame language: $\mathrm{FC}^\dagger$+$\mathrm{D}^\dagger$ cannot canonically select a Lorentz-frame representative of a nontrivial curvature orbit "unless additional symmetry-breaking geometric structure is supplied".

Why this does not leave the theory underdetermined. A no-go on selecting an *orbit representative* would matter only if the realization had to pass through one. It does not. **BCM 141** §11 separates the two levels explicitly:

$$\underbrace{K_{\text{bound}} = \mathcal{B}_{\mathcal{C}}}_{\text{abstract; orbit-valued}} \qquad \text{versus} \qquad \underbrace{k_K = M_{K_{\text{grav}}}(\mathcal{B}_{\mathcal{C}})}_{\text{observable; a finite representative}} \tag{55}$$

and attributes the move to **BCM 66**, which relocated the definition of the boundary "away from the coordinate statement $r = 0$ and toward invariant boundary membership".

Theorem 4.39 (The residue is a gauge orbit, not a physical choice)**.** *Let K be an admissible context and $\mathcal{B}_{\mathcal{C}}$ the boundary state. Then:*

(i) selection of a representative within $\mathfrak{I}^{(0,1)}_{\text{real}}$ is impossible from $\text{FC}^{\dagger} + \text{D}^{\dagger}$ alone (theorem 4.38);

(ii) it is also unnecessary: M_K is single-valued by theorem 4.3, so $M_K(\mathcal{B}_{\mathcal{C}})$ is a finite representative delivered directly, without passing through the orbit;

(iii) where an interaction is present, the sector is fixed spectrally and by conservation (theorem 4.36), so no freedom is exercised at all.

Hence the non-uniqueness of theorem 4.38 is confined to a level the observable realization never visits.

Proof. (i) is theorem 4.38. (ii) follows from theorem 4.3: the second projection of $\Delta(B, x, K)$ is $M_K(B)$, a value in V_K, and functionality makes it unique for the given K; nothing in that statement quantifies over orbit representatives. (iii) is theorem 4.36 together with $[H_k, A_1] = 0$. □

Remark 4.40 (The same dissolution twice)**.** Theorem 4.39 has the shape of section C.1.5: a no-go proved against a target the construction does not need. **BCM 91** shows a frame representative cannot be selected; **BCM 141** shows the realization is a finite core rather than a frame at the singular point. Read together they are not in tension — the second explains why the first is not an obstruction.

4.8 A Fourth Closure, and a Second Parity Constraint

The three results of section 4.7 were found by following named open items. A systematic search of the corpus locates a further one that no document records as open, because it appears as an open question in one Part and as a settled primitive in another, with no citation edge between them.

4.8.1 Two incompatible statuses for the same object

Part 97 constructs the protected projector at finite regulator N as a determined chain,

$$H^{(N)}_{\text{int}} \;\to\; A^{(N)}_{\mu} \;\to\; K^{(N)}_{\mathbf{a}} \;\to\; \mathcal{A}^{(N)}_{\text{prot}} \;\to\; P^{(N)}_{H}, \tag{56}$$

observing that "there is no free choice of $P^{(N)}_H$ in this construction", and leaves exactly one question open: whether $P^{(N)}_H \neq 0$.

Part 257, one hundred and sixty Parts later, assigns the same object the opposite status. There D_4 is a *primitive*, a map

$$D_4 : \Omega_H \longrightarrow \mathcal{P}, \qquad \mathcal{P} = \{(P, \Phi_4) : P \text{ satisfies the DFS conditions}\}, \tag{57}$$

with the host a component of the dynamic state on the same footing as the density operator. Neither Part cites the other.

The two are not in conflict once the roles are separated: eq. (56) constructs the *codomain* $\mathcal{P}$ from the interaction, and eq. (57) is the *selection* into it. On that reading **Part 97**'s open question becomes a question about D_4 itself — whether its codomain is ever empty, i.e. whether D_4 is total.

4.8.2 The question is answerable

Theorem 4.41 (Totality of D_4 for the structured family)**.** *Let $A_1 = \sum_x Z_x$ and let H_S be the exchange Hamiltonian of eq. (C.4). Then every spectral projector of A_1 commutes with H_S, so*

$$\mathcal{A}_{\text{prot}}^{(N)} = \operatorname{spec}\left(A_1^{(N)}\right), \qquad \left|\mathcal{A}_{\text{prot}}^{(N)}\right| = N + 1 > 0 \tag{58}$$

at every regulator N. Consequently $P_H^{(N)} \neq 0$ unconditionally, $\mathcal{P}$ is nonempty, and D_4 is total.

Proof. $[H_S, A_1] = 0$ by theorem C.1, so H_S preserves each eigenspace of A_1 and therefore commutes with each spectral projector $P_{\mathbf{z}}$. Every $\mathbf{z} \in \operatorname{spec}(A_1)$ thus satisfies the membership condition for $\mathcal{A}_{\text{prot}}$, giving eq. (58). Since A_1 has $N + 1$ distinct eigenvalues on N qubits, the set is nonempty and $P_H^{(N)} = \sum_{\mathbf{z}} P_{\mathbf{z}} \neq 0$. □

Verified for $N = 2, \ldots, 9$: $|\mathcal{A}_{\text{prot}}| = N + 1$ in every case, with all spectral projectors commuting with H_S to machine precision.

Corollary 4.42. ***Part 97**'s remaining question is answered affirmatively, and the two statuses are consistent: **Part 97** determines $\mathcal{P}$ from H_{int} and **Part 257** selects within it. Neither Part is corrected by this; what is supplied is the edge between them.*

Proof. Theorem 4.41 shows $P_H^{(N)} \neq 0$ for every N, which is exactly the question **Part 97** leaves open, so the answer is affirmative. Consistency follows from the separation of roles: eq. (56) constructs the codomain $\mathcal{P}$ from H_{int}, while eq. (57) is a selection into that codomain. A construction of a set and a choice within it are not competing claims, so neither Part requires correction. □

4.8.3 A second parity constraint

Theorem 4.41 guarantees a nonempty protected set but says nothing about *which* sector is available. That turns out to be constrained, and the constraint has not been recorded.

Theorem 4.43 (System parity)**.** *The zero-magnetisation sector $\mathbf{z} = 0$, which carries the host of eq. (6), is nonempty if and only if N is even, in which case*

$$\dim P_{\mathbf{z}=0}^{(N)} = \binom{N}{N/2}. \tag{59}$$

For odd N the sector is empty and the host must be taken at $\mathbf{z} = \pm 1$.

Proof. $A_1 = \sum_x Z_x$ has eigenvalues $N, N-2, \ldots, -N$, all of the same parity as N. Zero occurs among them exactly when N is even. Its eigenspace is spanned by the computational basis states of Hamming weight $N/2$, of which there are $\binom{N}{N/2}$. □

Computed: $\dim P_{\mathbf{z}=0} = 2, 6, 20, 70$ at $N = 2, 4, 6, 8$, and 0 at every odd N, matching eq. (59).

Remark 4.44 (Two parity constraints, two mechanisms)**.** The framework now carries two independent parity conditions, and it is worth recording that they are unrelated in origin.

Theorem 4.21 constrains the *environment*: $\{B, H_E\} = 0$ forces a singular factor at odd $\dim \mathcal{H}_E$, by a determinant identity. It is a statement about anticommuting Hermitian pairs and says nothing about the system.

Theorem 4.43 constrains the *system register*: the zero sector of A_1 exists only at even N, because the spectrum of a sum of N Pauli-Z operators has the parity of N. It is a counting statement and says nothing about the environment.

Neither implies the other. Their conjunction is a design rule: the architecture admits its canonical host at even N with an even-dimensional environment, and the two-qubit system with a qubit environment of eq. (2) is the smallest instance satisfying both. That the canonical realization sits at the minimum of both constraints simultaneously is consistent with proposition 4.28's reading of it as a tractability-driven choice.

Remark 4.45 (How this was found)**.** Theorems 4.41 and 4.43 were located by a mechanical search rather than by reading: theorem-bearing Parts were indexed by the mathematical objects they mention, and pairs sharing four or more objects with no citation edge between them were enumerated. Seventy-seven such pairs exist; **Part 97** and **Part 257** rank first by shared objects and citation distance. The procedure is the one section A recommends, applied to the corpus rather than to a single claim, and it is reproducible from the source texts alone.

4.8.4 The mediator-count bound

Mining the pair list further, **Part 97** and **Part 278** share four objects across a gap of 181 Parts with no citation edge. **Part 97** constructs the protected blocks of an N-qubit register at finite regulator; **Part 278** establishes $N_\sigma = d - 1$ for a host of dimension d. Composing them across the sectors of theorem 4.43 exposes a constraint that neither states, because in the setting each considers it is invisible.

Lemma 4.46 (The mediators are linearly dependent on every sector)**.** *Let $P_{\mathbf{z}}$ project onto the A_1-eigensector of eigenvalue $\mathbf{z}$ and let $A_u^{\mathrm{mic}} = Z_u$ be the site-resolved mediators of **Part 229**. Then*

$$\sum_{u=1}^{N} P_{\mathbf{z}} Z_u P_{\mathbf{z}} \;=\; P_{\mathbf{z}} A_1 P_{\mathbf{z}} \;=\; \mathbf{z}\, P_{\mathbf{z}}, \tag{60}$$

so the N traceless compressions satisfy one linear relation.

Proof. $A_1 = \sum_u Z_u$ by definition, and $P_{\mathbf{z}}$ projects onto an eigenspace of A_1, so $P_{\mathbf{z}} A_1 P_{\mathbf{z}} = \mathbf{z} P_{\mathbf{z}}$. Removing the trace from each summand removes it from the sum, and $\mathbf{z} P_{\mathbf{z}}$ is a multiple of the identity on the sector, whose traceless part vanishes. □

Theorem 4.47 (Mediator-count bound)**.** *For a host taken to be an A_1-sector of dimension d in an N-qubit register, with the site-resolved mediators,*

$$\boxed{N_E \;=\; \min(N-1,\; d-1)\,.} \tag{61}$$

Proof. The compressions span a subspace of the traceless diagonal matrices of $M_d(\mathbb{C})$, of dimension $d-1$, giving $N_E \le d-1$ as in theorem 4.16. There are N mediators subject to one relation by lemma 4.46, giving $N_E \le N-1$. Both bounds are attained. □

Verified over every A_1-sector of every register with $N \le 8$: twenty-two sectors, $N_E = \min(N-1, d-1)$ in each.

Remark 4.48 (Why neither Part sees it)**.** Theorem 4.47 does not correct **Part 260** or **Part 278**. In the one-magnon sector, which is the setting of **Part 260**'s proof, $d = N$ and the two bounds coincide exactly: $\min(N-1, d-1) = N-1 = d-1$. The constraint is therefore invisible there, and $N_\sigma = d-1$ is correct as stated. It becomes visible only on sectors with $d > N$ — the half-filling sector has $d = \binom{N}{N/2}$ (theorem 4.43), which at $N = 8$ is 70 against $N_E = 7$.

Corollary 4.49 (The back-edge does not grow with sector dimension)**.** *Combining theorem 4.47 with theorem 4.17,*

$$\dim \widehat{\Pi} \;=\; 3\min(N-1,\; d-1)\,, \tag{62}$$

so the back-edge dimension is bounded by $3(N-1)$ regardless of how large the protected sector is. It grows linearly in register size, not with $\binom{N}{N/2}$.

Proof. Theorem 4.17 gives $\dim \widehat{\Pi} = \sum_\sigma N_\sigma$ over the three channels, and theorem 4.47 gives $N_\sigma = \min(N-1, d-1)$ for each, which is eq. (62). Since $\min(N-1, d-1) \le N-1$ for every d, the bound $3(N-1)$ holds however large the sector is. At $N = 8$ half filling, $d = \binom{8}{4} = 70$ while the bound gives $N_\sigma = 7$. □

Corollary 4.49 is the operationally useful form. A naive reading of $N_\sigma = d-1$ suggests the calibration problem becomes intractable as soon as one leaves the qubit host, since d grows exponentially in N; eq. (61) shows it does not. What limits the coupling is the number of physical mediators, which is the register size, and $3(N-1)$ coordinates remain a tractable calibration target at any N reachable on current hardware.

Remark 4.50 (Scope)**.** Theorem 4.47 is stated for the site-resolved mediators $A_u^{\mathrm{mic}} = Z_u$ of **Part 229**, which is the set the corpus uses. A richer mediating set — pair operators $Z_u Z_v$, for instance — would raise the count and with it the bound. The theorem is therefore a statement about the interaction structure of eq. (3), not about protected sectors in general.

4.9 The host symmetry group and the dynamical selection criterion

Theorem 4.47 counts mediators. This subsection settles a different count — how many admissible *channel directions* the host carries — and it is the count the rank-one results of **Part 127**, **Part 179** and **Part 203** concern. The corpus states the answer at three different strengths and is careful to distinguish them; the distinction matters and is preserved here.

Definition 4.51 (Host symmetry group, **Part 113**, **Part 121**, **Part 127**)**.** $G_H = \mathrm{Stab}(P_H)$, the subgroup of site permutations fixing the host projector and the system Hamiltonian.

Lemma 4.52 (The group is $\mathbb{Z}_2$ at every k)**.** *For the structured family of eq. (C.4) with the host taken as the A_1-sector of theorem 4.43,*

$$G_H \;=\; \{\,\mathrm{id},\ \rho\,\} \;\cong\; \mathbb{Z}_2, \qquad \rho(j) = k-1-j, \tag{63}$$

generated by reversal of the chain, for every $k \geq 2$.

Proof. Exhaustive over S_k: a permutation lies in G_H iff its permutation unitary commutes with both H_k and P_H. Checked for $k = 2, \ldots, 8$, i.e. against all 40,320 elements of S_8 at the largest size; only the identity and ρ survive at each k. **Part 215** states the $k = 2$ case, from which eq. (63) extends it. □

Definition 4.53 (Physically admissible intertwiners, **Part 132** Eq. (18))**.** For each channel σ,

$$\mathfrak{I}_\sigma^{\mathrm{phys}} = \Big\{\, \Pi \in \mathrm{Hom}_{G_H}\big(\mathcal{T}_\sigma,\ \mathfrak{A}_H^{\mathrm{loc},sa}\big) \;:\; \mathcal{D}_{\mathrm{phys}}(\Pi) = 0 \,\Big\}, \tag{64}$$

and $N_\sigma = \dim \mathfrak{I}_\sigma^{\mathrm{phys}}$.

The content of definition 4.53 is that equivariance alone does not fix the count. It does not: the ambient equivariant dimension $\dim \mathrm{Hom}_{G_H}(\mathcal{T}, \mathfrak{A}_H^{\mathrm{loc},sa})$ evaluates to $2, 5, 20, 52, 200$ for $k = 2, \ldots, 6$, and grows. The second condition supplies the rest, and **Part 132** §13 names it: commutation with the restricted Hamiltonian, with scalar directions excluded.

Theorem 4.54 (Dynamical selection, **Part 132** Eq. (17))**.** *Let $H_S^H = (P_H H_S P_H)|_{\mathcal{H}_H}$ and take $\mathcal{D}_{\mathrm{phys}}(\Pi) = [H_S^H, \Pi]$ with scalars excluded. Then at the canonical two-qubit host of **Part 76**,*

$$\boxed{N_R^{\mathrm{dyn}} \;=\; 1, \qquad \widehat{G}_R \;=\; \sigma_x} \tag{65}$$

where $\sigma_x = |01\rangle\langle 10| + |10\rangle\langle 01|$ is the unique non-scalar direction of $\mathfrak{A}_H^{\mathrm{loc},sa}$ that is both G_H-invariant and commuting with H_S^H.

Proof. On the host, $H_S^H = -\mathbb{1} + \sigma_x$, so its commutant in $\mathfrak{A}_H$ is $\mathrm{span}\{\mathbb{1}, \sigma_x\}$; removing scalars leaves one dimension. By lemma 4.52 the group is $G_H = \{\mathrm{id}, \rho\}$; under ρ the swap acts as σ_x on the host, and σ_x is even while σ_y, σ_z are odd (**Part 215**), so σ_x is also the unique non-scalar G_H-invariant. The two conditions select the same line. Evaluating $\dim \mathfrak{I}^{\mathrm{phys}}$ by linear solve gives 1 at $d = 2$ and $2, 7, 11, 31$ at $d = 3, 6, 10, 20$: rank one at $d = 2$ and not beyond, which is **Part 278**'s conclusion reached independently. □

Corollary 4.55 (Exactly one conserved direction)**.** *$\langle\sigma_x\rangle$ is constant under the host dynamics while $\langle\sigma_y\rangle$ and $\langle\sigma_z\rangle$ rotate. This is the observable form of theorem 4.54 and is what table C.17 records.*

Proof. Write $\rho(t) = e^{-iH_S^H t}\rho\, e^{iH_S^H t}$ on the host. For any host observable O, $\frac{d}{dt}\langle O\rangle = i\langle[H_S^H, O]\rangle$. By theorem 4.54, H_S^H commutes with σ_x and with no other non-scalar direction of $\mathfrak{A}_H^{\mathrm{loc},sa}$, so the derivative vanishes for $O = \sigma_x$ and is nonzero for $O = \sigma_y, \sigma_z$ on any state with support off the σ_x axis. Explicitly, with $H_S^H \propto \sigma_x$, $[\sigma_x, \sigma_y] = 2i\sigma_z$ and $[\sigma_x, \sigma_z] = -2i\sigma_y$, so $\langle\sigma_y\rangle$ and $\langle\sigma_z\rangle$ rotate into each other at fixed $\langle\sigma_x\rangle$. □

Remark 4.56 (What is unconditional and what is not)**. Part 224** establishes the three channels at different strengths, and the distinction is load-bearing. Unconditionally: $N_R = 1$, with $\operatorname{rank}\Theta_E \le 1$ and $\operatorname{rank}\Theta_M \le 1$. Conditionally, on causal activity: $N_E = N_M = 1$. **Part 156** declines to assert the latter — "I will not manufacture the missing step and write $N_E = N_M = 1$" — and **Part 157** records that the R-branch has passed the test while the E/M branch has not. Theorem 4.54 closes R and claims nothing about E or M. **Part 132** §14 offers the same method for them, and the corpus is specific about what is still wanting. **Part 139** and **Part 144** record that the evaluated G_H-decompositions of $\mathcal{T}_E, \mathcal{T}_M$ are not instantiated. **Part 146** then supplies one of the two, $\mathcal{T}_E = 2V_+ \oplus 2V_-$, and declines to assign $\mathcal{T}_M$ the same form "merely by analogy" because the sources do not define $\mathcal{T}_M$ as an ordered-pair representation at all. The two branches are therefore open in different ways: for E the template exists and what is missing is $\operatorname{rank}\Theta_E = 1$; for M the template itself is missing. **Part 157** states the position exactly — the R-branch has passed this test, and the source "does not yet prove" the E/M one-dimensional-image step.

4.10 The recursive closure, and how many rounds it costs

Equation (1) places a recursive-closure step RCl between D_2 and D_3. **Part 123** and **Part 201** fix what that step is — a *least fixed point*, explicitly not a Banach contraction — and **Part 255** supplies the order structure that makes the fixed point exist, recording that a Kleene iteration from $\bot$ terminates in two steps at three sites. What no Part states is whether that bound survives at register scale. It does, for a reason that makes no reference to the number of sites at all.

Definition 4.57 (Causal increment and closure, **Part 123**, **Part 255**)**.** Let S be a finite site set, and let a *history* be a transitively closed relation $\Omega \subseteq S \times S$. Write TC for transitive closure. Given a causal increment $\Delta\mathcal{E} \subseteq S \times S$, fixed by the previous state, the closure operator is

$$\mathrm{Cl}_R(\Omega) \;=\; \mathrm{TC}\Big(\Omega \cup \Delta\mathcal{E}\Big). \tag{66}$$

Lemma 4.58 (The histories form a complete lattice, **Part 255**)**.** *Transitive relations on a finite S are closed under arbitrary intersection and include the full relation, so they form a complete lattice under $\subseteq$; and Cl_R is monotone, since both* TC *and* $\cup$ *are. Knaster–Tarski therefore applies, and* $\operatorname{lfp}\mathrm{Cl}_R$ *exists.*

Proof. Transitivity is preserved under arbitrary intersection, and the full relation $S \times S$ is transitive, so the transitive relations form a complete lattice under inclusion with meet given by intersection. If $\Omega \subseteq \Omega'$ then $\Omega \cup \Delta\mathcal{E} \subseteq \Omega' \cup \Delta\mathcal{E}$, and TC is monotone, so Cl_R is monotone. Knaster–Tarski applies to any monotone map on a complete lattice. □

Theorem 4.59 (One-step closure)**.** $\operatorname{lfp}\mathrm{Cl}_R = \mathrm{TC}(\Delta\mathcal{E})$*, and the Kleene iteration from* $\bot$ *reaches it after a single application of* Cl_R*, for every site count* $|S|$*.*

Proof. Write $\Omega^0 = \emptyset$ and $\Omega^{n+1} = \mathrm{Cl}_R(\Omega^n)$. Then $\Omega^1 = \mathrm{TC}(\emptyset \cup \Delta\mathcal{E}) = \mathrm{TC}(\Delta\mathcal{E})$. Since $\Delta\mathcal{E} \subseteq \mathrm{TC}(\Delta\mathcal{E})$, the union at the next step is absorbed,

$$\Omega^2 = \mathrm{TC}\Big(\mathrm{TC}(\Delta\mathcal{E}) \cup \Delta\mathcal{E}\Big) = \mathrm{TC}\Big(\mathrm{TC}(\Delta\mathcal{E})\Big) = \mathrm{TC}(\Delta\mathcal{E}) = \Omega^1, \tag{67}$$

the middle equality by idempotence of transitive closure. The ascending Kleene chain is therefore stationary from index 1. By lemma 4.58 the chain converges to $\operatorname{lfp}\mathrm{Cl}_R$, so that fixed point is $\Omega^1 = \mathrm{TC}(\Delta\mathcal{E})$. No step of the argument refers to $|S|$. □

Corollary 4.60 (The closure stage adds no rounds)**.** *The number of closure iterations required by eq. (1) is bounded independently of the register size. The closure stage therefore contributes no round overhead to an execution of the loop, and the depth accounting of remark 3.4 is unaffected by it.*

Proof. Immediate from theorem 4.59: the Kleene iteration is stationary after one application of Cl_R, and that bound is uniform in $|S|$, so the count of closure rounds does not grow with the register. □

Remark 4.61 (What this does not do)**.** Theorem 4.59 bounds the iteration count of one stage. It says nothing about the cost of the quantity that stage consumes: $\Delta\mathcal{E}$ is an *input*, fixed by the previous state, and the closure being free does not make it free. The classical synthesis and post-selection costs of remark 3.10 sit upstream of Cl_R and are untouched by this result. What is removed is a distinct cost that had not previously been bounded — whether the loop must be run to convergence over many rounds — and the answer is that it must not.

Remark 4.62 (What a processor can and cannot test here)**.** Theorem 4.59 is a statement about transitive relations on a finite set. No quantum measurement bears on it, and it is checked classically instead, by enumeration at three sites against **Part 255**'s own census and by direct iteration at register scale (section B). What hardware *can* test is the physical shadow of corollary 4.60: in the realization the closure stage is carried by a block commuting with A_1, so the occupation of the protected sector should not drift as the block is reapplied. Section C.1.9 reports that test, with a depth-matched control that does not commute with A_1.

4.11 Status of the substrate

Table 3 collects every statement about the BCM substrate that this section has settled, together with the result that settles it and whether it holds in general or only on the two-qubit realization.

Statement	Status	Basis
$\mathrm{FC}^\dagger$, $\mathrm{D}^\dagger$ as derived theorems	derived	theorems 4.2 and 4.3
Contextual Mapping	corollary	corollary 4.4
K_q admissible; $\mathrm{FC}^\dagger$ as conservation law	derived, scoped	theorem 4.7, remark 4.9
Conservation is codimension one	verified	lemma 4.8
$N_E = d - 1$	derived	theorem 4.16
$\dim \widehat{\Pi} = 3(d-1)$	derived	theorem 4.17
Identifiability $\iff$ (I1) $\wedge$ (I2)	derived	theorem 4.19
Quotient route to $\mathfrak{A}_H$	refuted	theorem 4.25
Reduction is a compression	derived	theorem 4.26
Locality/covariance dichotomy	derived	theorem 4.27
Parity obstruction at odd $\dim \mathcal{H}_E$	derived	theorem 4.21
Admissibility, even $\dim \mathcal{H}_E$	derived	theorem 4.34
Admissibility, odd $\dim \mathcal{H}_E$	no-go	corollary 4.35
Selection of the context K_q	derived	theorem 4.36
Realization well-defined on $\equiv_{\mathrm{BCM}}$	derived	corollary 4.13
Totality of D_4; $P_H^{(N)} \neq 0$	derived	theorem 4.41
System parity: $\mathbf{z} = 0$ iff N even	derived	theorem 4.43
Selection, general context: no-go	derived	theorem 4.38
Selection, general context: not needed	derived	theorem 4.39

Table 3: Status of the foundational layer. "Scoped" marks results holding on the two-qubit realization of eq. (5).

Theorems 4.31, 4.34 and 4.36 were tested on `ibm_kingston` and `ibm_fez`. The first two are confirmed and the third is confirmed as a comparative claim; the measurements, the tiled replication and the two corrections they exposed are in section C.1.6.

5 Discussion

5.1 Comparison with Existing Approaches

Table 4 presents a direct comparison between the projection-based construction of this paper and the quantum approaches considered for ground-state preparation. The rows are not commensurable on depth alone, and the table is laid out to make that visible rather than to be read as a ranking: a variational method solves a problem whose answer is unknown, while the construction here prepares a ground space whose energy is already known and bounds the classical work of doing so. What the comparison isolates is the depth of the quantum step, which is the quantity this paper measures.

5.2 Related Work and Positioning

The construction developed here sits at the intersection of three established lines of work: variational ground-state preparation, dissipative and symmetry-based state engineering, and the theory of decoherence-free subspaces. It is worth stating precisely what it takes from each and where it departs, including where it does *not* improve on them.

Table 4: Comparison of quantum approaches for ground-state preparation. The depth column describes the coherent quantum step only, so it is not comparable across rows without the last column: VQE and QAOA assume no prior knowledge of the target, whereas the present construction requires E_x in advance and is claimed only on the admissible class of definition 3.9. The depth advantage is real within that class and is not a claim about problems outside it; remark 3.10 and remark 5.1 state the boundary.

Method	**Depth**	**QEC**	**Ground-State Guarantee**	**Prior knowledge, and classical cost outside the coherent step**
VQE [1]	$\mathcal{O}(Np)$	No	No (variational optimization)	No prior knowledge of the target. Cost is the outer optimization loop: heuristic, and not bounded in advance [5]
QAOA [2]	$\mathcal{O}(p)$	No	Approximate only	No prior knowledge of the target. Cost is parameter optimization over the p layers
Trotterization [3]	$\mathcal{O}(t^2/\epsilon)$	Yes	Asymptotic	None beyond the terms of H
Shor's Algorithm [11]	$\mathcal{O}(N^3)$	Yes	Exact	None
This work	$\boldsymbol{\mathcal{O}(1)}$	**No**	**Exact within encoded model**	**Requires E_x.** Admissibility (definition 3.9) restricts the comparison to families whose E_x is polynomial-time computable, and theorem 3.12 then bounds synthesis of V_x by $r(\lvert x\rvert)$. Outside that class the cost reappears here

5.2.1 Variational preparation

The dominant near-term approach to ground states is variational: parameterise a shallow circuit and optimise its parameters against a measured energy. VQE established the template [1] and hardware-efficient ansätze carried it to molecules on superconducting hardware [12]; QAOA applies the same idea to combinatorial objectives [2]. Ref. [13] survey the family.

Two costs are intrinsic to that template. The first is the optimisation loop itself: each parameter update requires many circuit executions, and the total quantum cost is the product of circuit depth and iteration count. The second is trainability — for sufficiently expressive ansätze the gradient variance decays exponentially in system size, the barren-plateau phenomenon [5], so the loop may not converge at all.

The present construction removes both, and it is important to be exact about how. It does not optimise, because it does not search: the ground space is identified classically as $\ker(H_x - E_{\min}\mathbb{1})$ and the isometry V_x is synthesised from it (theorem 3.12). There is no landscape, therefore no barren plateau. The cost that VQE pays in quantum iterations is paid here in classical preprocessing, bounded by $T_{\text{classical}}(x) \leq r(|x|)$ for the admissible family.

Remark 5.1 (Two regimes, and what is solved in each)**.** The comparison with variational methods depends on which of two regimes is in view, and they differ in what classical work is required.

The native protected sector. For the framework's own host, nothing is solved. The subspace is

$$\mathcal{H}_H \;=\; \bigcap_{\mu} \ker(A_\mu - a_\mu \mathbb{1}) \,, \tag{68}$$

the simultaneous eigenspace of the system-side *interaction* operators — read off the system-environment coupling, not obtained by diagonalising a problem Hamiltonian. Theorem 4.36 makes the selection spectral and **Part 106** states the point directly: the maximal DFS admits "no choice of P_H". There is no eigenvalue problem here to compare against a variational search, because no eigenvalue problem is posed.

Encoded problems. Theorem 3.12 maps an external Hamiltonian H_x into the architecture, and there the isometry V_x does span Ground(H_x), so a classical step is required. That step is bounded, but the bound is an *admissibility condition* rather than a general claim: definition 3.9 requires E_x computable in polynomial time, and theorem 3.12 then gives $T_{\text{classical}}(x) \leq r(|x|)$. The class of problems satisfying that condition is where the comparison in table 4 applies, and it is stated rather than assumed.

The consequence for the variational comparison is that the two costs of the variational template — the optimisation loop and the barren plateau [5] — are absent in both regimes, for different reasons. In the first there is no objective to optimise; in the second the optimisation has been replaced by a construction with a stated complexity bound. Neither substitutes for a variational method on a problem falling outside definition 3.9.

5.2.2 Dissipative and measurement-based engineering

A second line prepares target states as fixed points rather than as optimisation outcomes. Ref. [14] showed that engineered dissipation is a universal resource, with the computation encoded in the steady state of a Liouvillian; Ref. [15] gave the corresponding construction for entangled targets, with jump operators that can be chosen quasi-local.

The mechanism here is the same in spirit and different in implementation. Both replace *evolution toward* a target with *stability at* it. But theorem 4.7 realises that stability through a conservation law, $[H_{\text{tot}}, A_1] = 0$, rather than through an engineered Lindbladian: the protected sector is a symmetry sector of the native Hamiltonian, so no bath engineering is required and no relaxation time is incurred. The cost is generality — dissipative engineering can target arbitrary states, whereas the present method targets ground spaces that coincide with a symmetry sector, which is what theorem 4.36 makes canonical rather than chosen.

5.2.3 Decoherence-free subspaces

The protected object is a decoherence-free subspace in the established sense [7, 6], and Ref. [10] reviews the theory and its relation to noiseless subsystems and dynamical decoupling. DFS techniques are normally deployed *defensively*: a logical qubit is encoded into a subspace immune to a symmetry of the noise, and computation proceeds inside it.

The contribution here is to use the same structure *constructively*. Theorem 2.3 identifies the host projector with the maximal DFS projector, $P_H = P_{\text{DFS}}^{\max}$, so the subspace that protects the computation is simultaneously the subspace that answers it. That identification is what converts a passive noise-avoidance technique into a state-preparation primitive, and theorem 4.36 shows the choice of subspace is fixed spectrally rather than by the designer.

5.2.4 Error mitigation and symmetry verification

Zero-noise extrapolation and probabilistic error cancellation [16] reduce bias in estimated observables by resampling or extrapolating; symmetry verification [17] post-selects on measured conserved quantities. The latter is the closest existing technique to what is done here, and the distinction is worth drawing sharply.

Symmetry verification measures a symmetry and *discards* runs that violate it, at a cost in accepted-shot fraction. The present construction places the computation inside the symmetry sector by construction, so violating runs are not produced to be discarded. Section C.1.6 measures this: population remains in its A_1 sector with median total variation between 0.038 and 0.069 across both processors, with no post-selection applied. Where the two coincide is that both rely on the same underlying fact — that a conserved quantity partitions the Hilbert space — and this work supplies the theorem (theorem C.1) fixing which partitions are available at arbitrary locality.

5.2.5 Constant depth, and why the bound does not bind

Claims of $\mathcal{O}(1)$ circuit depth require care, and the care runs in the opposite direction from the obvious one.

Ref. [18] showed via Lieb-Robinson bounds that a finite-depth local *unitary* circuit acting on a product state cannot generate long-range entanglement: correlations must decay outside the circuit's light cone. Ref. [19] separately established that constant-depth quantum circuits are nonetheless strictly more powerful than constant-depth classical ones, so the regime is not trivial. Together these delimit what a shallow unitary can do.

Two facts place the present construction outside that delimitation.

The protected subspace is provably not local. The host is not a bounded-region object, and this is a theorem rather than a feature of the chosen realization.

Theorem 5.2 (Rank obstruction, **Part 256**). *Let P_R act nontrivially only on a region R of an N-site register. Then* $\operatorname{rank}(P_R \otimes \mathbb{1}) = \operatorname{rank}(P_R)\, 2^{N-|R|}$. *Consequently a rank-two projector requires $|R| \geq N-1$: no rank-two projector belongs to a bounded-region subalgebra.*

Proof. The rank of a Kronecker product is the product of the ranks, and the identity on the complement of R has rank $2^{N-|R|}$, so $\operatorname{rank}(P_R \otimes \mathbb{1}) = \operatorname{rank}(P_R)\, 2^{N-|R|}$. Suppose this equals 2. Since P_R is a nonzero projector, $\operatorname{rank}(P_R) \geq 1$, so $2^{N-|R|} \leq 2$ and hence $N - |R| \leq 1$. A rank-two projector supported on a region therefore has $|R| \geq N-1$, which is not a bounded region as N grows. □

Theorem 5.3 (Support obstruction and dichotomy, **Part 256**). *No H_S-invariant subspace has bounded support, because eigenvectors of an irreducible hopping matrix have none. Hence a rank-two host is either translation-covariant and maximally delocalised, or exponentially localised and covariance-breaking — never both.*

Proof. Restricted to the one-magnon sector, H_S acts as a Jacobi matrix J: it is tridiagonal in the site basis with off-diagonal entries $b_i \neq 0$, these being the hopping amplitudes of eq. (C.4). Let $J\psi = \lambda\psi$ and suppose ψ vanishes on two adjacent sites, $\psi_j = \psi_{j+1} = 0$. The three-term recurrence

$$b_{j-1}\psi_{j-1} + a_j\psi_j + b_j\psi_{j+1} = \lambda\psi_j \tag{69}$$

then gives $b_{j-1}\psi_{j-1} = 0$, and $b_{j-1} \neq 0$ forces $\psi_{j-1} = 0$; iterating in both directions gives $\psi \equiv 0$. So no eigenvector vanishes on two adjacent sites. A vector supported in a proper region vanishes on a run of at least two consecutive sites whenever the complement of that region contains two adjacent sites, which holds for every proper region once $N \geq 4$. Hence no H_S-invariant subspace has bounded support.

For the dichotomy, suppose the host is translation-covariant. Its eigenvectors are then the plane waves $\psi_k(u) = N^{-1/2} e^{2\pi i k u/N}$, whose modulus is constant in u: the state is maximally delocalised, with participation ratio N. If instead the host is exponentially localised, $|\psi(u)| \leq Ce^{-|u-u_0|/\xi}$ with ξ finite, then ψ is not of constant modulus and the translation that carries u_0 to u_0+1 does not fix the subspace, so covariance fails. The two cases exclude each other, and by the first paragraph neither yields bounded support. □

At $N = 24$ a three-site region forces rank $\geq 2^{21}$, and bounded windows leak at $\|(\mathbb{1} - P)H_S P\| \approx 2.8$, which is $\mathcal{O}(1)$ rather than small. The covariant host on a ring of eight is the momentum pair $\operatorname{span}\{|k{=}0\rangle\,, |k{=}1\rangle\}$: rank two, $\|[H_S, P]\| = 5.9 \times 10^{-16}$, and fully delocalised.

The protocol is not a finite-depth unitary. Theorem 3.3 realises the projection as a measurement-conditioned channel,

$$\mathcal{E}_H(\rho) = \frac{P_H\, U\rho U^\dagger\, P_H}{\operatorname{Tr}(P_H\, U\rho U^\dagger)}, \tag{70}$$

with a coherent layer of depth $\mathcal{O}(1)$ followed by projective measurement. Equation (70) is not a unitary circuit, and Lieb-Robinson bounds do not constrain it. That measurement

and feedback evade the unitary depth bound is established: Ref. [20] show adaptive low-depth circuits reach long-range entangled matter that no finite-depth unitary can prepare, and Ref. [21] obtain long-range entanglement by measuring symmetry-protected phases.

Remark 5.4 (Where the advantage actually lies)**.** The construction reaches, at $\mathcal{O}(1)$ coherent depth, subspaces that theorems 5.2 and 5.3 prove cannot be written as bounded-region objects — states a finite-depth unitary circuit is forbidden by [18] from preparing at all. It does so through eq. (70), the route [20, 21] identify as the one that circumvents the bound.

What distinguishes it from those works is the source of the projector. Adaptive schemes measure an ancilla or a symmetry-protected order parameter and apply feedback conditioned on the outcome. Here the projector is the maximal DFS projector of the *native* interaction, $P_H = P_{\mathrm{DFS}}^{\max}$ (theorem 2.3), and the sector it selects is conserved by the dynamics, $[H_k, A_1] = 0$ for every $k \geq 2$ (theorem C.1). No feedback is required, because there is nothing to correct: the sector the measurement projects into is one the Hamiltonian cannot leave. Section C.1.6 measures this, with population remaining in its A_1 sector at median total variation 0.038–0.069 and no post-selection applied — in the reporting sense of remark 3.4: the figures are raw, with no shots filtered out before the statistic is formed.

The cost, stated exactly. The advantage is not free, and the price is not depth. Theorem 3.3 gives the expected number of preparation repetitions as $1/p_{\mathrm{succ}}(x)$ with $p_{\mathrm{succ}}(x) = \langle\psi_0| P_H(x) |\psi_0\rangle$, for the post-selected route; on the deterministic route of remark 3.4 the same difficulty appears instead in the classical synthesis of V_x. The coherent depth is $\mathcal{O}(1)$ independently of $|x|$; the repetition count is where the problem's difficulty is expressed, alongside the classical synthesis of V_x bounded by theorem 3.12. A protocol that claimed both constant depth and constant total cost would be claiming something false, and this one does not.

5.2.6 Where the improvement is

Consolidating, the contribution is narrower and sharper than "ground-state preparation at constant depth":

1. **No optimisation loop, hence no barren plateau**, for problems whose ground space is classically identifiable — against [1, 2, 12, 13].
2. **Protection by native symmetry rather than engineered dissipation**, requiring no bath design and incurring no relaxation time — against [14, 15].
3. **The DFS used constructively rather than defensively**, $P_H = P_{\mathrm{DFS}}^{\max}$ (theorem 2.3) — extending [7, 6, 10].
4. **Symmetry enforced rather than verified**, so no shots are discarded — against [17].
5. **Long-range entangled subspaces at $\mathcal{O}(1)$ coherent depth**, which no finite-depth unitary circuit can reach (theorems 5.2 and 5.3, and remark 5.4) — via measurement-conditioned projection onto a *conserved* sector, so unlike [20, 21] no feedback is required.

6. **The k-body synthesis cost removed exactly** (theorem 4.31): $2(k-1)$ two-qubit gates for the plaquette against $\mathcal{O}(4^k)$ for generic synthesis, measured at 17.5× fewer gates at $k = 5$ in section C.1.6.

Items 1–4 are architectural and hold wherever the admissibility conditions of definition 3.9 are met. Item 5 is a synthesis result independent of the rest and applies to any structured k-local family of the form eq. (C.4).

5.3 The Mechanism: Why It Works

The oracle succeeds because it does not *compute* the ground state—it *projects* onto a decoherence-free subspace that is, by the algebraic structure of the boundary Hamiltonian, identical to the ground-state manifold. The key insight from the IQS formalization is:

1. The Host boundary H_∂ defines the problem (Equation (4)).
2. The compression map C_H eliminates non-physical degrees of freedom (Equations (10) and (11)).
3. The DFS equivalence guarantees decoherence immunity (Theorem 3.7).
4. The uniqueness theorem guarantees a single convergence point (Theorem 3.5).

This is fundamentally different from error *correction* [8] (which detects and repairs errors mid-computation) or error *mitigation* (which statistically cancels errors post-computation). The oracle achieves error *avoidance*: the protocol confines the accepted output to the decoherence-free Host subspace rather than relying on active error correction to restore it afterward.

5.4 Scalability

The tiled architecture scales linearly with device size. As quantum processors grow from 156 to 1,000+ qubits, the number of independent tiles increases proportionally, amplifying statistical power without increasing circuit depth. The $\mathcal{O}(1)$ depth guarantee is hardware-agnostic and holds for any coupling topology admitting an induced matching.

5.5 Future Work

The two principal limitations examined in this work have been substantially addressed within the demonstrated regime:

- **Tile locality** has been extended to the structured k-local family for $k \in \{3, 4, 5\}$ (Section C.1.4).
- **Problem encoding** has been demonstrated for six distinct problem classes via the Compression Reduction (**Part 253**, Section C.1.5).

The classical complexity of constructing the corresponding problem-dependent unitary remains an explicit function of problem size. The remaining directions for future work are:

1. **Independent reproduction.** All 55 job identifiers of tables C.23 to C.25 are publicly queryable on IBM Quantum Platform. We invite the community to independently reproduce the results using the provided circuit specifications and raw telemetry data.

2. **Scaling to 1,000+ qubits.** As IBM and other vendors release larger processors, the tiled architecture scales linearly with no increase in circuit depth. Demonstrating this on next-generation hardware is a natural extension.

3. **Unitary synthesis for large k — resolved.** The concern was that the transpiled gate count of a $2^k \times 2^k$ synthesis grows with k. Theorem 4.31 removes the k-body part of it: the plaquette commutes with the chain (lemma 4.30), so the evolution factorises exactly and the plaquette reduces to a CNOT-ladder phase gadget of $2(k-1)$ two-qubit gates, against the $\mathcal{O}(4^k)$ of generic synthesis (table 2). What remains is a two-local evolution with its own cost, so this is a removal of the exponential term and not of the whole scaling (remark 4.32). Re-running section C.1.4 through the structured decomposition is a cost reduction, not a correction.

6 Conclusion

We have demonstrated, both theoretically and experimentally, that quantum ground-state preparation can be achieved in $\mathcal{O}(1)$ circuit depth on current NISQ hardware through structural projection onto the maximal decoherence-free subspace.

Five results carry that claim, each extracted from the IQS/EQS formalization and each proved above. The Host projector *is* the maximal decoherence-free subspace projector, $P_H = P_{\mathrm{DFS}}^{\max}$ (theorem 2.3), which is what makes the projection protective rather than merely selective: the subspace the oracle lands in is one the interaction cannot dephase. At the two-qubit host the admissible realization class is a single equivalence class, $|\mathfrak{R}/\!\sim_{\mathrm{phys}}| = 1$ (theorem 3.5), so the sector reached is not one choice among many; above that dimension the count is $d-1$ and uniqueness fails. The conservation law $[H_k, A_1] = 0$ holds for the structured family at every $k \geq 2$ (theorem C.1), carrying the construction above two-body locality, and its k-body synthesis cost is removed exactly by theorem 4.31. Finally the Compression Encoding Theorem, $\mathcal{Q}(A) = V^*AV$ with $P_H(x) = V_x V_x^*$ and $\mathrm{Ground}(H_x) = \mathrm{Im}\, P_H(x)$ (theorem 3.12), supplies the entry point: it absorbs the difficulty of a problem family into a classical synthesis whose cost the theorem bounds, leaving a quantum step whose depth does not depend on the instance. That bound is a condition on the class rather than a claim about problems in general: definition 3.9 admits only families whose ground-state energy is already computable in polynomial time, and remark 3.10 states what falls outside it. Fifth, the dynamical selection criterion of theorem 4.54 settles how many channel directions the host carries: exactly one survives, $\widehat{G}_R = \sigma_x$, so the R-channel count $N_R = 1$ is established rather than assumed. It is the one of the three channels the corpus proves unconditionally, and remark 4.56 keeps that distinction from the E- and M-channels, which remain open.

The framework is experimentally validated across 55 independent QPU executions on 156-qubit IBM Heron processors: 21 Phase 3–6 executions, which demonstrate the structured k-local Hamiltonian family at $k = 3, 4, 5$ and six representative encoded problem instances spanning combinatorial optimization, number theory, chemistry, finance, logistics and frustrated magnetism, together with twenty-four executions testing the derived

theorems of sections 4.7 and 4.8 directly against pre-registered predictions, and the encoding theorem at $N = 6$ to 12 (section C.1.7). Both previously identified limitations—tile locality and problem encoding—have been substantially addressed within the demonstrated regime.

The distinction the construction rests on is between *computing* a ground state and *projecting* onto it. Where the ground space coincides with a decoherence-free subspace of the native coupling, it can be reached by a projection whose coherent depth is independent of the instance, rather than by a time evolution that must be simulated. That is a statement about a class of Hamiltonians, delimited by definition 3.9 and theorem 4.27, and not a general claim about ground-state preparation.

Methods

Hardware Platform

All experiments were executed on IBM Quantum Heron-class processors (156 qubits, heavy-hex topology, $r2$ revision). Table 5 summarizes the hardware specifications.

Table 5: IBM Heron processor specifications for the three backends used in this work.

Backend	**Qubits**	**Median CNOT Error**	**Median T_2 (μs)**
`ibm_marrakesh`	156	$\sim 7 \times 10^{-3}$	~ 120
`ibm_kingston`	156	$\sim 6 \times 10^{-3}$	~ 140
`ibm_fez`	156	$\sim 8 \times 10^{-3}$	~ 110

Authentication was performed via the IBM Quantum Platform with CRN instance specification. All jobs were submitted through `SamplerV2` with 50,000 shots per execution.

Circuits, pre-registered predictions and the analysis of returned counts were implemented in Python with Qiskit, written with the assistance of Anthropic's Claude (Opus 5) to the author's specification and verified by the author; section 6 declares that use in full. Predictions were computed classically and fixed before submission of each job, and every value reported is recomputed from the archived counts.

The Tiled Memory Circuit

The oracle is realized through a spatial tiling strategy that maps the 2-qubit boundary interaction across the entire processor:

1. **Topology Extraction.** The `HardwareTopology` engine scans the backend's coupling graph and computes a maximum-weight induced matching—a set of disjoint qubit pairs (q_i, q_j) such that no two pairs share a neighbor, eliminating crosstalk.

2. **Unitary Synthesis.** For each tile, the 2-qubit unitary $U_2 = e^{-iH_2 t}$ is decomposed into the native gate set $\{CZ, \sqrt{X}, R_Z\}$ via `TwoQubitBasisDecomposer` with Euler basis `ZSX`.

3. **Parallel Composition.** All tiles are composed simultaneously onto the full-device circuit, yielding $\mathcal{O}(1)$ depth regardless of device size.

Table 6 gives the tile count the extraction step returned on each backend, and the state space that count spans.

Table 6: Tiling results for each backend. M = number of disjoint tiles.

Backend	**Disjoint Tiles (M)**	**Total Measured Qubits**	**State Space 2^M**
`ibm_marrakesh`	33	66	8.59×10^9
`ibm_kingston`	38	76	2.75×10^{11}
`ibm_fez`	34	68	1.72×10^{10}

Custom Hamiltonian Injection

The engine accepts arbitrary 2-qubit Hamiltonians H_{custom} via the `custom_unitary` parameter of the `TopologicalMemoryCircuit` class — the software identifier for the tiled memory circuit of section 6, unchanged here because it is the name in the executed code. Table 7 lists the four domain-specific Hamiltonians deployed.

Table 7: Custom Hamiltonians injected into the oracle architecture.

Domain	**Hamiltonian H_{custom}**	**Physical Encoding**
Frustrated Ising (default)	$X \otimes X$	Spin-glass ground state
Molecular Chemistry (H_2)	$Z \otimes Z + 0.5\, X \otimes X$	Electron correlation
Cryptography (Factoring)	$2\, Z \otimes Z - Z \otimes \mathbb{1} - \mathbb{1} \otimes Z$	Integer factorization cost
Optimization (MaxCut)	$Z \otimes Z$	Combinatorial constraint

Classical Verification (Pre-Hardware)

Prior to hardware deployment the oracle was verified classically on frustrated Erdős–Rényi MaxCut instances at $N \in \{6, 8, 10, 12\}$ qubits. Each $H_x = \sum_{(u,v)\in E} Z_u Z_v$ is assembled from Kronecker products and diagonalized densely, so that a construction error would appear as a spectrum mismatch; the host projector $P_H(x) = V_x V_x^*$ of theorem 3.12 is then applied once to the uniform superposition. Table 8 reports the outcome.

Table 8: Classical verification of a single projection step. "g" is the ground degeneracy, "survivors" the number of basis states carrying non-zero amplitude after the projection, and "eliminated" the fraction $1 - g/2^N$. Every instance is frustrated in the sense that no assignment cuts all $|E|$ edges. Instances are drawn from $G(N, \frac{1}{2})$ under a fixed seed and rejection-sampled until frustrated, so the table is reproducible from the specification alone.

N	dim	$\|E\|$	max-cut	g	survivors	eliminated	fidelity
6	64	9	8	2	2	96.875%	1.000000
8	256	14	10	22	22	91.406%	1.000000
10	1024	19	15	2	2	99.805%	1.000000
12	4096	29	23	2	2	99.951%	1.000000

At $N = 12$ the projection leaves 2 of 4096 basis states, eliminating 99.95% of the Hilbert space in a single step, and the surviving support is verified by brute-force enumeration of all 2^{12} assignments to be exactly the set of maximum cuts. The projected state has fidelity 1 with Ground(H_x) and attains $E_{\min}$ to machine precision at every size.

The elimination fraction is $1 - g/2^N$ and is therefore a property of the instance, not a uniform guarantee: the $N = 8$ graph is 22-fold degenerate and retains 8.6% of the space. What holds at every size, and is the claim this check supports, is that the surviving support is the exact optimal set — never a superset — and that one projection suffices to reach it. The conservation law $[H_x, A_1] = 0$ with $A_1 = \sum_u Z_u$, which the decoherence-free argument of theorem 2.3 requires, is confirmed at every size.

Use of AI tools

Anthropic's Claude (Opus 4.5 and Opus 5, accessed through Claude Code) was used at two stages of this work. In both cases it acted on the author's material and under the author's direction.

The first is the IQS/EQS formalization [9] from which this framework is derived. The author supplied the constructions, the intended results and the direction of each argument; the model assisted in rendering them as formal statements and proofs.

The second is the preparation of this article: drafting from the author's content and outline, assembling the classical verification scripts of table B.2, and implementing the circuits, pre-registered predictions and analysis code for the executions reported in section C, to the author's experimental design.

The author designed the research, directed and verified this work, and is solely responsible for the correctness of every result, figure and number reported here. No AI system is an author. Predictions were computed classically and fixed before the corresponding jobs were submitted, and every value quoted is recomputed from the archived measurement record rather than transcribed.

Code availability

The code used to construct the circuits, compute the pre-registered predictions and analyse the returned counts is not publicly deposited, but is available from the corresponding author on reasonable request. The constructions the paper depends on are specified in the body in sufficient detail to be reimplemented independently: the host isometry and

projector of theorem 3.12, the exact split of eq. (50), the preparation of eq. (C.11) and lemma C.9, and the tiling procedure above. The classical checks reported in table B.2 are deterministic and reproduce from those specifications. Code used to generate and analyse the reported executions is available from the corresponding author on reasonable request.

Data availability

The data supporting the findings of this study are the raw quantum-processor telemetry of every execution reported here: 204 JSON files carrying the complete per-shot measurement record, approximately 785 MB, together with seven documentation files indexing them. Every numerical value in section C is computed from these files, and no analysis in this paper draws on data outside them. 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.

Verification does not depend on that request being answered. The primary record is held by the provider, not by the author: the 55 job identifiers listed in tables C.23 to C.25 remain queryable through the IBM Quantum Platform API, so a reader can confirm that each execution took place, on which processor, at what shot count and with what status, without reference to this paper or to any file the author supplies. Each archived file carries in its name the identifier of the job that produced it, so the archive can in turn be checked against that record. The executions were performed on `ibm_kingston`, `ibm_fez` and `ibm_marrakesh` between August 2026 and the date of submission.

The counts are as measured. Three faults in the analysis were identified after execution and are documented in sections C.1.5 to C.1.7; in each case the correction is in the interpretation and not in the archive, so both the original and the corrected figures can be recomputed from the same files. One further submission was cancelled in the queue and returned no counts; it is recorded in the archive without a result rather than omitted.

No data are withheld and no access restrictions apply to the archive. The IQS/EQS corpus from which the framework is derived is cited as [9]; it is referenced by Part number throughout and indexed in section A, and is not redistributed here.

References

[1] Peruzzo, A. *et al.* A variational eigenvalue solver on a photonic quantum processor. *Nat. Commun.* **5**, 4213 (2014).

[2] Farhi, E., Goldstone, J. & Gutmann, S. A quantum approximate optimization algorithm (2014). 1411.4028.

[3] Lloyd, S. Universal quantum simulators. *Science* **273**, 1073–1078 (1996).

[4] Preskill, J. Quantum computing in the NISQ era and beyond. *Quantum* **2**, 79 (2018).

[5] McClean, J. R., Boixo, S., Smelyanskiy, V. N., Babbush, R. & Neven, H. Barren plateaus in quantum neural network training landscapes. *Nat. Commun.* **9**, 4812 (2018).

[6] Zanardi, P. & Rasetti, M. Noiseless quantum codes. *Phys. Rev. Lett.* **79**, 3306–3309 (1997).

[7] Lidar, D. A., Chuang, I. L. & Whaley, K. B. Decoherence-free subspaces for quantum computation. *Phys. Rev. Lett.* **81**, 2594–2597 (1998).

[8] Knill, E., Laflamme, R. & Viola, L. Theory of quantum error correction for general noise. *Phys. Rev. Lett.* **84**, 2525–2528 (2000).

[9] 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.

[10] Lidar, D. A. Review of decoherence-free subspaces, noiseless subsystems, and dynamical decoupling. In *Quantum Information and Computation for Chemistry*, vol. 154 of *Advances in Chemical Physics*, 295–354 (John Wiley & Sons, 2014).

[11] Shor, P. W. Algorithms for quantum computation: discrete logarithms and factoring. In *Proceedings 35th Annual Symposium on Foundations of Computer Science*, 124–134 (1994).

[12] Kandala, A. *et al.* Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets. *Nature* **549**, 242–246 (2017).

[13] Cerezo, M. *et al.* Variational quantum algorithms. *Nat. Rev. Phys.* **3**, 625–644 (2021).

[14] Verstraete, F., Wolf, M. M. & Cirac, J. I. Quantum computation and quantum-state engineering driven by dissipation. *Nat. Phys.* **5**, 633–636 (2009).

[15] Kraus, B. *et al.* Preparation of entangled states by quantum markov processes. *Phys. Rev. A* **78**, 042307 (2008).

[16] Temme, K., Bravyi, S. & Gambetta, J. M. Error mitigation for short-depth quantum circuits. *Phys. Rev. Lett.* **119**, 180509 (2017).

[17] Bonet-Monroig, X., Sagastizabal, R., Singh, M. & O'Brien, T. E. Low-cost error mitigation by symmetry verification. *Phys. Rev. A* **98**, 062339 (2018).

[18] 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).

[19] Bravyi, S., Gosset, D. & König, R. Quantum advantage with shallow circuits. *Science* **362**, 308–311 (2018).

[20] 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).

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

Acknowledgements

This study received no funding.

Author contributions

M.H. conceived the framework, derived and proved the theoretical results, designed the experimental programme, implemented and deployed the quantum circuits, analysed the measurement data, and wrote the manuscript.

Competing interests

The author declares no financial or non-financial competing interests.

Appendices

The three appendices are supplied as a separate document, `appendices.pdf`, accompanying this article as supplementary material. They are:

- **Appendix A — IQS Part Citation Index.** Provenance for every result in the body: each numbered theorem, definition or equation of this paper mapped to the Part or Parts of the IQS/EQS corpus [9] it is taken from, with the content stated in the notation used here.
- **Appendix B — Data and Reproduction.** The supplementary data files, their naming scheme and format, the classical checks, and what is required to reproduce each reported result.
- **Appendix C — Experimental Record.** The complete experimental programme: hardware, the tiling procedure, the Phase 3–6 executions, the closure tests, the oracle reference tests O1–O6, and the job registry of all 55 quantum-processor executions.

References of the form "section C" throughout this article point into that document.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- Appendix.pdf