DATED HISTORY

The Mercury Rung: Scalar Timing, Conditional Spin-2 Emergence, and Phase-1A Hardware Candidate Tests in QATNU

The Phase-1A hardware installment belongs to the experimental development record. The consolidated paper records the source-response checks and scoped null result.

Manuscript · p. 1, hardware candidate tests

Spin-2 search · null result

Evidence · Negative QATNU spin-2 search

The Mercury Rung

Scalar Timing, Conditional Spin-2 Emergence, and Phase-1A Hardware Candidate Tests in QATNU

Joshua Farrow
Independent Research

May 2026

Status. This paper gives a complete formulation of the QATNU "Mercury rung": the scalar timing sector, the spin-2 emergence target, the conditional derivation of a Pauli-Fierz-like tensor sector from bond-stress criticality and emergent Ward identities, corrected Phase-0A negative-control baselines, a four-circuit Phase-0A real-QPU hardware batch, and Phase-1A exploratory hardware candidate tests on embedded graph families. It does not claim that quantum gravity has been derived. The current results are narrower and deliberately stricter: scalar/bond signal is not gravity, tensor-looking raw anisotropy is not spin-2, and only a pre-registered tensor protocol with finite-size scaling, Ward tests, and the Mercury number R_g can pass. Phase-0A controls show active bond/scalar observables while still returning `FAIL_EXPECTED_NEGATIVE_CONTROL`; Phase-1A house, square, and triangular circuits show that larger hardware-readable graph families run on IBM hardware, while embedded 2D lattice cases are labeled only `PHASE1_TENSOR_CANDIDATE_NOT_CLAIMED`. The Kingston reruns preserve the boundary on a second IBM backend and include matched square and triangular pairs: `square_2x3` $\rightarrow$ `square_3x3` and `triangular_2x3` $\rightarrow$ `triangular_3x3`. The triangular pair is the strongest Phase-1A hardware-facing follow-up signal in this version: the 25-qubit `triangular_3x3` run increases mean bond occupation and scalar-load spread relative to `triangular_2x3`, while normalized raw anisotropy drops. This is not spin-2 evidence. It is an exploratory candidate result that identifies the next pre-registered tensor run.

Abstract

Emergent-gravity programs often succeed first at producing scalar or geometric analogies: clock slowing, graph distances, effective dimensions, or projection maps. The hard target is different. A substrate theory of gravity must produce the dynamical spin-2 content of weak-field general relativity without inserting it by hand. This paper formulates that target for the Quantum Automaton Tensor-Network Universe (QATNU), a finite local quantum substrate with matter qubits on vertices and bond registers on edges.

The paper has six parts. First, it separates the scalar timing sector from the tensor problem. QATNU naturally supports an internal scalar load Λ_i built from bond occupation and a clock law $\omega_{\text{eff}} = \omega/(1+\alpha\Lambda_i)$, but scalar timing is holonomy-free and cannot itself carry gravitational radiation. Second, it states the Mercury criterion for emergent gravity: the theory must target a known weak-field structure, avoid inserting the target by projection, and output at least one nontrivial number. For QATNU the primary number is $R_g = c_g/c_{\text{clock}}$, the ratio between the emergent tensor-wave speed and the scalar clock-sector speed. Third, it gives the conditional spin-2 derivation. The substrate bond-stress tensor M_{ab} defines a tensor susceptibility. If the tensor inverse kernel reaches a critical point $m_2^2 = \Delta_B^{(2)} - \lambda^2\Pi_2(0,0) = 0$ and the infrared Ward residuals vanish, then the long-wavelength correlator develops a transverse-traceless pole and the unique local quadratic effective action is Pauli-Fierz, with $c_g^2 = \rho_2/Z_2$. This is a derivation target, not an assumption.

Fourth, the paper reports corrected Phase-0A exact controls. The structural gate distinguishes null tensor channels from nonzero two-dimensional tensor controls: path, cycle, and star do not exercise a usable tensor channel; chorded diamond and bowtie controls exercise a nonzero anisotropy tensor but still fail because they are not three-dimensional infrared substrates. Fifth, the paper reports a full Phase-0A IBM hardware batch. Four controls were executed on real IBM Quantum backends and all returned `FAIL_EXPECTED_NEGATIVE_CONTROL`. Sixth, the paper reports the Phase-1A graph-family hardware pass: house, square, and triangular circuits ran on IBM hardware, with Kingston reruns for Phase-0A controls and matched Phase-1A square and triangular pairs. The 25-qubit `triangular_3x3` run is the largest and most angularly structured circuit in the current package. Relative to `triangular_2x3` on the same Kingston backend, it increases mean bond occupation and scalar-load spread while reducing normalized raw anisotropy. This is the best candidate-family follow-up signal so far, not positive evidence for spin-2 gravity. The square and triangular 2D lattices are only candidates, labeled `PHASE1_TENSOR_CANDIDATE_NOT_CLAIMED`, until a pre-registered tensor protocol supplies finite-size scaling, tensor susceptibility, TT fraction, Ward-residual checks, coarse-graining reproduction, and scalar/vector impostor rejection.

Contents

1 The problem **3**

2 The Mercury criterion **3**

3 The QATNU substrate **4**
3.1 Implemented exact-diagonalization form 4

4 The scalar timing sector **5**
4.1 Weighted graph closure 5
4.2 Why the scalar sector is not enough 6

5 The spin-2 target **6**

6 No-go constraints **7**

7 Conditional derivation of the QATNU spin-2 sector **7**
7.1 Tensor susceptibility 8
7.2 Ward residuals 9
7.3 Conditional theorem 9

8 Phase-0A: what has now been tested **9**

9 Hardware-readable Phase-0A Hamiltonian **10**

10 Structural spin-2 admissibility gate **10**

11 Phase-0A exact-simulation controls **11**

12 Real-QPU Phase-0A hardware batch **12**
12.1 Kingston backend-consistent Phase-0A rerun 14
12.2 Interpretation of the hardware batch 16

13 Phase-1A: exploratory embedded graph-family hardware candidate tests **16**
13.1 Phase-1A local Aer checks 17
13.2 Phase-1 IBM hardware runs 17
13.3 Kingston Phase-1A matched-family rerun 19
13.4 Hardened interpretation of the triangular pair 25
13.5 Interpretation of Phase 1A 26

14 Reproducibility package **27**

15 What remains to make this a positive spin-2 paper **27**

16 Claim registry **28**

17 Discussion **29**

18 Conclusion **29**

A Hardware run prompt **30**

B Notation **30**

C References **31**

1 The problem

The weak-field gravitational field is not a scalar clock lag. In linearized general relativity one writes

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu}, \qquad |h_{\mu\nu}| \ll 1, \tag{1}$$

and the propagating radiative degrees of freedom are the two transverse-traceless helicity-2 modes contained in h_{ij}^{TT}. A model that reproduces only a scalar redshift law has reproduced one observable face of weak gravity, not gravity itself.

This distinction is the central discipline of the QATNU program. The scalar sector is real enough to define, measure, and test. But the gravitational claim begins only when the substrate generates a protected spin-2 tensor sector whose equations of motion, propagator, and coupling structure match the Pauli-Fierz weak-field target.

The purpose of this paper is to put the whole story in one place. It defines the QATNU substrate, explains the scalar timing layer, derives the conditional spin-2 target, reports a first negative-control result, and gives the hardware route for reproducing the Phase-0 experiment on real quantum processors.

2 The Mercury criterion

A theory becomes more than a language when it produces an output it did not already put in. The historical role of Mercury's anomalous perihelion precession in general relativity is useful as a methodological benchmark: the decisive move was not aesthetic analogy, but a number that followed from equations constrained by other principles.

For QATNU, the corresponding requirement is:

> *derive the weak-field spin-2 sector from the substrate, including a nontrivial output number, without defining the tensor sector by the projection that one later claims to explain.*

This criterion has three parts.

1. **Known target.** The low-energy benchmark is a massless helicity-2 field with Pauli-Fierz kinetic structure and universal coupling to a conserved source.
2. **No insertion by hand.** A transverse-traceless projector may diagnose the result after the fact, but it may not be the reason the result exists.
3. **Output number.** The substrate must predict at least one dimensionless number not chosen to match the answer. The primary QATNU target is

$$R_g = \frac{c_g}{c_{\mathrm{clock}}}. \tag{2}$$

The rest of the paper is organized around this criterion.

3 The QATNU substrate

Let $G = (V, E)$ be a finite graph. Each vertex $i \in V$ carries a two-level matter system, and each edge $e = \langle i, j \rangle \in E$ carries a finite bond register. The total Hilbert space is

$$\mathcal{H} = \bigotimes_{i \in V} \mathcal{H}_i^{\text{mat}} \otimes \bigotimes_{\langle i,j \rangle \in E} \mathcal{H}_{ij}^{\text{bond}}, \qquad \dim \mathcal{H} = 2^{|V|} \chi_{\max}^{|E|}. \tag{3}$$

A minimal ladder Hamiltonian is

$$H = H_{\text{mat}} + H_{\text{bond}} + H_{\text{int}}, \tag{4}$$

with

$$H_{\text{mat}} = \frac{\omega}{2} \sum_{i \in V} X_i, \tag{5}$$

$$H_{\text{bond}} = \sum_{\langle i,j \rangle} \sum_{m=0}^{\chi_{\max}-1} \epsilon_m \Pi_{ij}^{(m)}, \tag{6}$$

$$H_{\text{int}} = \lambda \sum_{\langle i,j \rangle} F_{ij} \otimes \left(R_{ij} + R_{ij}^\dagger \right), \tag{7}$$

$$F_{ij} = \frac{1 - Z_i Z_j}{2}, \tag{8}$$

$$R_{ij} = \sum_{m=0}^{\chi_{\max}-2} \sqrt{m+1}\, |m+1\rangle\langle m|_{ij}. \tag{9}$$

The interaction is self-referential in the limited technical sense that local matter frustration promotes or demotes the bond register that defines the medium through which matter propagates. The model is finite, local, and unitary. It is therefore a testbed for emergence, not an appeal to nonlocal collapse or external measurement.

3.1 Implemented exact-diagonalization form

The current exact-diagonalization engine implements the same architecture with matter dimension 2^N, bond dimension $\chi_{\max}^{|E|}$, lookup tables for the matter Z values, explicit edge lists, and dense Hamiltonian construction. It includes transverse matter flips, diagonal $J_0 Z_i Z_j$ matter coupling, diagonal bond energy proportional to bond rung, an optional degree penalty, and frustration-controlled bond promotion/demotion. The code-level Hamiltonian is therefore slightly more concrete than Eq. (5)–(7):

$$H_{\text{ED}} = \frac{\omega}{2} \sum_i X_i + J_0 \sum_{\langle i,j \rangle} Z_i Z_j + \Delta_B \sum_e n_e + \kappa \sum_i (d_i - k_0)^2 \tag{10}$$

$$+ \lambda \sum_{e=\langle i,j \rangle} F_{ij} \left(T_e^+ + T_e^- \right), \tag{11}$$

where $T_e^\pm$ raise or lower the finite bond rung when allowed by the cutoff. Table 1 records the alignment between the analytic model, the exact-diagonalization implementation, and the Phase-0

hardware encoding.

Analytic object	Exact-diagonalization implementation	Hardware-readable Phase 0
Matter clocks $\frac{\omega}{2}\sum_i X_i$	Implemented as matter-state bit flips	Implemented as single-qubit X rotations
Bond ladder $\epsilon_m \Pi^{(m)}$	Linear rung energy $\Delta_B m$ with cutoff $\chi_{\max}$	One bond qubit per edge, $\chi_{\max} = 2$
Frustration tile F_{ij}	$F_{ij} = (1 - Z_i Z_j)/2$ from lookup table	Same expression compiled into controlled bond rotations
Promotion operator $R + R^\dagger$	Finite raise/lower moves of bond rung	Bond-qubit flip X_e controlled by frustration
Degree penalty	Optional $\kappa(d_i - k_0)^2$ term	Omitted in v0 to reduce depth
Tensor observable M_{ab}	Not part of legacy ED core	Raw 2D anisotropy diagnostic only

Table 1: Theory-to-code alignment. The hardware Phase-0 encoding is intentionally smaller than the full ladder theory. Its role is to test negative controls on real devices, not to complete the spin-2 derivation.

4 The scalar timing sector

The most immediate coarse observable is bond occupation. Let

$$n_{ij} = \sum_{m=0}^{\chi_{\max}-1} m \Pi_{ij}^{(m)}. \tag{12}$$

The scalar load at vertex i is

$$\Lambda_i = \sum_{j \in \partial i} \log\left(1 + \langle n_{ij} \rangle\right). \tag{13}$$

The clock law is

$$\omega_{\mathrm{eff},i} = \frac{\omega}{1 + \alpha \Lambda_i}, \tag{14}$$

so that in weak load

$$\frac{\delta \omega_i}{\omega} \approx -\alpha \Lambda_i. \tag{15}$$

This is a valid timing sector. It is also not gravity. A scalar field can slow clocks, but it has no intrinsic transverse-traceless helicity-2 content.

4.1 Weighted graph closure

From bond occupation or dressed couplings one can extract a weighted adjacency matrix $W = (w_{ij})$ and Laplacian

$$L_w = D - W, \qquad D_{ii} = \sum_j w_{ij}. \tag{16}$$

A leading scalar closure is

$$(\kappa L_w + \mu^2 I)\Lambda = A\rho, \tag{17}$$

with dynamical completion

$$\partial_t \Lambda + \Gamma(\kappa L_w + \mu^2 I)\Lambda = \Gamma A \rho(t). \tag{18}$$

If $L_w \phi_n = \lambda_n \phi_n$, the relaxation rates should satisfy

$$\tau_n^{-1} = \Gamma(\kappa \lambda_n + \mu^2). \tag{19}$$

The same operator defines a diffusion return probability,

$$P_{\text{return}}(t) = \frac{1}{|V|} \operatorname{Tr} e^{-tL_w}, \tag{20}$$

and an effective dimension,

$$d_{\text{eff}}^{(\text{diff})}(t) = -2 \frac{d \log P_{\text{return}}(t)}{d \log t}. \tag{21}$$

A shell profile gives an independent estimator,

$$d_{\text{eff}}^{(\Lambda)}(r) = 2 - \frac{d \log |\Lambda_{\text{shell}}(r)|}{d \log r}. \tag{22}$$

The severe scalar closure test is therefore

$$d_{\text{eff}}^{(\text{diff})} \approx d_{\text{eff}}^{(\Lambda)} \approx 3 \tag{23}$$

in the same window where Eq. (14) fits best.

4.2 Why the scalar sector is not enough

For a scalar load Λ_i, the discrete difference on an oriented edge is

$$(d\Lambda)_{ij} = \Lambda_j - \Lambda_i. \tag{24}$$

Around any closed loop $\gamma = (i_0, i_1, \ldots, i_m = i_0)$,

$$\sum_{a=0}^{m-1} (d\Lambda)_{i_a i_{a+1}} = \sum_{a=0}^{m-1} (\Lambda_{i_{a+1}} - \Lambda_{i_a}) = 0. \tag{25}$$

The scalar timing sector is holonomy-free in this sense. It cannot by itself represent gravitomagnetic circulation, transverse-traceless gravitational waves, or the constraint structure of linearized general relativity.

This is not a defect of the scalar sector. It is its boundary.

5 The spin-2 target

The weak-field benchmark is the Pauli-Fierz theory of a massless spin-2 field. For a symmetric perturbation $h_{\mu\nu}$ with trace $h = \eta^{\mu\nu} h_{\mu\nu}$, the quadratic Lagrangian density can be written as

$$\mathcal{L}_{\text{PF}} = \frac{1}{2} \partial_\lambda h_{\mu\nu} \partial^\lambda h^{\mu\nu} - \partial_\mu h^{\mu\nu} \partial^\lambda h_{\lambda\nu} + \partial_\mu h^{\mu\nu} \partial_\nu h - \frac{1}{2} \partial_\lambda h \partial^\lambda h. \tag{26}$$

It is invariant under the linearized gauge transformation

$$h_{\mu\nu} \mapsto h_{\mu\nu} + \partial_\mu \xi_\nu + \partial_\nu \xi_\mu. \tag{27}$$

In the source-free transverse-traceless sector,

$$\partial_i h_{ij}^{\mathrm{TT}} = 0, \qquad \delta^{ij} h_{ij}^{\mathrm{TT}} = 0, \tag{28}$$

and the wave equation is

$$\left(\partial_t^2 - c_g^2 \nabla^2\right) h_{ij}^{\mathrm{TT}} = 0. \tag{29}$$

The TT projector in momentum space is

$$\Pi_{ij}(\mathbf{k}) = \delta_{ij} - \frac{k_i k_j}{k^2}, \tag{30}$$

$$P_{ij,kl}^{\mathrm{TT}}(\mathbf{k}) = \frac{1}{2}\left(\Pi_{ik}\Pi_{jl} + \Pi_{il}\Pi_{jk} - \Pi_{ij}\Pi_{kl}\right). \tag{31}$$

This projector is allowed as a diagnostic. It is not allowed as the generator of the claimed physics.

6 No-go constraints

The spin-2 target is constrained by the Weinberg-Witten obstruction. An ordinary Lorentz-invariant quantum field theory with a Lorentz-covariant conserved stress tensor cannot contain a composite massless spin greater than one particle that carries that stress tensor in the usual way. The point is not that emergent gravity is impossible. The point is that any successful route must explain why the theorem's assumptions fail microscopically while the correct Lorentz and gauge structure reappears in the infrared.

QATNU has a possible route because its microscopic theory is not a continuum Lorentz-covariant field theory. It begins from a finite graph, local update rules, and bond registers. But this is not enough. A preferred microscopic graph only avoids one assumption. The model must still demonstrate:

1. an emergent limiting tensor-wave speed with decreasing anisotropy at long wavelength;
2. an emergent gauge or constraint structure that removes nonphysical polarizations;
3. an effective conserved source that appears in the same infrared limit.

Without these, a tensor mode in QATNU would be a medium excitation, not gravitational spin-2.

7 Conditional derivation of the QATNU spin-2 sector

The natural tensor candidate is not built from the scalar load itself. It is built from anisotropic bond stress. On an approximately embedded or charted graph, let each edge e have unit tangent $\hat{t}_e$ and centered occupation

$$q_e = n_e - \langle n \rangle_{\mathcal{N}(x)}. \tag{32}$$

In a coarse cell $\mathcal{N}(x)$ define the symmetric traceless tensor

$$M_{ab}(x) = \sum_{e \in \mathcal{N}(x)} q_e \left(\hat{t}_{e,a} \hat{t}_{e,b} - \frac{1}{d} \delta_{ab} \right). \tag{33}$$

This is a substrate observable. It is not made transverse-traceless by definition.

7.1 Tensor susceptibility

Couple a weak source J_{ab} to M_{ab}:

$$H_J = -\sum_x J_{ab}(x) M_{ab}(x). \tag{34}$$

Equivalently, at the edge level, the bond energy is biased as

$$\Delta_B(e) \mapsto \Delta_B - J_{ab}(x_e) \left(\hat{t}_{e,a} \hat{t}_{e,b} - \frac{1}{d} \delta_{ab} \right). \tag{35}$$

Define the connected susceptibility

$$\chi_{ab,cd}(\Omega, \mathbf{k}) = \frac{\delta \langle M_{ab}(\Omega, \mathbf{k}) \rangle}{\delta J_{cd}(\Omega, \mathbf{k})}. \tag{36}$$

The inverse tensor kernel is

$$K_{ab,cd}(\Omega, \mathbf{k}) = \chi^{-1}_{ab,cd}(\Omega, \mathbf{k}). \tag{37}$$

In the tensor channel the most general low-momentum form is

$$K_2(\Omega, k) = Z_2 \Omega^2 - \rho_2 k^2 - m_2^2 + O(k^4, \Omega^4, k^2 \Omega^2). \tag{38}$$

A perturbative fast-sector elimination gives the schematic tensor mass

$$m_2^2(\lambda) = \Delta_B^{(2)} - \lambda^2 \Pi_2(0,0) + O(\lambda^4), \tag{39}$$

where $\Pi_2(0,0)$ is the zero-frequency tensor polarization of the matter/frustration sector. Hence the tensor critical point is

$$m_2^2 = 0 \qquad \Longleftrightarrow \qquad \lambda_{c,2}^2 = \frac{\Delta_B^{(2)}}{\Pi_2(0,0)} + O(\lambda_{c,2}^4). \tag{40}$$

At criticality, the pole condition $K_2 = 0$ gives

$$\Omega^2 = c_g^2 k^2 + O(k^4), \qquad c_g^2 = \frac{\rho_2}{Z_2}. \tag{41}$$

This supplies the primary Mercury number:

$$R_g = \frac{c_g}{c_{\text{clock}}} = \frac{1}{c_{\text{clock}}} \sqrt{\frac{\rho_2}{Z_2}}. \tag{42}$$

7.2 Ward residuals

A massless tensor in a medium is not automatically a graviton. The nonphysical modes must decouple. Define the residuals

$$\mathcal{W}_{\rm div}(k) = \frac{\|k^a \chi_{ab,cd}\|}{\|\chi_{ab,cd}\|}, \qquad \mathcal{W}_{\rm tr}(k) = \frac{\|\delta^{ab}\chi_{ab,cd}\|}{\|\chi_{ab,cd}\|}. \tag{43}$$

The infrared Ward condition is

$$\mathcal{W}_{\rm div}(k) \to 0, \qquad \mathcal{W}_{\rm tr}(k) \to 0 \qquad (k \to 0). \tag{44}$$

If Eq. (39) reaches criticality and Eq. (43) vanishes, the long-wavelength correlator takes the form

$$\langle M_{ab}(\Omega,\mathbf{k}) M_{cd}(-\Omega,-\mathbf{k})\rangle_c \sim \frac{Z\, P^{\rm TT}_{ab,cd}(\mathbf{k})}{\Omega^2 - c_g^2 k^2 + i0} + \text{analytic terms}. \tag{45}$$

The TT projector appears here as the consequence of the infrared constraints, not as the definition of M_{ab}.

7.3 Conditional theorem

Conditional spin-2 theorem. *Suppose a QATNU graph family admits a long-wavelength chart and a tensor bond-stress observable M_{ab} as in Eq. (33). Suppose its inverse tensor kernel has the infrared form Eq. (38), reaches the tensor critical point Eq. (40), and satisfies vanishing Ward residuals Eq. (43). Then the radiative long-wavelength sector has two transverse-traceless propagating degrees of freedom, wave speed $c_g^2 = \rho_2/Z_2$, and a local quadratic effective action equivalent to Pauli-Fierz up to field normalization.*

This theorem is conditional, but it is not empty. It names the quantities that must be computed: $\Pi_2(0,0)$, Z_2, ρ_2, the Ward residuals, and R_g.

8 Phase-0A: what has now been tested

The full tensor theorem requires large embedded graph families and finite-size scaling. The present hardware-accessible test is smaller and more conservative. Phase 0A asks whether the QATNU circuit and analysis pipeline can do three things at once:

1. run a hardware-readable QATNU circuit;
2. extract nonzero bond and scalar observables;
3. refuse a spin-2 interpretation for a topology that is structurally inadmissible.

The third point is the key. This is a structural admissibility gate, not a dynamical tensor susceptibility measurement. It asks whether the graph/circuit is even eligible to host an infrared transverse-traceless sector. A one-dimensional path fails before any hardware counts are examined. The hardware counts then test a different claim: whether the QATNU encoding and scalar/bond readout are live on a real backend.

9 Hardware-readable Phase-0A Hamiltonian

For real quantum hardware, the bond ladder is reduced to one bond qubit per edge, so $\chi_{\max} = 2$. The Hilbert space is

$$\mathcal{H}_{\mathrm{v0}} = (\mathbb{C}^2)^{\otimes(|V|+|E|)}. \tag{46}$$

The Hamiltonian is

$$H_{\mathrm{v0}} = H_{\mathrm{mat}} + H_{ZZ} + H_{\mathrm{bond}} + H_{\mathrm{prom}}, \tag{47}$$

with

$$H_{\mathrm{mat}} = \frac{\omega}{2} \sum_{i \in V} X_i, \tag{48}$$

$$H_{ZZ} = J_0 \sum_{\langle i,j \rangle \in E} Z_i Z_j, \tag{49}$$

$$H_{\mathrm{bond}} = \Delta_B \sum_{e \in E} n_e, \qquad n_e = \frac{I - Z_e}{2}, \tag{50}$$

$$H_{\mathrm{prom}} = \lambda \sum_{e=\langle i,j \rangle} F_{ij} X_e, \qquad F_{ij} = \frac{1 - Z_i Z_j}{2}. \tag{51}$$

The initial state is

$$|\psi_0\rangle = |+\rangle^{\otimes|V|} \otimes |0\rangle^{\otimes|E|}. \tag{52}$$

The parameters used here are

$$\omega = 1.0, \qquad J_0 = 0.12, \qquad \Delta_B = 0.55, \qquad \lambda = 0.35, \qquad T = 2.0. \tag{53}$$

Bond occupation, scalar load, load spread, raw anisotropy, and the many-body ground gap are computed after exact statevector evolution. For embedded two-dimensional controls the raw anisotropy diagnostic is

$$M_{ab}^{(2D)} = \sum_e q_e \left(\hat{t}_{e,a} \hat{t}_{e,b} - \frac{1}{2} \delta_{ab} \right), \qquad q_e = \langle n_e \rangle - \frac{1}{|E|} \sum_f \langle n_f \rangle. \tag{54}$$

This diagnostic is not a TT derivation. It only checks whether a small control graph exercises a nonzero centered directional tensor channel.

10 Structural spin-2 admissibility gate

The corrected Phase-0A predicate has four stages.

1. It computes the centered bond pattern q_e and raw anisotropy M_{ab}.
2. It records whether the tensor channel is actually exercised, i.e. whether q_e and M_{ab} are nonzero above numerical tolerance.
3. It checks whether the assigned chart is three-dimensional.
4. It checks whether the graph is large enough to have a meaningful low-k infrared window.

This predicate can return different kinds of no. A path can fail because the tensor channel is

not exercised and because the chart has rank one. A chorded diamond or bowtie can exercise a nonzero two-dimensional anisotropy tensor and still fail because a two-dimensional single-cell graph has no three-dimensional helicity-2 infrared limit. This is better than a hardcoded failure label: it records why the graph is rejected.

11 Phase-0A exact-simulation controls

Table 2 gives the corrected exact statevector controls. The diamond control is the chorded plaquette, K_4 minus one edge, and is no longer identical to the cycle. The bowtie/kite control is included as a second nonzero two-dimensional tensor-channel negative control.

Graph	Edges	Qubits	$\langle n_e \rangle$	$\Delta\Lambda$	$\sum \lvert q_e \rvert$	Chart rank	Phase-0A result
Path	3	7	0.12878	0.12278	0.00371	1	null tensor channel
Cycle	4	8	0.13015	0.00000	0.00000	2	null tensor channel
Star	3	7	0.13050	0.24532	0.00000	2	null tensor channel
Diamond	5	9	0.12800	0.11933	0.00251	2	non-3D tensor control
Bowtie	5	9	0.12800	0.11933	0.00251	2	non-3D tensor control

Table 2: Corrected Phase-0A exact controls. Path, cycle, and star do not exercise a usable tensor channel. Chorded diamond and bowtie exercise a nonzero two-dimensional tensor channel but still fail the spin-2 admissibility gate because they are not three-dimensional infrared substrates.

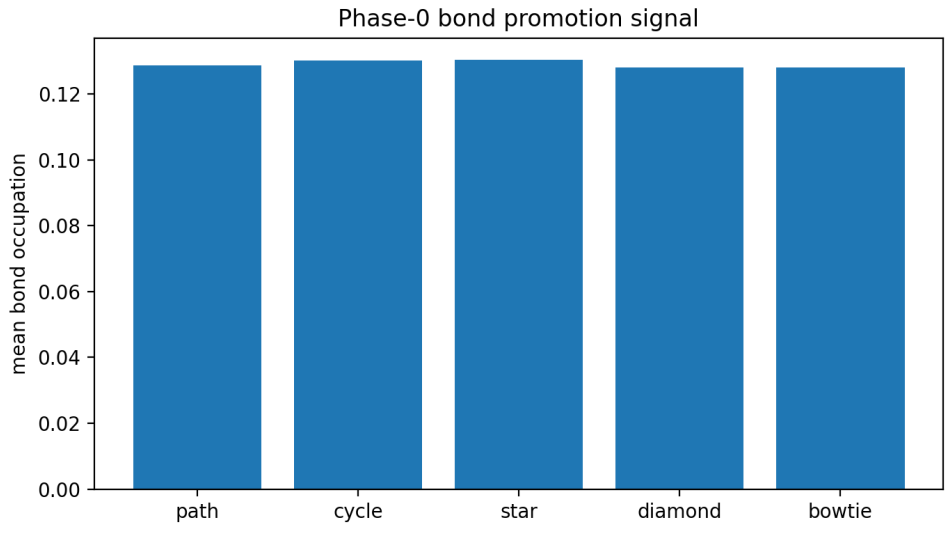

Figure 1: Mean bond occupation after exact evolution. The bond/scalar sector is active in all controls.

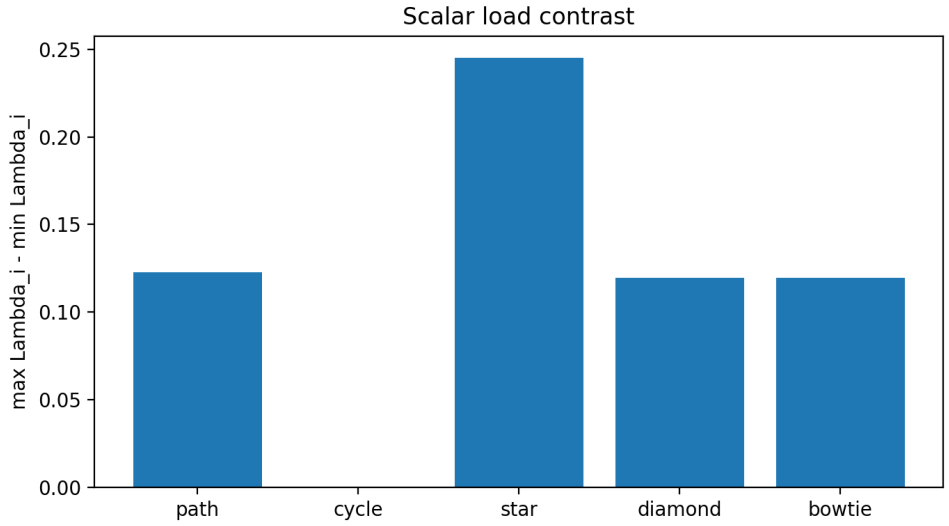

Figure 2: Scalar load spread for the corrected exact controls. The path, star, chorded diamond, and bowtie generate nonuniform scalar load; the symmetric cycle is uniform under the chosen state and parameters.

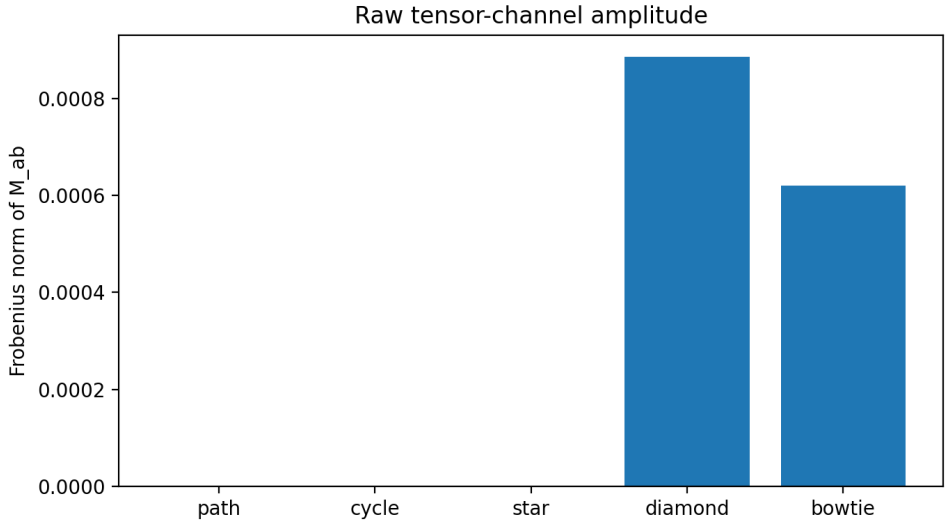

Figure 3: Raw anisotropy Frobenius norm. The nonzero values for chorded diamond and bowtie show that the tensor-channel diagnostic is being exercised, but these controls remain structurally inadmissible for spin-2 gravity.

12 Real-QPU Phase-0A hardware batch

The completed Phase-0A hardware batch ran four negative-control circuits on IBM Quantum hardware. The circuits are hardware-readable QATNU encodings with one matter qubit per vertex and one bond qubit per edge. The purpose of the batch is not to discover gravity. It is to test

whether active scalar/bond signals survive real backend execution while the spin-2 admissibility gate continues to refuse inadmissible topologies.

Topology	Backend	Job ID	Circuit role	$\langle n_e \rangle_{\mathrm{IBM}}$	Verdict
Path	`ibm_marrakesh`	`d8b6och59p8c73bm3fog`	1D null tensor channel	0.196833	fail expected
Star	`ibm_fez`	`d8b6tfke6kns73aalim0`	hub/no smooth IR limit	0.181000	fail expected
Cycle	`ibm_marrakesh`	`d8b6vm4e6kns73aallj0`	finite ring/no 3D IR	0.191875	fail expected
Diamond	`ibm_marrakesh`	`d8b6vrce6kns73aalls0`	finite plaquette control	0.156875	fail expected

Table 3: Completed IBM Phase-0A hardware batch. All four controls produced nonzero mean bond occupation and all four retained the expected negative-control verdict. The batch validates hardware execution and claim discipline, not a positive spin-2 sector.

The batch result can be summarized as

```
IBM Phase-0A batch:
active scalar/bond signal: yes, all four controls
spin-2 admissible substrate: no, all four controls
verdict: FAIL_EXPECTED_NEGATIVE_CONTROL, all four controls
```

For the path run, raw counts were also analyzed in detail. The path circuit used seven qubits, ran on `ibm_marrakesh`, completed with 2000 shots and 118 distinct measured bitstrings, and produced

$$\langle n_e \rangle_{\mathrm{IBM}} = [0.1905,\ 0.1585,\ 0.2415], \tag{55}$$

$$\Lambda_{\mathrm{IBM}} = [0.17437,\ 0.32150,\ 0.36345,\ 0.21632]. \tag{56}$$

The mean bond occupation was 0.19683 and the scalar-load spread was 0.18907. Thus the hardware circuit was not dead: it produced active bond occupation and nonuniform scalar load. The structural spin-2 gate still rejected it because a path has no transverse-traceless helicity-2 sector.

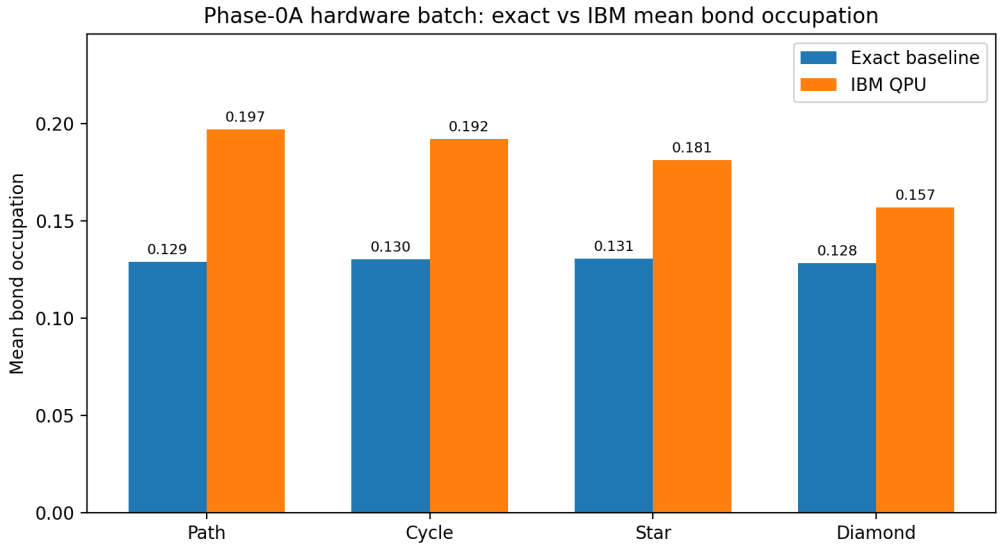

Figure 4: Exact baseline versus IBM mean bond occupation for the completed Phase-0A hardware batch. The IBM runs show active bond occupation across all four controls, while every topology remains structurally inadmissible for spin-2.

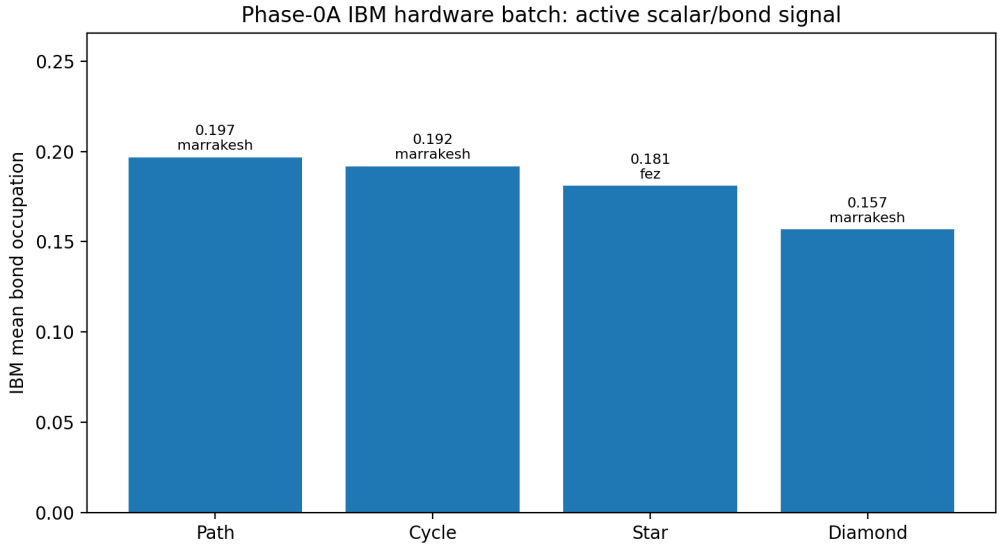

Figure 5: IBM-only mean bond occupation for the Phase-0A batch. Backend variation and hardware noise are treated as execution effects, not as new physics. The scientific result is the survival of the claim boundary.

12.1 Kingston backend-consistent Phase-0A rerun

After the first mixed-backend Phase-0A batch, the same four controls were rerun on `ibm_kingston`. This rerun is a backend-consistency check, not a new physics regime. It asks whether the Mercury-rung refusal layer survives when all four controls are executed on the same backend, with the same

shot count and result extractor.

Topology	Backend	Job ID	Shots	$\langle n_e \rangle$	$\Delta\Lambda$	Verdict
Path	ibm_kingston	d8b79k5mdd1s73b95ncg	2000	0.170500	0.171082	fail expected
Star	ibm_kingston	d8b79qh59p8c73bm46jg	2000	0.203500	0.398831	fail expected
Cycle	ibm_kingston	d8b7a1159p8c73bm46qg	2000	0.181500	0.076928	fail expected
Diamond	ibm_kingston	d8b7a6ce6kns73aam35g	2000	0.158500	0.025118	fail expected

Table 4: Backend-consistent Kingston rerun of the Phase-0A controls. The numeric observables shift relative to the first mixed-backend batch, but the verdict structure is unchanged: every control has active bond/scalar readout and every control remains structurally inadmissible for a spin-2 claim.

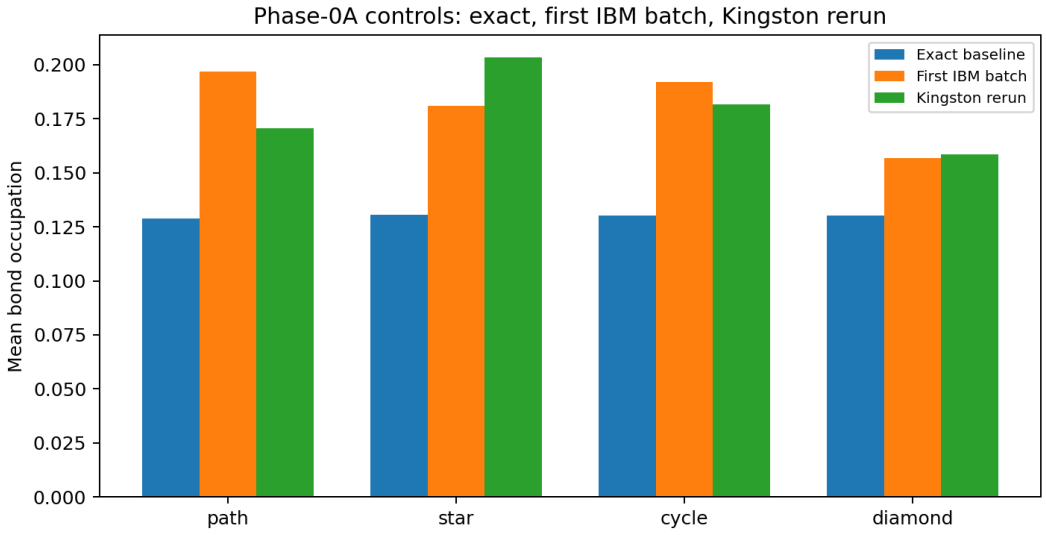

Figure 6: Phase-0A mean bond occupation across exact baselines, the first IBM batch, and the Kingston rerun. Backend-to-backend variation is visible, but the claim boundary is invariant.

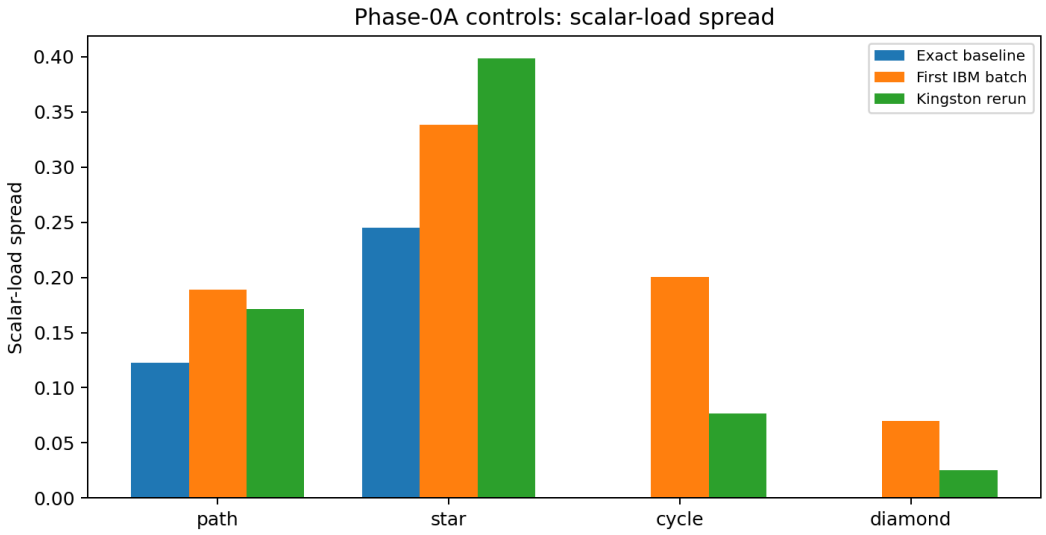

Figure 7: Phase-0A scalar-load spread across exact baselines, the first IBM batch, and the Kingston rerun. The hardware spread is device-sensitive; it is not interpreted as a tensor signal.

The Kingston rerun preserves the essential Phase-0A result:

```
active scalar/bond signal: yes, all controls
spin-2 admissible substrate: no, all controls
verdict: FAIL_EXPECTED_NEGATIVE_CONTROL, all controls
```

12.2 Interpretation of the hardware batch

The correct result statement is:

> *Four QATNU Phase-0A negative-control circuits ran on real IBM Quantum hardware. Each produced nonzero bond/scalar signal. Each was rejected by the structural spin-2 admissibility gate. The batch therefore validates the hardware-readable encoding, the bond/scalar readout, and the claim boundary, but it does not measure a dynamical tensor susceptibility or provide positive spin-2 evidence.*

This strengthens the Mercury-rung discipline. Before asking whether QATNU can derive a spin-2 sector, the protocol must show that it does not automatically hallucinate spin-2 from promoted-bond activity. The completed hardware batch provides that "no" on real devices.

13 Phase-1A: exploratory embedded graph-family hardware candidate tests

The Phase-1A graph-family pass extends the isolated hardware-prep package beyond the original Phase-0A negative controls. It still uses the hardware-readable v0 encoding: one matter qubit per vertex, one bond qubit per edge, four Trotter steps, and total evolution time $T = 2.0$. It does not yet couple these circuits back into the full exact-diagonalization pipeline, and it does not yet extract a tensor susceptibility.

The purpose of Phase 1A is narrower: test whether larger, embedded graph families can be compiled, locally simulated, and run on real IBM hardware while preserving the Mercury-rung claim boundary. The role labels are therefore deliberately conservative.

Topology	Vertices	Edges	Qubits	Role
`house`	5	6	11	small non-1D motif negative control
`square_2x3`	6	7	13	smallest square-lattice candidate
`triangular_2x3`	6	9	15	smallest triangular-lattice candidate
`square_3x3`	9	12	21	square-lattice scaling smoke test; hardware smoke test after smaller circuits
`triangular_3x3`	9	16	25	larger angular-structure smoke test; Kingston hardware candidate

Table 5: Phase-1A graph families. The square and triangular entries are admissibility candidates, not spin-2 claims. The 21-qubit `square_3x3` and 25-qubit `triangular_3x3` circuits extend the smoke tests into matched same-backend graph-family pairs.

13.1 Phase-1A local Aer checks

Local Aer checks show active bond/scalar observables for each family. The verdict labels are claim-boundary labels: house remains a negative control, while the embedded square and triangular families are candidate substrates requiring finite-size scaling and tensor-channel analysis.

Topology	Shots	$\langle n_e \rangle$	$\Delta\Lambda$	Raw anis.	Verdict
`house`	2000	0.133000	0.145126	0.278450	fail expected
`square_2x3`	2000	0.135000	0.146430	0.489535	candidate, not claimed
`triangular_2x3`	2000	0.139611	0.275658	0.449204	candidate, not claimed
`square_3x3`	512	0.134115	0.241870	0.294628	candidate, not claimed

Table 6: Phase-1A local Aer checks. Nonzero bond occupation, scalar-load spread, and raw anisotropy appear across the family, but these are still scalar/bond and raw geometry diagnostics, not a spin-2 susceptibility.

13.2 Phase-1 IBM hardware runs

The first Phase-1A hardware pass submitted the graph-family circuits to IBM Quantum with 2000 shots each. The initial smaller circuits established that non-1D and embedded 2D candidates could run; the 21-qubit `square_3x3` circuit was then submitted as a larger square-family hardware smoke test. All submitted jobs used Qiskit Runtime on the IBM Quantum Platform channel.

Topology	Backend	Job ID	$\langle n_e \rangle$	$\Delta\Lambda$	Raw anis.	Verdict
`house`	`ibm_marrakesh`	`d8b75h6honmc73cjjql0`	0.188917	0.193522	0.518695	fail expected
`square_2x3`	`ibm_marrakesh`	`d8b75b159p8c73bm4180`	0.165929	0.262930	0.311503	candidate, not claimed
`triangular_2x3`	`ibm_marrakesh`	`d8b75m159p8c73bm41l0`	0.177611	0.627958	0.431851	candidate, not claimed
`square_3x3`	`ibm_marrakesh`	`d8b77kehonmc73cjjt8g`	0.171500	0.530356	0.457432	candidate, not claimed

Table 7: Phase-1A IBM hardware pass. The embedded 2D graph families run and return active observables, but the results are not spin-2 evidence. They establish hardware reach for candidate graph families and identify the next analysis burden.

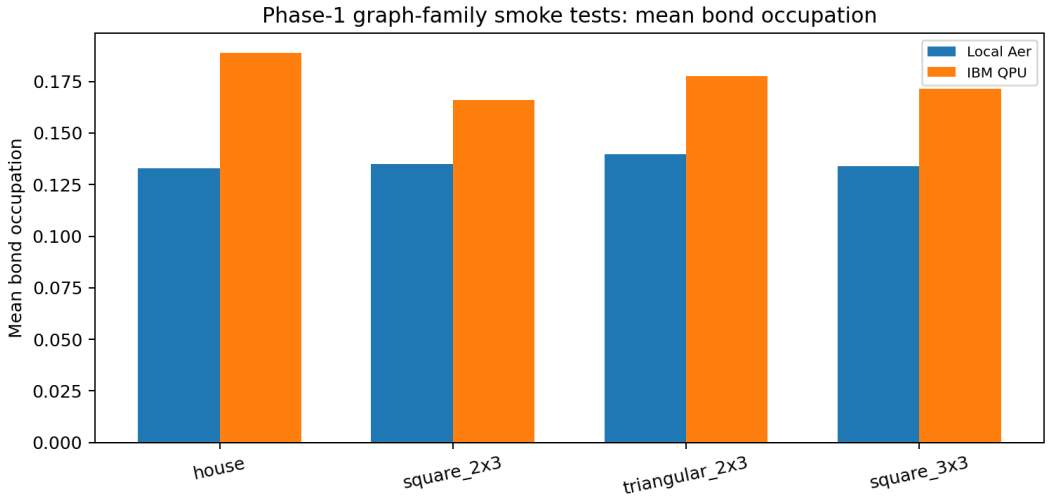

Figure 8: Phase-1A mean bond occupation. IBM values are higher than local Aer for the submitted circuits, consistent with hardware/transpilation/noise sensitivity rather than new physics.

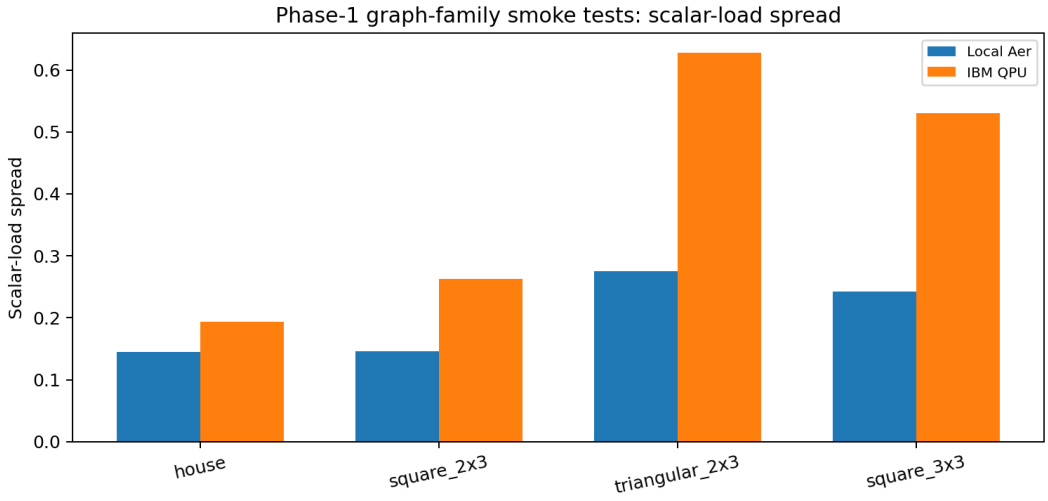

Figure 9: Phase-1A scalar-load spread. The triangular_2x3 hardware run shows a much larger scalar-load spread than local Aer. This should be treated as device- and layout-sensitive until repeated with calibration metadata and mitigation.

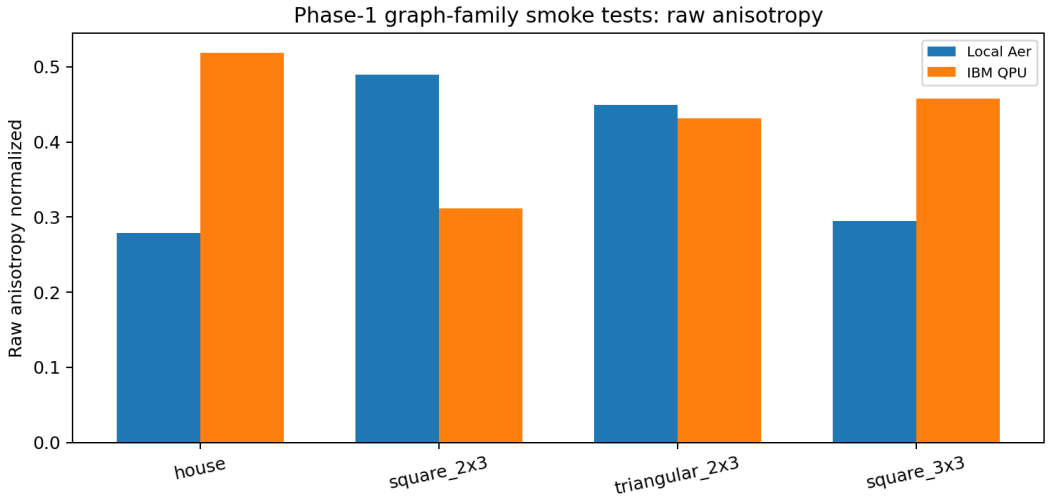

Figure 10: Phase-1A raw anisotropy diagnostic. Raw anisotropy is useful for checking that the graph-family readout is not trivial, but it is not a tensor susceptibility and cannot support a spin-2 claim on its own.

13.3 Kingston Phase-1A matched-family rerun

After the first Marrakesh pass, the Phase-1 house control, square family, and triangular family were checked on `ibm_kingston`. The house job reached `DONE` and was retrieved; no duplicate house job was submitted. This gives the cleanest Phase-1 comparison so far because `house`, `square_2x3`, `square_3x3`, `triangular_2x3`, and `triangular_3x3` share backend, shot count, Trotter settings, and result extractor. It is still not a finite-size spin-2 result: there is only one shot batch per size, no fixed physical layout policy, no mitigation, no backend-calibration metadata inside the result

object, and no tensor-channel observable. But it is a real hardware reach result: the same QPU path now executes the 21-qubit square and 25-qubit triangular candidates.

Topology	Backend	Job ID	Shots	$\langle n_e \rangle$	$\Delta\Lambda$	Raw anis.	Verdict
`house`	`ibm_kingston`	`d8b7bh5mdd1s73b95pg0`	2000	0.148167	0.176631	0.378877	fail expected
`square_2x3`	`ibm_kingston`	`d8b7acce6kns73aam3d0`	2000	0.179357	0.432255	0.485347	candidate, not claimed
`square_3x3`	`ibm_kingston`	`d8b7ahse6kns73aam3jg`	2000	0.178042	0.347370	0.066164	candidate, not claimed
`triangular_2x3`	`ibm_kingston`	`d8b7o4h59p8c73bm4ung`	2000	0.205333	0.645560	0.417111	candidate, not claimed
`triangular_3x3`	`ibm_kingston`	`d8b7ps4e6kns73aamsmg`	2000	0.256438	1.520218	0.133410	candidate, not claimed

Table 8: Kingston Phase-1A matched-family rerun. The house control remains a negative control on the same backend, while the square and triangular pairs execute and remain analyzable. None of these rows is promoted to spin-2 evidence; the result is hardware reach plus claim discipline, not a tensor pole.

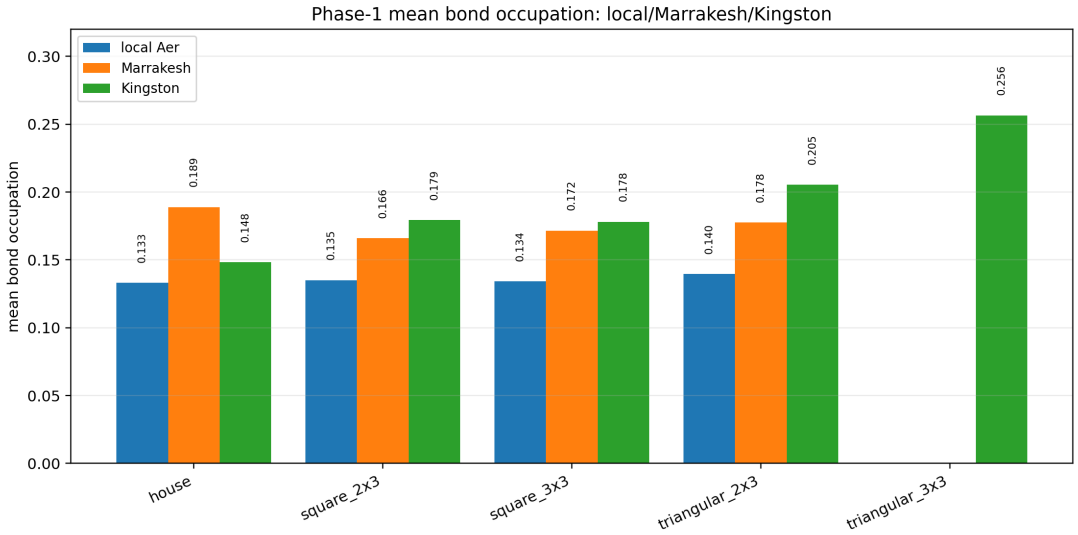

Figure 11: Phase-1A mean bond occupation across local Aer, the first Marrakesh pass, and the Kingston matched-family rerun. The Kingston data add cross-backend robustness, a same-backend house control, and same-backend square and triangular family reach.

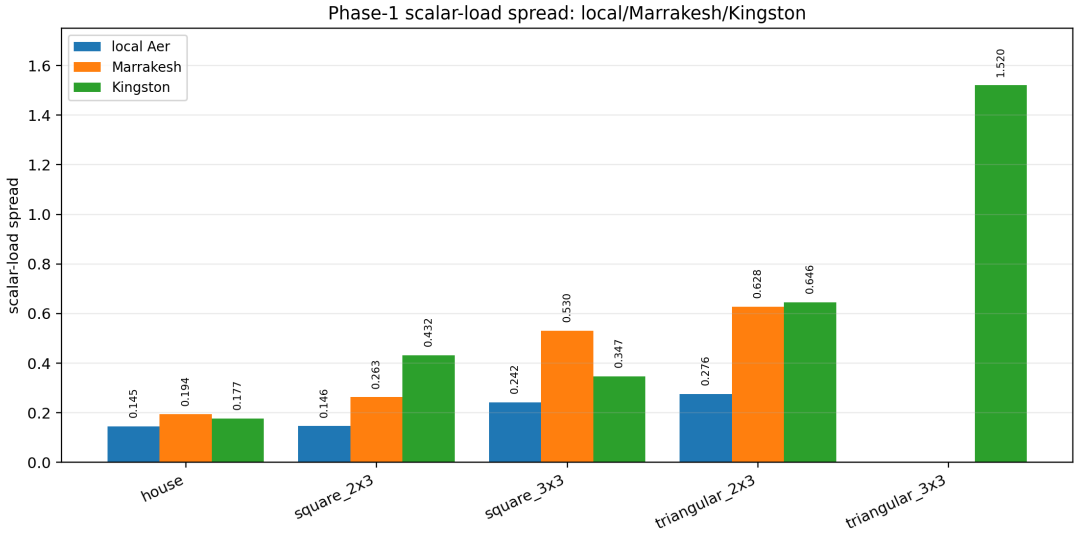

Figure 12: Phase-1A scalar-load spread for the house control, square family, and triangular family across local Aer, Marrakesh, and Kingston. The spread changes across hardware backends and is treated as device- and layout-sensitive, not as a tensor signal. The Kingston `triangular_3x3` point is the largest spread in the current hardware set.

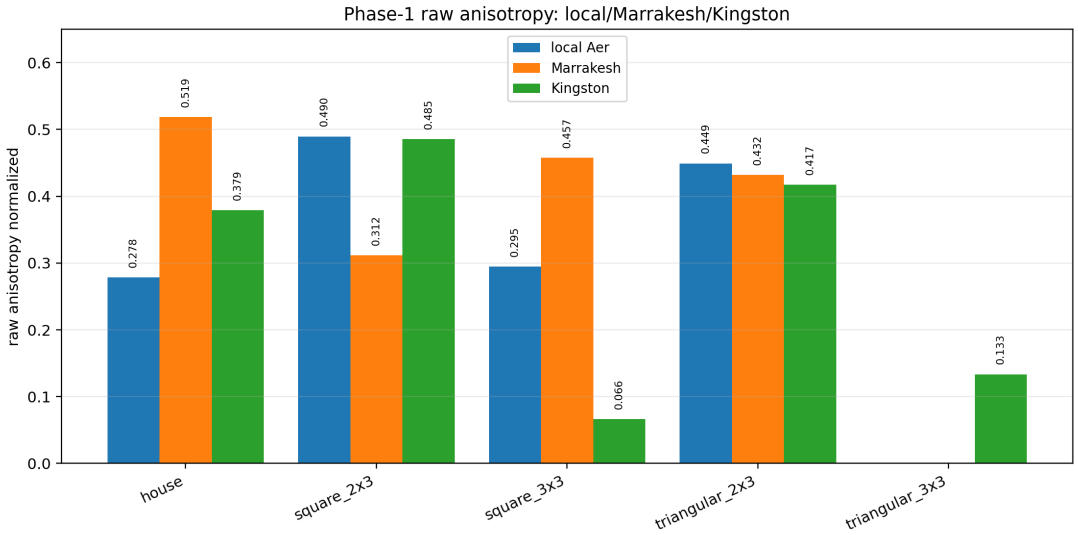

Figure 13: Phase-1A raw anisotropy across local Aer, Marrakesh, and Kingston. The triangular family shows the most interesting same-backend pattern: scalar-load spread grows from `triangular_2x3` to `triangular_3x3`, while normalized raw anisotropy decreases. This flags a follow-up target, not a spin-2 claim.

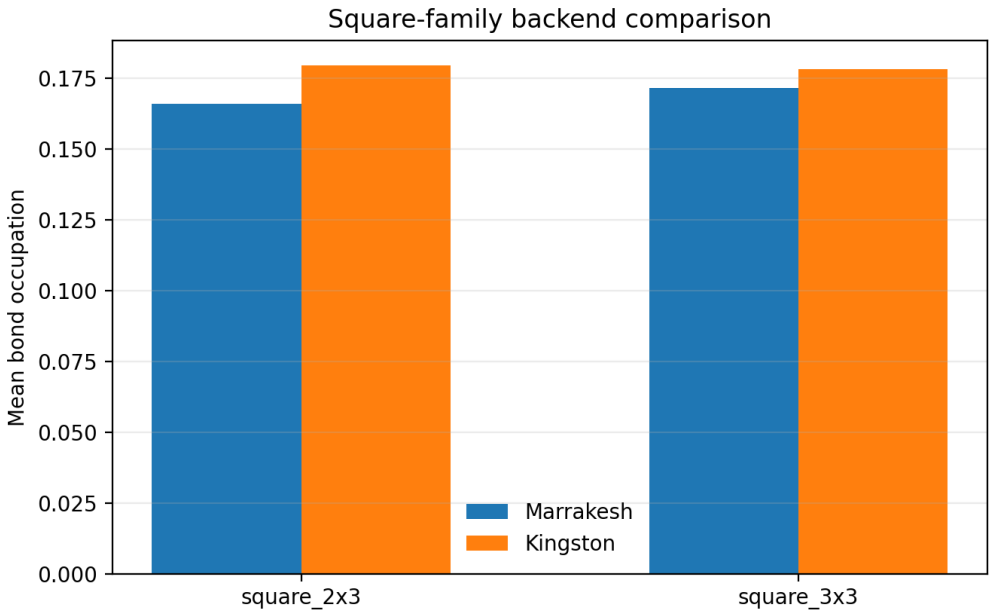

Figure 14: Square-family mean bond occupation for `square_2x3` and `square_3x3`. The near-equal Kingston means are an execution result only; they do not establish scaling or a tensor sector.

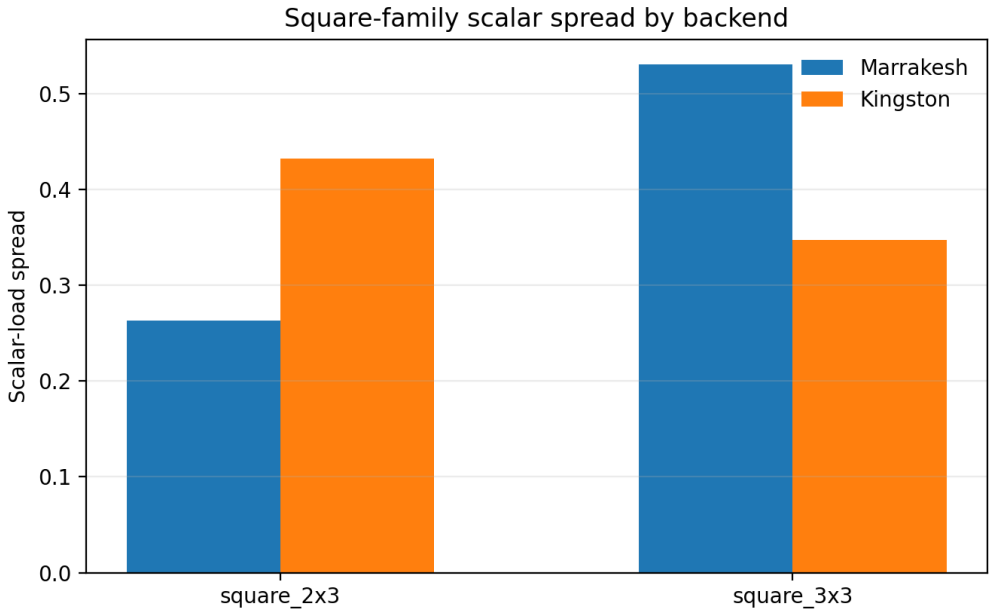

Figure 15: Square-family scalar-load spread. The square-family pair is now available on a single backend, with a house control in the same Kingston pass; repeated layouts and mitigation are still required before any finite-size conclusion can be drawn.

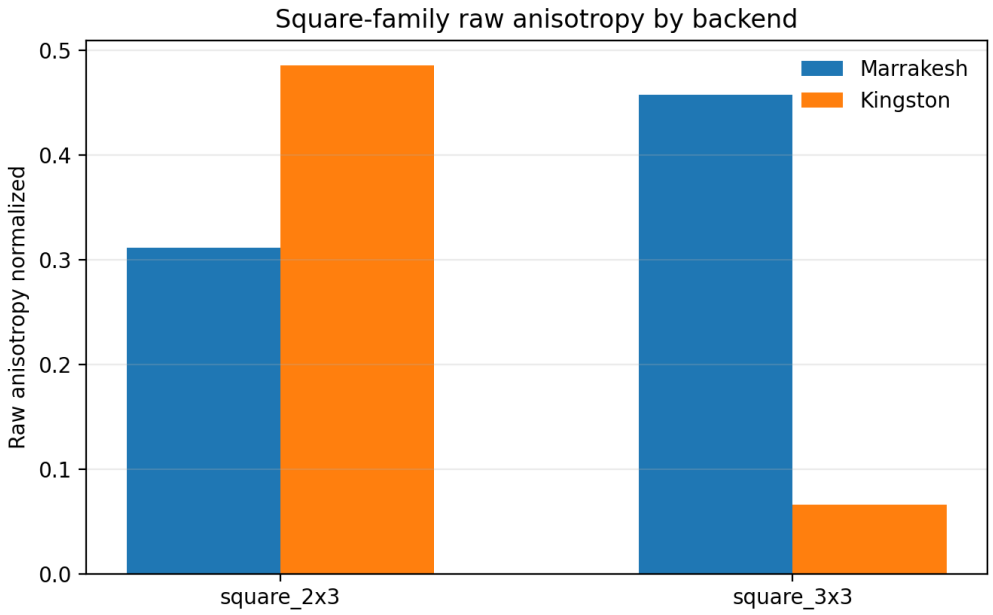

Figure 16: Square-family raw anisotropy. The large change in the Kingston raw anisotropy between `square_2x3` and `square_3x3` is not a spin-2 signal; it flags the need for fixed-layout repeats, tensor-channel observables, and scalar/vector impostor controls.

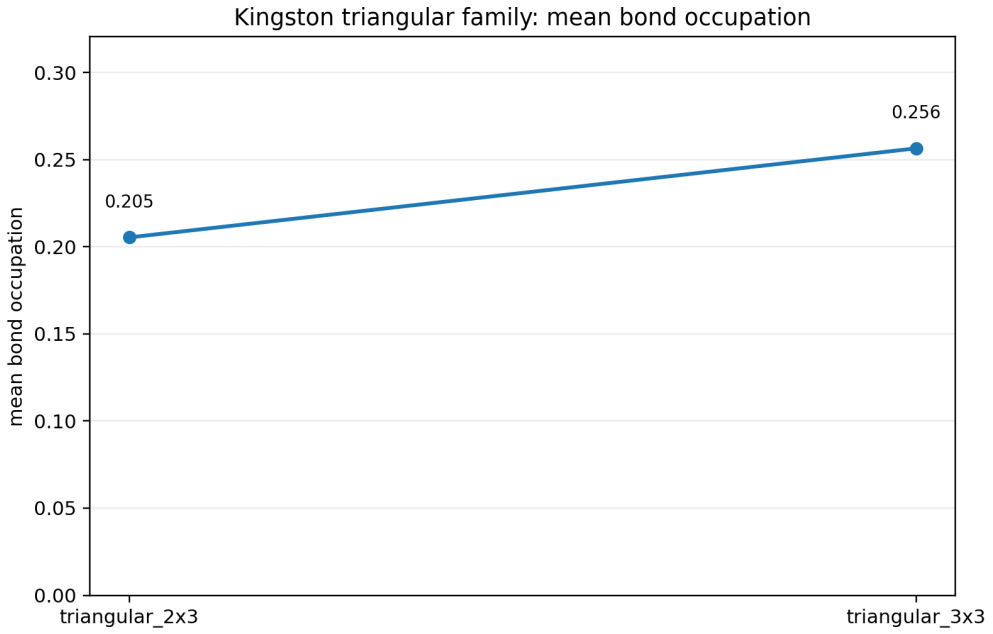

Figure 17: Kingston triangular-family mean bond occupation. The 25-qubit `triangular_3x3` run increases mean bond occupation relative to `triangular_2x3`.

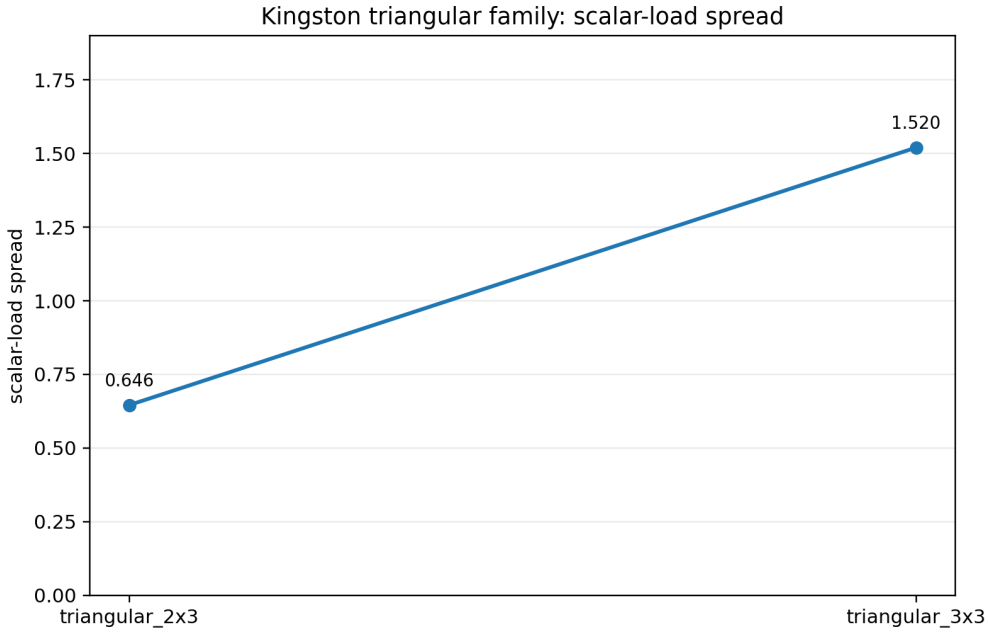

Figure 18: Kingston triangular-family scalar-load spread. The larger triangular graph produces a much larger scalar-load spread in this single hardware batch. The effect is hardware-facing and follow-up-worthy, but not yet interpretable as tensor scaling.

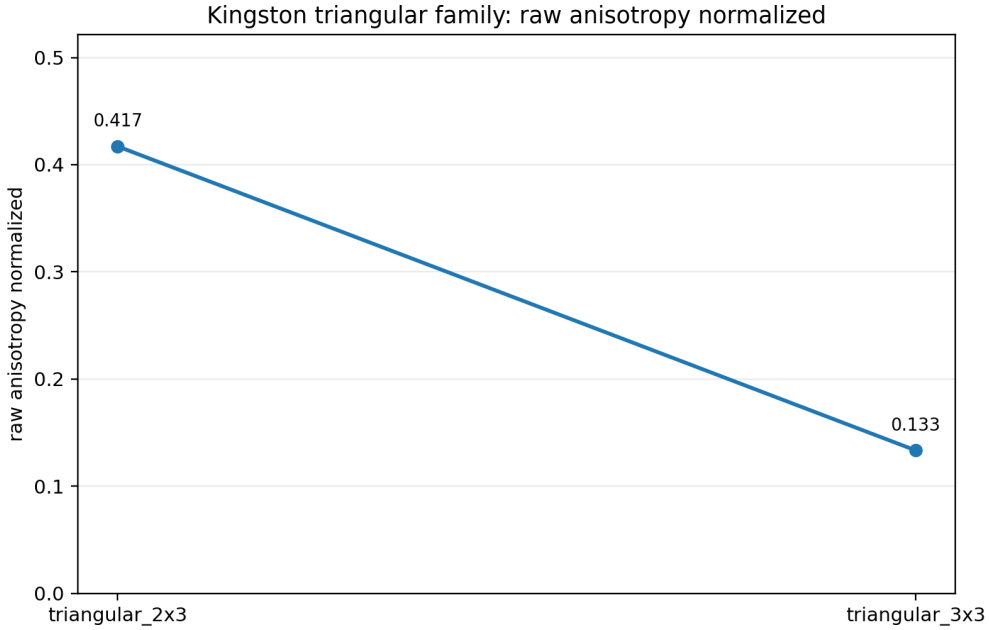

Figure 19: Kingston triangular-family raw anisotropy. Normalized raw anisotropy drops from `triangular_2x3` to `triangular_3x3`, even as scalar-load spread increases. This is the strongest "worth following up" pattern in the current hardware batch, but it remains a raw diagnostic rather than a spin-2 observable.

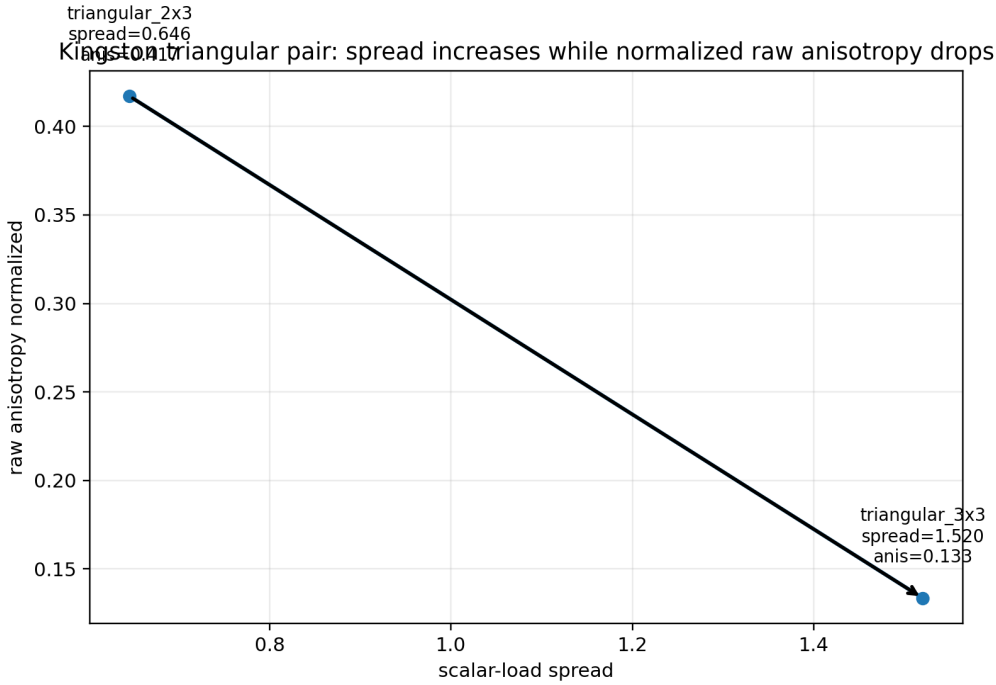

Figure 20: Same-backend triangular-pair summary. The move from `triangular_2x3` to `triangular_3x3` increases scalar-load spread and decreases normalized raw anisotropy. This is the clearest candidate-family follow-up signal, but a Mercury-rung claim still requires fixed-layout repeats, mitigation, tensor susceptibility, Ward-residual checks, and scalar/vector impostor rejection.

13.4 Hardened interpretation of the triangular pair

The Kingston triangular pair lands in the correct next lane, but it does not clear the Mercury rung. The hardened standard of this paper is

$$\text{scalar/bond signal} \neq \text{gravity}, \qquad \text{tensor-looking raw anisotropy} \neq \text{spin-2}. \tag{57}$$

Only a pre-registered tensor protocol with scaling, Ward tests, and R_g can pass.

The current triangular result supports four narrower statements. First, `triangular_3x3` is an embedded triangular patch, so the edge directions used for raw anisotropy are part of the circuit family rather than an arbitrary post-hoc chart. Second, it is part of a size pair, `triangular_2x3` $\rightarrow$ `triangular_3x3`. Third, it produced active bond marginals and structured scalar load on real IBM hardware. Fourth, it moved beyond Phase-0A null controls into an admissible embedded-lattice candidate family. These are engineering and protocol advances, not a spin-2 detection.

The run does not yet include the objects required by the tensor protocol:

- no pre-registered fit windows or acceptance thresholds;
- no tensor correlator $S_{ab,cd}(\mathbf{k}, t)$;
- no transverse-traceless fraction F_{TT};
- no effective tensor mass trend $m_g(L)$;
- no Ward/transversality residual $\mathcal{W}(k, L)$;

- no coarse-graining reproduction factor R_{TT};
- no Mercury number $R_g = c_g/c_{\mathrm{clock}}$;
- no fixed physical layout and calibration metadata inside the result record.

The most accurate label is therefore **Phase-1A exploratory hardware candidate result**. It is not "Mercury passed," and it is not "spin-2 observed." It is the first hardware-facing triangular-size pair that is interesting enough to justify a pre-registered tensor repeat.

A Phase-1B repeat should fix the backend, transpiler seed, initial layout, shot count, Trotter settings, readout-mitigation policy, and observables before execution. The first tensor observables should include a primitive orientation-cell tensor map M_{ab}, a structure factor $S_{ab,cd}(\mathbf{k})$, and the diagnostics

$$F_{\mathrm{TT}}(k,L) = \frac{\left\| P^{\mathrm{TT}} S(k,L) P^{\mathrm{TT}} \right\|_F}{\|S(k,L)\|_F}, \tag{58}$$

$$K_{\mathrm{TT}}(0,k;L) \simeq Z^{-1}\left(k^2 + m_g^2(L)\right), \tag{59}$$

$$\mathcal{W}(k,L) = \frac{\|k^a S_{ab,cd}(k,L)\|_F + \|\delta^{ab} S_{ab,cd}(k,L)\|_F}{\|S_{ab,cd}(k,L)\|_F}, \tag{60}$$

$$R_{\mathrm{TT}}(L \to bL) = \mathrm{corr}\Big(P^{\mathrm{TT}} M^{(L)},\, \mathcal{C}_b\left[P^{\mathrm{TT}} M^{(bL)}\right]\Big), \tag{61}$$

where $\mathcal{C}_b$ denotes the chosen coarse-graining map from the larger lattice back to the smaller comparison scale. These quantities are intentionally harder than raw anisotropy. They are the first checks that could distinguish a tensor channel from scalar/vector impostors and hardware layout artifacts.

13.5 Interpretation of Phase 1A

Phase 1A changes the status of the program, but only modestly. The hardware path is no longer restricted to four-qubit negative controls. It now reaches 11–25 qubit non-1D and embedded 2D graph circuits on IBM hardware, including a same-backend Kingston house control plus matched square and triangular graph-family pairs. This is genuine engineering progress. The triangular pair is now the strongest candidate-family target because `triangular_3x3` is the largest circuit in the package and it changes the scalar/bond observables in a nontrivial direction relative to `triangular_2x3`.

It is not yet physics evidence for an emergent graviton. A Mercury-rung spin-2 statement still requires at least four missing ingredients:

1. finite-size graph families with repeated same-backend and fixed-layout runs, now beginning from the observed `square_2x3` → `square_3x3` and `triangular_2x3` → `triangular_3x3` pairs;
2. a tensor-channel observable beyond scalar loads and raw two-dimensional anisotropy;
3. Ward-residual checks showing suppression of nonphysical polarizations in the infrared;
4. explicit scalar/vector impostor diagnostics, so ordinary bond activation cannot be mistaken for spin-2 emergence.

The correct Phase-1A result statement is therefore:

> *QATNU hardware-readable graph-family circuits now run beyond Phase-0A negative controls. The house motif remains a negative control. The square and triangular embedded*

lattices are tensor candidates only in the admissibility sense. The matched Kingston triangular pair is the strongest Phase-1A exploratory target because `triangular_3x3` *increases bond/scalar activation while reducing normalized raw anisotropy relative to* `triangular_2x3`*; nevertheless it is not spin-2 evidence until a pre-registered tensor protocol measures finite-size scaling,* F_{TT}*,* $m_g(L)$*, Ward residuals,* R_{TT}*, scalar/vector impostor rejection, and ultimately* R_g*.*

14 Reproducibility package

The code package contains exact-simulation utilities, OpenQASM circuits, count analyzers, and real-QPU runners. The Phase-0A exact result is reproduced by

```
cd qatnu_phase0_qpu
python src/run_phase0_exact.py
python src/generate_plots.py
```

The hardware batch is documented in

```
hardware/phase0_qpu/results/hardware_summary.md
```

with backend names, job ids, mean bond occupations, verdicts, claim boundary, and next Mercury-rung requirements. The Phase-1A graph-family pass, including the Kingston house control and matched square and triangular family reruns, is documented in `phase1_graph_families_summary.md` and `phase1_candidate_summary.json`. Individual raw IBM/local count JSON files are intentionally ignored in git; the reproducibility record is the summary plus the analyzer workflow.

A hardware count analysis has the form

```
python src/analyze_counts.py \
  --counts <counts.json> \
  --topology <path|star|cycle|diamond|house|square_2x3|square_3x3|triangular_2x3|triangular_3x
```

A hardware run command has the form

```
python src/run_ibm_qpu.py \
  --qasm qasm/qatnu_phase0_<topology>_N4.qasm \
  --backend <backend> \
  --shots 2000
```

The expected Phase-0A conclusion is not a positive spin-2 signal. It is an active scalar/bond signal plus structural refusal of the spin-2 claim. The expected Phase-1A candidate conclusion is likewise not a positive spin-2 signal: it is active scalar/bond and raw-anisotropy data on embedded graph families, with the verdict `PHASE1_TENSOR_CANDIDATE_NOT_CLAIMED` until finite-size tensor susceptibility and Ward-residual tests exist.

15 What remains to make this a positive spin-2 paper

The Phase-0A and Phase-1A results are gatekeepers and smoke tests, not the destination. A positive spin-2 paper must add the following.

1. **Finite-size graph families.** Repeat the matched families already reached in hardware, `square_2x3` → `square_3x3` and `triangular_2x3` → `triangular_3x3`, under fixed-layout and mitigation policies, then extend to larger sizes with a real low-k window.
2. **Tensor source response.** Add $J_{ab}M_{ab}$ perturbations and compute $\chi_{ab,cd}$.
3. **Critical scan.** Fit $m_2^2(\lambda) = \Delta_B^{(2)} - \lambda^2\Pi_2(0,0)$ and test whether a tensor critical point exists.
4. **Dispersion.** Extract Z_2, ρ_2, and $c_g^2 = \rho_2/Z_2$.
5. **Ward residuals.** Show that scalar/vector leakage and trace/divergence residuals decrease in the infrared.
6. **Mercury number.** Report $R_g = c_g/c_{\text{clock}}$ as a prediction, not a fit.

A null result at any step is informative. If m_2^2 never approaches zero, QATNU does not generate a massless tensor in that phase. If the Ward residuals do not vanish, the tensor is a medium mode, not spin-2 gravity. If R_g is not stable, the model has no emergent Lorentzian limit.

16 Claim registry

The current status is summarized in Table 9. The most important update is that real-QPU execution has moved in two stages: Phase-0A negative controls and Phase-1A embedded graph-family smoke tests. The claim remains narrow.

Table 9: Claim registry for the Mercury rung after Phase-0A and Phase-1 hardware runs.

Claim	Status	Evidence or next test
QATNU is a finite local quantum substrate	Implemented	ED and hardware-readable qubit encodings
Scalar load Λ_i is measurable	Implemented	Exact baselines, IBM Phase-0A batch, and Phase-1A graph families
Scalar timing can be represented by $\omega_{\text{eff}} = \omega/(1+\alpha\Lambda)$	Supported as effective law	Requires larger-system validation
Scalar timing is gravity	Rejected	Scalar sector is holonomy-free and lacks spin-2
Spin-2 target is Pauli-Fierz/TT	Standard benchmark	Known weak-field target
QATNU derives spin-2 unconditionally	Not claimed	Requires tensor criticality and Ward identities
Conditional spin-2 derivation	Formulated	Eqs. (39)–(45)
Negative-control exact simulations	Reported result	Corrected Phase-0A structural gate across path, cycle, star, diamond, bowtie

Claim	Status	Evidence or next test
Real-QPU Phase-0A execution	Reported result	IBM path, star, cycle, and diamond controls; all fail expected
Phase-1A graph-family execution	Reported result	IBM house, square_2x3, triangular_2x3, square_3x3, and triangular_3x3 runs; Phase-1A exploratory candidates not claimed
Path graph spin-2 claim	Rejected structurally	One-dimensional graph, null usable tensor channel, no TT sector
Positive spin-2 sector	Open	Requires admissible graph scaling and $\chi_{ab,cd}$

17 Discussion

The central scientific value of the Mercury rung is that it stops the program from treating every partial resemblance as success. Scalar timing is allowed to be real without being called gravity. Tensor projection is allowed as a diagnostic without being confused for derivation. Negative controls are treated as results rather than embarrassments.

The conditional spin-2 derivation is the bridge between the conceptual program and the next calculation. It says exactly what must happen: the bond-stress tensor susceptibility must reach a massless tensor critical point, the Ward residuals must vanish, and the resulting tensor wave speed must produce a stable R_g. If any of these fail, the gravitational claim narrows or stops.

The Phase-0A hardware batch is therefore not small in method, even though the graphs are small. It shows that the hardware-readable protocol can produce active scalar/bond observables on real QPUs while the structural gate still refuses false spin-2 interpretations across multiple controls. That is the first requirement for any credible emergence claim: the program must be able to say no before it is trusted when it says maybe.

The Phase-1A graph-family runs add the next boundary. They show that the hardware path can execute larger embedded families and produce active bond/scalar and raw-anisotropy data. The matched Kingston triangular pair is the most interesting hardware-facing result so far: the larger graph increases mean bond occupation and scalar-load spread while normalized raw anisotropy decreases. But the label is deliberately not `PASS`. It is `PHASE1_TENSOR_CANDIDATE_NOT_CLAIMED`. This is the correct intermediate status: the experiment has moved from negative controls into candidate graph families, while the spin-2 claim remains reserved for finite-size tensor susceptibility, Ward residuals, dispersion, and the Mercury number.

18 Conclusion

QATNU now has a complete Mercury-rung formulation, a four-control hardware-backed boundary result, a Kingston replication layer, and a Phase-1A embedded-graph hardware candidate test

reaching the 21-qubit square family and 25-qubit triangular family. The scalar timing sector is defined and bounded. The spin-2 target is stated without projection inflation. The conditional derivation identifies tensor criticality, Ward identities, and $R_g = c_g/c_{\text{clock}}$ as the quantities that decide whether the substrate can produce gravitational dynamics.

The first hardware result was Phase 0A: four QATNU negative-control circuits ran on IBM Quantum hardware. Path, star, cycle, and diamond/rhombus controls all produced active mean bond occupation and all returned `FAIL_EXPECTED_NEGATIVE_CONTROL`. The next hardware result is Phase 1A: house, square_2x3, triangular_2x3, square_3x3, and triangular_3x3 graph-family circuits also ran on IBM hardware, with Kingston reruns for Phase-0A controls and for matched Phase-1 square and triangular pairs. The house graph remains a negative control; the square and triangular lattices are candidates, not claims. These results are not evidence for a graviton, and they do not measure tensor susceptibility. They show that scalar/bond activity and raw anisotropy are not automatically promoted into a gravity claim, even after real hardware noise, finite shots, transpilation, backend execution, and backend replication.

The next paper should not repeat negative-control hardware runs unless they add a new admissibility class. It should move to Phase 1B: a pre-registered Kingston triangular repeat, beginning with `triangular_2x3` and `triangular_3x3`, with fixed backend, fixed transpiler seed, captured physical layout, calibration metadata, mitigation policy, and fixed observables. The first new layer should compute primitive M_{ab} maps, F_{TT}, $m_g(L)$, Ward residuals, and R_{TT} before any attempt at R_g. Only after that can tensor dispersion and the Mercury number R_g actually be tested.

A Hardware run prompt

> Run or reproduce the QATNU Mercury-rung hardware sequence. First reproduce the exact statevector baselines with `python src/run_phase0_exact.py`. Then consult `hardware/phase0_qpu/results/hardware_summary.md` and `hardware/phase0_qpu/results/phase1_graph_families_summary.md` for the completed IBM jobs. To repeat a Phase-0A run, submit the relevant `qasm/qatnu_phase0_<topology>_N4.qasm` circuit to IBM Quantum. To repeat a Phase-1A smoke test, submit `qasm/qatnu_phase1_<topology>.qasm`. Save raw bitstring counts with backend name, job id, shot count, layout/transpilation metadata, and calibration date if available, and analyze with `python src/analyze_counts.py --counts <counts.json> --topology <topology> --N <N>`. Report mean bond occupation, scalar-load spread, raw anisotropy, tensor-channel status, structural admissibility, and `spin2_phase0_result`. The expected Phase-0A conclusion is active scalar/bond signal plus structural spin-2 refusal. For Kingston replication, record the backend-consistent Phase-0A table and the Phase-1A Kingston comparisons `house`, `square_2x3` → `square_3x3`, and `triangular_2x3` → `triangular_3x3`. The expected Phase-1A conclusion is `PHASE1_TENSOR_CANDIDATE_NOT_CLAIMED` unless and until pre-registered finite-size tensor susceptibility, F_{TT}, $m_g(L)$, Ward-residual, R_{TT}, and scalar/vector-impostor tests are added.

B Notation

Symbol	Meaning	Status
H_{micro}	Microscopic QATNU Hamiltonian	Finite graph substrate
Λ_i	Internal scalar load at vertex i	Measured in exact, IBM Phase-0A, and Phase-1A graph-family runs where counts are available
L_w	Weighted graph Laplacian	Scalar closure operator
M_{ab}	Bond-stress tensor candidate	Spin-2 observable candidate
$K_{ab,cd}$	Inverse tensor susceptibility	Future tensor test
m_2^2	Tensor mass parameter	Criticality target
Z_2, ρ_2	Tensor kinetic and stiffness coefficients	Determine c_g
c_g	Emergent tensor-wave speed	Future output
c_{clock}	Scalar clock-sector speed	Future output
R_g	c_g/c_{clock}	Primary Mercury number
$\mathcal{W}_{\text{div}}, \mathcal{W}_{\text{tr}}$	Ward residuals	Gauge/constraint diagnostics

C References

1. A. Einstein, "Die Feldgleichungen der Gravitation," Sitzungsberichte der Preussischen Akademie der Wissenschaften, 1915.
2. M. Fierz and W. Pauli, "On relativistic wave equations for particles of arbitrary spin in an electromagnetic field," Proceedings of the Royal Society A, 1939.
3. S. Weinberg, "Photons and gravitons in perturbation theory: Derivation of Maxwell's and Einstein's equations," Physical Review, 1965.
4. S. Weinberg and E. Witten, "Limits on massless particles," Physics Letters B, 1980.
5. E. H. Lieb and D. W. Robinson, "The finite group velocity of quantum spin systems," Communications in Mathematical Physics, 1972.
6. C. W. Misner, K. S. Thorne, and J. A. Wheeler, *Gravitation*. W. H. Freeman, 1973.
7. R. M. Wald, *General Relativity*. University of Chicago Press, 1984.
8. IBM Quantum Documentation, Qiskit Runtime and IBM Quantum Platform documentation, accessed May 2026. https://quantum.cloud.ibm.com/docs
9. PyQuafu Documentation, "Welcome to PyQuafu's documentation," accessed May 2026. https://scq-cloud.github.io/
10. PyQuafu Documentation, "User guide," accessed May 2026. https://scq-cloud.github.io/0.2.x/index.html