The Mercury Rung

Full-Ladder QMC, Cube Hardware Validation, and the QATNU Spin-2 Null Test

Joshua Farrow
Independent Research

May 2026

Status. This paper gives the current Mercury-rung formulation for QATNU after the Phase-0A and Phase-1A hardware runs, the cubic-graph correction, the full-ladder stoquasticity theorem, the full-ladder worldline-QMC tensor null, and the first Phase-2A cube hardware source-response runs. The revision is stricter than the hardware-facing versions. It places QATNU inside the known lattice-graviton and rank-2 tensor-gauge literature; adds a binding spatial-dimension gate; demotes all two-dimensional square and triangular circuits to hardware pipeline validation; proves that the positive QATNU ladder is gauge-stoquastic; and reports a sign-free, full-ladder QMC test with frustration-gated bond-register histories and the rank-2 observable M_{ab}. In the natural parameter neighborhood scanned here, the tensor susceptibility does not grow with system size at converged $\chi_{\max}$; the channel is dominated by single-edge displaced-oscillator variance rather than a collective rank-2 critical mode. The new hardware update is narrower: the 20-qubit 2x2x2 cube ran on `ibm_kingston`, including a baseline and matched x-axis source pair. The runs validate the three-dimensional v0 encoding, source metadata, tensor-summary extraction, and an engineering optimization path; they do not establish a controlled tensor susceptibility, dispersion, Ward identity, or spin-2 sector. The next falsifiable modification is still mechanism-level: add an inter-edge bond–bond channel or spatially correlated frustration and rerun the $\chi_{\max}$-converged, L-scaled QMC test before asking hardware for thermodynamic evidence.

Abstract

Emergent-gravity programs often succeed first at producing scalar or geometric analogies: clock slowing, graph distances, effective dimensions, raw anisotropy maps, or projection diagnostics. 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, while also recovering the correct dimensional and Lorentzian limits.

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 revised contribution is now fivefold. First, it situates QATNU inside existing lattice-graviton and rank-2 tensor-gauge constructions. Prior work by Gu and Wen, Xu and Hořava, Xu, Pretko, and the fracton-elasticity literature already shows that qubit, bosonic, and tensor-gauge substrates can produce graviton-like helicity-2 or rank-2 tensor sectors, often with Lifshitz or fractonic rather than relativistic dynamics. QATNU is therefore not novel merely as a lattice route to tensor modes. Its distinct contribution is a hardware-readable bond-register substrate and a staged protocol that keeps scalar/bond activation, raw tensor-looking anisotropy, Lifshitz tensor phases, relaxational tensor response, and GR-like spin-2 recovery logically separate.

Second, the paper imposes a spatial-dimension gate. In D spacetime dimensions a massless graviton has $D(D-3)/2$ propagating polarizations. Thus $3+1$ dimensions give two propagating helicity-2 modes, while $2+1$ dimensions give none. This reclassifies the square and triangular two-dimensional hardware runs, including the matched Kingston triangular pair, as pipeline-validation and embedded-lattice smoke tests, not spin-2 admissible substrates for the target of this paper. They show active bond/scalar observables on real hardware, but they cannot host the desired $3+1$-dimensional radiative graviton even in principle.

Third, the conditional spin-2 derivation is formulated so that propagation itself is measured rather than assumed. Because the scalar timing sector is relaxational, the leading tensor null hypothesis is also relaxational, with an inverse response of the form $K_2^{\mathrm{diff}} \sim -iZ_D\Omega + D_T k^2 + m_2^2$. Only if the tensor response is propagating does it make sense to fit a dynamical exponent. The relativistic kernel $K_2 \sim Z_2\Omega^2 - \rho_2 k^2 - m_2^2$ is only the $z=1$ branch, while Lifshitz branches such as $K_2 \sim Z_z\Omega^2 - \rho_z k^{2z} - m_2^2$ are named partial-success/failure classes motivated by prior lattice-graviton constructions.

Fourth, the paper proves a sign-structure theorem for the full QATNU ladder. In the product basis, all diagonal terms are sign-trivial. The off-diagonal matter flip $(\omega/2)\sum_i X_i$ and the frustration-gated bond promotion $\lambda\sum_e F_e(R_e + R_e^\dagger)$ become non-positive after the diagonal gauge $U = (\prod_i Z_i) \otimes (\prod_e P_e)$ with $P_e = \mathrm{diag}((-1)^m)$. For $\omega, \lambda \geq 0$ and canonical positive ladder amplitudes, the gauged Hamiltonian is stoquastic for all $\chi_{\max}$. Hence the equilibrium imaginary-time path integral has non-negative weights. The paper then reports the full-ladder QMC instrument that this theorem permits: matter worldlines, frustration-gated bond-register worldlines, the rank-2 tensor observable M_{ab}, source-response measurements, Matsubara relaxation-versus-propagation diagnostics, and validation against exact diagonalization on graphs small enough to diagonalize.

Fifth, the roadmap is sharpened by a negative result. The full-ladder QMC reproduces ED diagonal observables at the percent/Trotter level for $\chi_{\max} = 3$; the source response on the 2x2x2 cube is linear; and the Matsubara fork is implemented, though too small/noisy here to classify propagation. The decisive finding is that at converged $\chi_{\max}$ the tensor susceptibility does not grow with system size in the scanned natural neighborhood ($\omega = 1, J_0 = 0.12, \Delta_B \in [0.05, 0.55], \lambda \in [0.35, 1.3], \beta \in [2, 8], L \leq 3$). At low Δ_B the susceptibility enhancement is a displaced-oscillator, single-edge variance effect, not a collective tensor critical mode. Thus the conditional spin-2 theorem faces a concrete antecedent failure for the current substrate family. The next Mercury test is not larger hardware alone; it is to add an inter-edge bond–bond channel or correlated frustration and rerun the identical $\chi_{\max}$-converged finite-size protocol.

Contents

1 The problem **4**

2 Related lattice-graviton and tensor-gauge constructions **4**

3 The Mercury criterion **5**

4 The QATNU substrate **5**
4.1 Implemented exact-diagonalization form 6

5 The scalar timing sector **7**
5.1 Weighted graph closure 7
5.2 Why the scalar sector is not enough 8

6 The spin-2 target **8**

7 No-go constraints **9**

8 Spatial-dimension gate **9**

9 Conditional derivation of the QATNU spin-2 sector **11**
9.1 Tensor susceptibility 11
9.2 Ward residuals 13
9.3 Conditional theorem 13
9.4 Massless limit and source-coupling warning 13

10 Stoquasticity and the native-3D QMC route **13**
10.1 Full-ladder gauge-stoquasticity theorem 14
10.2 What sign-free does and does not mean 15
10.3 Imaginary-time mapping and the tensor target 15
10.4 Prototype QMC instrument 16

11 Full-ladder QMC: bond-register worldlines and a converged null **16**
11.1 Imaginary-time mapping and gated bond worldlines 17
11.2 Validation against exact diagonalization 17
11.3 Tensor susceptibility, source response, and the Matsubara fork 18
11.4 $\chi_{\max}$ convergence and the displaced-oscillator structure 19
11.5 Finite-size scaling at converged truncation 20
11.6 Interpretation: a converged null for the nominal substrate family 21

12 Phase-0A: what has now been tested **21**

13 Hardware-readable Phase-0A Hamiltonian **21**

14 Structural spin-2 admissibility gate **22**

15 Phase-0A exact-simulation controls **23**

16 Real-QPU Phase-0A hardware batch **24**
16.1 Kingston backend-consistent Phase-0A rerun . . . 26
16.2 Interpretation of the hardware batch . . . 28

17 Phase-1A: embedded graph-family hardware pipeline validation **28**
17.1 Phase-1A local Aer checks . . . 29
17.2 Phase-1 IBM hardware runs . . . 29
17.3 Kingston Phase-1A matched-family rerun . . . 31
17.4 Hardened interpretation of the triangular pair . . . 37
17.5 Interpretation of Phase 1A . . . 38

18 Phase-2A: first 3D cube hardware source-response runs **39**
18.1 Trotter-2 engineering check . . . 40
18.2 Compact-source local checks . . . 41

19 Phase-2D mechanism triage: bond-bond QMC hook **41**

20 Reproducibility package **43**

21 What remains to make this a positive spin-2 paper **44**

22 Claim registry **45**

23 Discussion **48**

24 Conclusion **49**

A Hardware run prompt **50**

B Native-3D QMC next-test prompt **50**

C Notation **50**

D References **51**

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 the hardware negative controls, the Phase-2A cube source-response execution, and the full-ladder QMC tensor null.

2 Related lattice-graviton and tensor-gauge constructions

The target of this paper is not unprecedented. The relevant prior art is not merely general emergent-gravity philosophy, but explicit lattice and tensor-gauge constructions of graviton-like sectors.

Gu and Wen constructed qubit/quantum-spin models in which linearized diffeomorphism gauge symmetry and gapless helicity ± 2 excitations can emerge at low energy. In their controlled L-type model, the helicity-± 2 modes are gapless but have a nonrelativistic dispersion proportional to k^3; a second N-type route was argued to permit linear dispersion, but without the same controlled status. Xu and Hořava obtained emergent gravity at a Gaussian $z = 3$ Lifshitz fixed point from a lattice Bose liquid, with stability tied to gauge symmetry and self-duality. Xu's algebraic Bose liquid and the later rank-2 tensor-gauge/fracton literature likewise show that symmetric tensor fields and graviton-like excitations can arise in nonfundamental media, but often with restricted mobility, Lifshitz scaling, or non-GR source behavior.

This lineage changes the QATNU claim. QATNU is not claiming to be the first lattice or qubit path to helicity-2-like excitations. The open frontier is stricter: can a hardware-readable bond-register substrate produce a three-dimensional, Lorentz-recovering, gauge-protected spin-2 sector with the source-coupling structure of weak-field gravity, rather than a Lifshitz or fracton-like tensor medium?

The lessons for the Mercury rung are:

1. A gapless tensor mode is not enough. A Lifshitz spin-2 mode with $z \neq 1$ is a scientifically interesting partial success, but not a GR-like graviton.
2. Gauge protection is preferable to critical fine tuning. A massless tensor appearing only at an isolated $\lambda_{c,2}$ is weaker than a massless sector protected by an emergent constraint or gauge redundancy over a phase.
3. Three spatial dimensions are mandatory for the target used here. Two-dimensional spatial lattices can validate hardware, scalar observables, and raw anisotropy pipelines, but they cannot host two propagating graviton polarizations of 3+1-dimensional linearized gravity.

3 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 number or exponent not chosen to match the answer. The near-term outputs are the dispersion exponent z, the finite-size tensor mass/gap scaling, and a magnitude-bearing tensor stiffness or critical exponent in microscopic bond units. The ratio
$$R_g = \frac{c_g}{v_{\mathrm{LR}}} \tag{2}$$
is retained as a Lorentz-recovery diagnostic once a $z = 1$ tensor branch exists, not as the independent Mercury number.

The rest of the paper is organized around this criterion.

4 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

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

A minimal ladder Hamiltonian is

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

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.

4.1 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.

5 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

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

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.

5.1 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,

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

and an effective dimension,

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

A shell profile gives an independent estimator,

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

The severe scalar closure test is therefore

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

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

5.2 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.

6 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

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

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.

7 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.

8 Spatial-dimension gate

The spin-2 target is not merely a tensor algebra target. It is a dimensional target. In D spacetime dimensions a massless graviton has

$$N_{\mathrm{pol}}(D) = \frac{D(D-3)}{2} \tag{32}$$

propagating polarizations. For $D = 4$, $N_{\mathrm{pol}} = 2$. For $D = 3$, corresponding to a two-dimensional spatial lattice plus time, $N_{\mathrm{pol}} = 0$.

This is not a minor finite-size caveat. A square or triangular sheet may have meaningful scalar loads, bond anisotropy, and two-dimensional tensor algebra, but it is not an admissible host for the two propagating helicity-2 modes targeted by Eq. (31). In $2+1$ gravity the Einstein theory is

topological in the sense relevant here: there are no local gravitational waves. Therefore all currently executed two-dimensional square and triangular circuits must be classified as pipeline-validation runs, not as spin-2 admissible candidates for the $3+1$ Mercury target.

The graph-family gate is now:

1. **Phase-0A controls:** path, star, cycle, rhombus/diamond, and other small motifs. These validate rejection and hardware readout.
2. **Phase-1A two-dimensional validation:** house, square, and triangular patches. These validate embedded graph hardware, scalar/bond readout, and raw anisotropy maps, but fail the $3+1$ spin-2 admissibility gate.
3. **Phase-2A first admissible probe:** a 2x2x2 cubic graph in the same v0 encoding. This is a 3D tensor-susceptibility probe, not a continuum or dispersion measurement.
4. **Phase-2B/3 scaling families:** 2x2x3, 3x3x2, 3x3x3, and periodic 3D families, with the role of each fixed before interpreting any tensor signal.

In the current v0 hardware encoding, with one matter qubit per vertex and one bond qubit per edge, the first cubic patches are reachable in width, but not all of them can answer the same question:

3D patch	vertices	edges	hardware qubits	can measure z?	role
2x2x2 cubic	8	12	20	no	first 3D encoding and uniform tensor-susceptibility probe; $k=0$ gap/source response only
2x2x3 cubic	12	20	32	dirty single-axis	first single-axis dispersion probe; anisotropic and finite-size contaminated
3x3x2 cubic	18	33	51	partial	slab check with two long axes; not isotropic 3D scaling
3x3x3 cubic	27	54	81	yes	first isotropic v0 target for z and tensor finite-size scaling
3x3x3 periodic	27	81	108	yes	cleaner low-k torus once routing and noise budgets allow

Table 2: First three-dimensional QATNU targets in the v0 hardware encoding. The 2x2x2 cube is the first dimensionally admissible hardware probe and is cheaper than the previously run 25-qubit triangular 3x3 sheet, but it cannot measure a dispersion exponent. The first clean isotropic v0 z target is 3x3x3 at 81 qubits.

The earlier scalar closure standard also points in the same direction. The scalar sector requires agreement between diffusion-based and shell-profile effective dimensions in the same window. A two-dimensional shell response or a two-dimensional hardware sheet cannot be silently upgraded into a three-dimensional tensor continuum. The mirror warning is also important: a 2x2x2 cube has embedding dimension three, but it is not yet a three-dimensional continuum. Its open-cube Laplacian has only a tiny finite spectrum, and each axis supports only the uniform and zone-boundary modes. Thus the hierarchy is

$$\text{3D embedding at } 2\times 2\times 2 \quad \ll \quad d_{\text{eff}} \approx 3 \text{ scaling window} \quad \ll \quad \text{clean continuum/torus limit.} \tag{33}$$

The 2x2x2 cube is admissible as the first 3D tensor-susceptibility probe, not as a continuum pass.

9 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 admissible graph family, a tensor observable first requires a chart. For hardware fixtures this may be an explicit embedding; for an emergent-geometry claim it must also be reproduced by an intrinsic chart, such as a spectral or diffusion-map embedding derived from L_w. Let each edge e have unit tangent $\hat{t}_e$ and centered occupation

$$q_e = n_e - \langle n \rangle_{\mathcal{N}(x)}. \tag{34}$$

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

$$M_{ab}(x) = \sum_{e \in \mathcal{N}(x)} q_e \left(\hat{t}_{e,a} \hat{t}_{e,b} - \frac{1}{d} \delta_{ab} \right). \tag{35}$$

This is a substrate observable. It is not made transverse-traceless by definition. However, it is not yet an intrinsic geometric observable unless the result is stable under allowed chart choices. Chart stability is therefore a gate, not a decorative robustness check.

9.1 Tensor susceptibility

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

$$H_J = -\sum_x J_{ab}(x) M_{ab}(x). \tag{36}$$

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

$$\Delta_B(e) \mapsto \Delta_B - J_{ab}(x_e) \left(\hat{t}_{e,a} \hat{t}_{e,b} - \frac{1}{d} \delta_{ab} \right). \tag{37}$$

Define the connected susceptibility

$$\chi_{ab,cd}(\Omega, \mathbf{k}) = \frac{\delta \langle M_{ab}(\Omega, \mathbf{k}) \rangle}{\delta J_{cd}(\Omega, \mathbf{k})}. \tag{38}$$

The inverse tensor kernel is

$$K_{ab,cd}(\Omega, \mathbf{k}) = \chi^{-1}_{ab,cd}(\Omega, \mathbf{k}). \tag{39}$$

In the tensor channel one must not assume either propagation or the relativistic branch. Because the scalar timing closure is first-order and relaxational, the leading null hypothesis for the tensor channel is also relaxational:

$$K_2^{(\mathrm{diff})}(\Omega, k) = m_2^2 + D_T k^2 - i Z_D \Omega + O(k^4, \Omega k^2, \Omega^2). \tag{40}$$

A pole of this branch gives relaxation, not a propagating graviton. Only after a propagating response is established does it make sense to fit a dynamical exponent. The propagating branches

to distinguish are

$$K_2^{(z=1)}(\Omega,k) = Z_2\Omega^2 - \rho_2 k^2 - m_2^2 + O(k^4,\Omega^4,k^2\Omega^2), \tag{41}$$

$$K_2^{(z)}(\Omega,k) = Z_z\Omega^2 - \rho_z k^{2z} - m_2^2 + \cdots . \tag{42}$$

Here z is the dynamical exponent of a propagating branch. The logical fork is therefore: no tensor response; relaxational tensor response; propagating Lifshitz tensor response; and only finally relativistic $z = 1$ tensor response. The $z = 1$ form is the GR-like target; $z = 2$ or $z = 3$ would place QATNU in a Lifshitz/tensor-medium class rather than the relativistic Pauli-Fierz class.

A perturbative fast-sector elimination gives the schematic tensor mass

$$m_2^2(\lambda) = \Delta_B^{(2)} - \lambda^2\Pi_2(0,0) + O(\lambda^4), \tag{43}$$

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

$$m_2^2 = 0 \qquad \Longleftrightarrow \qquad \lambda_{c,2}^2 = \frac{\Delta_B^{(2)}}{\Pi_2(0,0)} + O(\lambda_{c,2}^4). \tag{44}$$

The stronger target is not fine tuning to one isolated critical point, but a gauge- or constraint-protected massless regime. The critical form is still useful as a local diagnostic: if the uniform susceptibility never grows under a λ or source scan, the tensor program fails before dispersion.

At criticality, a propagating branch gives

$$\Omega^2 \sim \frac{\rho_z}{Z_z} k^{2z}. \tag{45}$$

Only on the relativistic branch $z = 1$ can one define

$$\Omega^2 = c_g^2 k^2 + O(k^4), \qquad c_g^2 = \frac{\rho_2}{Z_2}. \tag{46}$$

The Lorentz-recovery diagnostic is then

$$R_g = \frac{c_g}{v_{\mathrm{LR}}} = \frac{1}{v_{\mathrm{LR}}}\sqrt{\frac{\rho_2}{Z_2}}, \tag{47}$$

where v_{LR} is the Lieb-Robinson or causal-cone velocity of the same substrate in the same microscopic units. This definition replaces the earlier informal idea of a scalar-clock speed: the scalar timing sector is relaxational in the present closure, so its natural speed scale is the substrate causal-cone velocity, not a wave speed extracted from Eq. (40). The first Mercury outputs are ordered: propagation versus relaxation, finite-size scaling of m_2 or χ_T^{-1}, then z on a graph large enough to fit it. R_g is meaningful only after a $z = 1$ branch has been established.

9.2 Ward residuals

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

$$\mathcal{W}_{\rm div}(k) = \frac{\|k^a \chi_{ab,cd}\|}{\|\chi_{ab,cd}\|}, \qquad \mathcal{W}_{\rm tr}(k) = \frac{\|\delta^{ab}\chi_{ab,cd}\|}{\|\chi_{ab,cd}\|}. \tag{48}$$

The infrared Ward condition is

$$\mathcal{W}_{\rm div}(k) \to 0, \qquad \mathcal{W}_{\rm tr}(k) \to 0 \qquad (k \to 0). \tag{49}$$

If Eq. (43) reaches a protected massless regime, Eq. (48) vanishes, and the measured exponent is $z = 1$, the long-wavelength correlator takes the form

$$\langle M_{ab}(\Omega, \mathbf{k}) M_{cd}(-\Omega, -\mathbf{k})\rangle_c \sim \frac{Z\, P^{\rm TT}_{ab,cd}(\mathbf{k})}{\Omega^2 - c_g^2 k^2 + i0} + \text{analytic terms}. \tag{50}$$

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

9.3 Conditional theorem

Conditional spin-2 theorem. *Suppose a QATNU graph family is three-dimensional in the infrared, admits both explicit and intrinsic charts for M_{ab}, has a tensor response that is propagating rather than relaxational, has a measured dynamical exponent $z = 1$, reaches a protected massless tensor regime, and satisfies vanishing Ward residuals Eq. (48). 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: the uniform tensor susceptibility, $\Pi_2(0,0)$ when a critical description is appropriate, the inverse response coefficients Z_D, D_T or Z_z, ρ_z, whether the mode is diffusive or propagating, the finite-size scaling of m_2 or χ_T^{-1}, chart-stability residuals, Ward residuals, and only then R_g. It also names the likely partial successes: a relaxational tensor channel is a medium response, not a graviton; a gapless gauge-protected tensor mode with $z \neq 1$ is a Lifshitz spin-2 medium, not a GR-like graviton.

9.4 Massless limit and source-coupling warning

A vanishing fitted tensor mass is not enough. Massive spin-2 responses can fail to approach the massless source-coupling structure smoothly; the extra scalar polarization can survive in the massless limit. The propagator pole test and the universal-source-coupling test are therefore logically independent. A QATNU tensor branch must test both.

10 Stoquasticity and the native-3D QMC route

The dimensional gate makes clear that thermodynamic scaling is needed. The hardware route can validate encodings and small source-response circuits, but it cannot efficiently reach the large

three-dimensional lattices needed to measure an infrared exponent. The sign structure of the QATNU ladder changes the computational strategy: the equilibrium ladder is gauge-stoquastic.

10.1 Full-ladder gauge-stoquasticity theorem

Work in the product basis in which matter qubits are Z-eigenstates and every bond register is in its rung basis $\{|m\rangle\}_{m=0}^{\chi_{\max}-1}$. The Hamiltonian has the schematic form

$$H = H_{\text{diag}} + \frac{\omega}{2}\sum_i X_i + \lambda \sum_{e=\langle i,j\rangle} F_{ij}(R_e + R_e^\dagger), \tag{51}$$

where H_{diag} contains rung energies, $J_0 Z_i Z_j$, $\Delta_B n_e$, degree penalties, and other diagonal terms. The frustration projector

$$F_{ij} = \frac{1 - Z_i Z_j}{2} \tag{52}$$

is diagonal and non-negative. The ladder operator has positive matrix elements

$$\langle m+1|R_e|m\rangle = \sqrt{m+1} > 0. \tag{53}$$

Thus the Hamiltonian is not stoquastic as written: both matter flips and bond promotions have positive off-diagonal matrix elements.

Define the diagonal gauge

$$U = \left(\prod_{i\in V} Z_i\right) \otimes \left(\prod_{e\in E} P_e\right), \qquad P_e = \sum_{m=0}^{\chi_{\max}-1} (-1)^m |m\rangle\langle m|. \tag{54}$$

All diagonal terms commute with U. For the matter sector,

$$Z_i X_i Z_i = -X_i, \tag{55}$$

so every matter-flip matrix element becomes $-\omega/2 \le 0$. For the bond ladder,

$$P_e R_e P_e = -R_e, \qquad P_e R_e^\dagger P_e = -R_e^\dagger, \tag{56}$$

so every allowed promotion or demotion matrix element becomes

$$-\lambda F_{ij}\sqrt{m+1} \le 0. \tag{57}$$

Therefore, for $\omega, \lambda \ge 0$ and canonical positive ladder amplitudes, $UHU^\dagger$ has non-positive off-diagonal elements in the product basis. The full $\chi_{\max}$ ladder is stoquastic, not merely the $\chi_{\max} = 2$ hardware truncation.

This theorem has a sharp domain. It assumes real positive matter-flip and promotion amplitudes. Chiral phases, sign-disordered couplings, complex hopping amplitudes, or additional non-diagonal bond terms would require a new sign analysis. Within this domain, however, the equilibrium imaginary-time weights are non-negative and the finite-temperature/statistical mechanics sector can be sampled without a sign problem.

10.2 What sign-free does and does not mean

Stoquasticity does not mean the ground state is classical, unentangled, or physically uninteresting. The transverse-field Ising model is the standard counterexample: it is stoquastic and sign-problem-free, yet it has a genuine quantum critical point with an emergent relativistic mode. The correct conclusion is narrower and stronger for computation:

> The equilibrium universality class of the QATNU ladder is accessible by sign-free imaginary-time Monte Carlo.

Thus the QPU/ED qubit wall is not the fundamental barrier for the Mercury numbers. It is a hardware/exact-diagonalization barrier. The thermodynamic questions–does a tensor susceptibility grow, does a tensor gap close, is the tensor response diffusive or propagating, and if propagating what is z–should be attacked first with native-3D worldline QMC.

10.3 Imaginary-time mapping and the tensor target

After the gauge, a Trotter or continuous-time expansion gives a positive-weight path integral. The matter transverse field produces ferromagnetic imaginary-time bonds. In a simple Ising-core discretization one obtains the usual anisotropic classical mapping with

$$K_\tau = -\frac{1}{2}\log\tanh\left(\frac{\omega\Delta\tau}{2}\right), \tag{58}$$

while spatial Z_iZ_j couplings contribute spatial Ising bonds. The QATNU-specific ingredient is the bond-register worldline: the promotion term contributes hopping vertices with amplitude

$$\lambda\sqrt{m+1}\,\Delta\tau \tag{59}$$

on frustrated spacetime plaquettes, together with the diagonal rung-energy and degree terms.

The scalar timing sector remains a useful observable, but the QMC Mercury target is tensorial. On a three-dimensional graph one should measure

$$M_{ab}(x,\tau) = \sum_{e\in\mathcal{N}(x)} q_e(\tau)\left(\hat{t}_{e,a}\hat{t}_{e,b} - \frac{1}{3}\delta_{ab}\right), \tag{60}$$

with chart-stability checks between explicit and intrinsic embeddings. The central QMC observables are the tensor susceptibility and connected correlator,

$$\chi_T(\lambda, L) = \sum_{x,\tau}\langle M_{ab}(x,\tau)M_{ab}(0,0)\rangle_c, \tag{61}$$

$$S_{ab,cd}(k,\omega_n) = \langle M_{ab}(k,\omega_n)M_{cd}(-k,-\omega_n)\rangle_c. \tag{62}$$

The mass/gap proxy is χ_T^{-1} only after the susceptibility-to-mass map is declared; the more robust statement is finite-size scaling of χ_T, correlation length, and the lowest tensor relaxation or propagation scale.

10.4 Prototype QMC instrument

A prototype Monte Carlo instrument was built for the matter-core sign-free mapping. It samples a $(d+1)$-dimensional anisotropic classical Ising model, measures spatial and imaginary-time correlators, estimates effective correlation lengths, and fits

$$\xi_\tau \sim \xi_x^z. \tag{63}$$

The instrument was first checked in $1+1$ dimensions, where the isotropic classical Ising benchmark should give $z = 1$. A representative run in the disordered regime gave spatial and temporal correlation-length ratios near unity and a log-slope $z \simeq 1.03$. An anisotropic control with $K_s \neq K_t$ produced a velocity ratio near 1.5 while still giving $z \simeq 1.06$, showing that the instrument distinguishes a nonuniversal velocity from the universal exponent. The same code then ran native three-spatial-dimensional samples with $L = 6, 8, 10$, corresponding to quantum matter Hilbert dimensions 2^{216}, 2^{512}, and 2^{1000}, which are out of reach for exact diagonalization or direct QPU width/depth scaling.

These results validate the method and the native-3D reach. They are not yet a QATNU tensor result. The prototype omits the full bond-register worldline update and the M_{ab} correlator. The next software result must add those two ingredients before reporting a QATNU-specific $m_2^2(\lambda)$, propagation-versus-relaxation classification, or z.

Layer	What was shown	What is not yet shown
Gauge-stoquasticity	Full positive ladder has a sign-free gauge for $\omega, \lambda \geq 0$	Robustness to complex/chiral couplings
QMC benchmark	Matter-core instrument recovers $z \approx 1$ on known anisotropic Ising tests	QATNU tensor exponent
Native-3D reach	$L = 6, 8, 10$ three-spatial-dimensional matter-core samples are routine	Full tensor scaling still requires full-ladder production statistics
Full-ladder QMC	ED validated and run on 2x2x2/3x3x3 cube fixtures	Returns a null for tensor criticality in the scanned region
Hardware cube	2x2x2 remains the first admissible 3D QPU source-response fixture	Thermodynamic finite-size scaling or z

Table 3: The new computational split. QPU hardware validates encodings and small source-response fixtures; sign-free QMC is the correct first tool for thermodynamic Mercury numbers.

11 Full-ladder QMC: bond-register worldlines and a converged null

The matter-core QMC prototype validated the measurement instrument, but it was not a QATNU tensor result. It omitted the bond-register worldlines and the rank-2 observable M_{ab}. This section reports the full-ladder version: a sign-free, imaginary-time worldline Monte Carlo with frustration-

gated bond registers, the cubic tensor observable, source-response probes, Matsubara-frequency diagnostics, and exact-diagonalization validation on small graphs. The headline result is a converged negative: in the natural parameter neighborhood of the nominal point, the rank-2 channel shows no emergent tensor criticality.

11.1 Imaginary-time mapping and gated bond worldlines

Work in the gauge-stoquastic frame. A discrete-time Trotter path integral with M slices and $\Delta\tau = \beta/M$ assigns to matter spins $s_i(\tau) \in \{\pm 1\}$ and bond rungs $m_e(\tau) \in \{0, \ldots, \chi_{\max} - 1\}$ the positive weight

$$W = \prod_{\tau} e^{-\Delta\tau E_{\mathrm{diag}}(\tau)} \prod_{i,\tau} \langle s_i(\tau+1)|e^{\Delta\tau(\omega/2)X_i}|s_i(\tau)\rangle \prod_{e,\tau} \langle m_e(\tau+1)|e^{\Delta\tau\lambda F_e(\tau)(R_e+R_e^\dagger)}|m_e(\tau)\rangle, \quad (64)$$

with

$$E_{\mathrm{diag}}(\tau) = J_0 \sum_{\langle ij\rangle} s_i(\tau)s_j(\tau) + \Delta_B \sum_e m_e(\tau). \quad (65)$$

The matter time bonds are the standard transverse-field factors: $\cosh(\Delta\tau\omega/2)$ if a spin is unchanged across the time link and $\sinh(\Delta\tau\omega/2)$ if it flips. The QATNU-specific term is the bond time bond. A rung may hop $m \to m \pm 1$ only on a slice where its two matter endpoints are frustrated,

$$F_e(\tau) = \frac{1 - s_i(\tau)s_j(\tau)}{2} = 1, \quad (66)$$

with amplitude $\Delta\tau\lambda\sqrt{\min(m, m') + 1}$; a hop on an unfrustrated slice has zero weight. All measured quantities below, including $\langle n_e\rangle$, Λ_i, and M_{ab}, are diagonal. Since the stoquastic gauge is also diagonal, these thermal expectations are gauge-invariant and can be validated against ED in the original frame.

Matter spins are updated by a checkerboard Metropolis sweep on bipartite lattices, with a serial fallback for non-bipartite validation graphs. The bond sector is updated by exact heat-bath resampling. Conditioned on the matter gate pattern $\{F_e(\tau)\}$, each bond worldline is a periodic one-dimensional chain in imaginary time. It is sampled exactly by transfer matrices

$$[A_e^{(\tau)}]_{m'm} = K_{m'm}^{(\tau)}\, e^{-\Delta\tau\Delta_B^{(e)}m}, \qquad K_{m\pm1,m}^{(\tau)} = \Delta\tau\lambda\sqrt{m+1} \quad \text{iff} \quad F_e(\tau) = 1. \quad (67)$$

A traceless tensor source $J_{ab} = J\,\mathrm{diag}(1, -1/2, -1/2)$ enters through a per-edge rung-energy bias, $\Delta_B^{(e)} \to \Delta_B - J_{aa}$ for an axis-a edge.

11.2 Validation against exact diagonalization

With $\chi_{\max} = 3$, both ladder amplitudes $\sqrt{1}$ and $\sqrt{2}$ are exercised. ED agreement is therefore a test of the gated promotion rule, not only of the matter core. Table 4 reports diagonal observables on small graphs; the QMC reproduces the ED edge occupations and scalar load at the percent/Trotter-bias level, including edge ordering on an inequivalent chorded diamond. A wrong gate or ladder amplitude would miss by far more than this.

test ($\chi_{\max}=3$, $\beta=2$, $M=24$)	ED	QMC	note
triangle $\langle n_e\rangle$	0.4457	0.438(3)	$O(\Delta\tau)$ residual
triangle scalar load Λ	0.7372	0.726(4)	consistent
chorded diamond, rim edges	0.447	~ 0.44	per-edge agreement
chorded diamond, chord edge	0.442	0.43	reproduces ordering
square, checkerboard update	0.4515	0.445	vectorized matter
cube $\lambda=0$	0.384	0.386	pure-diagonal limit

Table 4: Full-ladder QMC against exact diagonalization. Agreement is at the expected $\Delta\tau$-limited percent level on graphs with inequivalent edges.

11.3 Tensor susceptibility, source response, and the Matsubara fork

On the $2\times2\times2$ cube, the axis-aligned tensor observable is especially direct. Let $Q_a=\sum_{e\in a} q_e$ with $q_e=n_e-\langle n\rangle$ over the edges in a time slice. Then

$$M_a(\tau)=Q_a(\tau)-\frac{1}{3}\sum_b Q_b(\tau),\qquad a=x,y,z. \tag{68}$$

The calculation uses two independent uniform-susceptibility estimators. The equilibrium estimator is the connected fluctuation of the zero-Matsubara tensor,

$$\chi_T=\frac{\beta}{|E|}\langle \overline{M}_a\overline{M}_a\rangle_c, \tag{69}$$

and the explicit source response is the finite difference

$$\chi_T^{\mathrm{src}}=\frac{M_{xx}(+J)-M_{xx}(-J)}{2J}. \tag{70}$$

The baseline anisotropy vanishes within noise, $\langle M_{xx}\rangle\simeq-0.01$. The source response is linear: $\chi_T^{\mathrm{src}}=3.16$ at $J=0.3$ and 3.19 at $J=0.6$, of the same order as the fluctuation estimator after accounting for single-component versus summed normalization.

The propagation-versus-relaxation diagnostic is the small-Matsubara-frequency analytic structure of $\chi^{-1}(i\omega_n)$. A relaxational branch gives $\chi^{-1}\sim|\omega_n|$, while a propagating branch gives $\chi^{-1}\sim\omega_n^2$. On the smallest cube the fork is implemented but ambiguous: representative low-frequency fits give $R^2(\omega_n^2)=0.995$ versus $R^2(|\omega_n|)=0.982$, both high and separated by less than the prototype uncertainty. This is method validation, not a propagation verdict.

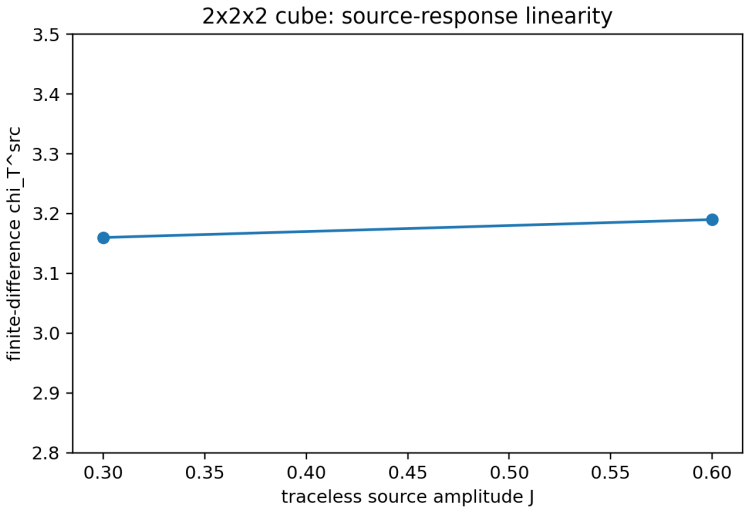

Figure 1: The $2 \times 2 \times 2$ cube source response is linear in the tested range. This validates the tensor-source channel, but it does not establish a spin-2 mode or a dispersion law.

11.4 $\chi_{\max}$ convergence and the displaced-oscillator structure

The operator $R_e = \sum_m \sqrt{m+1}|m+1\rangle\langle m|$ is a truncated bosonic creation operator. On a frustrated edge the bond sector behaves like a displaced oscillator, $\Delta_B n + \lambda(R + R^\dagger)$, so lowering Δ_B increases occupation and makes truncation convergence nontrivial. Table 5 and Fig. 2 show why the ladder cutoff must be converged before interpreting any tensor number.

$\chi_{\max}$	3	4	5	6	10	14
nominal $\chi_T/\lvert E\rvert$ ($\Delta_B = 0.55$)	0.51	0.73	0.82	0.84	—	—
$\langle n\rangle$, nominal	0.27	0.31	0.34	0.33	—	—
low-Δ_B $\chi_T/\lvert E\rvert$ ($\Delta_B = 0.20$)	1.1	2.1	3.3	4.8	11.2	15.5
$\langle n\rangle$, low-Δ_B	0.75	1.15	1.53	1.85	2.64	2.89

Table 5: Truncation convergence on the $2 \times 2 \times 2$ cube at $\beta = 4$, $\lambda = 0.35$. The nominal point converges by $\chi_{\max} \sim 5$–6. The low-Δ_B point requires $\chi_{\max} \sim 14$–16, and the apparent enhancement at small cutoff is not yet physics.

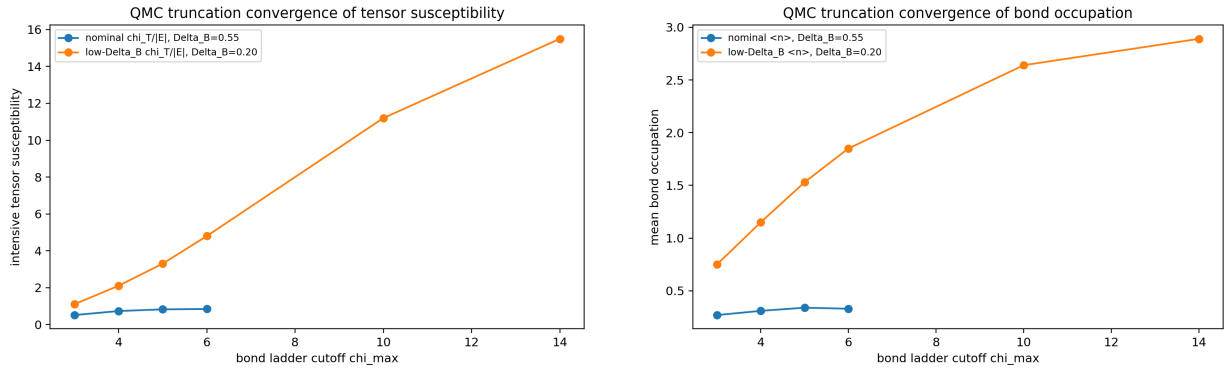

Figure 2: Cutoff convergence for tensor susceptibility and mean bond occupation. The low-Δ_B point is a strongly populated displaced-oscillator regime; the hardware v0 encoding with $\chi_{\max} = 2$ cannot represent this regime.

11.5 Finite-size scaling at converged truncation

A tensor critical point is a claim about correlations growing with system size, not about a large one-size susceptibility. At converged truncation, the finite-size result is negative. At the displaced point $(\Delta_B = 0.20, \lambda = 0.35, \beta = 4, \chi_{\max} = 12)$,

$$\left.\frac{\chi_T}{|E|}\right|_{L=2} = 14.1, \qquad \left.\frac{\chi_T}{|E|}\right|_{L=3} = 13.0. \tag{71}$$

The intensive susceptibility is flat to slightly decreasing. The disconnected single-edge occupation-variance baseline is about 20.8 at both sizes, larger than the full $\chi_T/|E|$. Thus the tensor channel is accounted for by single-edge displaced-oscillator variance, with no constructive inter-edge correlation and no L-growth. The nominal point is likewise flat, $0.55 \to 0.53$.

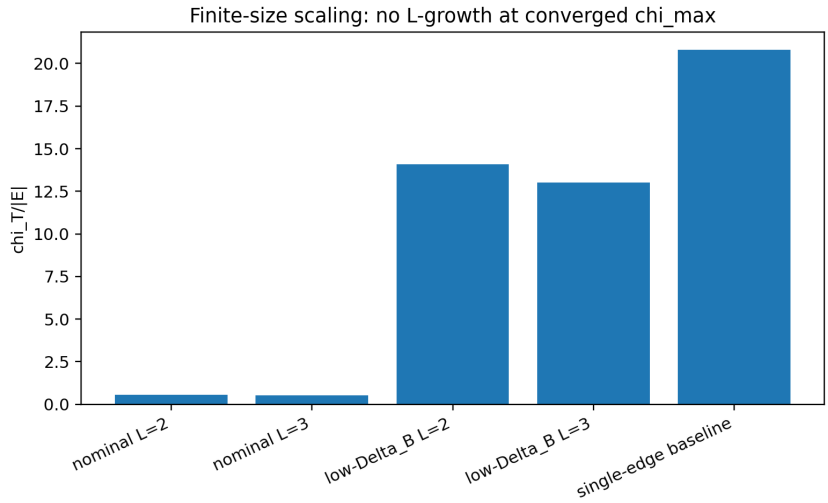

Figure 3: Finite-size scaling at converged truncation gives the central null: $\chi_T/|E|$ does not grow from $L = 2$ to $L = 3$ at either the nominal or low-Δ_B point. The low-Δ_B enhancement is dominated by single-edge variance, not collective tensor criticality.

11.6 Interpretation: a converged null for the nominal substrate family

In the scanned natural neighborhood

$$\omega = 1, \quad J_0 = 0.12, \quad \Delta_B \in [0.05, 0.55], \quad \lambda \in [0.35, 1.3], \quad \beta \in [2, 8], \quad L \leq 3, \tag{72}$$

with converged $\chi_{\max}$, QATNU shows no emergent tensor criticality. The susceptibility does not grow with L, and its magnitude is set by trivial per-edge bond fluctuations rather than a collective rank-2 mode. This is a clean negative for the region a reader would probe first. It does not exclude criticality elsewhere, and it carries prototype caveats: few seeds, no formal autocorrelation analysis, and only $L \leq 3$.

The result also identifies the missing mechanism. Since R_e is an independent per-edge bosonic operator, a single frustrated edge cannot by itself produce a collective critical tensor mode. Long-range tensor order can only come from inter-edge correlations transmitted through the shared matter/frustration sector, and the present interaction does not generate them strongly enough at any scanned point. The conditional spin-2 theorem therefore now faces a concrete antecedent failure. Reaching its hypotheses likely requires more than larger λ.

The sharpened next test is to add a direct bond–bond coupling or engineer spatially correlated frustration, then rerun the identical $\chi_{\max}$-converged, L-scaled χ_T protocol. If that inter-edge channel produces an L-growing tensor susceptibility, the program has its first genuine candidate critical surface. If it does not, the substrate as currently defined does not host the target spin-2 sector in three dimensions.

12 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.

13 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

$$\mathcal{H}_{\mathrm{v0}} = (\mathbb{C}^2)^{\otimes(|V|+|E|)}. \tag{73}$$

The Hamiltonian is

$$H_{\mathrm{v0}} = H_{\mathrm{mat}} + H_{ZZ} + H_{\mathrm{bond}} + H_{\mathrm{prom}}, \tag{74}$$

with

$$H_{\mathrm{mat}} = \frac{\omega}{2} \sum_{i \in V} X_i, \tag{75}$$

$$H_{ZZ} = J_0 \sum_{\langle i,j \rangle \in E} Z_i Z_j, \tag{76}$$

$$H_{\mathrm{bond}} = \Delta_B \sum_{e \in E} n_e, \qquad n_e = \frac{I - Z_e}{2}, \tag{77}$$

$$H_{\mathrm{prom}} = \lambda \sum_{e = \langle i,j \rangle} F_{ij} X_e, \qquad F_{ij} = \frac{1 - Z_i Z_j}{2}. \tag{78}$$

The initial state is

$$|\psi_0\rangle = |+\rangle^{\otimes |V|} \otimes |0\rangle^{\otimes |E|}. \tag{79}$$

The parameters used here are

$$\omega = 1.0, \qquad J_0 = 0.12, \qquad \Delta_B = 0.55, \qquad \lambda = 0.35, \qquad T = 2.0. \tag{80}$$

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

$$M_{ab}^{(2D)} = \sum_e q_e \left(\hat{t}_{e,a} \hat{t}_{e,b} - \frac{1}{2} \delta_{ab} \right), \qquad q_e = \langle n_e \rangle - \frac{1}{|E|} \sum_f \langle n_f \rangle. \tag{81}$$

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

14 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 graph family has a three-dimensional spatial infrared limit, not merely a two-dimensional embedding.
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, bowtie, square sheet, or triangular sheet can exercise a nonzero two-dimensional anisotropy tensor and still fail because a two-dimensional spatial graph has no 3+1-dimensional helicity-2 infrared limit. This is better than a hardcoded failure label: it records why the graph is rejected.

15 Phase-0A exact-simulation controls

Table 6 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 6: 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.

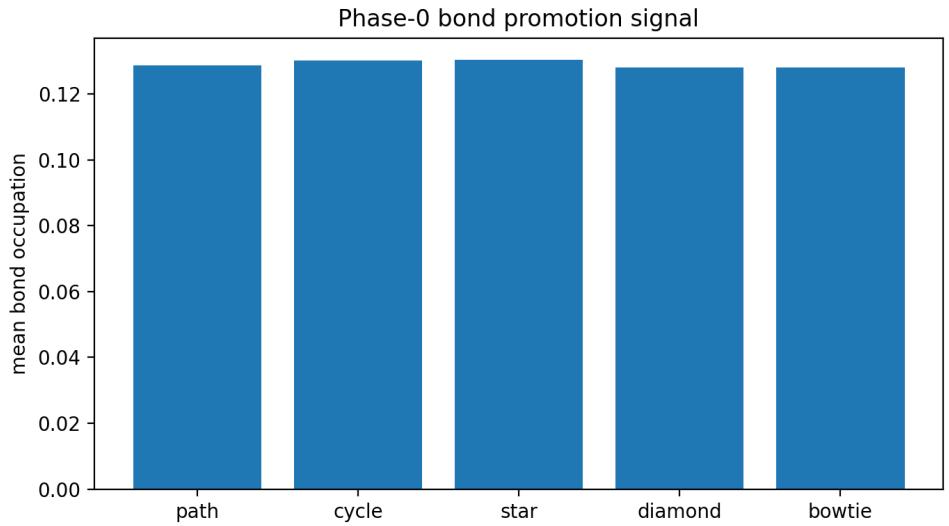

Figure 4: Mean bond occupation after exact evolution. The bond/scalar sector is active in all controls.

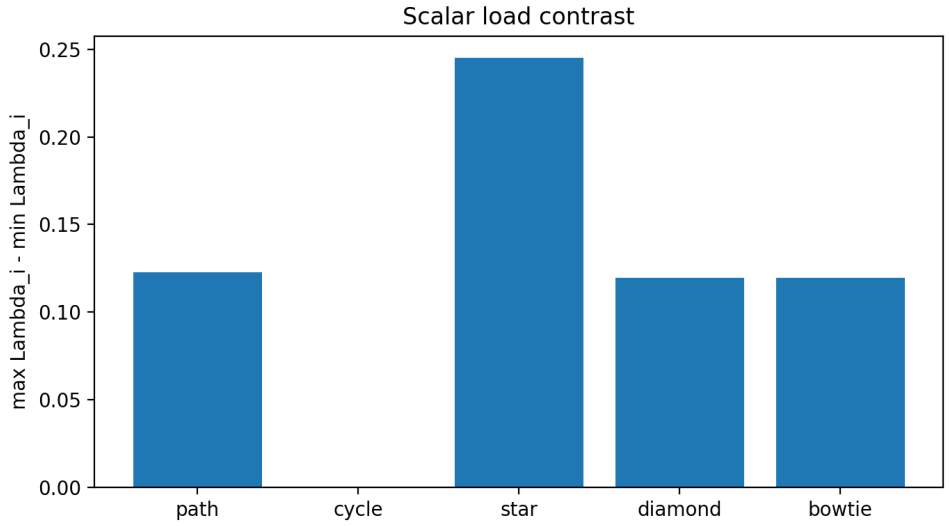

Figure 5: Scalar load spread for the corrected exact controls. The path, star, chorded diamond, and bowtie generate nonuniform scalar load; the symmetric cycle is uniform under the chosen state and parameters.

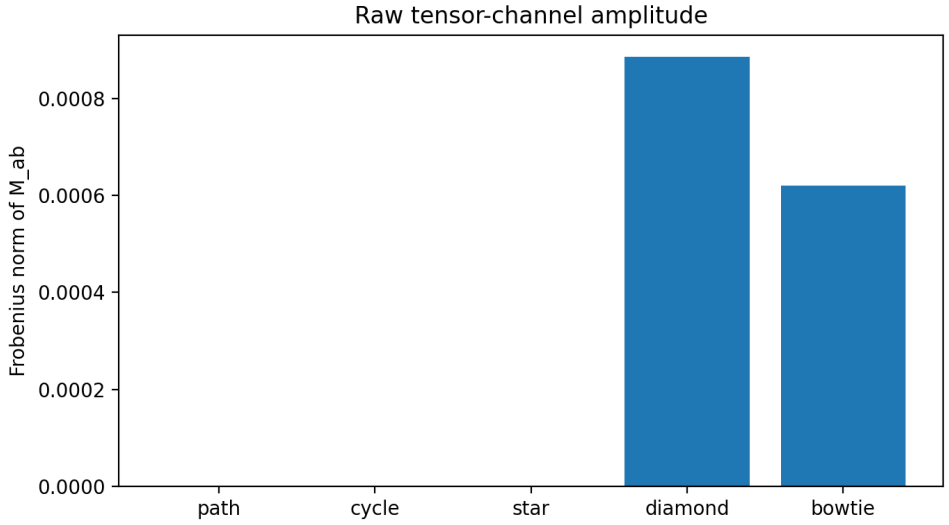

Figure 6: Raw anisotropy Frobenius norm. The nonzero values for chorded diamond and bowtie show that the tensor-channel diagnostic is being exercised, but these controls remain structurally inadmissible for spin-2 gravity.

16 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 7: 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

$$\langle n_e \rangle_{\mathrm{IBM}} = [0.1905,\ 0.1585,\ 0.2415], \tag{82}$$

$$\Lambda_{\mathrm{IBM}} = [0.17437,\ 0.32150,\ 0.36345,\ 0.21632]. \tag{83}$$

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.

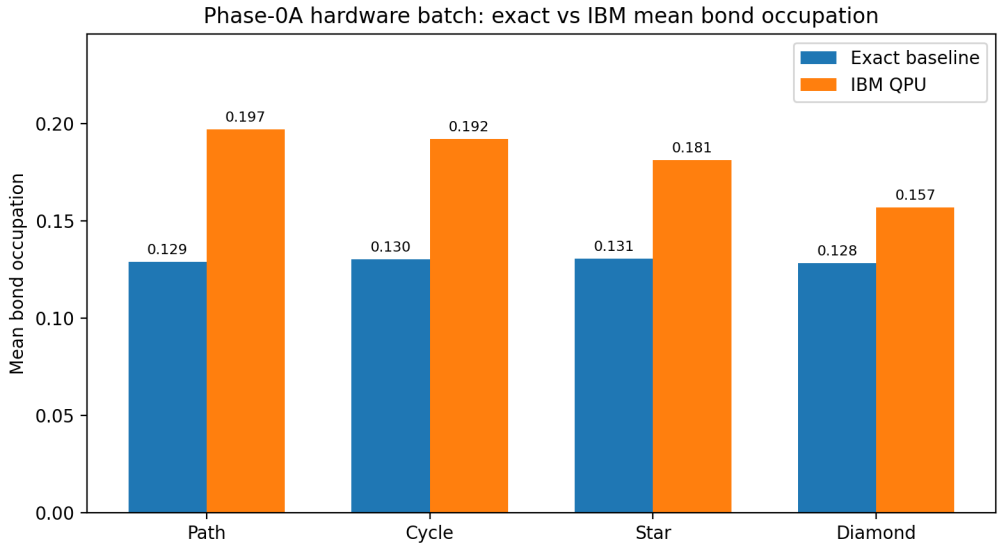

Figure 7: Exact baseline versus IBM mean bond occupation for the completed Phase-0A hardware batch. The IBM runs show active bond occupation across all four controls, while every topology remains structurally inadmissible for spin-2.

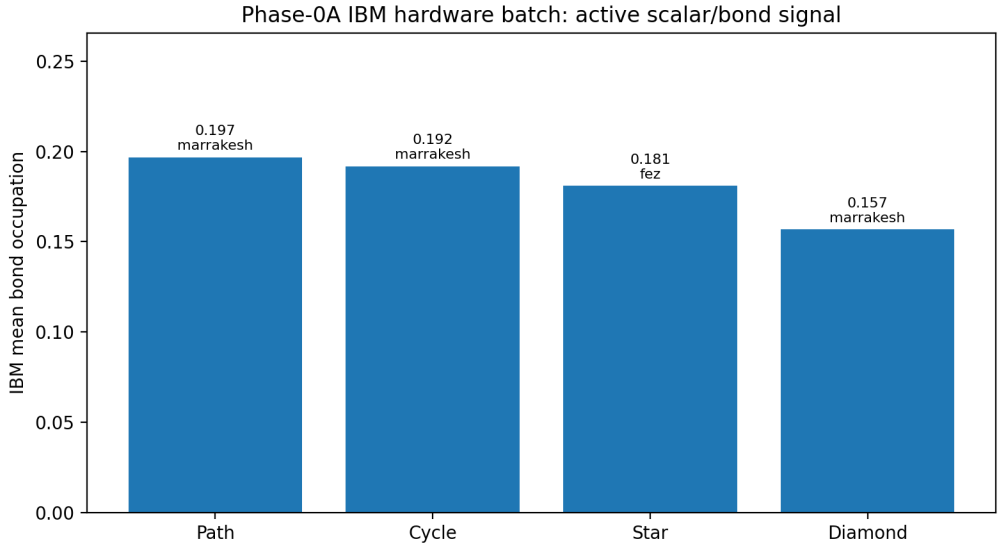

Figure 8: IBM-only mean bond occupation for the Phase-0A batch. Backend variation and hardware noise are treated as execution effects, not as new physics. The scientific result is the survival of the claim boundary.

16.1 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 8: 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.

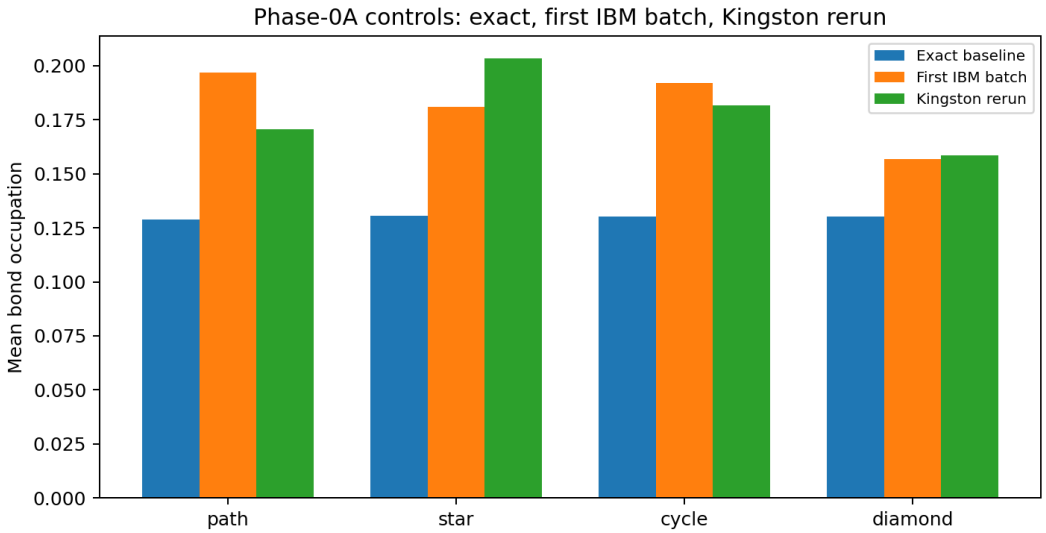

Figure 9: Phase-0A mean bond occupation across exact baselines, the first IBM batch, and the Kingston rerun. Backend-to-backend variation is visible, but the claim boundary is invariant.

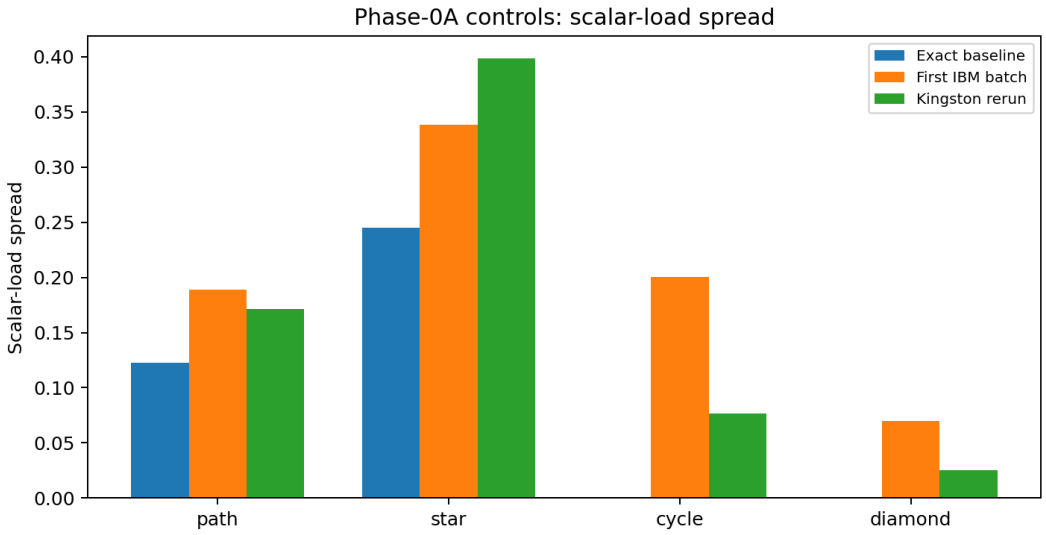

Figure 10: Phase-0A scalar-load spread across exact baselines, the first IBM batch, and the Kingston rerun. The hardware spread is device-sensitive; it is not interpreted as a tensor signal.

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
```

16.2 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.

17 Phase-1A: embedded graph-family hardware pipeline validation

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. Under the spatial-dimension gate, these two-dimensional graph families are now classified as hardware pipeline-validation runs rather than spin-2 admissible candidates.

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	2D square-lattice validation
`triangular_2x3`	6	9	15	2D triangular-lattice validation
`square_3x3`	9	12	21	2D square-lattice scaling smoke test; hardware smoke test after smaller circuits
`triangular_3x3`	9	16	25	2D triangular-lattice smoke test; not 3D-admissible

Table 9: Phase-1A graph families. The square and triangular entries are embedded 2D validation families, not 3D spin-2 admissible substrates. The 21-qubit `square_3x3` and 25-qubit `triangular_3x3` circuits extend the smoke tests into matched same-backend graph-family pairs.

17.1 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 10: 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.

17.2 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 11: 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.

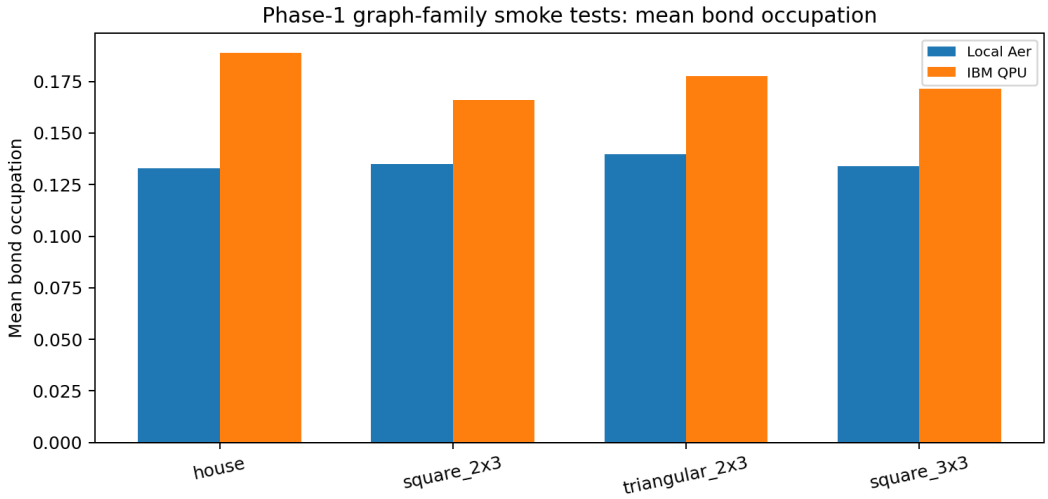

Figure 11: Phase-1A mean bond occupation. IBM values are higher than local Aer for the submitted circuits, consistent with hardware/transpilation/noise sensitivity rather than new physics.

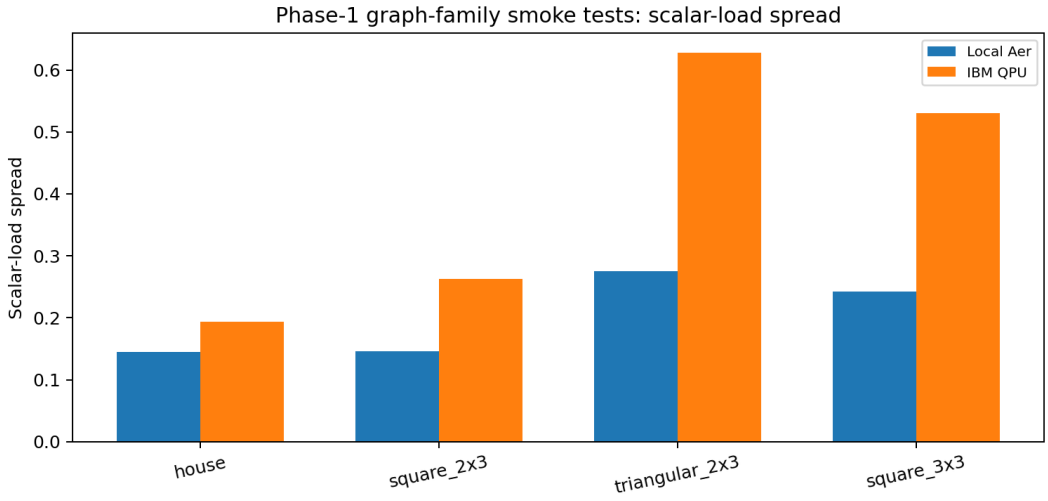

Figure 12: Phase-1A scalar-load spread. The triangular_2x3 hardware run shows a much larger scalar-load spread than local Aer. This should be treated as device- and layout-sensitive until repeated with calibration metadata and mitigation.

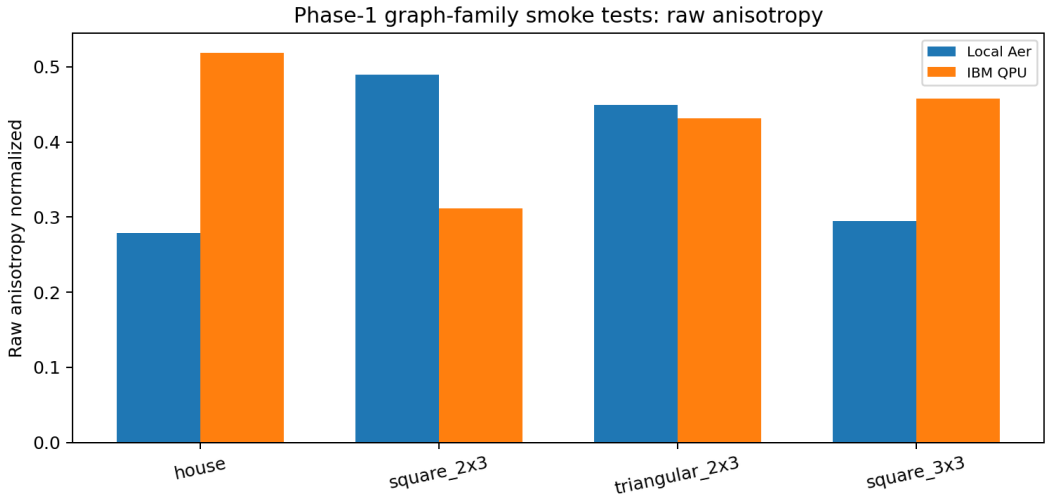

Figure 13: Phase-1A raw anisotropy diagnostic. Raw anisotropy is useful for checking that the graph-family readout is not trivial, but it is not a tensor susceptibility and cannot support a spin-2 claim on its own.

17.3 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 12: 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.

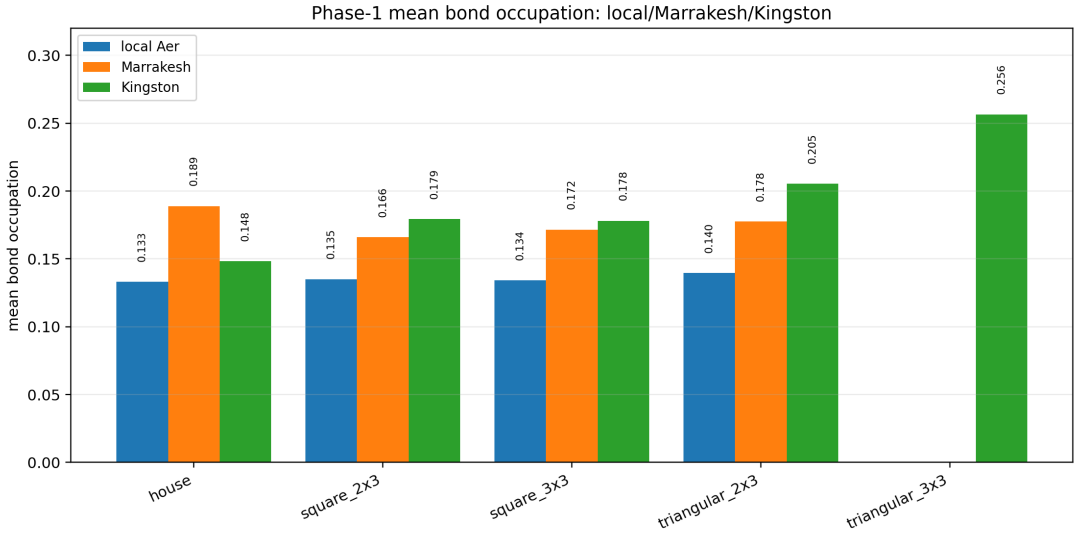

Figure 14: Phase-1A mean bond occupation across local Aer, the first Marrakesh pass, and the Kingston matched-family rerun. The Kingston data add cross-backend robustness, a same-backend house control, and same-backend square and triangular family reach.

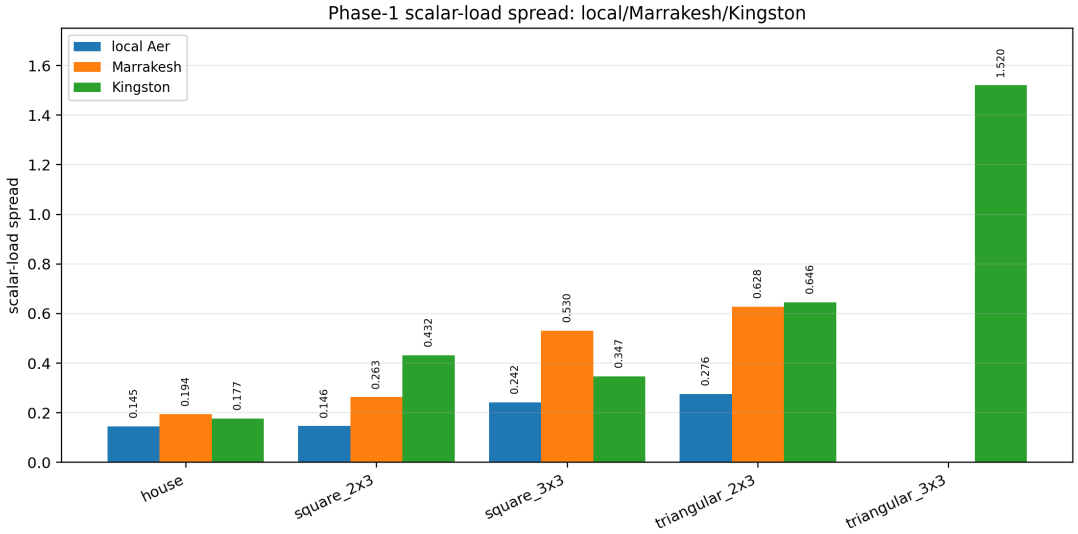

Figure 15: Phase-1A scalar-load spread for the house control, square family, and triangular family across local Aer, Marrakesh, and Kingston. The spread changes across hardware backends and is treated as device- and layout-sensitive, not as a tensor signal. The Kingston `triangular_3x3` point is the largest spread in the current hardware set.

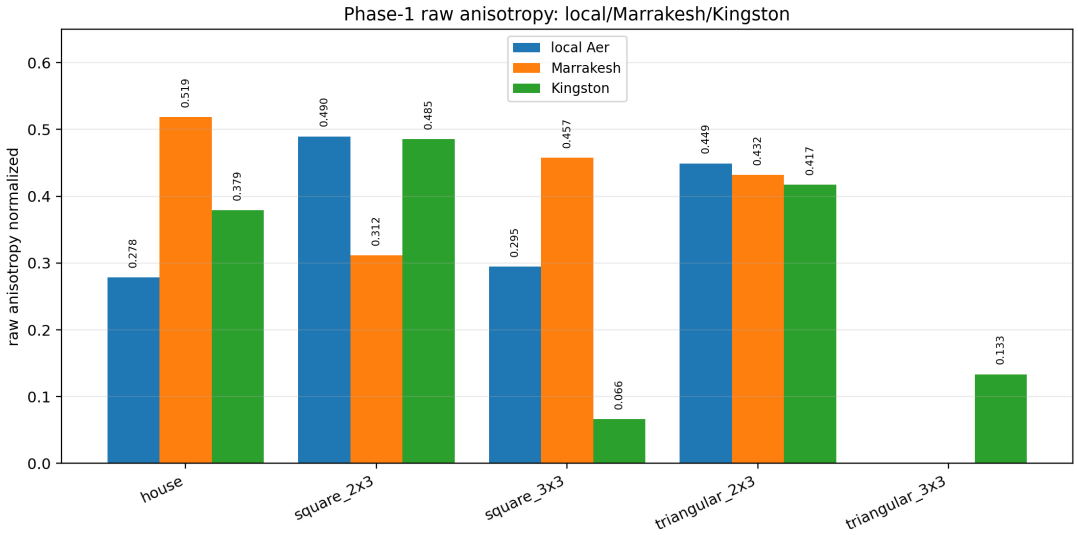

Figure 16: Phase-1A raw anisotropy across local Aer, Marrakesh, and Kingston. The triangular family shows the most interesting same-backend pattern: scalar-load spread grows from `triangular_2x3` to `triangular_3x3`, while normalized raw anisotropy decreases. This flags a follow-up target, not a spin-2 claim.

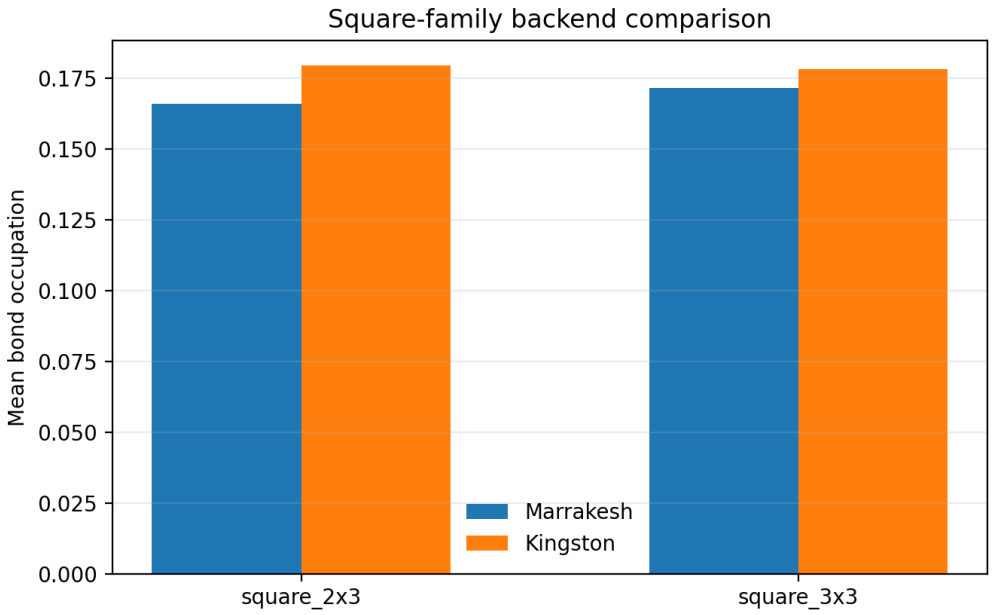

Figure 17: Square-family mean bond occupation for `square_2x3` and `square_3x3`. The near-equal Kingston means are an execution result only; they do not establish scaling or a tensor sector.

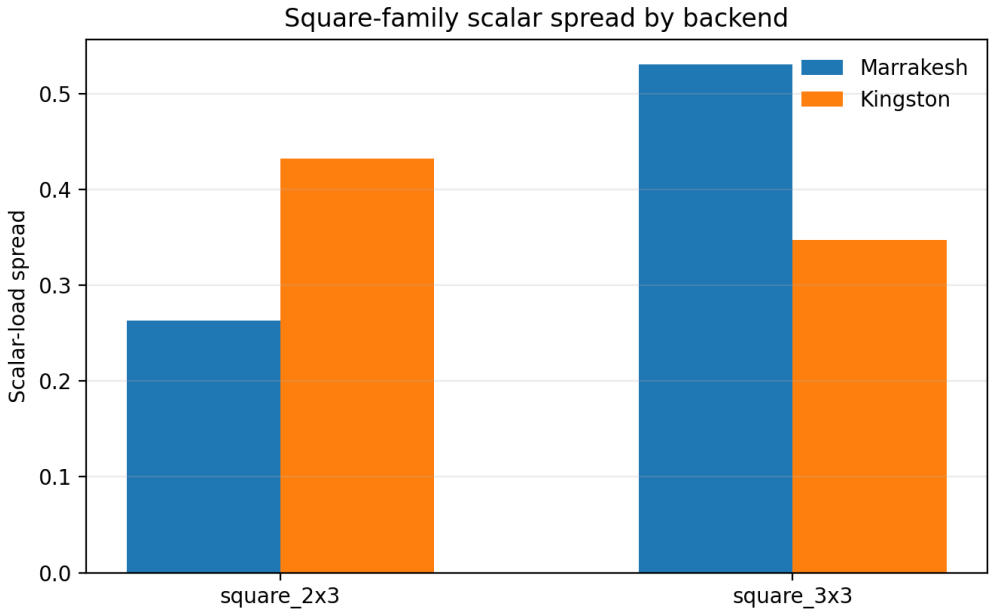

Figure 18: Square-family scalar-load spread. The square-family pair is now available on a single backend, with a house control in the same Kingston pass; repeated layouts and mitigation are still required before any finite-size conclusion can be drawn.

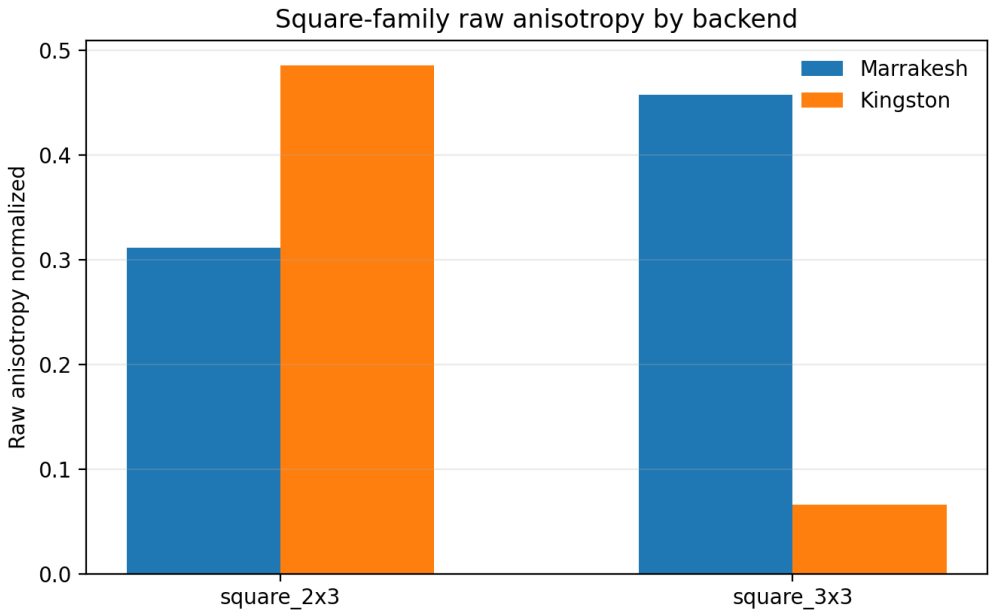

Figure 19: Square-family raw anisotropy. The large change in the Kingston raw anisotropy between `square_2x3` and `square_3x3` is not a spin-2 signal; it flags the need for fixed-layout repeats, tensor-channel observables, and scalar/vector impostor controls.

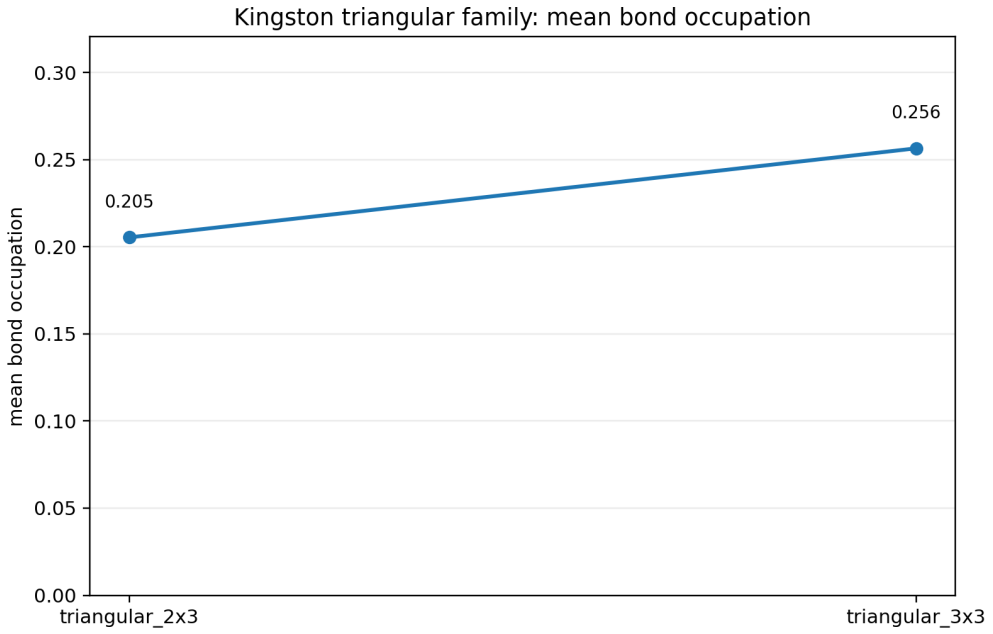

Figure 20: Kingston triangular-family mean bond occupation. The 25-qubit `triangular_3x3` run increases mean bond occupation relative to `triangular_2x3`.

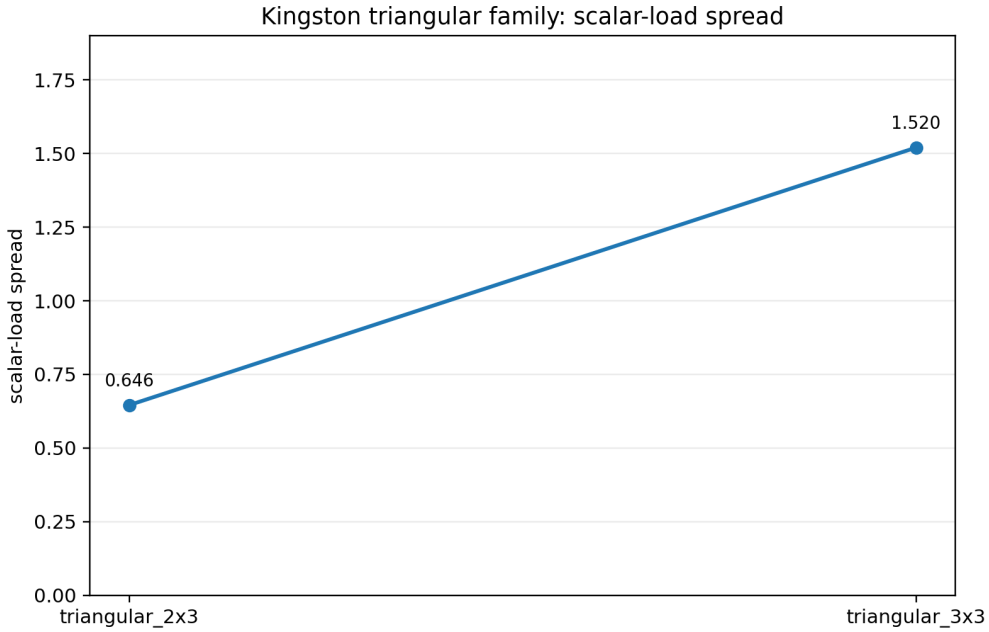

Figure 21: Kingston triangular-family scalar-load spread. The larger triangular graph produces a much larger scalar-load spread in this single hardware batch. The effect is hardware-facing and follow-up-worthy, but not yet interpretable as tensor scaling.

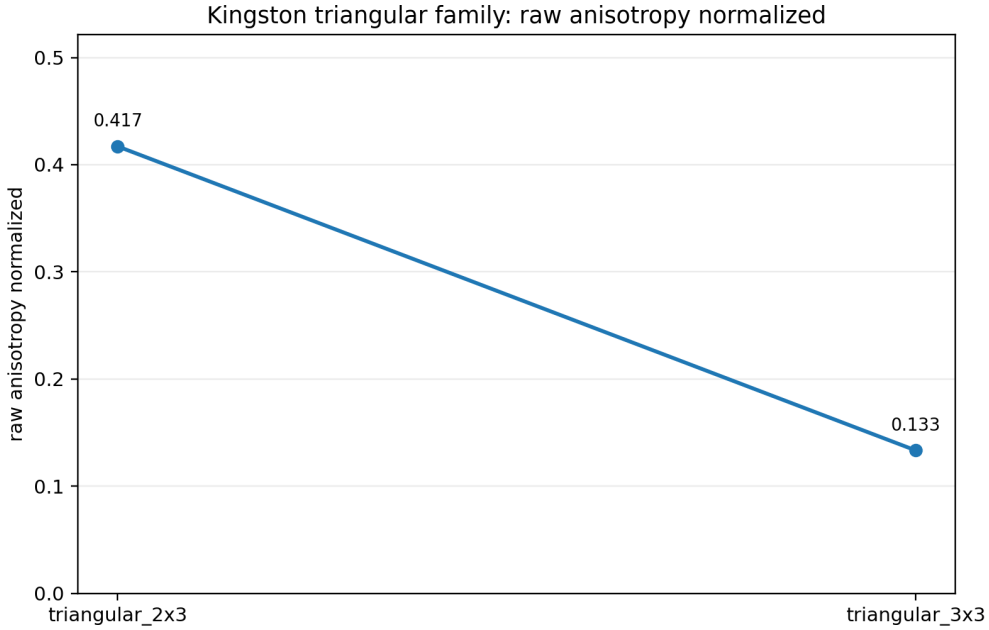

Figure 22: Kingston triangular-family raw anisotropy. Normalized raw anisotropy drops from `triangular_2x3` to `triangular_3x3`, even as scalar-load spread increases. This is the strongest "worth following up" pattern in the current hardware batch, but it remains a raw diagnostic rather than a spin-2 observable.

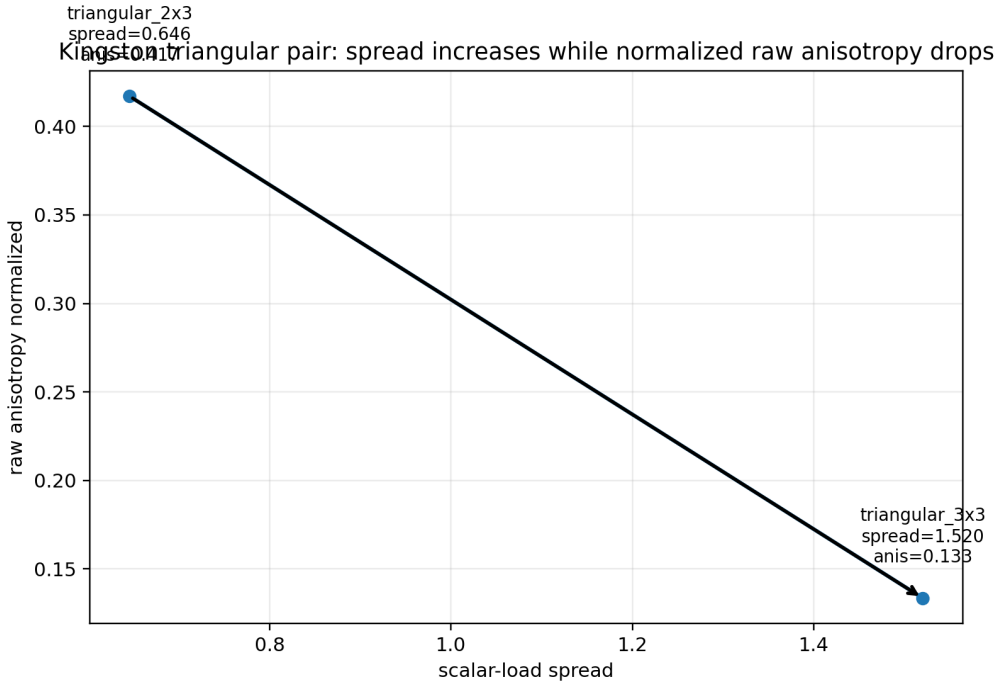

Figure 23: Same-backend triangular-pair summary. The move from `triangular_2x3` to `triangular_3x3` increases scalar-load spread and decreases normalized raw anisotropy. This is the clearest candidate-family follow-up signal, but a Mercury-rung claim still requires fixed-layout repeats, mitigation, tensor susceptibility, Ward-residual checks, and scalar/vector impostor rejection.

17.4 Hardened interpretation of the triangular pair

The Kingston triangular pair is useful, but the dimensional gate changes its status. It is not the correct next spin-2 lane; it is the strongest two-dimensional hardware pipeline-validation result. The hardened standard of this paper is

$$\text{scalar/bond signal} \neq \text{gravity}, \qquad \text{tensor-looking raw anisotropy} \neq \text{spin-2}. \tag{84}$$

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

$$F_{\mathrm{TT}}(k, L) = \frac{\left\| P^{\mathrm{TT}} S(k, L) P^{\mathrm{TT}} \right\|_F}{\|S(k, L)\|_F}, \tag{85}$$

$$K_{\mathrm{TT}}(0, k; L) \simeq Z^{-1} \left(k^2 + m_g^2(L)\right), \tag{86}$$

$$\mathcal{W}(k, L) = \frac{\|k^a S_{ab,cd}(k, L)\|_F + \|\delta^{ab} S_{ab,cd}(k, L)\|_F}{\|S_{ab,cd}(k, L)\|_F}, \tag{87}$$

$$R_{\mathrm{TT}}(L \to bL) = \mathrm{corr}\Big(P^{\mathrm{TT}} M^{(L)},\, \mathcal{C}_b \left[P^{\mathrm{TT}} M^{(bL)}\right]\Big), \tag{88}$$

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.

17.5 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 dimensional gate, however, prevents promoting those embedded sheets into 3+1 spin-2 candidates.

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

1. a three-dimensional graph family, beginning with a 2x2x2 cubic patch rather than another 2D sheet;
2. a tensor-channel observable beyond scalar loads and raw two-dimensional anisotropy;
3. a measured dynamical exponent z, with $z = 1$ required before any GR-like claim;
4. Ward-residual checks showing suppression of nonphysical polarizations in the infrared;
5. 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 em-*

bedded lattices are two-dimensional hardware validation families, not three-dimensional spin-2 admissible substrates. The matched Kingston triangular pair is the strongest two-dimensional hardware-facing signal because `triangular_3x3` *increases bond/scalar activation while reducing normalized raw anisotropy relative to* `triangular_2x3`*; nevertheless it is not spin-2 evidence. The next admissible run is a pre-registered 2x2x2 cubic graph.*

18 Phase-2A: first 3D cube hardware source-response runs

The dimensional gate identified the 2x2x2 cube as the first hardware fixture that is spatially three-dimensional in the v0 encoding. It has eight matter qubits, twelve bond qubits, and twenty hardware qubits total. It is a $k = 0$ source-response fixture. It cannot measure z, c_g, R_g, a Ward residual, or a transverse-traceless propagator. Its proper role is narrower: validate the three-dimensional encoding, the axis-aligned tensor observable, source metadata, and finite-difference tensor-summary extraction on real hardware.

The Phase-2A Kingston package executes exactly that role. The baseline cube and a matched x-axis traceless source pair ran on `ibm_kingston` with 2000 shots each. The source convention is

$$J_{ab} = J\,\mathrm{diag}(1, -1/2, -1/2), \qquad J = 0.05, \tag{89}$$

implemented as an axis-dependent bond-energy shift. The finite-difference proxy is

$$\chi^{\mathrm{src}}_{T,xx} \approx \frac{M_{xx}(+J_x) - M_{xx}(-J_x)}{2J}. \tag{90}$$

This is a uniform tensor-source proxy, not a continuum tensor susceptibility.

Run	Backend	Job ID	Submitted UTC	Shots	Distinct bitstrings	Depth/CZ
Baseline cube	`ibm_kingston`	`d8dm8gtmdsks73d44g70`	2026-05-30 22:27:22	2000	1922	1671/942
$+J_x$ source	`ibm_kingston`	`d8donca4gq0s73aqjlq0`	2026-05-31 01:15:36	2000	1938	1671/942
$-J_x$ source	`ibm_kingston`	`d8dong24gq0s73aqjlug`	2026-05-31 01:15:49	2000	1919	1671/942

Table 13: Phase-2A Kingston cube jobs. All three 4-step circuits used the same reported footprint: 156 physical qubits in the transpiled object, 20 classical bits, depth 1671, size 3602, and operations `cz=942`, `sx=1739`, `rz=901`, `measure=20`. The nearly maximal bitstring diversity is a warning that the physics readout is noise-limited.

Observable	Baseline	$+J_x$	$-J_x$	$\Delta(+,-)$	Derived
Mean bond occupation	0.246708	0.226667	0.230375	-0.003708	–
M_{xx}	-0.126833	-0.201667	-0.179000	-0.022667	$\chi^{\mathrm{src}}_{T,xx} = -0.226667$
M_{yy}	-0.206833	-0.042667	-0.040000	-0.002667	–
M_{zz}	0.333667	0.244333	0.219000	0.025333	–
$\|M\|_F$	0.412553	0.319670	0.285661	0.034009	–

Table 14: Phase-2A cube tensor summaries from the 4-step Kingston run. The finite-difference result is reported as a hardware observable, not as a physical negative susceptibility. The 4-step circuits are too deep and too bitstring-spread for a controlled tensor interpretation.

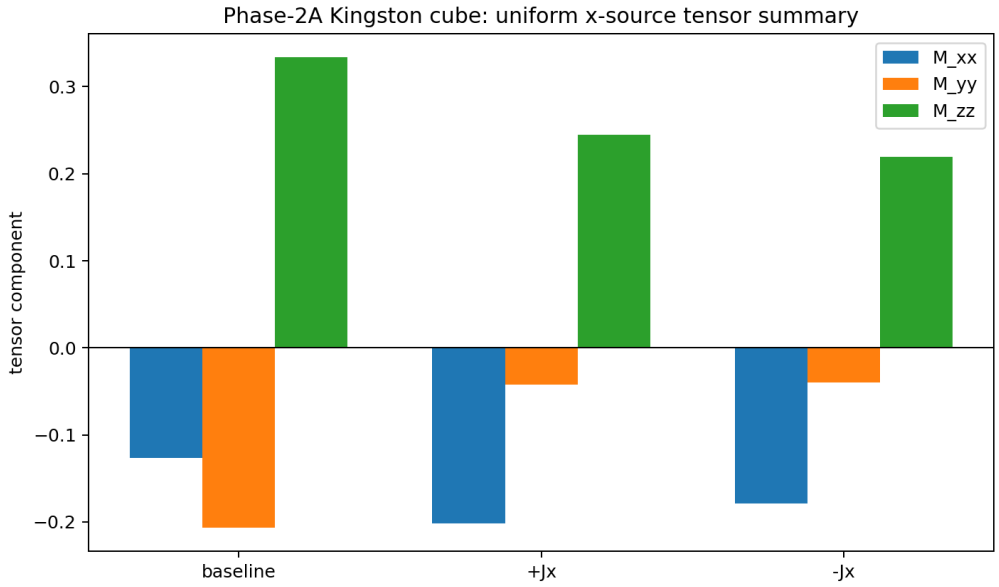

Figure 24: Axis-aligned tensor components for the Phase-2A 4-step Kingston cube baseline and $\pm J_x$ source pair. The components are nonzero, so the tensor-summary pipeline is active, but this is not a continuum tensor response or a spin-2 signal.

The operational result is positive: the 20-qubit cube fixture, source metadata, count capture, and tensor-summary analyzer all work end to end on an IBM backend. The physics interpretation is negative or, more precisely, noncommittal: the measured finite-difference signal is smaller than the noise/plumbing scale and cannot be used as a stable tensor susceptibility. The correct result statement is

> *Phase-2A hardware validates the dimensionally admissible cube source-response workflow. It does not establish spin-2 emergence, dispersion, Ward identities, z, c_g, or a Mercury number.*

18.1 Trotter-2 engineering check

Because the 4-step cube circuits were too deep for a clean open-plan physics readout, a shallower trotter2 source-pair fixture was generated and submitted on `ibm_kingston`. This was an engineering check, not a physics replacement. It cut the logical two-qubit gate count in half locally and reduced the transpiled Kingston depth from 1671 to 813 and the CZ count from 942 to 411.

Fixture	Depth	CZ count	Distinct bitstrings	$\chi^{\mathrm{src}}_{T,xx}$	Interpretation
4-step uniform source	1671	942	1919–1938	-0.226667	too deep; noisy hardware observation
2-step uniform source	813	411	1799–1802	0.278333	shallower; sign flips; engineering only

Table 15: Trotter-depth engineering comparison. The 2-step circuit is a better hardware substrate, but the sign flip relative to the 4-step pair shows that the uniform-source observable is not yet a stable hardware physics quantity.

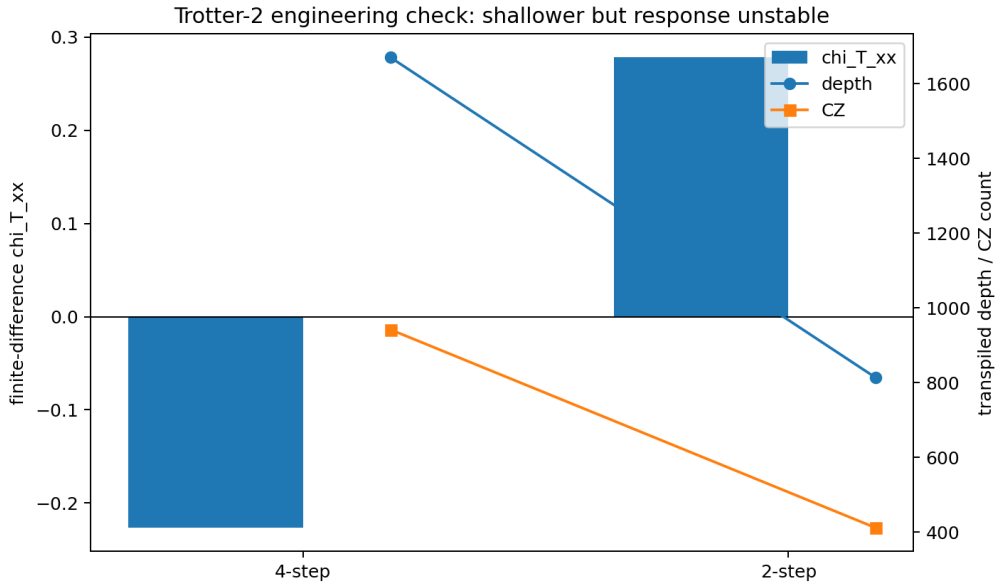

Figure 25: The trotter2 fixture reduces depth and CZ count, but the finite-difference response changes sign. This pushes the next hardware design toward source-off subtraction, amplitude scans, fixed layouts, and mitigation rather than repetition of the same uniform-source pair.

18.2 Compact-source local checks

A compact source is closer to the desired localized-source protocol than a uniform tensor source, but a 2x2x2 cube has no interior site. The first local compact-source fixtures therefore apply a traceless profile only to the three bonds incident to vertex 0:

$$x\text{-edge} : +J, \qquad y\text{-edge} : -J/2, \qquad z\text{-edge} : -J/2. \tag{91}$$

The local ideal checks show source-on/source-off mechanics and amplitude linearity, but also show that the response is small and boundary-dominated:

Compact source	ΔM_{xx}	$\chi^{\mathrm{src}}_{T,xx}$	$\|M(+J)\|_F$	$\|M(-J)\|_F$
$J = 0.05$	-0.001809	-0.018087	0.001144	0.001071
$J = 0.20$	-0.007201	-0.018002	0.004997	0.003822

Table 16: Local ideal compact-source checks on the 2x2x2 cube. The raw deformation scales with source amplitude, but the normalized response is small. The fixture validates circuit mechanics, not a continuum localized-mass response.

19 Phase-2D mechanism triage: bond-bond QMC hook

The full-ladder QMC null result closed the nominal per-edge promotion dynamics in the scanned region. The next mechanism-level question is whether an inter-edge channel can create a collective tensor response. The Phase-2 package adds a diagonal line-graph bond-bond coupling hook and runs a deliberately small triage on open 2x2x2 and 3x3x3 cubes. The run uses low statistics, small

$\chi_{\max} = 3$, $\beta = 1.0$, four imaginary-time slices, two seeds, and only three couplings. It is a machinery test, not a finite-size-scaling result.

Topology	K_{bb}	Mean bond occupation	$\chi_T/\lvert E\rvert$
2x2x2 cube	0.00	0.673958	0.276923
2x2x2 cube	0.02	0.718229	0.374405
2x2x2 cube	0.05	0.755208	0.388199
3x3x3 cube	0.00	0.674537	0.368208
3x3x3 cube	0.02	0.679630	0.390949
3x3x3 cube	0.05	0.795602	0.403963

Table 17: Phase-2D bond-bond mechanism triage, aggregated over two seeds. The hook can increase occupation and sometimes $\chi_T/|E|$, but seed spread is still large and the run is not a controlled scaling result.

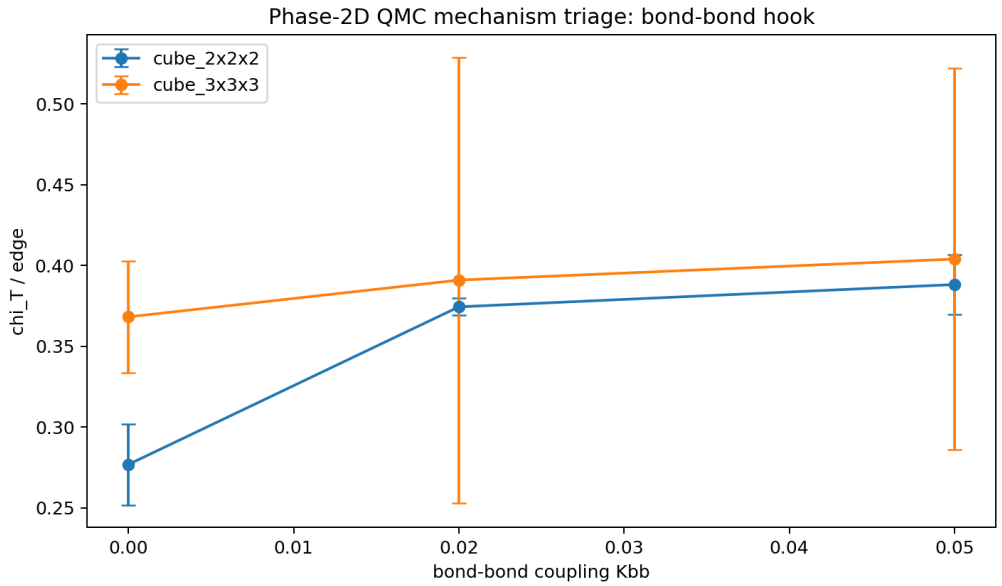

Figure 26: Bond-bond QMC mechanism triage for $\chi_T/|E|$. The small run suggests that the mechanism hook is active, but it is not yet evidence for collective tensor criticality.

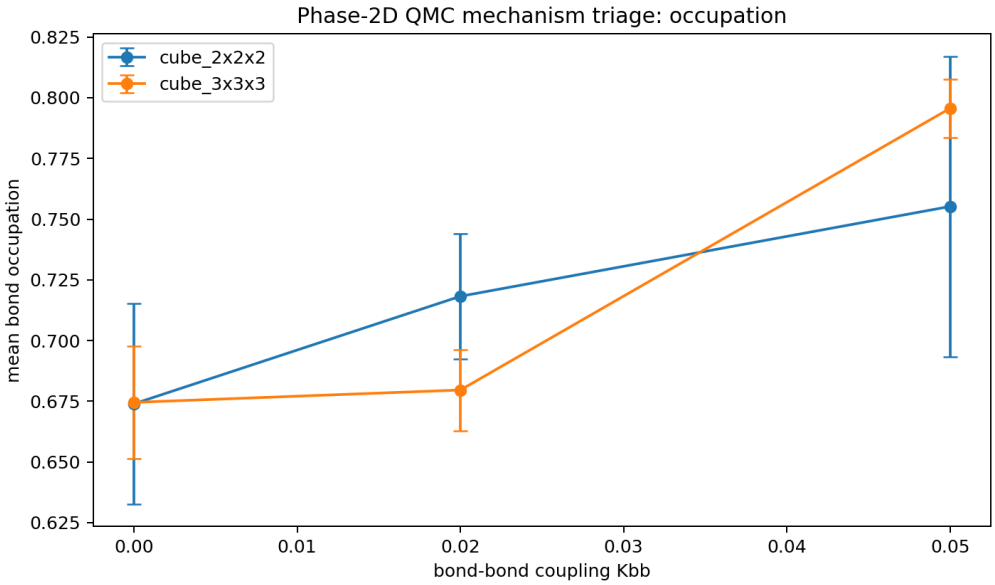

Figure 27: Mean bond occupation in the bond-bond QMC triage. The occupation response confirms that the hook changes the sampled ensemble; higher statistics, cutoff convergence, and a real L comparison are still required.

The triage changes the next-step recommendation. The next QPU job should not be another 4-step uniform-source cube pair. The better hardware direction is a compact-source or source-off subtraction design with fixed layout, amplitude linearity, and mitigation. The better theory/computation direction is a higher-statistics QMC convergence pass on the bond-bond mechanism at $K_{bb} = 0.02$ and 0.05, with more seeds and controlled $\chi_{\max}$, β, and sampling before any large-QPU geometry is attempted.

20 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`. The Phase-2A cube hardware and optimization update is documented in `phase2_kingston_2026-05-30_summary.md`, `phase2_cube_2x2x2_readiness.md`,

`phase2_cube_x_kingston_source_response.json`, and `phase2_optimization_qmc_triage.md`. Individual raw IBM/local count JSON files are intentionally ignored in git in the working repository; they are included in the local reproducibility package for this paper.

A hardware count analysis has the form

```
python src/analyze_counts.py \
  --counts <counts.json> \
  --topology <topology> --N <N>
```

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 conclusion is likewise not a positive spin-2 signal: it is active scalar/bond and raw-anisotropy data on two-dimensional embedded graph families, useful for pipeline validation but not admissible for the 3+1 spin-2 target. The next admissible topology is a three-dimensional cubic patch.

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

The Phase-0A and Phase-1A hardware results are gatekeepers and smoke tests, not the destination. The stoquasticity theorem changes the fastest route to the destination. The QPU path should not be asked to perform the thermodynamic finite-size scaling job first. Its near-term role is controlled source-response validation on the first three-dimensional fixture, while QMC carries the thermodynamic Mercury numbers.

1. **Phase-2A-QPU: 2x2x2 cube, 20 qubits.** This stage is now executed on `ibm_kingston`. It validates the 3D v0 encoding, defines the axis-aligned tensor observable, and exercises a uniform source-response proxy at $k = 0$. It cannot measure z, c_g, or R_g, and the first 4-step and 2-step hardware source pairs are too noisy/unstable for a physics interpretation.
2. **QPU source-pair tensor response.** The first uniform source pair used $J_{ab} = J\,\mathrm{diag}(1, -1/2, -1/2)$ and estimated

$$\chi_T \approx \frac{M_{xx}(+J) - M_{xx}(-J)}{2J}. \tag{92}$$

   The 4-step pair gave $\chi_{T,xx}^{\mathrm{src}} = -0.226667$; the shallower 2-step pair gave $+0.278333$. The sign flip prevents interpreting the uniform-source observable as a stable hardware physics quantity.
3. **Phase-2B-QMC: matter-core instrument validation.** The prototype sign-free QMC instrument has already validated the z-extraction logic on an anisotropic Ising benchmark. This is not a QATNU tensor result; it is the calibration of the measuring device.
4. **Phase-2C-QMC: full QATNU bond-register worldlines.** This stage is implemented and ED-validated. It measures rung worldlines, M_{ab}, source response, and Matsubara-frequency structure. In the scanned nominal neighborhood it returns a null: $\chi_T/|E|$ does not grow with L at converged $\chi_{\max}$.

5. **Phase-2D-QMC: mechanism modification.** A diagonal line-graph bond–bond coupling hook is now implemented and smoke-tested on 2x2x2 and 3x3x3 cubes. The hook changes occupation and sometimes $\chi_T/|E|$, but the run is low-statistics and not a scaling result. The next step is a higher-statistics, more-seed, cutoff-controlled QMC pass at $K_{bb} = 0.02$ and 0.05.
6. **Phase-2E-QMC: dispersion only after criticality.** If and only if Phase-2D finds an L-growing tensor channel, measure relaxation versus propagation and then z. Without a critical tensor channel, fitting a dispersion exponent would be empty.
7. **Hardware scaling after QMC triage.** The next hardware run should not repeat the deep 4-step uniform-source cube pair. It should use compact-source or source-off subtraction fixtures, amplitude-linearity checks, fixed layouts, and mitigation, and should be guided by the QMC-identified mechanism.
8. **Bond-reduction validation.** A bond-reduced encoding may reduce a 3x3x3 target from 81 to 27 qubits, but only if explicit-bond observables are shown to be recoverable from matter-only proxies such as $f(\langle Z_iZ_j\rangle, \langle X_iX_j\rangle, \ldots)$. If $\langle n_e\rangle$ is not recoverable, bond reduction is a different model rather than a compression.
9. **Ward residuals and source coupling.** Only after an admissible propagating tensor branch is seen should one test infrared Ward residuals, scalar/vector leakage, and universal source coupling.
10. **Lorentz recovery.** Only if $z = 1$, extract c_g and report $R_g = c_g/v_{\mathrm{LR}}$ as a Lorentz-recovery diagnostic.

A null result at any step is informative. If the 2x2x2 cube has no controlled tensor-source response, the hardware tensor-source program fails before dispersion. If QMC finds only relaxational tensor response, QATNU has a medium response rather than a graviton. If the three-dimensional tensor gap never approaches zero, QATNU does not generate a massless tensor in that phase. If the Ward residuals do not vanish, the tensor is not spin-2 gravity. If $z \neq 1$, the result belongs to the Lifshitz lineage rather than to GR. If R_g is not stable after $z = 1$ is established, the model has no emergent Lorentzian limit.

22 Claim registry

The current status is summarized in Table 18. The two most important updates are that the full-ladder QMC route has been implemented and used to produce a converged null for tensor criticality in the scanned nominal region, and that the first dimensionally admissible 2x2x2 cube hardware source-response workflow has now run on `ibm_kingston`. Both updates narrow the claim rather than inflate it.

Table 18: Claim registry for the Mercury rung after the full-ladder QMC null result and Phase-2A cube hardware update.

Claim	Status	Evidence or next test
QATNU is a finite local quantum substrate	Implemented	ED and hardware-readable qubit encodings

Claim	Status	Evidence or next test
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 or gauge protection and Ward identities
Conditional spin-2 derivation	Revised	Requires 3D infrared limit, propagating response, measured $z = 1$, chart stability, Ward residuals, and source coupling
Diffusive tensor branch	Added as leading null	$K_2 \sim m_2^2 + D_T k^2 - iZ_D\Omega$; medium response, not graviton
Lifshitz spin-2 outcome	Named partial success/failure	Gapless gauge-protected tensor mode with $z \neq 1$; not GR-like gravity
Negative-control exact simulations	Reported result	Corrected Phase-0A structural gate across path, cycle, star, diamond, bowtie
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, and triangular runs; 2D pipeline validation, not 3D spin-2 admissibility
Gauge-stoquasticity of full ladder	Proven under real positive amplitudes	Gauge $U = (\prod_i Z_i) \otimes (\prod_e P_e)$ makes off-diagonal elements non-positive

Claim	Status	Evidence or next test
Sign-problem-free equilibrium QMC	New route	Applies to thermal/imaginary-time observables; not a real-time claim
QMC matter-core instrument	Prototype validation	Recovers $z \approx 1$ on known anisotropic Ising benchmark; not full QATNU tensor result
Full QATNU tensor QMC	Implemented and ED-validated	Includes bond-register worldlines, M_{ab}, source response, and Matsubara fork
Tensor criticality in nominal neighborhood	Rejected in scanned region	At converged $\chi_{\max}$, $\chi_T/\lvert E\rvert$ is flat from $L = 2$ to $L = 3$ and dominated by single-edge variance
$\chi_{\max}$ convergence requirement	Established	Nominal converges by $\chi_{\max} \sim 5$–6; low-Δ_B requires $\chi_{\max} \sim 14$–16; hardware $\chi_{\max} = 2$ is insufficient in populated regimes
Inter-edge bond–bond coupling	Implemented as mechanism hook	Low-statistics 2x2x2/3x3x3 triage shows hook activity; requires higher-statistics, cutoff-controlled L scaling
First admissible 3D graph	Run as Phase-2A hardware workflow	2x2x2 cubic patch in v0 encoding: 8 vertices, 12 edges, 20 qubits; Kingston baseline and $\pm J_x$ source pair validate plumbing, not spin-2
2x2x2 uniform source response	Hardware engineering result	4-step pair gives $\chi^{\mathrm{src}}_{T,xx} = -0.226667$; 2-step pair gives $+0.278333$; sign flip blocks physics interpretation
2x2x2 as continuum	Rejected	3D embedding is not a $d_{\mathrm{eff}} \approx 3$ scaling window; no dispersion exponent

Claim	Status	Evidence or next test
Dispersion exponent z	Not yet measurable in hardware runs	Requires at least 2x2x3 for dirty single-axis estimate and 3x3x3 for first isotropic v0 target
Bond-reduced encoding	Required milestone	Validate explicit-bond $\langle n_e \rangle$ against matter-only proxies before reducing 3x3x3 from 81 to 27 qubits
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, propagation rather than relaxation, $\chi_{ab,cd}$, Ward residuals, and source coupling

23 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 related-constructions literature makes this discipline more important, not less. It shows that gapless tensor modes can arise in lattice and tensor-gauge settings, but also that the generic controlled outcome need not be relativistic GR. QATNU therefore has to distinguish five levels: scalar/bond activation, raw anisotropy in an embedded chart, relaxational tensor response, a Lifshitz or fracton-like tensor medium, and a true $3+1$ Pauli-Fierz sector.

The conditional spin-2 derivation is the bridge between the conceptual program and the next calculation. It now says exactly what must happen: the admissible graph family must be three-dimensional in the infrared; the tensor susceptibility must show a controlled source response; the response must be propagating rather than merely relaxational; the dynamical exponent must be measured rather than assumed; chart dependence must be controlled; and Ward residuals and source coupling must pass independently. If the branch is diffusive, it is a medium relaxation mode. If $z \neq 1$, the result may still be valuable, but it belongs to the Lifshitz/tensor-medium family rather than to GR-like gravity. If $z = 1$ but source coupling fails, the result is still not gravity.

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 a hardware boundary, not a spin-2 boundary. They show that the hardware path can execute larger embedded two-dimensional families and produce active bond/scalar and raw-anisotropy data. The matched Kingston triangular pair is the most interesting two-dimensional hardware-facing result so far: the larger graph increases mean bond occupation and scalar-load spread while normalized raw anisotropy decreases. But the dimensional gate prevents calling it a spin-2 candidate for the target of this paper.

The stoquasticity theorem changed the computational priority, and the full-ladder QMC run has now changed the scientific status. The next thermodynamic boundary is no longer implementation of bond-register worldlines; that has been done and validated. The boundary is mechanism: the present per-edge promotion dynamics does not generate an L-growing collective tensor susceptibility in the scanned region. The next test must add or engineer inter-edge correlations and ask whether the same QMC protocol then develops a tensor critical surface.

24 Conclusion

QATNU now has a Mercury-rung formulation, a real IBM hardware validation layer, and a stricter admissibility standard. The scalar timing sector is defined and bounded. The spin-2 target is stated without projection inflation. The conditional derivation now distinguishes four outcomes that must not be confused: scalar/bond activation, relaxational tensor response, Lifshitz or fracton-like tensor media, and a true 3+1-dimensional Lorentz-recovering Pauli-Fierz sector.

The hardware results remain useful, but their status is narrower than before. Phase 0A shows that small negative controls run on real IBM QPUs while remaining structurally rejected. Phase 1A shows that larger house, square, and triangular circuits also run and produce active bond/scalar and raw-anisotropy observables. The matched Kingston triangular pair is the strongest hardware-facing two-dimensional signal so far: mean bond occupation and scalar-load spread increase from `triangular_2x3` to `triangular_3x3`, while normalized raw anisotropy drops. But the spatial-dimension gate forbids promoting that result into a spin-2 candidate for 3+1-dimensional gravity. A two-dimensional sheet has zero local graviton polarizations for the target theory.

The next paper should not repeat two-dimensional hardware smoke tests unless they are used as noise or pipeline controls. The hardware instrument has now moved to Phase-2A: a 20-qubit 2x2x2 cube baseline and matched x-axis source pairs were run on `ibm_kingston`. This is the first dimensionally admissible QPU source-response execution, but it is not a controlled tensor susceptibility. The 4-step circuit is too deep and bitstring-spread for interpretation; the 2-step engineering fixture is shallower but flips the finite-difference sign. The next hardware step should be compact-source/source-off subtraction, amplitude scans, fixed-layout repeats, and mitigation, not a larger unmitigated uniform-source circuit.

The thermodynamic instrument has also moved: full-ladder, sign-free, native-3D QMC is implemented and ED-validated. It carries bond-register worldlines and the M_{ab} tensor observable. The result is a negative one for the nominal per-edge substrate. At converged $\chi_{\max}$ in the natural parameter neighborhood, $\chi_T/|E|$ does not grow with L, and the apparent enhancement at low Δ_B is a displaced-oscillator, single-edge variance effect rather than a collective rank-2 critical mode. The first bond–bond mechanism hook is now implemented and smoke-tested, but only as a triage. It suggests the hook is active; it does not yet show tensor criticality. This does not end the QATNU program, but it shifts the next falsifiable test to mechanism design: find an inter-edge or correlated-frustration channel that creates an L-growing tensor response under the same cutoff-converged

QMC protocol. If such a mechanism appears, the next gates are relaxation versus propagation, z, chart stability, Ward residuals, source coupling, and only then $R_g = c_g/v_{\mathrm{LR}}$.

A Hardware run prompt

Run or reproduce the QATNU Mercury-rung hardware sequence with the corrected dimensional gate. Phase-0A path, star, cycle, and diamond circuits are negative controls. Phase-1A house, square, and triangular circuits are two-dimensional pipeline-validation tests, not 3D spin-2 candidates. Phase-2A is the 2x2x2 cubic fixture: 8 matter qubits plus 12 bond qubits, 20 hardware qubits total. Save raw bitstring counts with backend name, job id, shot count, layout/transpilation metadata, calibration date if available, and source metadata.

The first Phase-2A Kingston run has already executed the isotropic cube baseline and a uniform x-axis traceless source pair. To reproduce it, run the baseline and $\pm J_x$ fixtures with the same source convention $J_{ab} = J\,\mathrm{diag}(1, -1/2, -1/2)$ and report axis sums Q_x, Q_y, Q_z, the diagonal traceless M_{ab}, $\|M\|_F$, and the finite-difference proxy $\chi_T = [M_{xx}(+J) - M_{xx}(-J)]/(2J)$. Do not claim spin-2: 2x2x2 does not measure z, c_g, R_g, Ward residuals, or a TT propagator.

Do not submit another deep 4-step uniform-source pair as the next physics run. The current 4-step result is too depth/noise dominated and the 2-step engineering pair flips sign. The next hardware design should use source-off subtraction, compact-source or amplitude-scan fixtures, fixed physical layout, readout mitigation, and a local/noise-calibrated preflight before interpretation.

B Native-3D QMC next-test prompt

Extend the ED-validated full-ladder QATNU worldline QMC instrument. Keep the gauge-stoquastic Hamiltonian in the matter-Z and bond-rung product basis, with positive imaginary-time weights. Preserve the validated ingredients: matter worldlines, diagonal bond-rung energies, frustration-gated bond-rung hopping vertices with weights proportional to $\lambda\sqrt{m+1}\Delta\tau$, exact bond-worldline heat-bath resampling, the M_{ab} observable, source response, and Matsubara diagnostics.

Validate the next model change in three stages. First, add either a direct bond–bond coupling or spatially correlated frustration, with the null model retained as the control. Second, reproduce the ED validation table and the $\chi_{\max}$ convergence table; do not interpret any tensor number at an unconverged cutoff. Third, rerun the L-scaled intensive susceptibility test. Only if $\chi_T/|E|$ grows with L should the code proceed to propagation-versus-relaxation and z extraction. If the susceptibility remains flat, record the result as a Mercury-rung null for that modified substrate.

C 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	Implemented in full-ladder QMC; not yet a spin-2 pass
$K_{ab,cd}$	Inverse tensor susceptibility	Matsubara fork implemented; classification ambiguous on smallest cube
m_2^2	Tensor mass parameter	No gap-closing inferred in scanned nominal region
Z_2, ρ_2	Tensor kinetic and stiffness coefficients	Determine c_g
c_g	Emergent tensor-wave speed	Requires a propagating critical tensor branch; not reached
v_{LR}	Lieb-Robinson / substrate causal-cone velocity	Defines Lorentz-recovery diagnostic
R_g	c_g/v_{LR}	Lorentz-recovery diagnostic after $z = 1$ is established
$\mathcal{W}_{\text{div}}, \mathcal{W}_{\text{tr}}$	Ward residuals	Gauge/constraint diagnostics

D 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. Z.-C. Gu and X.-G. Wen, "Emergence of helicity ±2 modes (gravitons) from qubit models," Nuclear Physics B 863, 90–129, 2012; arXiv:0907.1203.
7. Z.-C. Gu and X.-G. Wen, "A lattice bosonic model as a quantum theory of gravity," arXiv:gr-qc/0606100.
8. C. Xu, "Novel algebraic boson liquid phase with soft graviton excitations," Physical Review B 74, 224433, 2006.
9. C. Xu and P. Hořava, "Emergent gravity at a Lifshitz point from a Bose liquid on the lattice," Physical Review D 81, 104033, 2010.
10. M. Pretko, "Emergent gravity of fractons: Mach's principle revisited," Physical Review D 96,

024051, 2017.

11. M. Pretko and L. Radzihovsky, "Fracton-elasticity duality," Physical Review Letters 120, 195301, 2018.
12. S. Deser, R. Jackiw, and G. 't Hooft, "Three-dimensional Einstein gravity: dynamics of flat space," Annals of Physics 152, 220–235, 1984.
13. E. Witten, "(2+1)-dimensional gravity as an exactly soluble system," Nuclear Physics B 311, 46–78, 1988.
14. H. van Dam and M. Veltman, "Massive and massless Yang-Mills and gravitational fields," Nuclear Physics B 22, 397–411, 1970.
15. V. I. Zakharov, "Linearized gravitation theory and the graviton mass," JETP Letters 12, 312, 1970.
16. C. W. Misner, K. S. Thorne, and J. A. Wheeler, *Gravitation*. W. H. Freeman, 1973.
17. R. M. Wald, *General Relativity*. University of Chicago Press, 1984.
18. IBM Quantum Documentation, Qiskit Runtime and IBM Quantum Platform documentation, accessed May 2026. https://quantum.cloud.ibm.com/docs
19. PyQuafu Documentation, "Welcome to PyQuafu's documentation," accessed May 2026. https://scq-cloud.github.io/
20. PyQuafu Documentation, "User guide," accessed May 2026. https://scq-cloud.github.io/0.2.x/index.html