DATED HISTORY

The Mercury Rung: Scalar Timing, Spin-2 Emergence, and Phase-0 Negative Controls in QATNU

Phase-0 negative-control installment. The consolidated search combines the controls with the candidate and source-response tests.

Manuscript · p. 1, negative controls

Spin-2 search · null result

Evidence · Negative QATNU spin-2 search

The Mercury Rung

Scalar Timing, Spin-2 Emergence, and Phase-0 Negative Controls in QATNU

Joshua Farrow
Independent Research

May 2026

Status. This paper gives a complete formulation of the QATNU "Mercury rung": the scalar timing sector, the spin-2 emergence target, the conditional derivation of a Pauli-Fierz-like tensor sector from bond-stress criticality and emergent Ward identities, and a Phase-0 exact-simulation result on small negative-control graphs. It does not claim that quantum gravity has been derived. It reports the first result that the diagnostic can reject inadmissible substrates while measuring active bond and scalar observables, and it supplies hardware-ready circuits for repetition on real quantum processors.

Abstract

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

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

Fourth, the paper reports a Phase-0 exact-simulation result. A hardware-readable qubit encoding with one matter qubit per vertex and one bond qubit per edge was run on four negative-control graphs: a path, a cycle, a star, and a diamond plaquette. All four show nonzero bond occupation and scalar load, but all four correctly fail the spin-2 admissibility test. This is the intended first result: the diagnostic can say no. The accompanying code package emits OpenQASM 2 circuits and includes IBM Quantum and Quafu/BAQIS runners, so the same controls can be executed on real superconducting quantum processors once account tokens are supplied. The next admissible experiment is a finite-size scaling study on embedded graph families with a genuine low-momentum tensor sector.

Contents

1 The problem **3**

2 The Mercury criterion **3**

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

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

5 The spin-2 target **6**

6 No-go constraints **7**

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

8 Phase-0: what can be tested now **9**

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

10 Phase-0 exact-simulation result **10**
10.1 Interpretation of the result 12

11 Real quantum hardware path **12**
11.1 IBM Quantum 12
11.2 Quafu/BAQIS 13
11.3 Hardware claim boundary 13

12 Reproducibility package **13**

13 What remains to make this a positive spin-2 paper **14**

14 Claim registry **14**

15 Discussion **15**

16 Conclusion **15**

A Hardware run prompt **15**

B Notation **16**

C References **16**

1 The problem

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

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

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

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

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

2 The Mercury criterion

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

For QATNU, the corresponding requirement is:

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

This criterion has three parts.

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

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

The rest of the paper is organized around this criterion.

3 The QATNU substrate

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

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

A minimal ladder Hamiltonian is

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

with

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

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

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

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

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

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

3.1 Implemented exact-diagonalization form

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

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

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

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

hardware encoding.

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

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

4 The scalar timing sector

The most immediate coarse observable is bond occupation. Let

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

The scalar load at vertex i is

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

The clock law is

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

so that in weak load

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

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

4.1 Weighted graph closure

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

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

A leading scalar closure is

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

with dynamical completion

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

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

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

The same operator defines a diffusion return probability,

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

and an effective dimension,

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

A shell profile gives an independent estimator,

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

The severe scalar closure test is therefore

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

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

4.2 Why the scalar sector is not enough

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

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

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

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

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

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

5 The spin-2 target

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

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

It is invariant under the linearized gauge transformation

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

In the source-free transverse-traceless sector,

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

and the wave equation is

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

The TT projector in momentum space is

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

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

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

6 No-go constraints

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

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

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

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

7 Conditional derivation of the QATNU spin-2 sector

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

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

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

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

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

7.1 Tensor susceptibility

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

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

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

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

Define the connected susceptibility

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

The inverse tensor kernel is

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

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

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

A perturbative fast-sector elimination gives the schematic tensor mass

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

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

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

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

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

This supplies the primary Mercury number:

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

7.2 Ward residuals

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

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

The infrared Ward condition is

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

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

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

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

7.3 Conditional theorem

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

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

8 Phase-0: what can be tested now

The full tensor theorem requires large embedded graph families and finite-size scaling. The current near-term test is smaller: verify that the diagnostic says no on graph families that cannot support spin-2 in the first place.

The Phase-0 negative controls are:

- a one-dimensional path;
- a finite cycle;
- a hub-dominated star;
- a single diamond plaquette.

The test is successful only if these fail the spin-2 admissibility check while still producing nonzero scalar and bond observables. In other words, Phase 0 asks whether the exam has a working "no" answer.

9 Hardware-readable Phase-0 Hamiltonian

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

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

The Hamiltonian is

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

with

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

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

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

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

The initial state is

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

The parameters used here are

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

Bond occupation, scalar load, load spread, raw anisotropy, and the many-body ground gap are computed after exact statevector evolution. The raw anisotropy diagnostic in the assigned two-dimensional chart is

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

This is not a spin-2 claim. It is a minimal centered directional diagnostic for the negative controls.

10 Phase-0 exact-simulation result

Table 2 gives the exact statevector results. All four graph families return the intended status:

$$\texttt{FAIL_EXPECTED_NEGATIVE_CONTROL}. \tag{55}$$

Graph	Matter	Edges	Qubits	$\langle n_e \rangle$	$\Delta\Lambda$	Gap	Spin-2 result
Path	4	3	7	0.1288	0.1228	0.6491	fail expected
Cycle	4	4	8	0.1301	0.0000	0.5485	fail expected
Star	4	3	7	0.1305	0.2453	0.6279	fail expected
Diamond	4	4	8	0.1301	0.0000	0.5485	fail expected

Table 2: Phase-0 exact-simulation results for the hardware-readable QATNU encoding. The scalar/bond sector is active, but none of the graph families is admissible as an infrared spin-2 substrate.

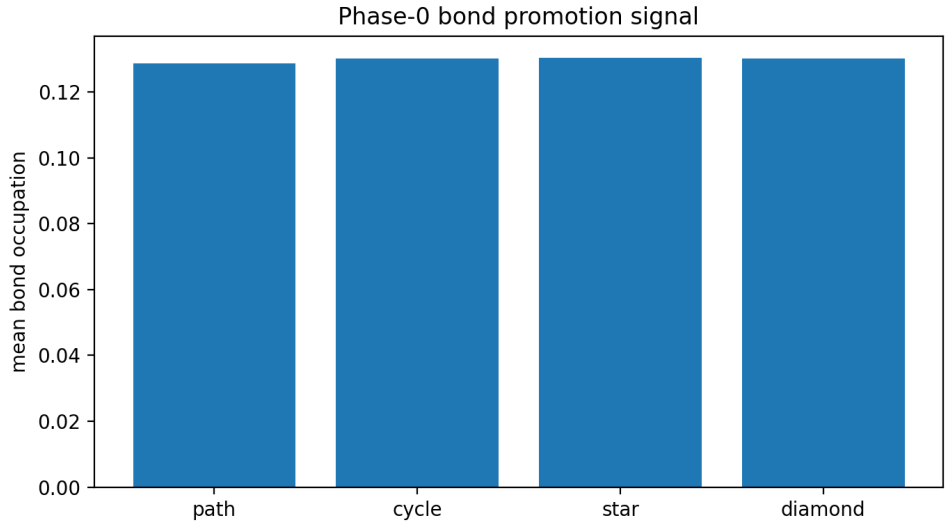

Figure 1: Mean bond occupation after evolution. All tested graphs show nonzero bond activation.

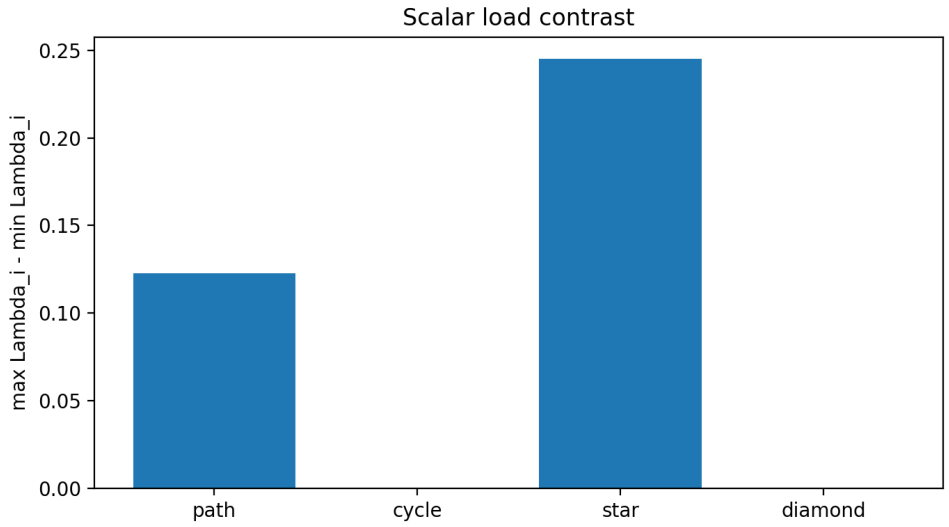

Figure 2: Scalar load spread. The path and star have nonuniform load; the symmetric cycle and diamond are uniform under this state and parameter choice.

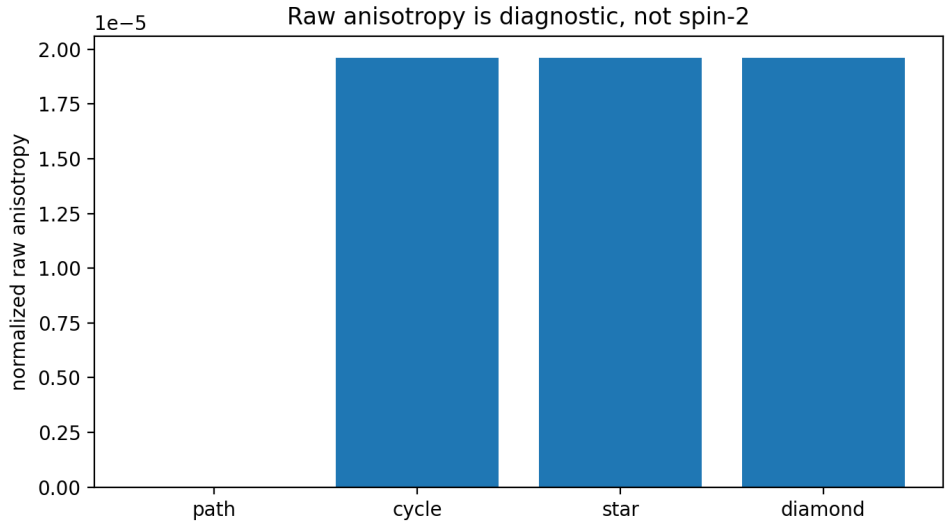

Figure 3: Raw centered anisotropy diagnostic in the assigned two-dimensional charts. Values near numerical zero are expected for the symmetric controls. This diagnostic is not a TT derivation.

10.1 Interpretation of the result

The result is deliberately modest. It proves three things and only three things.

1. The QATNU qubit encoding is active: bond occupation and scalar load are nonzero.
2. The spin-2 diagnostic is not automatically positive: inadmissible graph families are rejected.
3. The next experiment must move to admissible embedded graph families with a low-momentum window and enough angular structure for tensor susceptibility extraction.

This is why the Phase-0 result belongs in the Mercury rung paper. Before asking whether QATNU can generate a spin-2 sector, the protocol must show that it can refuse one.

11 Real quantum hardware path

The accompanying code package emits OpenQASM 2 circuits for the Phase-0 graphs and includes two runners:

- an IBM Quantum runner using Qiskit Runtime;
- a Quafu/BAQIS runner using PyQuafu.

The hardware circuits are not intended to discover gravity. Their purpose is to test whether the scalar/bond observables and negative-control classification survive compilation, noise, finite shot count, and real backend constraints.

11.1 IBM Quantum

IBM's Open Plan is the most straightforward general route for a first reproduction. IBM currently advertises free Open Plan access to real QPUs with up to 10 minutes of quantum runtime per 28-day rolling window, and its documentation describes Open Plan instances as a free route for running circuits on QPUs. The package runner expects an environment variable

`IBM_QUANTUM_TOKEN`

and submits an OpenQASM circuit through Qiskit Runtime.

11.2 Quafu/BAQIS

Quafu is the China-based route. The PyQuafu documentation states that Quafu shares open access to superconducting quantum computing resources with public users and provides experimental backends through its Python SDK. The package runner expects

`QUAFU_API_TOKEN`

and can list available backends before submission.

11.3 Hardware claim boundary

No real-QPU counts are reported in this paper. That is intentional: hardware submission requires a user account, API token, backend availability, and queue completion. The scientific claim here is that the circuits and runners are ready, and that the exact-simulation baseline is fixed before hardware noise enters.

12 Reproducibility package

The code package contains:

- exact statevector simulation: `src/phase0_core.py`;
- result reproduction: `src/run_phase0_exact.py`;
- figure generation: `src/generate_plots.py`;
- hardware analysis: `src/analyze_counts.py`;
- IBM runner: `src/run_ibm_qpu.py`;
- Quafu runner: `src/run_quafu_qpu.py`;
- circuits: `qasm/qatnu_phase0_*_N4.qasm`;
- exact results: `results/phase0_summary.csv` and `results/phase0_summary.json`.

The local baseline is reproduced by

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

The first hardware circuit to run is

`qasm/qatnu_phase0_path_N4.qasm`

The expected scientific result is not a positive spin-2 signal. It is

`FAIL_EXPECTED_NEGATIVE_CONTROL`

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

The Phase-0 result is a gatekeeper, not the destination. A positive spin-2 paper must add the following.

1. **Admissible graph families.** Use embedded two- and three-dimensional graph families with increasing size and a real low-k window.
2. **Tensor source response.** Add $J_{ab}M_{ab}$ perturbations and compute $\chi_{ab,cd}$.
3. **Critical scan.** Fit $m_2^2(\lambda) = \Delta_B^{(2)} - \lambda^2 \Pi_2(0,0)$ and test whether a tensor critical point exists.
4. **Dispersion.** Extract Z_2, ρ_2, and $c_g^2 = \rho_2/Z_2$.
5. **Ward residuals.** Show that scalar/vector leakage and trace/divergence residuals decrease in the infrared.
6. **Mercury number.** Report $R_g = c_g/c_{\text{clock}}$ as a prediction, not a fit.

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

14 Claim registry

The current status is summarized in Table 3.

Table 3: Claim registry for the Mercury rung.

Claim	Status	Evidence or next test
QATNU is a finite local quantum substrate	Implemented	ED and hardware-readable qubit encodings
Scalar load Λ_i is measurable	Implemented	Bond occupation and scalar load in Phase 0
Scalar timing can be represented by $\omega_{\text{eff}} = \omega/(1+\alpha\Lambda)$	Supported as effective law	Requires larger-system validation
Scalar timing is gravity	Rejected	Scalar sector is holonomy-free and lacks spin-2
Spin-2 target is Pauli-Fierz/TT	Standard benchmark	Known weak-field target
QATNU derives spin-2 unconditionally	Not claimed	Requires tensor criticality and Ward identities
Conditional spin-2 derivation	Formulated	Eqs. (39)–(45)
Negative controls fail	Reported result	Phase-0 exact simulations
Real-QPU execution	Prepared, not yet run	OpenQASM plus IBM and Quafu runners

Claim	Status	Evidence or next test
Positive spin-2 sector	Open	Requires admissible graph scaling and $\chi_{ab,cd}$

15 Discussion

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

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

The Phase-0 result is therefore not small in method, even though the graphs are small. It shows that the software/hardware protocol has a way to refuse false positives. That is the first requirement for any credible emergence claim.

16 Conclusion

QATNU now has a complete Mercury-rung formulation. The scalar timing sector is defined and bounded. The spin-2 target is stated without projection inflation. The conditional derivation identifies tensor criticality, Ward identities, and $R_g = c_g/c_{\text{clock}}$ as the quantities that decide whether the substrate can produce gravitational dynamics. The first executable result has been run: small negative-control graphs show active bond/scalar observables and correctly fail the spin-2 admissibility test.

The honest conclusion is not that QATNU has derived gravity. It is stronger than another proposal and weaker than a discovery: QATNU now has a result-bearing test. The next paper should run the tensor susceptibility program on admissible embedded graph families and report either the first nontrivial approach to a spin-2 sector or the first serious null result.

A Hardware run prompt

> Run the QATNU Phase-0 negative-control experiment on a real quantum processor. First reproduce the exact statevector baseline with `python src/run_phase0_exact.py`. Then submit the smallest OpenQASM circuit, `qasm/qatnu_phase0_path_N4.qasm`, to either IBM Quantum or Quafu/BAQIS. Save raw bitstring counts as JSON with backend name, job id, shot count, and calibration date if available. Analyze counts with `python src/analyze_counts.py --counts <counts.json> --topology path --N 4`. Report mean bond occupation, scalar loads, scalar load spread, raw anisotropy, and `spin2_phase0_result`. The expected result is `FAIL_EXPECTED_NEGATIVE_CONTROL`. A positive spin-2 claim from this graph means the diagnostic is not severe enough.

B Notation

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

C References

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