The Holographic Quantum Stack Machine

A finite boundary-encoded architecture for particle-like states

David Elliman

5 November 2026

Abstract

The Holographic Quantum Stack Machine (HQSM) is a finite quantum architecture in which distributed registers carry internal degrees of freedom and pairs of boundary addresses define an interior spatial representation. We give a mathematical specification of an eight-site prototype and distinguish its exact properties from its proposed interpretation in particle physics. Explicit $[[8, 2, 2]]_2$ and $[[8, 2, 2]]_3$ stabilizer encodings hide logical information from any single ledger site. A 64-dimensional chord register contains an isometric 32-dimensional shell–anchor representation with an exactly computable radial hopping operator. Integer registers provide reversible bookkeeping for selected quantum numbers. We replace an unrestricted pair of controlled-NOT gates by a chirality-selected exchange that conserves electric charge on the entire state space, and describe auxiliary-sector elimination using the Schur complement. A three-component mass ansatz reproduces the Koide identity only on the appropriate positive-root branch; its proposed fixed phase is incompatible with the precision muon–electron mass ratio. An explicitly calibrated version gives a testable, conditional tau-mass estimate. Positive boundary weights yield a bounded barycentric map, but do not supply a spacetime metric or gravitational field equations. Reproducible calculations separate these finite algebraic results from the additional dynamics, symmetries and continuum limits required for a physical theory. The construction is presented as a computational model and research framework, with derivations accessible to readers familiar with undergraduate linear algebra and quantum mechanics.

Keywords: quantum architecture; stabilizer codes; boundary encoding; effective Hamiltonians; quantum simulation; lepton mass relations.

1 Introduction and scope

Computer architecture offers a useful language for separating the definition of an object, its stored state, its address and the operations acting on it. In quantum physics the corresponding separation is between a species of excitation, its Hilbert space, a spatial representation and an evolution law. The HQSM explores this analogy with explicit finite matrices. Its starting question is computational: can a small boundary register encode internal quantum information and an interior position representation, while maintaining reversible dynamics and specified conservation laws?

This question is related to quantum simulation, whose purpose is to reproduce the dynamics and measurement statistics of a quantum model [1]. It also draws inspiration from holographic descriptions and the relation between bulk reconstruction and quantum error correction [2, 3]. Those theories provide motivation rather than a derivation of the present architecture. Here, "holographic" means that an interior description is encoded in designated boundary-address and ledger registers. It does not assert an AdS/CFT duality, an entropy–area law or a theory of quantum gravity.

The identity of electrons does not require a common mutable memory location. In quantum field theory, electrons are excitations of the same field, with common couplings and fermionic statistics [4]. Likewise, sharing a class definition in software does not force different instances to share their runtime state. We therefore separate a shared classical species table from independent logical quantum degrees of freedom. A pointer selects a logical subsystem or spatial address; it does not clone an unknown quantum state [5].

The paper has three aims. First, it gives definitions and short proofs for the parts of HQSM that can be made exact: finite encodings, register operations, conservation tests and auxiliary-state elimination. Second, it provides a corrected reference implementation for two problematic extensions: the weak-interaction gate analogy and the charged-lepton mass ansatz. Third, it identifies the additional results needed to connect a finite architecture to measured particle physics. Neither an exact matrix identity nor a passing software test, by itself, establishes that the matrix describes nature.

1.1 How to read the model

We use three levels of statement throughout. A *definition* specifies a chosen model. A *finite-model result* follows from that definition and can be proved or checked. A *physical hypothesis* identifies a model quantity with an observable and must be tested independently. For example, a negative eigenvalue shift of a finite Hamiltonian is a result; calling its absolute value an electron mass is an additional hypothesis.

Table 1: The scope of the construction. The third column states what the finite result does not establish.

Component	Established within the specified model	Additional physical requirement
Boundary ledger	Two logical qudits protected against one known-site erasure	Scalable storage, preparation and physically local access
Chord pointers	Four exact chord-radius values and an isometric shell representation	A rotationally and relativistically consistent continuum limit
Register transactions	Reversible shifts and selected conserved quantities	Gauge dynamics, spin representations and scattering amplitudes
Auxiliary sectors	Exact stationary Schur reduction	Controlled dynamical approximation and physical mass matching
Generation ansatz	A conditional Koide relation	A mechanism fixing parameters and radiative corrections
Barycentric map	Convexity and an antipodal tanh identity	Metric, causal structure and gravitational field equations

1.2 Notation and prerequisites

A ket $|\psi\rangle$ is a column vector and its adjoint $\langle\psi|$ is the conjugate row vector. A normalized pure state satisfies $\langle\psi \mid \psi\rangle = 1$. A density matrix ρ is positive and has unit trace; it represents pure or mixed states. A tensor product describes jointly available degrees of freedom. An *isometry* V satisfies $V^\dagger V = \mathbb{I}$ and preserves inner products while embedding a smaller space in a larger one. A projector $P = P^2 = P^\dagger$ selects a subspace. An observable A is conserved by a time-independent Hamiltonian H when $[H, A] = HA - AH = 0$.

Except when quoting particle masses in MeV, numerical energies and lengths in the finite prototype are in model units. We use $\hbar = c = 1$ in dynamical equations and the $+---$ metric convention when discussing four-momentum. The physical local dimension $d \in \{2, 3\}$ is distinct from the code distance, which is two in both constructions. A qubit has dimension two and a qutrit dimension three. The cube graph is denoted Q_3; it is not a qutrit.

2 Architecture and resource accounting

Let $\mathcal{V} = \{0,1\}^3$ be the eight vertices of a cube. Vertices are connected when their binary labels differ in one coordinate. The graph has 12 edges. Adding its six square faces gives a closed cell complex homeomorphic to a two-sphere. The graph alone is a one-dimensional complex; it is the boundary cell complex that justifies the term "surface". Figure 1 gives the motivating architectural overview; figure 2 separates the two encodings used in the formal model.

For each local dimension d define a physical ledger

$$\mathcal{H}_{\mathrm{led},d} = (\mathbb{C}^d)^{\otimes 8}, \qquad V_d : \mathcal{L}_d = (\mathbb{C}^d)^{\otimes 2} \longrightarrow \mathcal{H}_{\mathrm{led},d}. \tag{2.1}$$

The image of V_d is a stabilizer code. The qubit and qutrit ledgers may be used together, in which case the physical ledger dimension is 6^8 and its logical dimension is $2^2 3^2 = 36$. This counts internal storage only.

A separate pointer register is

$$\mathcal{P} = \mathrm{span}\{|v, w\rangle : v, w \in \mathcal{V}\} \simeq \mathbb{C}^8 \otimes \mathbb{C}^8. \tag{2.2}$$

It needs six qubits in a binary address implementation. For N labelled pointer registers the address space has dimension 64^N. Allocating another pointer therefore requires additional capacity; it is not supplied by a two-logical-qudit code. A closed quantum circuit has a fixed Hilbert space, so allocation must mean activating pre-existing blank registers or moving between occupation sectors of a specified Fock space. For identical fermions, the appropriate many-particle construction is the antisymmetric Fock space over the one-particle space. Antisymmetry is an imposed kinematic structure, not a consequence of shared addressing.

The address v names a site but does not by itself specify where the bits representing that address are physically stored. Thus a claim that all hardware is confined to the boundary still requires a placement and routing scheme for the pointer, stack and controller registers. The present specification treats their storage as additional boundary resources. Error correction of the ledger does not automatically protect these resources.

This distinction also sets a locality requirement. If a logical operator has support on several sites, it needs a physical implementation using available couplings, gates or measurements. The absence of a bulk distance in the encoding does not make such an operation instantaneous or cost-free.

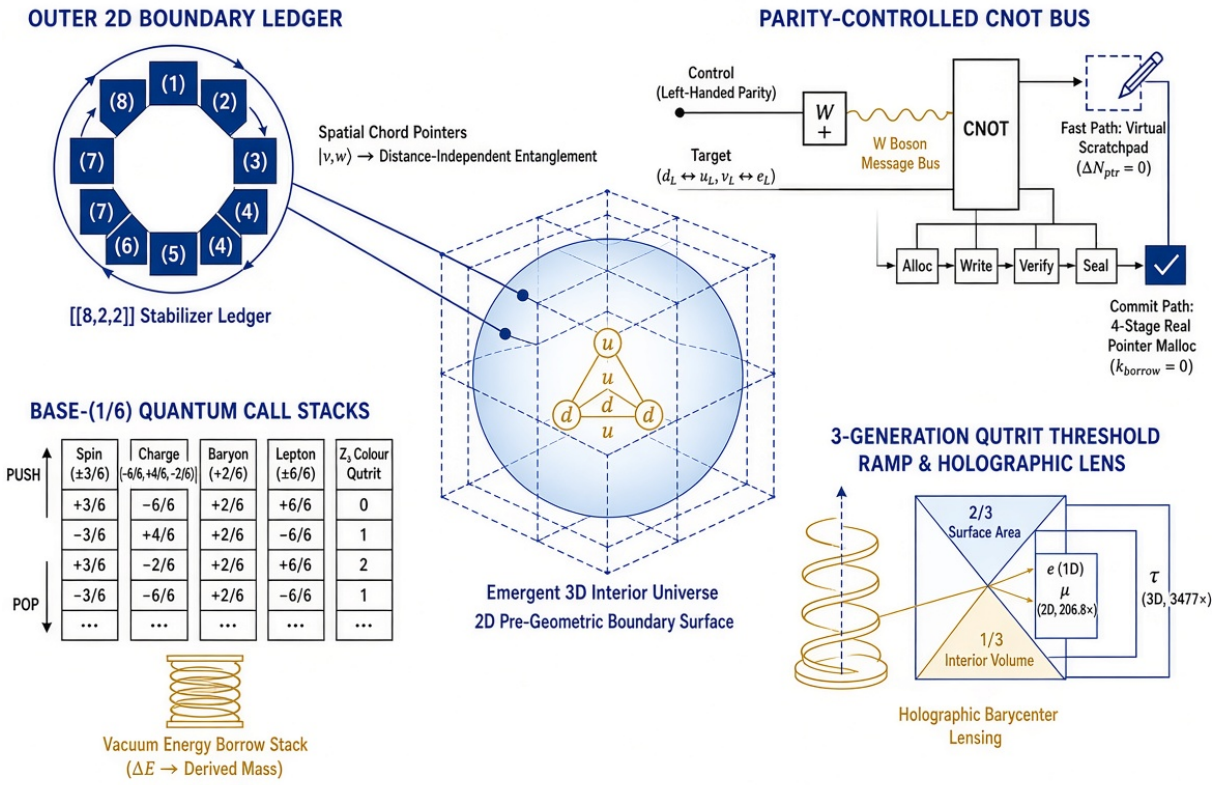


Figure 1: Conceptual overview supplied with the HQSM proposal, retained as an architectural guide. The register ring is schematic and repeats site labels; the mathematical model has exactly eight sites. "Left-handed parity" is interpreted as chirality, and the corrected operation is the charge-preserving exchange in section 6. "Borrow" denotes an auxiliary excitation cost, not a violation of energy conservation. The mass and area–volume labels describe the motivating ansatz rather than derived physical laws. The additional diagrams below specify the mathematical construction precisely.

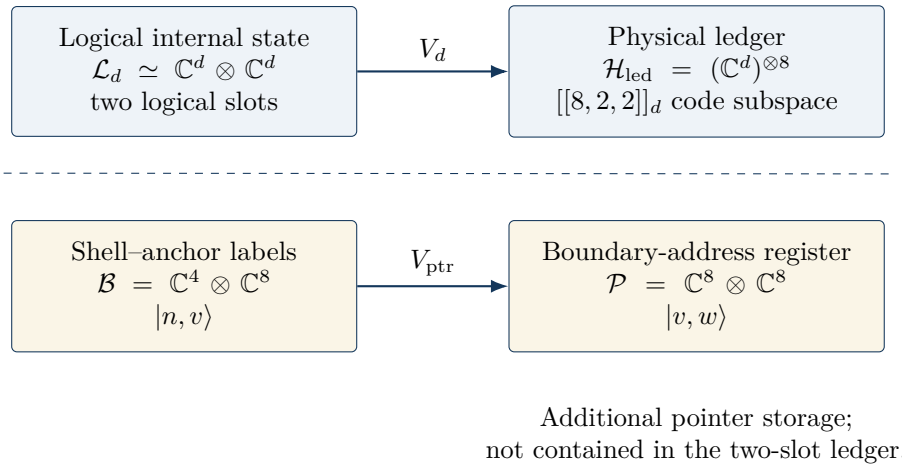


Figure 2: Two distinct encodings. The upper map distributes internal logical information over physical ledger sites. The lower map encodes shell–anchor labels into a separate boundary-address register. The registers can be combined by tensor product; neither is identified with the other.

3 Chord geometry and the interior representation

3.1 Four exact radii from eight boundary vertices

Embed the vertices in the unit sphere by

$$\mathbf{r}(v) = \frac{1}{\sqrt{3}}(2v_1 - 1, 2v_2 - 1, 2v_3 - 1), \qquad \mathbf{x}(v, w) = \frac{\mathbf{r}(v) + \mathbf{r}(w)}{2}. \tag{3.1}$$

The embedding is prescribed; it is not dynamically derived. Let $m = d_{Q_3}(v, w)$ be the Hamming distance, namely the number of differing binary coordinates, and put $n = 3 - m$.

Proposition 3.1 (Chord-radius identity)**.** *For every ordered pair* (v, w)*,*

$$\|\mathbf{x}(v, w)\|^2 = \frac{n}{3}, \qquad R_n = \sqrt{n/3} \in \{0, 1/\sqrt{3}, \sqrt{2/3}, 1\}. \tag{3.2}$$

Proof. In each of the m differing coordinates, the two signs cancel. In each of the n matching coordinates, the midpoint is $\pm 1/\sqrt{3}$. Squaring and summing gives $n/3$. □

The 64 ordered chord labels give only 27 distinct midpoint positions, with shell populations $1, 6, 12, 8$ from the centre outward. Several different register states therefore have the same position. The finite construction does not fill a continuous ball. Its chord midpoints lie in the inscribed cube, and even arbitrary positive endpoint weights remain in that cube's convex hull.

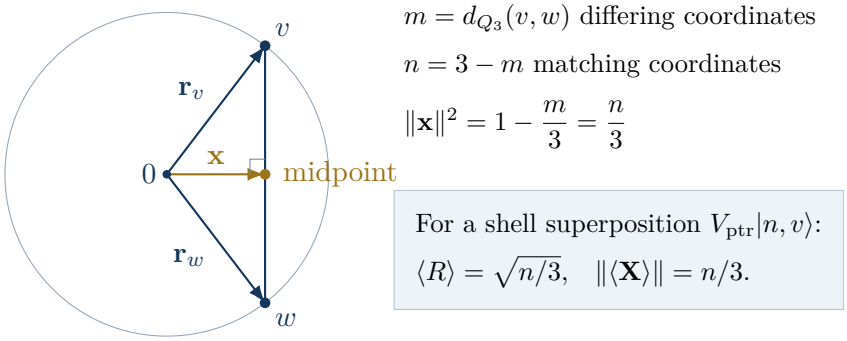


Figure 3: A cross-section illustrating the chord midpoint. The right-hand relations distinguish the radius of each constituent chord from the norm of the mean position of a shell superposition. The cross-section is schematic; the exact allowed radii follow from the cube's binary coordinates.

3.2 An isometric shell–anchor code

Define $\mathcal{B} = \text{span}\{|n, v\rangle : n = 0, 1, 2, 3;\ v \in \mathcal{V}\}$ and

$$V_{\text{ptr}}|n, v\rangle = \binom{3}{3-n}^{-1/2} \sum_{w : d_{Q_3}(v,w) = 3-n} |v, w\rangle. \tag{3.3}$$

Each column has unit norm. Different anchors have different first addresses, and different shells have disjoint Hamming-distance supports. Thus $V_{\text{ptr}}^\dagger V_{\text{ptr}} = \mathbb{I}_{32}$. The label v is an *anchor*, not an independent continuous angular coordinate. In particular, the eight states with $n = 0$ remain mutually orthogonal despite having the same midpoint at the origin.

Let $\mathbf{X}$ and R be diagonal operators on $\mathcal{P}$ with values $\mathbf{x}(v, w)$ and $\|\mathbf{x}(v, w)\|$, respectively. Then

$$V_{\text{ptr}}^\dagger R V_{\text{ptr}}|n, v\rangle = R_n|n, v\rangle, \tag{3.4}$$

$$\langle n, v|V_{\text{ptr}}^\dagger \mathbf{X} V_{\text{ptr}}|n, v\rangle = \frac{n}{3}\mathbf{r}(v). \tag{3.5}$$

To obtain the second formula, observe that each coordinate is flipped in a fraction $(3-n)/3$ of the allowed partners w. Hence its average midpoint component is $n/3$ times the corresponding component of $\mathbf{r}(v)$. For $n = 1, 2$, $\langle R\rangle \neq \|\langle \mathbf{X}\rangle\|$. Consequently it would be incorrect to assign every encoded state the sharp position $R_n\mathbf{r}(v)$. It has a sharp chord radius and a distribution of positions on that shell.

3.3 Exactly compressed hopping and its limits

Let A be the adjacency matrix of Q_3 and $L = 3\mathbb{I} - A$ its graph Laplacian. Hopping only the second endpoint gives $H_r = J_r\mathbb{I}_8 \otimes L$. The encoded subspace is invariant under this operator, and

$$H_r V_{\mathrm{ptr}} = V_{\mathrm{ptr}}(J_r L_r \otimes \mathbb{I}_8), \qquad L_r = \begin{pmatrix} 3 & -\sqrt{3} & 0 & 0 \\ -\sqrt{3} & 3 & -2 & 0 \\ 0 & -2 & 3 & -\sqrt{3} \\ 0 & 0 & -\sqrt{3} & 3 \end{pmatrix}. \tag{3.6}$$

This is an exact finite intertwining relation: evolving and then decoding agrees with evolving in the encoded space. It follows by counting the edges between adjacent Hamming shells; the off-diagonal magnitude is $\sqrt{(n+1)(3-n)}$. A derivation is given in section B.

Let F_i flip coordinate i of a vertex label. Parallel endpoint flips $F_i \otimes F_i$ preserve n and change the anchor. Consequently,

$$\sum_{i=1}^{3}(\mathbb{I} - F_i \otimes F_i)V_{\mathrm{ptr}} = V_{\mathrm{ptr}}(\mathbb{I}_4 \otimes L). \tag{3.7}$$

The eigenvalues of L are $0, 2, 4, 6$, with multiplicities $1, 3, 3, 1$. They are the spectrum of a cube graph, not the complete spectrum $\ell(\ell+1)$ of the spherical Laplacian. Cubic symmetry is exact; continuous rotational symmetry has not been established.

A convenient family of encoded Hamiltonians is

$$H_{\mathcal{B}} = J_r L_r \otimes \mathbb{I}_8 + \sum_{n=0}^{3} |n\rangle\langle n| \otimes \left[\frac{J_\Omega}{(n+\eta)^2}L + U_n\right], \qquad \eta > 0, \tag{3.8}$$

where U_n is a diagonal anchor potential. This form follows from boundary operators diagonal in the Hamming distance and the first address. In the prototype, $\eta = 1/2$ and part of U_n is proportional to $-1/(n+1/2)$. The regularized radial variable $n+\eta$ is distinct from the geometrical radius $R_n = \sqrt{n/3}$. Thus this potential must be described as a chosen finite attractive interaction, not an exact Coulomb law in the midpoint geometry. A physical hydrogen calculation would additionally require a common continuum regulator, calibrated units and converged spectral and transition observables.

4 Distributed qubit and qutrit ledgers

4.1 Explicit stabilizer constructions

A stabilizer code stores a logical state in the joint eigenspace of commuting operators [6]. For prime d, let

$$X|j\rangle = |j+1 \bmod d\rangle, \qquad Z|j\rangle = \omega^j|j\rangle, \qquad \omega = e^{2\pi i/d}. \tag{4.1}$$

For $a, b \in \mathbb{F}_d^8$, write $X(a) = \bigotimes_u X^{a_u}$ and $Z(b) = \bigotimes_u Z^{b_u}$. They commute when $a \cdot b = 0$ modulo d. Choose the row spaces C_d of

$$G_2 = \begin{pmatrix} 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 1 & 1 & 0 & 0 \end{pmatrix}, \tag{4.2}$$

$$G_3 = \begin{pmatrix} 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 1 & 0 & 0 \\ 1 & 2 & 0 & 1 & 2 & 0 & 1 & 2 \end{pmatrix}. \tag{4.3}$$

Both have rank three over their respective fields and satisfy $G_d G_d^T = 0$. The six generators $X(G_{d,i})$ and $Z(G_{d,i})$ therefore commute and are independent. Their common $+1$ eigenspace has dimension $d^{8-6} = d^2$.

For representatives $\ell \in C_d^\perp / C_d$, an explicit orthonormal basis is

$$|\ell\rangle_L = d^{-3/2} \sum_{c \in C_d} |c + \ell\rangle. \tag{4.4}$$

The nine qutrit codewords and four qubit codewords arise from disjoint cosets. The minimum weight of a vector in $C_d^\perp \setminus C_d$ is two, so these are $[[8, 2, 2]]_d$ codes. Explicit weight-two witnesses and the single-site argument are given in section A. The ternary construction is a CSS code placed on eight sites; its supports need not be cube faces. Neither construction has been shown to be a scalable topological surface code.

4.2 What a single site can reveal

For these codes every nonidentity one-site Weyl operator obeys

$$V_d^\dagger X_u^a Z_u^b V_d = 0, \qquad (a, b) \neq (0, 0). \tag{4.5}$$

Since Weyl operators form a basis of all one-site matrices, every encoded density matrix has

$$\rho_u = \mathrm{tr}_{\mathcal{V} \setminus \{u\}}(V_d \rho_L V_d^\dagger) = \mathbb{I}_d / d. \tag{4.6}$$

No measurement on one ledger site distinguishes logical inputs. The Knill–Laflamme conditions then imply correctability of an erasure at a known single site [7].

Two qualifications matter. Code distance two implies that a single-site reduced state is independent of the logical state, but does not by itself require that reduced state to be maximally mixed. Equation (4.6) uses the explicit generators above. Also, one erasure is not one arbitrary error at an unknown location: distance two detects all single-site errors but does not generally correct them without location information. The statement concerns the ledger factor, not other memory stored at a site.

4.3 Entanglement, logical access and no-signalling

The two logical qubits can be prepared in the singlet

$$|\Psi^-\rangle_L = (|01\rangle_L - |10\rangle_L)/\sqrt{2}. \tag{4.7}$$

This is a choice of state, not a consequence of having two pointers. Let $A_0 = Z$, $A_1 = X$, $B_0 = -(Z + X)/\sqrt{2}$ and $B_1 = -(Z - X)/\sqrt{2}$ act on the respective logical slots. The Clauser–Horne–Shimony–Holt expression [8] is

$$S = \langle A_0 B_0 + A_0 B_1 + A_1 B_0 - A_1 B_1 \rangle = 2\sqrt{2}. \tag{4.8}$$

If pointer motion acts only on spatial factors, it leaves this correlation unchanged at any assigned separation. This reproduces the familiar distance independence of a preserved Bell state. It does not explain why that state was prepared or how its physical distribution occurred.

For any trace-preserving operation $\mathcal{E}_A$ on one logical factor,

$$\mathrm{tr}_A[(\mathcal{E}_A \otimes \mathrm{id}_B)(\rho_{AB})] = \rho_B. \tag{4.9}$$

The proof uses the completeness relation for the operation's Kraus operators; see section C. Conditional states can change when a measurement outcome is selected, but that outcome must be communicated. This is logical no-signalling. To infer relativistic locality in an emergent space one must additionally identify physically localized observables and causal operations. Nonlocal support of a logical readout does not supply that identification.

5 Quantum-number registers and reversible transactions

5.1 The meaning of the one-sixth convention

The original "base-(1/6)" terminology is best understood as an integer rescaling, not a numeral base. For dimensionless charge $q = Q/e$, spin projection $j_z = J_z/\hbar$, baryon number B and lepton number L, define

$$J_6 = 6j_z, \qquad Q_6 = 6q, \qquad B_6 = 6B, \qquad L_6 = 6L. \tag{5.1}$$

The factor six makes the listed half-integer and third-integer quantities integral. It does not identify angular momentum, charge and baryon number as the same physical dimension. The physical labels follow the standard particle assignments [9]. In the convention $q = T_3 + Y$, the left-handed quark doublet has $Y = 1/6$. This is a normalization of hypercharge within the Standard Model, not evidence that the four operators in (5.1) have a common dynamical origin [10].

Table 2: Selected integer labels. The spin column refers to a chosen projection or, for massless gauge bosons, helicity. Colour triality $t \in \mathbb{Z}_3$ records a representation's centre charge; it is not a colour basis label. Antiparticles reverse additive charges and triality.

State	J_6	Q_6	B_6	L_6	$SU(3)$ representation	t
u	± 3	4	2	0	**3**	1
d	± 3	-2	2	0	**3**	1
e^-	± 3	-6	0	6	**1**	0
ν_e	± 3	0	0	6	**1**	0
γ	± 6	0	0	0	**1**	0
g	± 6	0	0	0	**8**	0
$W^\pm$	$-6, 0, 6$	± 6	0	0	**1**	0

The neutrino row is a bookkeeping space for spin-1/2 states, not an assertion that both chiralities participate in the Standard Model charged current. Lepton number is used for the selected beta-decay processes; Majorana mass terms and anomalous baryon-plus-lepton-number violation would require an extended treatment. A spin magnitude $s = 1/2$ also should not be confused with its projection $j_z = \pm 1/2$. General rotations require all three noncommuting angular-momentum operators, not just a scalar register.

5.2 Registers, rather than literal call stacks

On the unbounded counter space $\ell^2(\mathbb{Z})$ define

$$T_k|n\rangle = |n + k\rangle, \qquad T_k^\dagger = T_{-k}, \qquad T_k^\dagger N T_k = N + k\mathbb{I}. \tag{5.2}$$

One may call T_k a `PUSH` and its inverse a `POP`. A finite register with $D = 2M + 1$ levels can instead use modular shifts

$$T_k^{(D)}|n\rangle = |((n + M + k) \bmod D) - M\rangle. \tag{5.3}$$

These are unitary, but integer conservation becomes modular conservation whenever a boundary is crossed. A simulation of an unbounded charge must avoid wraparound, introduce an overflow register, or bound its truncation error. The reference implementation labels these operations as cyclic register shifts. A literal last-in–first-out call stack would additionally store an ordered sequence, a stack pointer and reversible control logic; those structures are not present in the simple counter model.

A transaction is a term $g|f\rangle\langle i| + g^*|i\rangle\langle f|$ in a Hamiltonian. For a diagonal conserved label A, it preserves A exactly if $A_f = A_i$. The general and basis-independent check is

$$[H, A] = 0 \quad \Longrightarrow \quad e^{iHt} A e^{-iHt} = A \quad \text{for all } t. \tag{5.4}$$

This operator test is stronger than checking the expectation value in one chosen input state. It is central to the gate correction below.

5.3 Colour, composite states and neutrinos

A qutrit can hold the three basis states of the fundamental colour representation. Its cyclic shift is only one possible operation; the continuous $SU(3)$ group has eight generators and a gluon transforms in the eight-dimensional adjoint representation [11]. The centre group $\mathbb{Z}_3$ retains triality but discards most of this structure. Zero triality is necessary for a colour singlet but not sufficient:

$$\mathbf{3} \otimes \mathbf{3} \otimes \mathbf{3} = \mathbf{1} \oplus \mathbf{8} \oplus \mathbf{8} \oplus \mathbf{10}. \tag{5.5}$$

Every term on the right has zero triality. A baryon singlet therefore needs an explicit projection, such as the antisymmetric colour tensor $\epsilon_{abc}/\sqrt{6}$. The proton's valence content uud has $Q_6 = 4 + 4 - 2 = 6$ and $B_6 = 2 + 2 + 2 = 6$, but a fixed list of three spin projections does not describe its complete spin structure. Spin, orbital motion and gluon contributions must be treated in a dynamical hadron model.

At the level of weak quantum-number bookkeeping,

$$d \longrightarrow u + e^- + \bar{\nu}_e \tag{5.6}$$

conserves Q_6, B_6 and L_6. The antineutrino supplies $L_6 = -6$. An initial half-integer total angular momentum cannot turn into a state containing only two spin-1/2 fermions and integer orbital angular momentum. Adding a third spin-1/2 particle permits half-integer total angular momentum. This is a useful consistency argument, but does not determine the neutrino's measured properties or a decay rate. The prototype's count of six balanced projection assignments out of sixteen ignores orbital amplitudes and helicity selection; it is not a branching probability.

For the selected particle labels in a generation,

$$\sum Q_6 = 3(4 - 2) + (0 - 6) = 0, \qquad \sum (B_6 - L_6) = 3(2 + 2) - (6 + 6) = 0. \tag{5.7}$$

These are arithmetic sums of assigned charges. Gauge anomaly cancellation requires a different sum over left-handed Weyl representations, including right-handed fields as their conjugates. For example,

$$6(1/6)^3 + 3(-2/3)^3 + 3(1/3)^3 + 2(-1/2)^3 + 1^3 = 0 \tag{5.8}$$

is the familiar $U(1)_Y^3$ cancellation. A finite charge table can check this identity; it does not derive the representation content or establish all anomaly constraints of a quantum field theory.

6 A charge-preserving replacement for the CNOT bus

6.1 Chirality and the charged current

Parity reverses spatial coordinates. Chirality is an algebraic property of a relativistic spinor, selected by $P_L = (1-\gamma^5)/2$ and $P_R = (1+\gamma^5)/2$. Helicity is spin projected along momentum. These concepts cannot be represented by a single interchangeable "parity bit"; chirality and helicity coincide for massless particles with the appropriate particle/antiparticle convention. Weak parity violation is experimental physics [12]. Its Standard Model description is a chiral gauge interaction [10, 13].

Suppressing generation mixing, the charged-current structure is

$$\mathcal{L}_{\mathrm{cc}} = -\frac{g_W}{\sqrt{2}} \left(\bar{u}_L \gamma^\mu d_L + \bar{\nu}_L \gamma^\mu e_L\right) W_\mu^+ + \mathrm{h.c.} \tag{6.1}$$

The gamma matrices encode the spinor and spacetime structure. Quark mixing introduces the CKM matrix and lepton mixing the PMNS matrix. The photon, gluon and Z have other couplings; they do not all perform this charged flavour transition. Equation (6.1) motivates a conditional transaction, not a deterministic elementary CNOT with a universal flip probability.

6.2 The failure of independent flips

Take local basis labels $|\chi, t\rangle$, with $\chi = 0, 1$ denoting right and left chirality. For quarks $t = 0, 1$ denotes u, d and for leptons it denotes ν, e. The right-chiral neutrino label is a formal spectator in this bookkeeping space, not a claim that an additional neutrino species has been observed. The original gate analogy uses

$$U_C = |0\rangle\langle 0| \otimes \mathbb{I} + |1\rangle\langle 1| \otimes X. \tag{6.2}$$

Applying $U_C \otimes U_C$ correctly maps $|d_L \nu_L\rangle$ to $|u_L e_L\rangle$, both of charge $Q_6 = -2$. However, it also maps $|u_L \nu_L\rangle$, with $Q_6 = 4$, to $|d_L e_L\rangle$, with $Q_6 = -8$. The full commutator has

$$\|[U_C \otimes U_C, Q_6^{\mathrm{tot}}]\|_2 = 12. \tag{6.3}$$

Here $\|\cdot\|_2$ is the spectral operator norm. A charge check on only the first input misses this failure. A single flavour flip can be used if an explicit boson or reservoir carries the compensating charge; without that resource it is not a charge-conserving closed operation.

6.3 Corrected exchange dynamics

Define $T_q^+ = |u_L\rangle\langle d_L|$ and $T_\ell^- = |e_L\rangle\langle \nu_L|$. On the full 16-dimensional two-register space choose

$$H_{\mathrm{ex}} = g\, T_q^+ \otimes T_\ell^- + g^*\, T_q^- \otimes T_\ell^+. \tag{6.4}$$

All states outside $\mathrm{span}\{|d_L \nu_L\rangle, |u_L e_L\rangle\}$ are spectators. The first transition raises the quark charge by six units and lowers the lepton charge by six units, so

$$[H_{\mathrm{ex}}, Q_6^{\mathrm{tot}}] = 0, \qquad [e^{-itH_{\mathrm{ex}}}, Q_6^{\mathrm{tot}}] = 0 \tag{6.5}$$

for every input and every time. Baryon and lepton number are unchanged as well. This provides a consistent finite exchange operation.

For real g in a degenerate two-state block,

$$e^{-itH_{\mathrm{ex}}} |d_L \nu_L\rangle = \cos(gt)|d_L \nu_L\rangle - i \sin(gt)|u_L e_L\rangle. \tag{6.6}$$

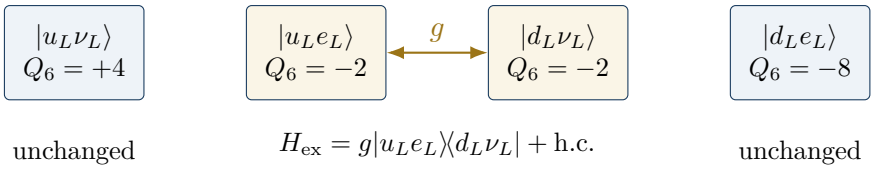


Figure 4: The corrected exchange acts only between the equal-charge states. An unrestricted pair of CNOT gates would also couple the two outer states, violating charge conservation. The displayed diagram is the left-chiral block of the full 16-state operation.

Unit transfer occurs only for a selected pulse area $gt = \pi/2$ modulo π. With detuning Δ, the transition probability becomes

$$P(t) = \frac{4|g|^2}{\Delta^2 + 4|g|^2} \sin^2\left(\frac{t}{2}\sqrt{\Delta^2 + 4|g|^2}\right). \tag{6.7}$$

The repaired gate captures a selection rule and conserved exchange. It still lacks momentum dependence, gauge fields and the spinor matrix elements needed to reproduce weak scattering or decay rates.

7 Auxiliary states, effective interactions and particle creation

7.1 A precise meaning for a virtual scratchpad

In perturbative quantum field theory, a virtual particle is represented by an internal propagator, not an independently observable asymptotic particle. Four-momentum is conserved at interaction vertices; an internal momentum need not obey the external-particle mass-shell relation [4]. Calling this a temporary violation of energy conservation is misleading.

A useful finite analogue splits a Hilbert space with projectors P and $Q = \mathbb{I} - P$ into retained and auxiliary sectors:

$$H = \begin{pmatrix} H_{PP} & H_{PQ} \\ H_{QP} & H_{QQ} \end{pmatrix}, \qquad |\psi\rangle = \begin{pmatrix} \psi_P \\ \psi_Q \end{pmatrix}. \tag{7.1}$$

Solving the auxiliary row of $H\psi = E\psi$ gives $\psi_Q = (E - H_{QQ})^{-1} H_{QP}\psi_P$ whenever the inverse exists. Substitution into the retained row gives the exact stationary equation

$$H_{\rm eff}(E)\psi_P = E\psi_P, \qquad H_{\rm eff}(E) = H_{PP} + H_{PQ}(E - H_{QQ})^{-1} H_{QP}. \tag{7.2}$$

This is the Feshbach–Schur reduction [14]. The self-energy term contains the effect of auxiliary amplitudes without assigning them independent external pointers. It is generally energy dependent; replacing it by a constant Hamiltonian for all times needs an approximation or a controlled block diagonalization [15].

If the auxiliary block is separated by a gap ε and the coupling is small, the leading correction near $E = 0$ is

$$H_{\rm eff} \simeq H_{PP} - H_{PQ} H_{QQ}^{-1} H_{QP}. \tag{7.3}$$

A positive excitation penalty suppresses auxiliary occupation. The term "borrow stack" can denote the level counter for that penalty, but total energy is always governed by the full time-independent Hamiltonian. Auxiliary amplitudes can affect measured interference and scattering even though they are not independent detector events. Defining a detector to vanish on Q merely specifies an observable; it does not prove that all effects of Q are unobservable.

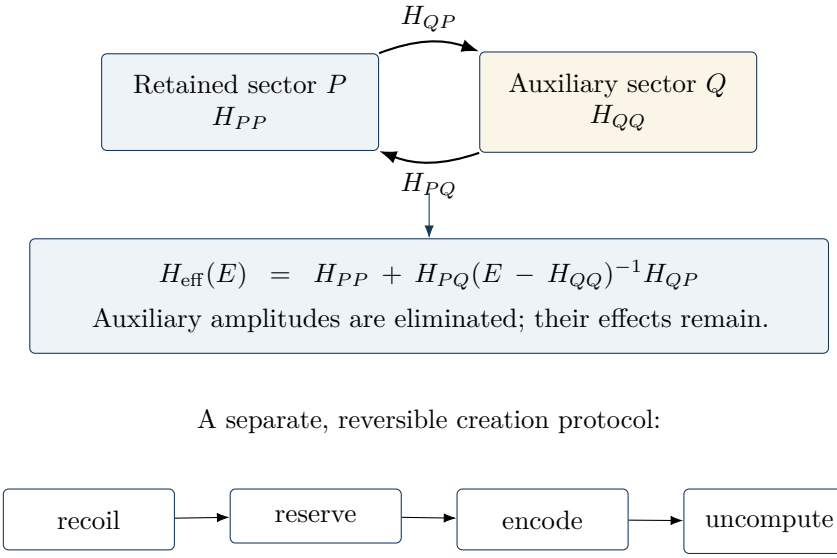


Figure 5: Effective interaction by auxiliary-sector elimination. The four-step creation protocol is a separate computational choice using preallocated resources. It is not implied by the Schur identity or by the number of vertices in a Feynman diagram.

7.2 Why a self-energy shift is not yet a mass

Consider the smallest nontrivial example,

$$H = \begin{pmatrix} 0 & g \\ g^* & \varepsilon \end{pmatrix}, \quad E_- = \frac{\varepsilon - \sqrt{\varepsilon^2 + 4|g|^2}}{2} = -\frac{|g|^2}{\varepsilon} + O(|g|^4/\varepsilon^3), \quad \varepsilon > 0. \tag{7.4}$$

The lower eigenvalue is negative. The original five-state electron sector has the same issue: its default lowest eigenvalue is approximately -0.073405 in model units, and the prototype reports its absolute value as a positive "derived mass". Taking an absolute value is not a derivation of a rest mass.

A physical excitation energy must be measured relative to a consistently defined interacting vacuum. A relativistic mass additionally requires a dispersion relation, such as $E(\mathbf{p})^2 - \mathbf{p}^2 = m^2$, or an appropriately defined propagator pole. For confined quarks, mass definitions also depend on renormalization conventions. In the Standard Model charged fermion masses are associated with Yukawa couplings and electroweak symmetry breaking [10, 13]. The finite Schur model does not replace that mechanism. The corrected reference code returns energy shifts with their signs and makes no mass identification.

7.3 Real emission and reversible allocation

An isolated on-shell massive particle cannot radiate a nonzero real photon while remaining the same on-shell particle in vacuum. If $p = p' + k$, $p^2 = p'^2 = m^2$ and $k^2 = 0$, then $p' \cdot k = 0$. In the rest frame of p', this would require $mk^0 = 0$, contradicting a positive-energy photon. This argument is general; a census on the finite momentum grid $\{-1, 0, 1\}^3$ is only an illustration.

Radiation can occur when another system absorbs recoil, when an initial state has excitation energy, or through other allowed scattering and decay channels. A third body is therefore one mechanism, not a universal requirement for real-particle creation. The following protocol is suitable for a recoil-assisted finite example:

1. Include the recoil and energy source in the total state.
2. Reserve a blank internal register in a preallocated pool.

3. Encode the outgoing pointer and entangle it with the participating systems.
4. Uncompute work registers when the constraints permit a completed transition.

Allocation and deallocation must retain reversibility; uncomputation is possible only if the ancillary information has been coherently accounted for. A projection onto a successful "commit" subspace is a measurement with a success probability, not a deterministic unitary reset. One may define a committed sector by $k_{\text{aux}} = 0$ and an increased occupation counter, but these are protocol conventions rather than universal criteria for real particles. The memory cost and gate depth depend on the chosen encoding and controller, not solely on the number of conceptual stages.

8 The generation qutrit and the charged-lepton mass ansatz

8.1 An algebraic identity with a positivity condition

For three positive masses define

$$K = \frac{m_e + m_\mu + m_\tau}{(\sqrt{m_e} + \sqrt{m_\mu} + \sqrt{m_\tau})^2}. \tag{8.1}$$

Koide's phenomenological relation is $K = 2/3$ [16]. Its geometric interpretation and phase parametrizations predate HQSM [17, 18]; neither the relation nor the phase near $2/9$ is claimed as a new discovery here.

Let $\mathbf{v} = (\sqrt{m_e}, \sqrt{m_\mu}, \sqrt{m_\tau})^T$ and $\hat{\mathbf{u}} = (1,1,1)^T/\sqrt{3}$. Decompose $\mathbf{v} = \mathbf{v}_\parallel + \mathbf{v}_\perp$ relative to this axis. Then

$$\|\mathbf{v}_\parallel\|^2 = \frac{1}{3}\left(\sum_g \sqrt{m_g}\right)^2, \qquad K = \frac{1}{3}\left(1 + \frac{\|\mathbf{v}_\perp\|^2}{\|\mathbf{v}_\parallel\|^2}\right). \tag{8.2}$$

Thus $K = 2/3$ is equivalent to equal squared norms in the one-dimensional symmetric subspace and its two-dimensional orthogonal complement. It is also equivalent to an angle $\pi/4$ between $\mathbf{v}$ and $\hat{\mathbf{u}}$. Equal norms are an extra constraint; they do not follow from counting one radial and two tangential spatial directions. In fact, equal variance per orthonormal component would give a 1:2 norm ratio, not equal norms. Areas and volumes also have different physical dimensions, so literal fractions of area and volume cannot be equated without a separate dimensionally consistent model.

A qutrit clock supplies the parametrization

$$r_g = a\left[1 + \sqrt{2}\cos(2\pi g/3 + \delta)\right], \qquad m_g = r_g^2, \quad g = 0,1,2, \tag{8.3}$$

where $a > 0$ has units of the square root of mass. The sums of the three cosines and their squares are zero and $3/2$, respectively, so

$$\sum_g r_g = 3a, \qquad \sum_g r_g^2 = 6a^2. \tag{8.4}$$

However, $\sqrt{m_g} = |r_g|$. Equations (8.4) therefore imply the physical Koide identity only when the three amplitudes are nonnegative. One positive branch is $|\delta| \le \pi/12$, with the other branches related by permutations. For $0 < \delta < \pi/12$, the ordering (e, μ, τ) corresponds to $g = (1,2,0)$.

At $\delta = \pi/3$, one amplitude is negative. The signed ratio $\sum r_g^2/(\sum r_g)^2$ remains $2/3$, but the physical ratio is

$$K = \frac{6}{(1+2\sqrt{2})^2} \simeq 0.4094. \tag{8.5}$$

The original all-phase test used signed amplitudes in the denominator and therefore tested a different identity. The corrected implementation uses the principal nonnegative square roots and includes this counterexample.

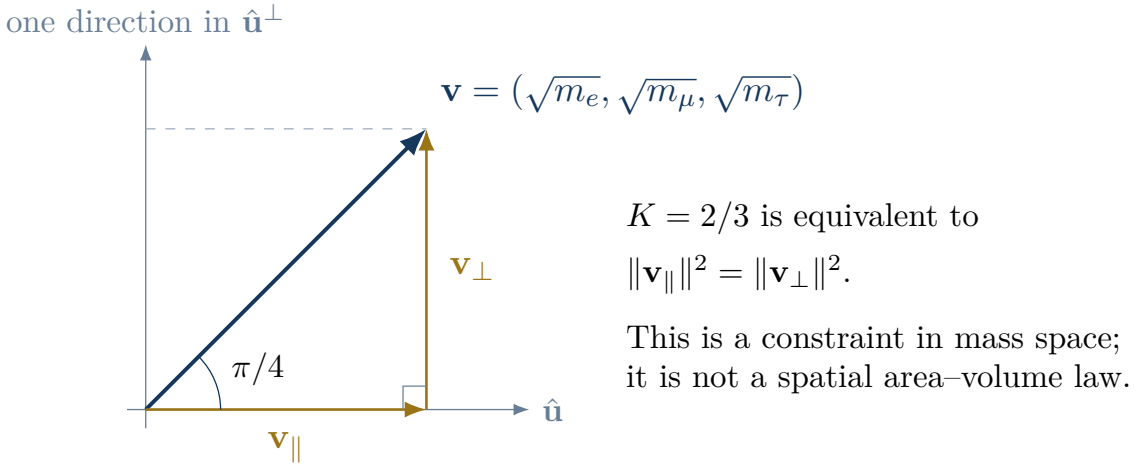


Figure 6: Koide geometry in the plane spanned by $\mathbf{v}$ and the symmetric axis. Equality of the two projected norms is equivalent to the mass relation. The two-dimensional orthogonal subspace of the full three-component mass space is represented by one direction in this drawing.

8.2 Testing the fixed phase and making calibration explicit

The numerical proposal $\delta = 2/9$ radians is not fixed by the ratio of two spatial dimensions to nine two-qutrit states. No Hamiltonian, symmetry or variational principle has been supplied that selects this angle. Moreover, taking it as an exact prediction gives

$$\left.\frac{m_\mu}{m_e}\right|_{\delta=2/9} = 206.7703159727, \tag{8.6}$$

whereas the 2022 CODATA ratio is $206.7682827 \pm 0.0000046$ [19]. The relative residual is 9.83×10^{-6}, about 442 quoted standard uncertainties. This is a discrepancy, not precision agreement. The count assumes the fixed-phase formula has no theory uncertainty; introducing unspecified corrections removes its status as an exact prediction.

A consistent phenomenological repair treats both a and δ as calibrated parameters. We fix $m_e = 0.51099895000$ MeV from the specified PDG electron listing [20] and solve for the phase using the independent muon–electron ratio. On the positive branch this gives

$$\delta_{\text{fit}} = 0.2222220470731, \qquad a_{\text{fit}} = 17.7159927708\ \text{MeV}^{1/2}. \tag{8.7}$$

This is a two-input fit, including the overall scale. The remaining conditional prediction is $m_\tau = 1776.9690248$ MeV. It is precisely the large-root Koide estimate from the two lighter masses, as can be seen without introducing a clock:

$$\sqrt{m_\tau} = 2(\sqrt{m_e} + \sqrt{m_\mu}) + \sqrt{3\,(m_e + m_\mu + 4\sqrt{m_e m_\mu})}. \tag{8.8}$$

The tau datum is the PDG 2025 average [21]. Propagating the ratio uncertainty with m_e fixed gives about 0.0000361 MeV uncertainty on the conditional tau prediction, much smaller than the 0.09 MeV experimental uncertainty. This propagation does not assign a theory uncertainty to the Koide ansatz. The experimental central values give $K \simeq 0.66666446$. Agreement of the calibrated ansatz with the tau mass is an interesting empirical property; it is not a derivation of lepton masses, of exactly three generations, or of the HQSM dynamics. The distinction between pole masses and scale-dependent running masses must also be specified before testing radiative stability.

Table 3: Fixed-phase and calibrated versions of the ansatz. Data are frozen to the named sources for reproducibility; they are not inferred from the schematic figure. "Input" means the quantity calibrates a parameter.

Quantity	$\delta = 2/9$	Calibrated phase	Reference value
m_e (MeV)	Input	Input	0.51099895000(15)
m_μ/m_e	206.77031597	Input	206.7682827(46)
m_τ (MeV)	1776.9849684	1776.9690248	1776.93(9)
Muon-ratio residual / σ	442.02	Fitted	—
Tau residual / σ	0.61	0.43	—

8.3 What a neutrino qutrit must reproduce

A three-level Hamiltonian conjugated by the discrete Fourier matrix can produce flavour transitions. Unitarity guarantees that its probabilities sum to one, but does not make it a neutrino model. In vacuum the ultrarelativistic oscillation amplitude has the form

$$\mathcal{A}_{\alpha\to\beta}(L,E) = \sum_{j=1}^{3} U_{\beta j} U_{\alpha j}^{*} \exp\left(-i\frac{m_j^2 L}{2E}\right), \tag{8.9}$$

with mixing matrix U, measured mass-squared differences and appropriate matter corrections [22]. A Fourier matrix and three arbitrarily chosen clock frequencies are illustrative choices, not a determination of these parameters. Electrical neutrality alone does not imply oscillations or prevent decay. A proposed neutrino sector must recover the observed L/E dependence, mixing angles and mass-squared splittings.

9 Weighted chord maps and the proposed gravitational interpretation

Let ϕ_v and $\mathbf{b}$ be dimensionless potential and frame-bias parameters, and define

$$s_v = -\beta\phi_v + \mathbf{b}\cdot\mathbf{r}(v), \quad W_v = e^{s_v} > 0, \qquad \mathbf{x}_W(v,w) = \frac{W_v\mathbf{r}(v) + W_w\mathbf{r}(w)}{W_v + W_w}. \tag{9.1}$$

All exponents must be dimensionless. If ϕ or a physical acceleration is inserted, appropriate energy, length and time scales must be supplied; an acceleration cannot simply be added to a dimensionless exponent. The weights are taken as prescribed classical parameters in this section. Making them depend on an unknown quantum state's expectation values would require a separate mean-field or fully quantum dynamical specification.

Proposition 9.1 (Convexity and the antipodal identity)**.** *For unit endpoint vectors and finite positive weights, $\|\mathbf{x}_W\| \le 1$, with strict inequality for distinct endpoints. For $w = \bar{v}$ antipodal to v,*

$$\mathbf{x}_W(v,\bar{v}) = \tanh\xi\,\mathbf{r}(v), \qquad \xi = \tfrac{1}{2}(s_v - s_{\bar{v}}). \tag{9.2}$$

Proof. Put $\lambda = W_v/(W_v + W_w) \in (0,1)$. Then

$$\|\lambda\mathbf{r}(v) + (1-\lambda)\mathbf{r}(w)\|^2 = 1 - \lambda(1-\lambda)\|\mathbf{r}(v) - \mathbf{r}(w)\|^2. \tag{9.3}$$

For antipodal endpoints the coefficient of $\mathbf{r}(v)$ is $(e^{s_v} - e^{s_{\bar{v}}})/(e^{s_v} + e^{s_{\bar{v}}}) = \tanh\xi$. □

Coincident endpoints remain on the boundary; they are an exception to any claim of a strictly interior position for every chord. A bounded spatial coordinate is also not a causal horizon.

Event horizons, Rindler horizons and de Sitter horizons require a spacetime metric and a causal structure [23]. Equation (9.2) alone supplies neither.

The prototype chooses a discrete potential

$$\phi = -L^{+}(\mathcal{E} - \bar{\mathcal{E}}\mathbf{1}), \tag{9.4}$$

where L^{+} is the Moore–Penrose inverse of the cube Laplacian. Subtracting the mean makes the source orthogonal to the constant zero mode. This is a finite graph Poisson prescription. Its sign, normalization and connection to physical energy are choices to be justified by a subsequent dynamical theory.

Only differences of the s_v matter. Therefore two choices satisfying

$$-\beta\phi_v = \mathbf{b} \cdot \mathbf{r}(v) + c \tag{9.5}$$

give identical weighted positions. This explains why a matched potential and frame bias agree to machine precision in the prototype. It is an identity of the map, not a derivation of the equivalence principle. Likewise, changing the coordinates assigned to a prescribed path can bend its plotted trajectory without demonstrating gravitational light bending. A physical prediction requires a propagation law, a metric, an operational definition of angles and clocks, and a comparison with observations such as redshift, time delay and deflection. No Einstein equation follows from the present convexity argument.

10 Reproducibility and validation

The archived software release [24] contains the four original verification scripts, their captured outputs and SHA-256 source hashes, plus a separate corrected reference implementation. The original scripts were run without modification on 5 November 2026 and reported 43 passing checks in total. These counts reproduce the programs' own assertions; they do not endorse all interpretations in their comments or output labels.

Table 4: Original suite results and their evidential scope. Short labels map to the complete filenames in the software archive's manifest.

Suite	Result	Scope of the passing checks
Distributed ledger	9/9	Specified code, logical operators and selected encoded Bell experiments
Stack machine	10/10	Finite register algebra, chosen self-energy matrices and weighted maps
Virtual/real compilation	12/12	Selected transitions, algebraic sums and originally implemented phase tests
Boundary atom benchmark	12/12	Finite interacting matrices, flux enumeration and encoding identities

The revised `verification/hqsm_reference.py` implements the mathematical definitions used here. The independent driver `verification/verify_paper.py` reports 23 passing checks. It enumerates the qubit and qutrit codes, tests every nonidentity one-site Weyl operator, checks all 64 chord radii and the mean-position correction, verifies radial intertwining, and tests the repaired exchange on the full state space. It also retains the nonconserving CNOT pair and the signed-root phase example as regression counterexamples. The mass inputs, fitted parameters and conditional residuals are saved in machine-readable JSON.

The recorded environment is Python 3.13.7, NumPy 2.4.4 and SciPy 1.17.1. Finite-field code calculations are exact integer arithmetic; matrix and trigonometric checks use double precision

with tolerance 10^{-12} where appropriate. The accompanying README gives commands and pinned versions. Numbers in the mass table use explicit frozen data inputs, not values silently downloaded at execution time.

A passing original mass test, for example, only means its residual met its programmed tolerance. It did not compare that residual to the experimental uncertainty. The new tests make this distinction explicit. Similarly, testing one charge-balanced transition does not replace the full commutator check, and radial refinement of one chosen finite potential does not establish convergence to continuum hydrogen.

11 Research programme and conclusion

The finite HQSM construction has a coherent computational core: independent internal and spatial encodings, explicit erasure-correcting ledgers, reversible register operations, an exactly compressed hopping law and charge-conserving exchange dynamics. Effective interactions can be derived within a stated auxiliary-sector model. These are useful, reproducible objects for studying quantum information and simulation.

Connecting this architecture to particle physics requires several further results. A growing family of boundary graphs must have a specified continuum limit, resource scaling and physical locality. An interacting Hamiltonian must implement gauge constraints, spin and statistics, and recover relativistic dispersion and scattering. Finite gauge encodings can be studied in established quantum-link frameworks [25], but replacing $SU(3)$ by triality alone is insufficient. A mass model must define its vacuum, symmetry-breaking structure, renormalization scheme and parameters before confronting independent data. A gravitational extension must supply dynamics and a metric rather than only a coordinate map.

The immediate falsifiable targets are correspondingly concrete: test operator conservation on every sector; measure encoding and truncation errors; demonstrate convergence of spectra and propagation on refined graphs; predict an observable not used for calibration; and verify that the prediction survives changes of regulator. The exact fixed phase $2/9$ already fails the precision muon-ratio test. The calibrated Koide ansatz remains an empirical relation whose origin is unexplained.

HQSM is therefore presented as a finite boundary-encoded quantum model, not as an established replacement for the Standard Model or general relativity. Its most productive feature is the ability to state proposed mechanisms as operators and test them. The corrections developed here preserve that architectural idea while making its mathematical content, physical assumptions and outstanding obligations explicit.

Code and data availability. The supporting Python code, verification outputs and unmodified prototype snapshots are archived on Zenodo under the concept DOI 10.5281/zenodo.23991264 (all versions). The calculations reported here use release 1.0.0, 10.5281/zenodo.23991265, distributed under the MIT License [24]. The manuscript, LaTeX source, editable TikZ diagrams and supplied overview image are archived at 10.5281/zenodo.23991267 under CC BY 4.0. Particle-data inputs are identified in section 8 and the bibliography.

A Details of the stabilizer codes

For either matrix G_d, its rank is three and its row space C_d obeys $C_d \subseteq C_d^{\perp}$. Thus $\dim C_d^{\perp} = 8 - 3 = 5$ and the quotient $C_d^{\perp}/C_d$ has dimension two over $\mathbb{F}_d$, giving d^2 cosets. The states in (4.4) are invariant under $X(c)$ for every $c \in C_d$, because adding c permutes the summation. They

are invariant under $Z(c)$ because $c \cdot (c' + \ell) = 0$ for $c' \in C_d$ and $\ell \in C_d^\perp$. Distinct cosets have disjoint computational-basis supports, establishing orthonormality.

Every column of G_d is nonzero. Consider a nonidentity one-site Weyl operator $E_u = X_u^a Z_u^b$. If $a \neq 0$, some $Z(G_{d,i})$ commutes with E_u only up to a phase $\omega^k \neq 1$. If $a = 0$ then $b \neq 0$ and an X-type generator has the same property. For that stabilizer S,

$$V_d^\dagger E_u V_d = V_d^\dagger S^\dagger E_u S V_d = \omega^k V_d^\dagger E_u V_d, \tag{A.1}$$

forcing the compression to vanish. The argument works for every logical matrix element, not only for a sample of encoded pure states.

In the binary code, the vector $(1,1,0,0,0,0,0,0)$ belongs to $C_2^\perp \setminus C_2$. Every nonzero word of C_2 has weight four or eight, so this weight-two operator is logical. In the ternary code, $(0,0,0,0,0,0,1,1)$ is orthogonal to all three rows but is not in C_3. Indeed, a nonzero combination of the first two rows alone has weight at least three; a combination with a nonzero third-row coefficient has weight six. These witnesses, together with the absence of weight-one logical operators, establish distance exactly two.

For a known-site erasure, the error products $E_i^\dagger E_j$ span all operators on that site. The Knill–Laflamme condition is

$$V^\dagger E_i^\dagger E_j V = \alpha_{ij} \mathbb{I}_L. \tag{A.2}$$

Equation (4.5) proves this condition for the specified codes. In a general erasure-correcting code, it guarantees a logical-state-independent local density matrix, which may be nonuniform. For instance, appending a fixed ancillary qubit $|0\rangle$ to a code leaves that ancilla's reduced state pure while providing no information about the logical state. Maximal mixing therefore requires the additional property proved above.

B Hamming-shell hopping and positional moments

Fix the anchor v and write the normalized uniform shell state as

$$|m; v\rangle = \binom{3}{m}^{-1/2} \sum_{w : d(v,w)=m} |w\rangle. \tag{B.1}$$

Each vertex of the m shell has $3 - m$ neighbours in shell $m + 1$ and m neighbours in shell $m - 1$. Counting edges gives

$$\langle m+1; v | A | m; v \rangle = \frac{\binom{3}{m}(3-m)}{\sqrt{\binom{3}{m}\binom{3}{m+1}}} = \sqrt{(m+1)(3-m)}. \tag{B.2}$$

There are no edges within a fixed shell. Substituting $n = 3 - m$ and $L = 3\mathbb{I} - A$ gives (3.6). Uniform amplitudes within each shell remain uniform under hopping, so the encoded subspace is invariant, not merely a space onto which one may project after evolution.

The variance of the full midpoint vector in an encoded basis state is

$$\left\langle \|\mathbf{X} - \langle \mathbf{X} \rangle\|^2 \right\rangle = \langle R^2 \rangle - \|\langle \mathbf{X} \rangle\|^2 = \frac{n}{3} - \frac{n^2}{9} = \frac{n(3-n)}{9}. \tag{B.3}$$

It is nonzero on the two intermediate shells. This explains how every constituent chord can have the same radius while their mean position has a smaller norm. Compression also need not preserve products of observables: $V^\dagger X^2 V$ generally differs from $(V^\dagger X V)^2$ when X maps states outside the code subspace. Care is therefore needed when constructing a spatial metric from compressed operators.

The cube spectrum follows directly from the binary characters $f_k(v) = (-1)^{k\cdot v}$ with $k \in \{0,1\}^3$. Flipping coordinate i multiplies f_k by $(-1)^{k_i}$, so $Lf_k = 2|k|f_k$, where $|k|$ is the Hamming weight of k. This gives eigenvalues $2|k|$ and multiplicities $\binom{3}{|k|}$. It supplies an exact finite angular spectrum while making its difference from continuum spherical harmonics explicit.

C Logical no-signalling and the Bell bound

Let $\mathcal{E}_A(\rho) = \sum_r K_r \rho K_r^\dagger$ with $\sum_r K_r^\dagger K_r = \mathbb{I}_A$. For any observable B of the second logical subsystem,

$$\begin{aligned}\operatorname{tr}\left[(\mathbb{I}_A \otimes B)\sum_r (K_r \otimes \mathbb{I})\rho_{AB}(K_r^\dagger \otimes \mathbb{I})\right] &= \operatorname{tr}\left[\left(\sum_r K_r^\dagger K_r \otimes B\right)\rho_{AB}\right] \\ &= \operatorname{tr}[(\mathbb{I}_A \otimes B)\rho_{AB}]. \end{aligned} \tag{C.1}$$

All expectation values of B are unchanged, proving (4.9). Selecting a single r breaks the completeness sum and describes conditioning, not a signalling channel without the outcome record.

For dichotomic observables $A_i^2 = B_j^2 = \mathbb{I}$ on separate factors, define $\mathcal{S} = A_0 \otimes (B_0 + B_1) + A_1 \otimes (B_0 - B_1)$. Expanding gives

$$\mathcal{S}^2 = 4\mathbb{I} - [A_0, A_1] \otimes [B_0, B_1]. \tag{C.2}$$

Each commutator has operator norm at most two, hence $\|\mathcal{S}\| \le 2\sqrt{2}$. The singlet and observables in section 4 attain this bound. An isometric encoding preserves their logical expectation values. These statements are consequences of the assumed quantum tensor-product structure, not of the pointer geometry.

D Controlled elimination and error estimates

Suppose $H_{QQ} \ge \varepsilon\mathbb{I}$ with $\varepsilon > 0$. For real $E < \varepsilon$, the inverse in (7.2) is negative definite. Thus the self-energy is negative semidefinite on the retained subspace. This sign check rules out interpreting every downward shift from a zero diagonal entry as a positive rest energy.

In the two-state example, rationalizing the square root gives

$$E_- = -\frac{2|g|^2}{\varepsilon + \sqrt{\varepsilon^2 + 4|g|^2}}, \qquad 0 \le E_- + \frac{|g|^2}{\varepsilon} = \frac{4|g|^4}{\varepsilon(\varepsilon + \sqrt{\varepsilon^2 + 4|g|^2})^2} \le \frac{|g|^4}{\varepsilon^3}. \tag{D.1}$$

The last inequality is a concrete remainder bound for the leading effective interaction. It is stronger than describing the correction as "small" without a scale.

For a more general dynamical approximation, let H and H_B be self-adjoint and V an isometry. If $\|HV - VH_B\| \le \epsilon$, differentiating the product of propagators and integrating gives

$$\|e^{-itH}V - Ve^{-itH_B}\| \le |t|\epsilon. \tag{D.2}$$

This identifies the residual and time interval that a numerical encoding test must control. It does not imply that an energy-dependent Schur complement is an exact time-local Hamiltonian at every energy.

E Mass calibration and interpretation of residuals

For the ordering $(e, \mu, \tau) = (g = 1, 2, 0)$, define

$$F(\delta) = \left[\frac{1 + \sqrt{2}\cos(4\pi/3 + \delta)}{1 + \sqrt{2}\cos(2\pi/3 + \delta)}\right]^2. \tag{E.1}$$

On $0 < \delta < \pi/12$, the numerator increases and the positive denominator decreases, so F is strictly increasing. It starts at one and diverges at the upper endpoint. There is therefore one solution for every measured ratio greater than one. The reference implementation brackets this root; the electron mass then fixes a. There is no unreported phase-selection freedom within this chosen ordering and branch.

Near the fitted phase, $\mathrm{d}\log F/\mathrm{d}\delta$ is about 56 per radian. The difference between $2/9$ and the fitted phase is only 1.75×10^{-7} radians, yet it produces a roughly 9.83 parts-per-million mass-ratio residual. A small phase difference is not small compared with the measurement precision. The 442-standard-uncertainty comparison in section 8 tests the fixed formula; it is not a goodness-of-fit claim for a model with unknown correction terms. Once δ is calibrated, the muon datum cannot also count as an independent successful prediction.

F Computer-science and physics terminology

Table 5: Terminology used in this paper. Each analogy has a specific mathematical meaning.

Architectural term	Mathematical object	Boundary of the analogy
Species definition	Fixed classical table of representations and couplings	Does not imply a shared mutable quantum state
Ledger	Encoded logical subsystem $V_d\mathcal{L}_d$	Finite capacity; local marginal protection applies only to this subsystem
Pointer	Boundary-address state $\lvert v, w\rangle$	Distinct addresses can have the same midpoint
Stack	Integer or modular counter with reversible shifts	Does not store a last-in–first-out history by itself
Message bus	A charge-conserving coupling between sectors	Does not reproduce a gauge field merely by being spin one
Scratchpad	Auxiliary subspace $Q\mathcal{H}$	Its effects survive elimination; energy is conserved in the full system
Commit	A designated output sector with cleared work registers	A postselected success is not a deterministic reset
Generation clock	Three-level spectral parametrization	Three generations and parameter values remain assumptions
Holographic lens	Positive-weight chord barycenter	A coordinate map is not a spacetime metric or a gravitational dynamics

References

[1] R. P. Feynman. “Simulating physics with computers”. In: *International Journal of Theoretical Physics* 21 (1982), pp. 467–488. DOI: 10.1007/BF02650179.

[2] J. M. Maldacena. “The Large N Limit of Superconformal Field Theories and Supergravity”. In: *Advances in Theoretical and Mathematical Physics* 2 (1998), pp. 231–252. DOI: 10.4310/ATMP.1998.v2.n2.a1. arXiv: hep-th/9711200.

[3] A. Almheiri, X. Dong, and D. Harlow. "Bulk locality and quantum error correction in AdS/CFT". In: *Journal of High Energy Physics* 2015.4, 163 (2015). DOI: 10.1007/JHEP04(2015)163. arXiv: 1411.7041.

[4] D. Tong. *Lectures on Quantum Field Theory*. Author's lecture notes; especially sections 2–6. University of Cambridge. URL: https://www.damtp.cam.ac.uk/user/tong/qft.htm (visited on 11/05/2026).

[5] W. K. Wootters and W. H. Zurek. "A single quantum cannot be cloned". In: *Nature* 299 (1982), pp. 802–803. DOI: 10.1038/299802a0.

[6] D. Gottesman. "Stabilizer Codes and Quantum Error Correction". PhD thesis. California Institute of Technology, 1997. arXiv: quant-ph/9705052.

[7] E. Knill and R. Laflamme. "Theory of quantum error-correcting codes". In: *Physical Review A* 55 (1997), pp. 900–911. DOI: 10.1103/PhysRevA.55.900. arXiv: quant-ph/9604034.

[8] J. F. Clauser, M. A. Horne, A. Shimony, and R. A. Holt. "Proposed Experiment to Test Local Hidden-Variable Theories". In: *Physical Review Letters* 23 (1969), pp. 880–884. DOI: 10.1103/PhysRevLett.23.880.

[9] S. Navas et al. "Review of Particle Physics". In: *Physical Review D* 110, 030001 (2024). With the 2025 update; numerical inputs are specified separately. DOI: 10.1103/PhysRevD.110.030001. URL: https://pdg.lbl.gov/2025/.

[10] J. de Blas, S. Dittmaier, and R. Kogler. *Electroweak Model and Constraints on New Physics*. Review of Particle Physics, chapter 10. Particle Data Group. 2025. URL: https://pdg.lbl.gov/2025/reviews/rpp2025-rev-standard-model.pdf (visited on 11/05/2026).

[11] Particle Data Group. *Quantum Chromodynamics*. Review of Particle Physics, chapter 9. 2025. URL: https://pdg.lbl.gov/2025/reviews/rpp2025-rev-qcd.pdf (visited on 11/05/2026).

[12] C. S. Wu, E. Ambler, R. W. Hayward, D. D. Hoppes, and R. P. Hudson. "Experimental Test of Parity Conservation in Beta Decay". In: *Physical Review* 105 (1957), pp. 1413–1415. DOI: 10.1103/PhysRev.105.1413.

[13] S. Weinberg. "A Model of Leptons". In: *Physical Review Letters* 19 (1967), pp. 1264–1266. DOI: 10.1103/PhysRevLett.19.1264.

[14] G. Dusson, I. M. Sigal, and B. Stamm. *The Feshbach–Schur map and perturbation theory*. 2021. arXiv: 2105.02058.

[15] S. Bravyi, D. P. DiVincenzo, and D. Loss. "Schrieffer–Wolff transformation for quantum many-body systems". In: *Annals of Physics* 326 (2011), pp. 2793–2826. DOI: 10.1016/j.aop.2011.06.004. arXiv: 1105.0675.

[16] Y. Koide. "New view of quark and lepton mass hierarchy". In: *Physical Review D* 28 (1983). Erratum: Phys. Rev. D 29, 1544 (1984), pp. 252–254. DOI: 10.1103/PhysRevD.28.252.

[17] R. Foot. *A note on Koide's lepton mass relation*. 1994. arXiv: hep-ph/9402242.

[18] P. Żenczykowski. "A Note on Koide's Doubly Special Parametrization of Quark Masses". In: *Open Physics* 16 (2018), pp. 427–429. DOI: 10.1515/phys-2018-0058.

[19] National Institute of Standards and Technology. *Muon–electron mass ratio: 2022 CODATA recommended value*. 206.7682827(46). URL: https://physics.nist.gov/cgi-bin/cuu/Value?mmusme (visited on 11/05/2026).

[20] Particle Data Group. *Electron particle listing*. Electron mass input, page 1. 2025. URL: https://pdg.lbl.gov/2025/listings/rpp2025-list-electron.pdf (visited on 11/05/2026).

[21] Particle Data Group. *Tau particle listing*. Mass average: 1776.93(9) MeV, page 1. 2025. URL: https://pdg.lbl.gov/2025/listings/rpp2025-list-tau.pdf (visited on 11/05/2026).

[22] M. C. Gonzalez-Garcia and R. Wendell. *Neutrino Masses, Mixing, and Oscillations*. Review of Particle Physics, chapter 14. Particle Data Group. 2025. URL: https://pdg.lbl.gov/2025/reviews/rpp2025-rev-neutrino-mixing.pdf (visited on 11/05/2026).

[23] D. Tong. *Lectures on General Relativity*. Author's lecture notes. University of Cambridge. URL: https://www.damtp.cam.ac.uk/user/tong/gr/gr.pdf (visited on 11/05/2026).

[24] D. Elliman. *Supporting code and verification for The Holographic Quantum Stack Machine*. Version 1.0.0. Nov. 5, 2026. DOI: 10.5281/zenodo.23991265.

[25] S. Chandrasekharan and U.-J. Wiese. "Quantum Link Models: A Discrete Approach to Gauge Theories". In: *Nuclear Physics B* 492 (1997), pp. 455–471. DOI: 10.1016/S0550-3213(97)80041-7. arXiv: hep-lat/9609042.