Fermions and Colour

Transport and its limits in qutrit–qubit networks

David Elliman
Neuro-Symbolic Ltd
dave@neusym.ai

4 November 2026

Abstract

The quantum state of two fermions, such as quarks, changes sign when they are exchanged, and quarks also carry a three-valued internal label called colour. Can networks of small quantum registers reproduce both features? We build each site from four qutrits (three-level quantum systems). Its lowest-energy states record whether the site is occupied and carry a three-component SU(3) colour. Moving the occupation, carrying the colour and obtaining the fermion sign turn out to be separate requirements.

Interactions that swap pairs of qutrits move a particle only indirectly, through brief virtual visits to higher-energy states; the same processes also create particle pairs and unwanted colour-dependent returns. A cancellation that removes these on one isolated link fails where links meet, because different links share virtual states. The fermion sign can instead come from preparing the empty sites in an antisymmetric colour state, using a known correspondence between internal symmetry and particle statistics. This needs shared entanglement, allows at most three empty sites, and is not singled out by the hopping Hamiltonian.

Adding colour-neutral auxiliary registers keeps the virtual processes on different links apart. Six existing Hamiltonian parameters are then fixed from single-link equations, without fitting masses, mixing angles or propagation scores. At an operating point chosen in advance, a simplified (effective) model of a four-site junction gives minimum fidelity 0.99499 with ideal fermion hopping and maximum leakage out of the encoded states of 0.207%, over sixteen hop times on 102 tested states. A separate implementation reproduces these results. A point with a smaller energy penalty on pair creation fails the same sixteen-hop test. Virtual creation and removal of particle pairs explain, to leading order, the remaining coherent error.

A residual calculation bounds the difference between this effective model and the full four-site Hamiltonian. For the same sampled test it gives microscopic fidelity at least 0.99375 and leakage at most 0.326%. The bound is evaluated numerically, without interval certification. This is a finite-model result: the interactions still require supplied three-register controls, precisely calibrated parameters and a prepared background state. Scalable preparation, robustness and a complete microscopic theory of fermions remain open.

Finite-model research draft with independent computational checks. The microscopic bounds are numerical and restricted to the stated experiment; human peer review remains pending.

Keywords: qutrits; fermionic encodings; colour transport; effective Hamiltonians; exchange statistics

Contents

1 What must a coloured particle do? 3

2 A block with occupation and colour 4

3 What pair exchange actually generates 7

4 Two hopping channels 9

5 Testing the fermion sign at a junction 12

6 The state of the empty blocks 14

7 A local dressed junction 16

8 Conclusions 27

A Spectrum and boundary algebra of the block 28

B Exact second-order cancellation certificate 29

C Permutation classification of the junction 30

D Link phases, rotation classes and their scope 31

E Proof and evaluation of the residual bound 33

F Verification and reproducibility 34

1 What must a coloured particle do?

This paper asks a deliberately limited question. Suppose we build a small network out of quantum registers (finite-dimensional quantum systems, the quantum analogue of memory cells). Can a pattern in that network behave like a moving particle that carries an internal "colour" label and obeys the exchange rule of fermions (particles such as electrons and quarks, whose joint quantum state changes sign when two of them are swapped)? These questions can be studied in finite quantum systems without assuming that the systems already describe elementary particles.

A useful microscopic model of a moving particle must do more than transfer an excitation between two registers. It must specify three things: what counts as a particle, which internal state moves with it, and how two particles behave when their paths exchange.

Here a *qutrit* is a three-dimensional quantum register, that is, a quantum system with three basis states. When we call it a fundamental colour register, we mean that a colour rotation $U \in \mathrm{SU}(3)$ acts on it as the three-dimensional representation **3**. (The group SU(3) is the set of 3×3 unitary matrices with unit determinant; it is the symmetry group of colour in the strong interaction. A *representation* is a rule assigning a matrix to each group element, and **3** is the basic three-dimensional one, also called the *fundamental* or *triplet* representation.) Colour is an internal quantum label. It is not a spatial direction, and global SU(3) covariance alone is not a local gauge theory of quantum chromodynamics: we only require that the model treats all colour rotations, applied everywhere at once, as symmetries. A qubit is a two-dimensional register; it can supply the binary grading (the even/odd bookkeeping) used in a fermionic encoding. Neither register's dimension, by itself, determines the statistics of a collective excitation.

The model studied below distinguishes three operations, each of which can succeed or fail independently:

1. **Occupation transport:** a two-valued label changes from 10 to 01 on neighbouring blocks. In words, "particle here, none there" becomes "none here, particle there".
2. **Colour transport:** the colour associated with the occupied block follows that change.
3. **Fermionic exchange:** two nonzero sequences of legal hops differ by a minus sign after their final internal states are identified consistently.

The third condition concerns an interference experiment: two different routes lead to the same final arrangement, and the fermion sign is the relative sign between them. Multiplying every hopping amplitude by -1 does not impose it: two histories containing the same three links acquire the same three scalar factors, so their relative sign is unchanged.

Finite-register descriptions of gauge fields and quantum information have a substantial history. Qutrit codes provide useful finite state spaces[1]; Hamiltonian lattice gauge theory and gauge magnets specify how local constraints relate matter and links[2, 3]. Our question is narrower: what does a particular static, SU(3)-invariant pair-exchange Hamiltonian actually generate in a selected low-energy subspace? Here "pair exchange" means an interaction built only from operators that swap the states of two registers, and "static" means that the Hamiltonian is fixed rather than switched on and off in time. The use of representation multiplicities (repeated copies of the same representation) for encoding is related to exchange-only quantum computation[4], but a sequence of controlled gates is not the same construction as one fixed Hamiltonian.

The statistics benchmarks likewise use established ideas. Jordan–Wigner, Bravyi–Kitaev and related encodings represent fermions by spin operators[5–9]. Levin and Wen formulate

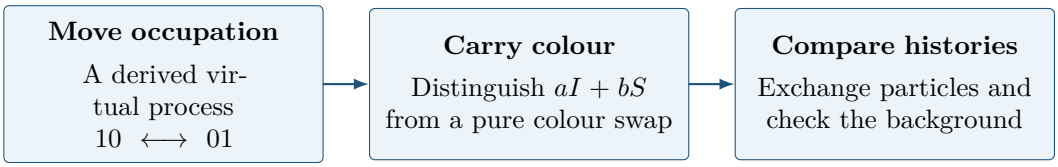


Figure 1: Three tests that must be satisfied together. The pair-exchange construction derives occupation motion. A colour swap and a statistics sign require additional checks, including the state left in nominally empty blocks.

statistics directly in terms of hopping operators[10]. Gauge constraints and string operators can produce fermionic excitations in bosonic microscopic systems[11–13]. We use these results as reference constructions; the existence of a minus sign after a prescribed string dressing (an extra operator attached along the particle's path) is not a new derivation of fermions from qutrit exchange.

The contribution of this paper is the analysis of one particular pair-exchange law: its virtual hopping coefficients, an isolated-link cancelling family, the failure of that cancellation at a junction, and a tuned auxiliary extension. The exchange sign obtained from an antisymmetric colour background is an application of known colour–orbital symmetry relations, not a new statistics mechanism. The corresponding occupancy limit is also known[14]; entangled-photon experiments have simulated the associated bosonic and fermionic interference[15, 16].

The presentation follows the physical construction. Sections 2–4 build the basic unit (a "block" of four qutrits) and work out which couplings between blocks make a particle move. Sections 5–6 explain the exchange test and show that the state of the empty blocks decides whether it can be passed. Section 7 realizes the relational code on fixed site registers, explains its preparation and scaling costs, and reports the qualified auxiliary-junction result, including numerical error bounds. Appendix D keeps the finite rotation diagnostics separate from dynamical exchange.

2 A block with occupation and colour

This section builds the basic unit of the model. The aim is a small group of qutrits whose lowest-energy states can be read as "empty" or "occupied", while also carrying one complete colour triplet. We first find such states using symmetry, then fix a convention for which of them counts as occupied, and finally give a Hamiltonian that makes them the ground states.

2.1 Two copies of a colour triplet

Each block contains four fundamental qutrits, labelled $0, 1, 2, 3$. Its Hilbert space (the space of all its quantum states) is

$$\mathcal{H}_{\text{block}} = (\mathbb{C}^3)^{\otimes 4}, \qquad \dim \mathcal{H}_{\text{block}} = 81. \tag{1}$$

The operator S_{ij} exchanges the complete states of qutrits i and j. It is Hermitian, satisfies $S_{ij}^2 = I$, and commutes with $U^{\otimes 4}$ for every $U \in \mathrm{SU}(3)$. In words, swapping two registers does not care which colour basis is used. Thus a real linear combination of pair swaps is an SU(3)-invariant Hamiltonian.

Schur–Weyl duality separates the colour representation from the multiplicity with which it occurs[17]. This theorem says that the 81 states of four qutrits split into pieces, each of which is a colour representation of SU(3) paired with a representation of the permutations of the four qutrits; the permutation part counts how many equivalent copies of that colour representation

there are. To avoid ambiguous names for the two 15-dimensional representations, we label colour representations by their Dynkin indices (the standard pair of integers that identifies an SU(3) representation):

$$(\mathbf{3})^{\otimes 4} \cong R_{(4,0)} \otimes V_{(4)} \oplus R_{(2,1)} \otimes V_{(3,1)} \oplus \bar{\mathbf{6}} \otimes V_{(2,2)} \oplus \mathbf{3} \otimes V_{(2,1,1)}. \tag{2}$$

Here V_λ is a representation of the permutation group S_4. The dimensions are $15{+}45{+}12{+}9 = 81$. In particular, the last summand contains three equivalent colour triplets. Choosing between these copies can encode a neutral logical degree of freedom without deleting any of the three colours. This is the key idea of the block: *which* copy of the triplet the block is in can serve as an occupation label, while the colour itself is carried in full by every copy.

An explicit triplet embedding makes this statement concrete. Define a normalized antisymmetric pair by

$$|\bar{a}\rangle_{ij} = \frac{1}{\sqrt{2}} \sum_{b,c=0}^{2} \epsilon_{abc} |b\rangle_i |c\rangle_j, \qquad W|c\rangle = \frac{1}{\sqrt{2}} \sum_{a,b=0}^{2} \epsilon_{cab} |\bar{a}\rangle_{01} |\bar{b}\rangle_{23}. \tag{3}$$

The Levi-Civita tensor ϵ_{abc} is $+1$ on even permutations of $(0,1,2)$, -1 on odd permutations, and zero when indices repeat. Each first-stage pair transforms as an antitriplet (the conjugate representation $\bar{\mathbf{3}}$); combining two antitriplets antisymmetrically in the second stage selects a fundamental triplet. The three maps

$$W_0 = W, \qquad W_1 = S_{02} W, \qquad W_2 = S_{03} W \tag{4}$$

have mutually orthogonal images and satisfy $W_i^\dagger W_j = \delta_{ij} I_3$, so they place a colour state into three non-overlapping copies inside the block. We retain the last two copies:

$$E = (W_1, W_2), \qquad E^\dagger E = I_6, \qquad \mathcal{H}_{\text{code}} \cong \mathbb{C}^2 \otimes \mathbb{C}^3. \tag{5}$$

The resulting six-dimensional *code space* is a two-state label times a colour triplet. The first factor distinguishes equivalent colour triplets. It is an occupation *label*, not yet a physical particle number.

2.2 Fixing the occupation convention

The two-state label can be described in several bases, and choosing the wrong one later leads to a wrong hopping operator. This subsection fixes the convention once and for all. Let X, Y, Z denote Pauli matrices in the copy basis (W_1, W_2). Direct permutation of the qutrits gives

$$E^\dagger S_{01} E = X \otimes I_3, \qquad E^\dagger S_{23} E = -X \otimes I_3, \qquad E^\dagger (S_{02} S_{13}) E = -Z \otimes I_3. \tag{6}$$

(These are the actions of the physical swaps when restricted to the code space.) We define occupation by the eigenvalue of X. The normalized states are

$$|n, c\rangle_E = \frac{W_1|c\rangle + (-1)^n W_2|c\rangle}{\sqrt{2}}, \qquad n = 0, 1, \quad c = 0, 1, 2. \tag{7}$$

Here n is the occupation (0 for empty, 1 for occupied) and c the colour. It is helpful to introduce a second set of Pauli symbols, τ_x, τ_y, τ_z, specifically for this occupation basis:

$$\tau_z = X, \qquad \tau_x = Z, \qquad \tau_y = -Y, \qquad \hat{n} = \frac{I - \tau_z}{2}. \tag{8}$$

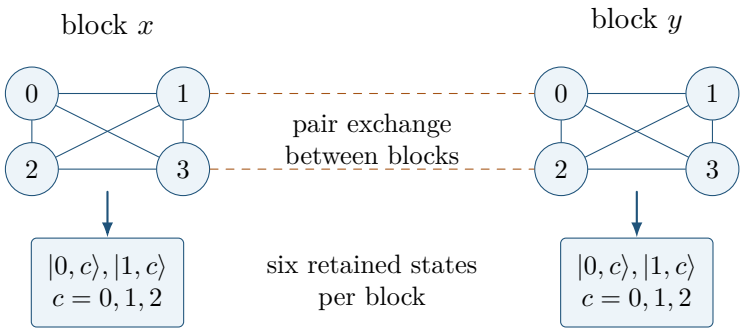


Figure 2: The physical and encoded descriptions. Each block has four fundamental qutrits and six internal pair bonds. The selected subspace carries a binary multiplicity label and a complete colour triplet. Dashed inter-block lines indicate available pair couplings, not distances or an assumed spatial tiling.

The distinction prevents an important basis error. A number-preserving two-block hop is proportional to

$$T = (|01\rangle\langle 10| + |10\rangle\langle 01|) \otimes S_{\mathrm{c}} = \frac{\tau_x\tau_x + \tau_y\tau_y}{2} \otimes S_{\mathrm{c}} = \frac{ZZ + YY}{2} \otimes S_{\mathrm{c}}. \tag{9}$$

Here S_{c} swaps the two colour triplets. So T moves the particle from one block to the other and moves its colour with it; this is the target operation that the rest of the paper tries to generate. The expression $(XX + YY)/2$ would instead conserve occupation defined along Z, which is not our convention.

Both $|0, c\rangle_E$ and $|1, c\rangle_E$ transform as colour triplets. Calling $n = 0$ "empty" therefore does not make it a colour-singlet vacuum: an empty block still carries a colour. This distinction will be decisive in Section 6.

2.3 An isolated six-state ground space

An energy gap is needed before virtual-transition calculations are meaningful. We therefore want a Hamiltonian on each block, built from pair swaps, whose lowest-energy states are exactly the six code states and which separates them from all other states by a finite energy. Define the class sum $\Sigma_4 = \sum_{i<j} S_{ij}$ and the pair operator

$$M = 4I + 2(S_{01} + S_{23}) + S_{02} + S_{03} + S_{12} + S_{13}. \tag{10}$$

The kernel of M is $\operatorname{im} W_0$. The desired six states have $M = 2$, but so do six unwanted states in the $\bar{\mathbf{6}}$ sector. Consequently, the old choice $(G/2)M$ does not isolate the desired excited manifold. Some cross-block couplings connect it resonantly (that is, at equal energy) to the unwanted states; this is not true of every coupling.

The fixed selector used here is

$$\begin{aligned} L &= I - \frac{1}{2}M + 2(\Sigma_4 + 2I) \\ &= 3I + S_{01} + S_{23} + \frac{3}{2}(S_{02} + S_{03} + S_{12} + S_{13}), \qquad H_{\mathrm{local}} = GL, \quad G > 0. \end{aligned} \tag{11}$$

All its nonconstant terms are pair swaps. Its spectrum, including multiplicities, is

$$\operatorname{spec} L = \{0^{(6)}, 1^{(3)}, 2^{(6)}, 4^{(6)}, 5^{(15)}, 6^{(30)}, 11^{(15)}\}. \tag{12}$$

Thus $\ker L = \operatorname{im} E$, with gap G: the six code states have zero energy and every other state costs at least G. Appendix A derives the spectrum by resolving Σ_4 and the two disjoint internal swaps. The gap is an isolation mechanism. It is not, by itself, an autonomous error-correction process or protection against every perturbation.

3 What pair exchange actually generates

We now couple two blocks with pair swaps between their qutrits and ask what effective motion results inside the code space. The short answer is: at first order in the coupling, colours can be exchanged but the particle cannot move; at second order, through brief visits to higher-energy states, it can move, but other unwanted processes (such as creating two particles at once) appear at the same order.

3.1 First order: colour exchange without occupation motion

Take two blocks, x and y, with $H_0 = G(L_x + L_y)$. Let $\mathcal{E} = E_x \otimes E_y$, with the encoded factors reordered as occupation–occupation–colour–colour, and let $\mathcal{P} = \mathcal{E}\mathcal{E}^\dagger$, $\mathcal{Q} = I - \mathcal{P}$. Here $\mathcal{P}$ projects onto the code space of the two blocks and $\mathcal{Q}$ onto everything else. We write encoded matrices as $\mathcal{E}^\dagger O \mathcal{E}$ rather than confusing them with physical-space projections.

To compute what a physical operator does inside the code, we need its *compression* $E^\dagger O E$, the part of O that acts within the code space. The local colour matrix unit $e_{ab}^{(i)} = |a\rangle\langle b|_i$ (which changes colour b to colour a on qutrit i) has compression

$$E^\dagger e_{ab}^{(i)} E = A_i \otimes |a\rangle\langle b| + \delta_{ab} \frac{I_2 - A_i}{3} \otimes I_3, \quad A_{0,1} = \frac{2I_2 + 3X}{8}, \quad A_{2,3} = \frac{2I_2 - 3X}{8}. \tag{13}$$

The two terms respectively act on the retained colour and on its scalar part. Expanding a swap as $S_{xi,yj} = \sum_{a,b} e_{ab}^{(xi)} e_{ba}^{(yj)}$ gives

$$\mathcal{E}^\dagger S_{xi,yj} \mathcal{E} = A_i \otimes A_j \otimes \left(S_{\mathrm{c}} - \frac{I_9}{3} \right) + \frac{I_{36}}{3}. \tag{14}$$

Because every A_i is diagonal in the occupation basis, a first-order cross-block swap cannot change either occupation. It can still change colour.

For example, take $r = (1, 1, -1, -1)$ and

$$V_r = g \sum_{i,j=0}^{3} r_i r_j S_{xi,yj}. \tag{15}$$

Then $\sum_i r_i A_i = 3X/2$ and $\sum_i r_i = 0$, so

$$\mathcal{E}^\dagger V_r \mathcal{E} = \frac{9g}{4} \tau_z^x \tau_z^y \otimes \left(S_{\mathrm{c}} - \frac{I_9}{3} \right). \tag{16}$$

The coefficient depends on the two occupations, but neither occupation moves. This is exact beyond the first-order projection: the physical Hamiltonian $H_0 + V_r$ commutes with S_{01}^x and S_{01}^y separately, so the occupation of each block is conserved exactly.

There is also a quantitative first-order limitation, which shows that no choice of pair couplings can produce the target hop at first order. The identity and all 28 pair swaps on eight physical qutrits compress to a nine-dimensional operator space. Adding the target T in Eq. (9) raises the rank to ten. The matrix $J = \tau_y^x \tau_y^y \otimes S_{\mathrm{c}}$ is orthogonal to that entire pair span under the trace inner product, whereas

$$\mathrm{Tr}(J^\dagger T) = 18, \qquad \|J\|_F = 6. \tag{17}$$

Consequently every first-order encoded pair law A obeys

$$\|A - T\|_F \geq 3. \tag{18}$$

The Frobenius norm is the square root of the sum of squared matrix entries. So every first-order law stays a fixed, finite distance away from the target. This bound excludes the specified target from the linear pair span. It says nothing against higher-order virtual transport.

Occupation block of $GH^{(2)}$	V_+	V_-
$\langle 01\vert\cdot\vert 10\rangle$	$\frac{404}{2025}I - \frac{12}{25}S_c$	$-\frac{491}{3500}I + \frac{3}{500}S_c$
$\langle 00\vert\cdot\vert 11\rangle$	$-\frac{491}{3500}I + \frac{3}{500}S_c$	$\frac{404}{2025}I - \frac{12}{25}S_c$
$\langle 00\vert\cdot\vert 00\rangle$, $\langle 11\vert\cdot\vert 11\rangle$	$-\frac{6017}{10500}I + \frac{67}{500}S_c$	$-\frac{1052}{2025}I + \frac{4}{25}S_c$
$\langle 01\vert\cdot\vert 01\rangle$, $\langle 10\vert\cdot\vert 10\rangle$	$-\frac{1052}{2025}I + \frac{4}{25}S_c$	$-\frac{6017}{10500}I + \frac{67}{500}S_c$
Single-occupation flips	0	0

Table 1: Exact second-order coefficients, including $g_0^2 = 16/225$. Hermitian conjugates supply the reverse off-diagonal blocks. $I = I_9$. All blocks are retained; in particular, pair creation is of the same perturbative order as hopping.

3.2 Second order: virtual excursions enable motion

Higher-order processes work through *virtual excursions*: the coupling briefly takes the system out of the code space into a higher-energy state, and a second application brings it back, possibly in a different code state. Such two-step processes are suppressed by the energy cost of the intermediate state but can do things a single step cannot.

We measure energies in one fixed reference unit and set $\hbar = 1$. The inherited numerical coupling is $g_0 = -4/15$. This is an input, not a mass prediction. To distinguish perturbative order from an energy scale, write $H(\lambda) = H_0 + \lambda V$, with dimensionless expansion parameter λ. For the first-order-silent couplings below (those with zero first-order effect in the code),

$$H_{\text{eff}}(\lambda) = \lambda^2 H^{(2)} + O(\lambda^3), \qquad H^{(2)} = -\mathcal{E}^\dagger V \mathcal{Q}(\mathcal{Q}H_0\mathcal{Q})^{-1}\mathcal{Q}V\mathcal{E}. \tag{19}$$

The inverse is taken only on the excited subspace. The formula describes a transition out of the code, propagation through an excited state with its actual energy denominator, and a return to the code. It is the usual second-order degenerate-perturbation expression[18]; it must not be obtained by inverting the singular H_0 on its ground space.

For a real port vector v, define $V_v = g_0 \sum_{ij} v_i v_j S_{xi,yj}$. A *port vector* is simply a set of four weights saying how strongly each qutrit of a block takes part in the coupling. The vectors

$$s = (1,-1,1,-1), \qquad t = (1,-1,-1,1), \qquad V_+ = V_s + V_t, \quad V_- = V_s - V_t \tag{20}$$

give zero first-order code matrix. At second order each occupation block is $aI_9 + bS_c$, since these are the two invariant operators on $\mathbf{3}\otimes\mathbf{3} = \mathbf{6}\oplus\bar{\mathbf{3}}$. The identity term leaves the colours at their sites; the swap term exchanges them. So for a hop, a measures the unwanted part (the particle moves but leaves its colour behind) and b the wanted part (the colour travels with the particle).

Table 1 gives the complete block structure for $V_\pm$. Its rows are, in order: hopping ($01 \leftrightarrow 10$), pair creation or annihilation ($00 \leftrightarrow 11$), energy shifts of configurations with equal occupations, energy shifts with unequal occupations, and processes that change only one occupation. The coefficients were calculated by an exact rational solve in the 560-dimensional physical colour-weight sector $(3,3,2)$ (the states containing three qutrits of each of two colours and two of the third). This sector contains both colour arrangements needed to determine a and b; SU(3) covariance then determines the nine-dimensional colour operator. It is a symmetry reduction of the two-block problem, not a claim that a dense 6561-state time evolution was run.

For V_s or V_t alone, both the hop and pair block equal $(16789/1134000)I - (237/2000)S_c$, but there are also eight nonzero single-flip blocks. Their coefficients have magnitudes corresponding to $-101I/1200 + 9S_c/400$, with signs depending on the block. Those terms cannot be discarded when checking occupation constraints.

It is essential to distinguish two parity statements. *Matter parity* is whether the total number of particles is even or odd; fermion models normally require it to be conserved. The question is whether that conservation holds only at the order computed, or exactly. Define the physical involution $\Pi = S^x_{01} S^y_{01}$, whose code action is $(-1)^{\hat{n}_x + \hat{n}_y}$. Then

$$[H_0, \Pi] = 0, \qquad \Pi V_+ \Pi = V_+, \qquad \Pi V_- \Pi = -V_-. \tag{21}$$

Thus $H_0 + \lambda V_+$ conserves Π exactly. The Hamiltonian $H_0 + \lambda V_-$ does not conserve Π, but matter parity is defined on the code, and its extension to excited states is a choice. The involution $\Pi' = S^x_{01} S^y_{23}$ satisfies

$$[H_0, \Pi'] = 0, \qquad \Pi' V_- \Pi' = V_-, \qquad \Pi' V_+ \Pi' = -V_+, \tag{22}$$

and Eq. (6) makes its code action $-(-1)^{\hat{n}_x + \hat{n}_y}$. So $H_0 + \lambda V_-$ also conserves matter parity on its low-energy band, at every order and at finite coupling. Neither involution commutes with a mixture of V_+ and V_-. For a mixture, a zero entry at second order must not be described as an all-order microscopic symmetry.

3.3 Finite gaps and the limits of the effective law

The second-order formula is only the first term of a series. The exact energy-dependent return operator, which includes every order, is

$$H_{\mathrm{ret}}(z) = \mathcal{P}H\mathcal{P} + \mathcal{P}H\mathcal{Q}(z - \mathcal{Q}H\mathcal{Q})^{-1}\mathcal{Q}H\mathcal{P}. \tag{23}$$

This Schur-complement or Feshbach formula[19] retains every excited-state return. Its expansion around $\lambda = 0$ gives Eq. (19), but replacing it by its second-order part is an approximation. In particular, cancellations between second-order coefficients need not persist at finite coupling or long propagation times.

The transport calculations retain the declared colour-weight sectors. The clock Hamiltonian and its three coherent spectator blocks from Ref. [20] are not part of the junction Hamiltonian studied here. The spectator particles tested below are a different extension.

4 Two hopping channels

The previous section showed that second-order hopping exists but, for the simple couplings tried, always drags along an unwanted identity term a. This section explains why, using symmetry, and then shows how a more elaborate coupling can cancel a on a single link, and why that cancellation breaks down once several links meet at a junction.

4.1 A symmetry explanation of the fixed ratios

Four real vectors form a basis for the site weights on a block:

$$e_+ = (1, -1, 0, 0), \quad e_- = (0, 0, 1, -1), \quad r = (1, 1, -1, -1), \quad u = (1, 1, 1, 1). \tag{24}$$

Write $B_{pq} = \sum_{ij} p_i q_j S_{xi,yj}$ and keep the coupling normalization explicit:

$$V(c) = g_0 \sum_{p,q \in \{e_+, e_-, r, u\}} c_{pq} B_{pq}, \qquad c_{pq} \in \mathbb{R}. \tag{25}$$

There are 16 independent real coefficients. So the most general pair coupling between two blocks is described by a 4×4 real matrix c. The symbols $p \otimes q$ sometimes used for a port pattern refer to these weights; they are not additional microscopic tensor factors.

The selector commutes with both S_{01} and S_{23}. Each occupation state has definite eigenvalues under these swaps: $(+,-)$ for $n=0$ and $(-,+)$ for $n=1$. Changing occupation therefore reverses *both* signs. Under conjugation, the port vectors carry the characters shown in Table 2.

Port	e_+	e_-	r	u
Character under (S_{01}, S_{23})	$(-,+)$	$(+,-)$	$(+,+)$	$(+,+)$

Table 2: A character specifies whether the port changes sign under each internal swap. The two odd ports must occur once each on a block to reverse both of its code characters.

Since the excited-state resolvent preserves these characters, a second-order hop can contain only e_+ followed by e_-, or the reverse, on each block. This is a *selection rule*: symmetry alone forbids most processes. Of the 256 ordered pairs of port products, 192 cannot connect any two code states and vanish exactly. Of the remaining 64, 55 are nonzero in at least one occupation block. The selection rule fixes which processes can occur; their nonzero sizes still depend on the virtual energies.

For every coefficient matrix c, the hopping block is

$$\begin{aligned} G\langle 01|H^{(2)}|10\rangle &= a(c) I_9 + b(c) S_{\mathrm{c}}, \\ a(c) &= \frac{101}{2025} c_{++}c_{--} - \frac{491}{14000} c_{+-}c_{-+}, \\ b(c) &= -\frac{3}{25} c_{++}c_{--} + \frac{3}{2000} c_{+-}c_{-+}. \end{aligned} \tag{26}$$

All fractions include g_0^2. Only two combinations of couplings matter, $c_{++}c_{--}$ and $c_{+-}c_{-+}$; we call these the two *hopping channels*. The pair-creation block has the two coefficient pairs interchanged. Ports involving r or u cannot contribute to these two-flip blocks, although they can contribute to diagonal and single-flip blocks.

If only the first product in Eq. (26) is present, $b/a = -243/101$. If only the second is present, $b/a = -21/491$. Neither channel alone permits $a = 0$ with $b \neq 0$. In other words, each channel on its own has a locked ratio of wanted to unwanted hopping, and neither ratio is the pure colour-carrying hop. These are exact restrictions on the fixed selector and the stated perturbative order. An antisymmetric inter-block pattern need not have zero hop: $B_{+,-} - B_{-,+}$ gives $(a, b) = (491/14000, -3/2000)$ and obeys the second ratio.

4.2 Cancellation on one isolated link

Mixing the two channels can make a vanish, but at a price. With $e_\pm$ alone, combining the two channels always generates single-occupation flips: $a = 0$ with $b \neq 0$ needs all four of $c_{\pm\pm}$ to be nonzero, while zero flips would require $c_{++}c_{-+} = c_{+-}c_{--} = 0$. That does not extend to the full family. The ports r and u contribute to the single-flip blocks and can cancel them, even though they do not alter Eq. (26).

The following explicit real coefficient matrix cancels the identity part of the second-order hop. Order its rows and columns as (e_+, e_-, r, u), and define

$$\rho = \sqrt{\frac{56560}{39771}}, \qquad \alpha = \frac{5621}{404}, \qquad \beta = \frac{136411}{3232}, \qquad d = \sqrt{\frac{\beta}{\alpha}}. \tag{27}$$

Then take

$$c = \begin{pmatrix} 1 & \rho & 1 & d \\ \rho & 1 & \alpha\rho & -d\alpha\rho \\ 1 & \alpha\rho & 0 & 0 \\ d & -d\alpha\rho & 0 & 0 \end{pmatrix}. \tag{28}$$

On one isolated link, this symmetric inter-block pattern has zero first-order code matrix, every second-order single-flip block vanishes, and

$$G\langle 01|H^{(2)}|10\rangle = -\frac{39064}{331425} S_{\mathrm{c}}. \tag{29}$$

That is, at second order the hop is a pure colour-carrying swap. Appendix B gives an algebraic cancellation certificate. A separate construction of the 560-state resolvent agrees within 9.3×10^{-14} in the checked entries.

This isolated-link cancellation does not supply the required protected fermion model. Its coefficients are selected to satisfy algebraic cancellation conditions, its microscopic interaction contains both parities, and its pair-creation and diagonal returns remain. An overall small λ in $H(\lambda)$ is needed for a controlled perturbative interpretation. No higher-order identity cancellation or all-order matter-parity conservation follows from Eq. (29).

Appendix B gives the complete real first-order-silent cancellation family, up to scale, discrete signs and two continuous ratios in the r, u entries. Hop and pair blocks depend on the same two products through an invertible coefficient matrix, and vanish only together: one cannot keep the hop while removing pair creation. At $a = 0$, the pair block is approximately $0.0359I - 0.1692S_{\mathrm{c}}$, compared with the hop $-0.1179S_{\mathrm{c}}$. The printed pattern also has a diagonal colour coefficient about 832 times the hop. Thus an isolated pure-swap hop block is not an isolated transport process.

The diagonal colour term also commutes with the hop. Such a term keeps rotating the colour state while the particle hops, which degrades colour transport. For a normalized input colour state $|\chi\rangle$, its fidelity with ideal colour exchange at the hop time $T = \pi/(2|b|)$ is

$$F_{\mathrm{colour}} = \cos^2(d_b T) + \langle\chi|S_{\mathrm{c}}|\chi\rangle^2 \sin^2(d_b T). \tag{30}$$

For distinct basis colours this reduces to $\cos^2(d_b T)$; for a swap eigenstate it equals one. Thus a vanishing identity coefficient in the hop block alone does not guarantee colour transport for arbitrary inputs.

The restriction that does survive is useful: *with this selector, a coupling of definite matter-parity character has one of the locked second-order ratios.* A general mixture can evade the ratio obstruction through cancellation. Exact symmetry protection and cancellation at one perturbative order are distinct design requirements.

4.3 Why the isolated-link cancellation fails at a junction

A network needs junctions, where one block is linked to several others. There, a virtual excursion can start on one link and return through a different link sharing the same block. For links e and f, the additional second-order operator is

$$H_{ef}^{(2)} = -\mathcal{P}\left[V_e \mathcal{Q}(\mathcal{Q}H_0\mathcal{Q})^{-1}\mathcal{Q}V_f + V_f \mathcal{Q}(\mathcal{Q}H_0\mathcal{Q})^{-1}\mathcal{Q}V_e\right]\mathcal{P}. \tag{31}$$

It is absent from a sum of isolated-link calculations. For $V_\pm$, each vertex excites both endpoints and this return vanishes. The additional r, u ports can excite only the shared block, so the cancelling pattern does not have that protection.

For the entire real cancellation family in Appendix B, an unavoidable centre-only flip has magnitude relative to the hop

$$R = \frac{\rho}{39064/331425}\left(\frac{297\alpha}{1000} + \frac{33\beta}{250} + \frac{3}{10}\sqrt{\alpha\beta}\right) = 171.73082281125\ldots. \tag{32}$$

So at a junction, an unwanted process that flips only the centre's occupation has a matrix element more than 170 times the size of the desired hop. This ratio is independent of the family's two continuous off-core ratios. For the printed pattern the entry is 20.2413604 in the $G = 1$ normalization. The maximum entry over the full three-site code is 237.873036; the smaller value 212.153293 applies only to the all-distinct-colour weight sector. Once the between-link terms are retained, the tested star laws have minimum sampled fidelity (overlap with the ideal evolution) near zero and leakage (probability of leaving the code space) around 0.96 for $\lambda = 10^{-2}$ down to 10^{-4}. A three-block low-band calculation in a 34,650-state symmetry sector confirms the return terms and their order. These results reject the tested junction implementations; they do not retract the isolated-link identity in Eq. (29). They show why that identity cannot be applied link by link to obtain a network law.

5 Testing the fermion sign at a junction

This section sets up the experiment that decides whether particles behave as fermions. On a small star-shaped network, two particles can reach the same final arrangement by two different sequences of hops. For fermions, the two routes must differ by exactly a minus sign. We describe the test, explain what an added layer of qubits can and cannot supply, and review one standard way of producing the sign.

5.1 The two paths must be physically legal

We use a central site 0 with leaves $1, 2, 3$ (a four-site "star"). Let t_{ij} move a particle from j to i, and initially occupy leaves 1 and 3. The two histories are

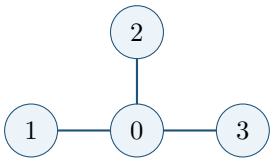
$$\mathscr{L} = t_{03}t_{20}t_{01}, \qquad \mathscr{R} = t_{01}t_{20}t_{03}, \tag{33}$$

with the rightmost hop acting first. In $\mathscr{L}$, the particle at leaf 1 moves through the centre to leaf 2, then the particle at leaf 3 moves to the centre; in $\mathscr{R}$, the particle at leaf 3 goes to leaf 2 first and the particle at leaf 1 then moves to the centre. Both end with sites 0 and 2 occupied, but which original particle ends at which site is interchanged between the two histories. The intermediate configurations are shown in Fig. 3. This order agrees with the directed-hop convention in the calculation and never attempts to hop into an already occupied site.

For a single internal species, the Levin–Wen criterion is $\mathscr{L} = -\mathscr{R}$ on the relevant nonzero two-particle sector[10]. Directed hops are generally singular operators: a hop annihilates a state (gives zero) if the source is empty or the destination is occupied. They should not be replaced indiscriminately by globally invertible toggles, that is, by operators that simply flip occupations regardless of the current state.

5.2 What the qubit layer supplies

Qubit strings can realize the required anticommutation algebra, the defining algebraic rule of fermion operators. On the eight-vertex face-incidence graph of an octahedron, the checked reference encodings use either eight face qubits or twelve edge qubits. The local face construction



Initially occupied: $1, 3$
Finally occupied: $0, 2$

$\mathscr{L}$: $\quad 1 \to 0, \quad 0 \to 2, \quad 3 \to 0$

occupied sets: $\{1,3\} \to \{0,3\} \to \{2,3\} \to \{0,2\}$

$\mathscr{R}$: $\quad 3 \to 0, \quad 0 \to 2, \quad 1 \to 0$

occupied sets: $\{1,3\} \to \{0,1\} \to \{1,2\} \to \{0,2\}$

Figure 3: The directed junction comparison. The two particles traverse the same three links in different orders. A common scalar hopping sign cancels from the comparison. Internal colours and any state left in the empty sites must also be compared.

has hop support at most four qubits; the edge construction has support at most five. (The *support* is the number of qubits a single hop operator acts on.) These are finite-graph examples of established fermion-to-spin encoding methods, particularly the superfast encoding[6, 9].

Adding $\mathbb{Z}_3$ qutrit-link shifts preserves the qubit exchange word and moves a three-valued detector charge. To describe a particle with a fixed charge, one must also restrict the toggle to an occupied-to-empty hop and impose the appropriate matter–link matching constraint. A bare endpoint toggle can create a pair, and is not automatically a number-preserving, gauge-invariant charged-particle hop. Nor does $\mathbb{Z}_3$ detector charge establish transport of an SU(3) triplet: a three-valued label is not the same as a full colour state.

There is a simple reason not to overinterpret the failure of a pure-qutrit unitary ansatz. For invertible operators on one common D-dimensional space,

$$\mathscr{U} = T_2 T_3^{-1} T_1 T_2^{-1} T_3 T_1^{-1} \quad \Longrightarrow \quad \det \mathscr{U} = 1. \tag{34}$$

Therefore $\mathscr{U} = -I_D$ requires even D, because $\det(-I_D) = (-1)^D$. This rules out that global unitary identity on the full space of finitely many qutrits, whose dimension is odd. It does not rule out fermionic excitations on constrained sectors or the singular directed-hop relation in Eq. (33). Indeed, ordinary two-fermion hopping on six sites acts in an odd-dimensional space of dimension $\binom{6}{2} = 15$ and still satisfies the directed relation.

5.3 Planar charge–flux dressing

One supplied route to the statistics sign uses planar $\mathbb{Z}_2$ charge–flux attachment. The idea, familiar from the toric code, is to bind a "charge" and a "flux" together: neither alone is a fermion, but the composite is. Put a qubit on each link and define commuting star and plaquette operators

$$A_v = \prod_{e \ni v} \sigma_e^z, \qquad B_f = \prod_{e \in \partial f} \sigma_e^x. \tag{35}$$

For a chosen framing face $f(v)$ adjacent to each site, impose

$$A_v \tau_z^v = 1, \qquad B_{f(v)} \tau_z^v = 1. \tag{36}$$

An occupied block then carries both an electric defect and a magnetic defect. A ribbon moving this composite has an X-type charge string and a Z-type dual flux string. Their odd crossing gives the minus sign. This is the standard fermionic charge–flux composite of the $\mathbb{Z}_2$ quantum double[11]; the constraints and routing rule are inputs here.

For the declared shortest, non-enclosing ribbons, the monochrome sector $|cccc\rangle$ (all four sites with the same colour) gives $\mathscr{L} = -\mathscr{R}$. With V_+, its nonzero scalar hop is

$$t_{\text{mono}} = \frac{1}{G}\left(\frac{404}{2025} - \frac{12}{25}\right) = -\frac{568}{2025G}. \tag{37}$$

Charge strings alone give the plus sign. The result is stable under the tested changes of adjacent framing face, but a ribbon wound around another occupied site can alter the comparison. Such a winding is a physically different path in the presence of flux, not an innocuous convention.

The diagonal terms in Table 1 need no ribbon. Both hopping and pair creation flip two endpoints and can be dressed to commute with Eq. (36). A term flipping only one endpoint is incompatible with this particular two-endpoint ribbon prescription. Gauge invariance of the pair term does not establish its full mixed pair–hop statistics algebra; only the two-particle hopping comparison has been tested.

6 The state of the empty blocks

So far we have ignored a subtle point from Section 2.2: an empty block still carries a colour. This section shows that this leftover "background" colour can record which path the particles took, and so spoil the interference that defines the fermion sign. We then ask what background state avoids the problem, and find that the answer is tightly constrained by representation theory.

6.1 The exact four-register comparison

For an arbitrary colour state on the four blocks, write the colour part of one hop as $K_{ij} = aI + bS_{ij}$. Temporarily separate the ribbon sign from the matter matrices. With the histories of Eq. (33), their matter factors are

$$L_{\mathrm{c}} = K_{03}K_{20}K_{01}, \qquad R_{\mathrm{c}} = K_{01}K_{20}K_{03}. \tag{38}$$

The dressed histories have the additional relative minus sign from the ribbon. A scalar-times-permutation comparison asks whether $L_{\mathrm{c}} = \vartheta P R_{\mathrm{c}}$ on all $3^4 = 81$ colour states. Here P is a rearrangement of the four colour registers and ϑ a number; the question is whether the two histories end in the same state up to relabelling and an overall factor.

Proposition 6.1. *For the same coefficients a, b on the three links, and $(a, b) \neq (0, 0)$, Eq. (38) satisfies a scalar-times-register-permutation identity only in the following two cases:*

$$b = 0: \quad P = I, \quad \vartheta = 1, \tag{39}$$

$$a = 0: \quad P = S_{02}S_{13}, \quad \vartheta = 1. \tag{40}$$

In particular, the particle-colour permutation S_{02} alone does not work for any nonzero choice of a, b.

In words: the two histories can only be compared cleanly if the hop either never moves colour ($b = 0$) or always moves it ($a = 0$), and in the second case the comparison necessarily swaps the colours of the empty sites as well as those of the particles. The proof is a finite exact polynomial classification, described in Appendix C. It covers all 24 register permutations and takes account of the one-dimensional kernel of the S_4 group algebra on four triplets. It does not classify more general state-dependent intertwining operators. Independent coefficients on the three links give further solutions, including $P = S_{02}$, but that case occurs only where two links carry no colour and the third is a singular projector ($b_i/a_i = \pm 1$), so that $L_{\mathrm{c}} = R_{\mathrm{c}}$ and no comparison of paths remains. If every link has $b_i \neq 0$, the only solution is still $P = S_{02}S_{13}$, with all $a_i = 0$.

For the derived V_+ coefficients, the best scalar fit using the double transposition has relative Frobenius residual about 0.589; the identity comparison has residual about 0.989.

These numbers are diagnostics of failure, not statistical phases. The fitted matter multiplier near 0.808 is not a unit-modulus exchange phase, and the ribbon sign must still be applied separately.

6.2 What a pure swap does to the background

The $b = 0$ branch moves occupation while leaving colour on the sites. It therefore fails the intended colour-carrying interpretation. The $a = 0$ branch transports colour, but swaps the colour of the empty block in the opposite direction: when a particle's colour moves forward, the empty block's colour moves back to take its place.

Under a pure swap, the two histories in Eq. (38) induce inverse four-cycles. Acting on an initial register ordering (c_0, c_1, c_2, c_3), they give

$$L_{\mathrm{c}} : (c_3, c_0, c_1, c_2), \qquad R_{\mathrm{c}} : (c_1, c_2, c_3, c_0). \tag{41}$$

Initially the particle colours are c_1, c_3 and the backgrounds are c_0, c_2. Finally the particles occupy $0, 2$, while the backgrounds occupy $1, 3$. The ratio of the two register permutations is exactly $S_{02}S_{13}$: it exchanges *both* pairs.

For example, take identical red particles and distinct green and blue backgrounds:

$$(g, r, b, r) \longrightarrow \begin{cases} (r, g, r, b), & L_{\mathrm{c}}, \\ (r, b, r, g), & R_{\mathrm{c}}. \end{cases} \tag{42}$$

The particle colours now agree, but the two background states are orthogonal. Tracing over that background removes the interference between the histories. The full evolution remains coherent; the failure is that it leaves a record of the path in degrees of freedom that were being called vacuum. This is the same reason that "which-path" information destroys interference in a two-slit experiment.

One possible restriction is to demand that every interchange of empty registers act as a scalar on the background, so that it cannot record anything. If all empty pairs can be interchanged, the transpositions generate their full permutation group. Its one-dimensional representations are the trivial and sign representations, so the background must be totally symmetric or totally antisymmetric. For m empty triplets,

$$\dim \mathrm{Sym}^m(\mathbb{C}^3) = \binom{m+2}{2}, \qquad \dim \Lambda^m(\mathbb{C}^3) = \binom{3}{m}. \tag{43}$$

At $m = 3$ the antisymmetric space is a singlet (a single, colour-neutral state). For $m \geq 4$ it vanishes, leaving the symmetric space. The latter is the irreducible SU(3) representation with Dynkin labels $(m, 0)$, so for $m > 0$ it contains no singlet. The conclusion uses irreducibility, not merely the fact that the dimension exceeds one.

6.3 A short representation-theory obstruction

The antisymmetric option is the familiar relation between internal symmetry and orbital statistics: for bosons, a state that is antisymmetric in colour must also be antisymmetric in position, so its motion carries a fermion-like sign. Herviou gives the lattice correspondence using colour and orbital representations: an antisymmetric colour sector of bosons supports the orbital sector associated with colour-symmetric fermions, with the bound that the number of antisymmetrized colours cannot exceed their dimension[14]. The limit of three empties here is that same restriction. Photonic demonstrations of statistics-dependent interference provide experimental precedents[15, 16]; they do not prepare the state or Hamiltonian of this model.

For the stronger demand that *every* empty-pair swap act as $-I$, the mean swap expectation must be -1. The lowest class-sum eigenvalues give the explicit bounds

$$\frac{1}{\binom{m}{2}}\sum_{i<j}\langle S_{ij}\rangle \geq \begin{cases} -1, & m=3, \\ -1/3, & m=4, \\ -1/5, & m=5,6. \end{cases} \tag{44}$$

These values are also instances of the complete-graph antisymmetric-marginal bounds of Allerstorfer et al.[21], using $\langle S\rangle = 1 - 2q_W$. They are standard representation-theory consequences, not a new exclusion principle. They exclude this exact universal-background sign for four or more empties. They do not exclude every fermionic encoding made from qutrits.

6.4 The standard centre-charge restriction

Could a different local code make an empty block a true singlet while retaining a triplet particle? The natural single-occupancy local space would be $\mathbf{1} \oplus \mathbf{3}$ (empty is colourless, occupied carries a triplet), rather than the present $\mathbf{3} \oplus \mathbf{3}$.

The following is a standard consequence of the action of the centre on tensor-product representations[17]. (The *centre* of SU(3) consists of the three matrices that commute with every group element, namely the identity times a cube root of unity; its eigenvalue on a representation records the *triality*.)

Proposition 6.2 (Fixed-product triality obstruction)**.** *Under the diagonal product action of* SU(3)*, no fixed tensor product*

$$\mathcal{H} = (\mathbf{3})^{\otimes N_q} \otimes (\bar{\mathbf{3}})^{\otimes N_{\bar{q}}} \tag{45}$$

contains both a singlet and a fundamental triplet as invariant subspaces.

Proof. For $\omega = e^{2\pi i/3}$, the central element ωI_3 acts on all of $\mathcal{H}$ as the same scalar,

$$\rho(\omega I_3) = \omega^{N_q - N_{\bar{q}}} I_{\mathcal{H}}. \tag{46}$$

It must therefore have that scalar on every invariant subspace. A singlet requires eigenvalue 1, whereas a fundamental triplet requires ω. These are different. □

Put simply, with this symmetry action a fixed collection of triplet and antitriplet registers cannot contain both a colourless state and a triplet. This proposition fixes the symmetry action as well as the microscopic carrier. It is not a prohibition on all qutrit-based descriptions of charged matter. A relational encoding can distribute compensating charge outside the region called a site; a gauge-dressed encoding can include boundary flux; and a direct-sum local carrier can have more than one triality. Each changes a stated premise. Simply adding a colour-neutral control bit, or switching off a coupling to an empty fundamental qutrit, does not change its centre eigenvalue.

7 A local dressed junction

This section turns the lessons of Sections 4 and 6 into a working construction, and then measures how well it works. In outline: we place the empty sites in a shared antisymmetric (singlet) colour state, which supplies the fermion sign; we show that a simple local hop then reproduces fermionic hopping exactly; we discuss what this costs in entanglement and scalability; and we add auxiliary two-state systems that keep virtual processes on different links apart. Finally, we calibrate the adjustable parameters, simulate the resulting effective model, and bound its error relative to the full microscopic Hamiltonian.

7.1 Realizing the relational code on fixed registers

A relational description need not imply that one shared hardware register must be moved from site to site. Here "relational" means that the colours are recorded relative to which sites are occupied, rather than in permanently assigned registers. Keep the occupation and colour registers at fixed sites, and encode their relation in the state. For four sites and two particles, add one boundary colour qutrit, denoted ∂, which never hops. If the occupied sites are $i < j$ and the empty sites are $k < l$, define the isometry

$$\mathcal{V}|i,j;c_i,c_j\rangle = |n(i,j)\rangle|c_i\rangle_i|c_j\rangle_j \otimes \frac{1}{\sqrt{6}} \sum_{a,b,c=0}^{2} \epsilon_{abc}|a\rangle_k|b\rangle_l|c\rangle_\partial. \tag{47}$$

The occupation string $n(i,j)$ has ones at i,j. The colour factors on the right are inserted into their physical site slots. The last factor is a singlet of the two empty registers and the boundary qutrit; it is not a product of three colourless sites. There are $\binom{4}{2}3^2 = 54$ code states (six choices of occupied pair times nine particle colours), giving a fixed-register realization of

$$\mathcal{H}_{\text{rel}} = \text{span}\{|i,j\rangle : i < j\} \otimes \mathbb{C}^3 \otimes \mathbb{C}^3. \tag{48}$$

The two colour entries refer to occupied positions in their current order, not to permanently distinguishable particles.

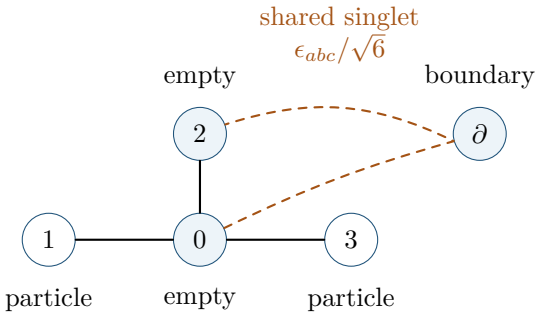


Figure 4: One code configuration: particles at sites 1 and 3; empty sites 0 and 2 share an antisymmetric singlet with the boundary register. Solid lines are allowed star hops. The dashed curve denotes a correlation in the prepared state, not an extra Hamiltonian link. Local hopping operators therefore coexist with a nonlocal preparation resource.

Register	Colour representation	Function
Site occupation	**1**	Neutral copy label in each four-qutrit block.
Site colour	**3**	Particle colour when occupied; background colour when empty.
Boundary colour	**3**	Completes the shared singlet and does not hop.
Two auxiliaries per link	**1**	Distinguish virtual channels in the auxiliary law.

Table 3: Registers in the junction construction. The auxiliaries are absent from the exact pure-hop identity.

Each site still transforms as a triplet in either occupation state. Four physical four-qutrit blocks and one boundary qutrit contain 17 fundamentals, with centre phase $\omega^{17} = \omega^2$, exactly the phase of the two particle colours. The construction respects Proposition 6.2; it relocates the neutral vacuum to a shared sector rather than finding a local $\mathbf{1} \oplus \mathbf{3}$ inside one fixed block.

7.2 A local hop and its known statistics mechanism

Let T_{vu} move a particle from u to v. The operator

$$T_{vu} = |01\rangle\langle 10|_{n_u n_v} \otimes S_{uv} \tag{49}$$

changes the two occupations and exchanges their colours. It acts on two encoded sites, with no operator string (no extra operators on the sites in between). This is locality at the block level; Eq. (49) has not thereby been derived as a single exchange of two elementary qutrits.

Reordering the empty registers after a hop changes the singlet by $(-1)^{N^{\text{empty}}_{(u,v)}}$, because the singlet is antisymmetric. Since the number of empty and occupied sites strictly between the endpoints sums to $|u-v|-1$,

$$\mathcal{V}^\dagger T_{vu} \mathcal{V} = (-1)^{|u-v|-1} t^{\mathrm{F}}_{vu}, \qquad (I - \mathcal{V}\mathcal{V}^\dagger) T_{vu} \mathcal{V} = 0. \tag{50}$$

The first identity says that, inside the code, the local hop equals the standard fermionic hop up to a fixed sign per link; the second says that the hop never takes the state out of the code. Here t^{F}_{vu} is the independently constructed fermionic hop, including relocation of the particle's colour and the occupation sign $(-1)^{N^{\text{occupied}}_{(u,v)}}$. The remaining factor depends only on the link. It can be removed by site phases on a tree; on graphs with cycles its products around loops must be retained. The identity was checked for several particle numbers with at most three empties, including spectator particles. These are not the coherent clock-sector spectators of Ref. [20].

For the three-hop histories, the particle colours acquire their exchange and the two empty colours contribute $S_{13} = -I$. Thus, without a ribbon,

$$\mathscr{L} = -\mathsf{S}_{\text{particles}} \mathscr{R} \tag{51}$$

on the full code. This is the fermionic exchange relation: the two histories differ by a minus sign once the particle colours are matched. A symmetric background gives the plus sign instead. On the four-site star with identical particle colours and unit hopping magnitude, the two-particle spectra are respectively

$$\{-\sqrt{3}, -\sqrt{3}, 0, 0, \sqrt{3}, \sqrt{3}\}, \qquad \{-2, -1, -1, 1, 1, 2\}. \tag{52}$$

The distinction is visible at a junction, whereas the tested open-chain spectra coincide.

The connection to prior work can be stated as an evaluated map. Regard empties as hard-core bosons B carrying colour (bosons that cannot share a site), and take identical particle colours. A particle hop $u \to v$ is then a background-boson hop $v \to u$:

$$T_{vu} = E^B_{uv} \quad \text{as matrix elements under this identification.} \tag{53}$$

All 72 directed comparisons on the four-site, two-particle sector agree, including the forbidden zero hops. This identifies the sign with the known colour–orbital correspondence[14]. With arbitrary particle colours the map additionally transports their colour; this is a short application, not a claim of a new general theorem. Local fermion representations themselves have longstanding precedents[7, 8, 10].

7.3 Preparation and scaling are physical costs

The local hop is not free: the sign comes from the prepared background, and that background has costs of its own. Equation (49) contains no string operator, but Eq. (47) requires

entanglement shared by the empties and the boundary, however far apart they become. Operator locality has not supplied local state preparation. Moradi Kalarde et al. explain this distinction in their operational comparison of fermions and bosons: suitable pre-shared entanglement can enable a bosonic simulation while changing the state-preparation resources[22]. Our finite construction illustrates that cost; it does not establish their general operational theorem.

Nor does the pure hop prepare the singlet by energy minimization; that is, the desired background is not simply the lowest-energy state. With two identical-colour particles and two mobile empty registers plus one boundary qutrit, the lowest energies in the background colour sectors are:

Graph	Hop coefficient	Singlet	Mixed	Symmetric
Star	± 1	-1.7321	-2.0000	-2.0000
Star plus link 1–2	$+1$	-2.4812	-2.4812	-2.0000
Star plus link 1–2	-1	-2.4812	-2.7321	-2.7321

Table 4: Pure-hop reference energies in units of the hopping magnitude, resolved by the total colour of the empties and boundary. These are not ground-state energies of the tuned auxiliary Hamiltonian.

For identical particle colours the singlet-background sector is conserved but excited on the star. Allowing all particle colours gives a state at -2 inside that sector, degenerate with other background sectors. Thus the hopping term does not uniquely prepare the desired background. In the identical-colour test, the added frustrated link only produces a tie for lowest energy. Singlet-cluster selection on frustrated bosonic lattices[14] and hole-induced singlets in particular $U = \infty$ fermionic Hubbard models[23] are leads for a preparation mechanism, not evidence that one exists here.

For $m \leq 3$ empties the singlet construction adds $3-m$ pinned boundary qutrits. For $m > 3$, the required totally antisymmetric background is absent. The construction can accommodate more particles but cannot describe a dilute network with arbitrarily many empties. One may instead use a symmetric coloured background and a separate ribbon sign, but that changes the vacuum and statistics resources. Multiplying the independently established background and ribbon signs predicts a minus sign for symmetric background plus ribbon and a plus sign for antisymmetric background plus ribbon; that composition has not been simulated in this junction study.

7.4 Separate colour symmetries explain the necessary tuning

Before building the full Hamiltonian, it helps to know which unwanted terms actually threaten the singlet background. A symmetry argument answers this. Let $E_i^{ab} = |a\rangle\langle b|$ act on the colour at site i. There are separate generators for particle colour and background colour:

$$Q_A^{ab} = \sum_i n_i E_i^{ab}, \tag{54}$$

$$Q_B^{ab} = \sum_i (1-n_i) E_i^{ab} + \sum_{r\in\partial} E_r^{ab}. \tag{55}$$

The traceless combinations generate two independent SU(3) rotations: one rotates only the particle colours, the other only the background colours. A hop moves each species together with its colour, so $[T_{vu}, Q_A^{ab}] = [T_{vu}, Q_B^{ab}] = 0$. In particular, the background singlet sector is invariant. This is an operator identity, not a condition on one initial colour assignment.

Write the occupation projectors on a link as

$$P_= = n_u n_v + (1-n_u)(1-n_v), \qquad P_{\neq} = n_u(1-n_v) + (1-n_u)n_v. \tag{56}$$

The occupation-diagonal colour swap $P_{=}S_{uv}$ exchanges two colours within one species and preserves both rotations. In contrast, $P_{\neq}S_{uv}$ exchanges a particle colour with a background colour without moving occupation. It preserves total colour but breaks the separate symmetries and can leave the singlet code. An exhaustive two-site matrix check confirms these identities.

Consequently, *among these number-conserving diagonal colour terms*, only the unlike-occupation term needs cancellation for code invariance. This does not remove pair creation, guarantee the desired orbital energies, or protect the cancellation against higher-order returns. The numerical auxiliary study below cancelled all diagonal colour coefficients; the symmetry analysis suggests a less restrictive future design, which has not been tested dynamically.

7.5 The auxiliary Hamiltonian

Section 4.3 showed that virtual excursions returning through a different link spoil the junction. The remedy studied here is to give each link its own "flag": two extra two-state systems (auxiliaries), one for each hopping channel, which are flipped whenever that channel's coupling acts on that link. A process that starts on one link and returns on another would leave a flag raised, so at second order it cannot return to the low-energy space. Two neutral two-state auxiliaries on every star link distinguish its virtual channels. Let $n_{e,\pm}$ be their excitation projectors and $X_{e,\pm}$ their flips. In units of the fixed energy scale G, the microscopic law to be approximated is

$$\begin{aligned} H_\star/G = &\sum_{i=0}^{3} L_i + \sum_e (n_{e,+} + n_{e,-}) \\ &+ \lambda \sum_e \left(\sqrt{w}\, V_{e,+} X_{e,+} + V_{e,-} X_{e,-}\right) + \lambda^2 H_{\mathrm{ct}}, \end{aligned} \tag{57}$$

where $g_0 = -4/15$ is included in the dimensionless port operators $V_{e,\pm}$ of Section 3. The auxiliary excitation energy is G; GL_i is the block selector. A second-order return must restore the auxiliary it excited. This eliminates second-order returns through a different link without discarding any allowed term.

The remaining controls are physical pair couplings and internal swaps:

$$H_{\mathrm{ct}} = g_0 \sum_{e=(0,v)} \sum_{p,q\in\{r,u\}} c_{pq} B^e_{pq} + h_0 S^{(0)}_{01} + h \sum_{v=1}^{3} S^{(v)}_{01}. \tag{58}$$

("ct" stands for counterterm: an extra term added to cancel unwanted effects.) The first term adjusts colour-dependent returns; the fields h_0 and h detune pair creation, meaning that they raise the energy cost of creating a pair so that it happens less. They do not remove that operator. Each λ vertex in Eq. (57) is a pair swap multiplied by an auxiliary flip, hence a supplied three-register interaction. No pair-only implementation of these controlled vertices is asserted.

At second order the two isolated channels have coefficients

$$(a_+, b_+) = \left(\frac{218}{1575}, -\frac{66}{175}\right), \qquad (a_-, b_-) = \left(-\frac{3062}{37125}, -\frac{452}{37125}\right). \tag{59}$$

Thus the perturbative choice $w_0 = -a_-/a_+ = 10717/17985$ gives

$$a_{\mathrm{d}} = 0, \qquad b_{\mathrm{d}} = -\frac{958682}{4046625}. \tag{60}$$

The weight w is chosen so that the unwanted identity parts of the two channels cancel, leaving a pure colour-carrying hop. At that order the diagonal colour terms cancel with $h_0 = h$ and

$$(c_{rr}, c_{ru}, c_{ur}, c_{uu}) \simeq (0.0028172, 0, 0, 0.590524). \tag{61}$$

These are four independent nonzero settings, w_0, c_{rr}, c_{uu}, h. Their cancellation is perturbative; using them unchanged at finite λ leaves residual unwanted operators.

One protection does hold exactly:

$$\widetilde{\Pi} = \left(\prod_{i=0}^{3} S_{01}^{(i)}\right)\left(\prod_{e} Z_{e,-}\right), \qquad [H_\star, \widetilde{\Pi}] = 0. \tag{62}$$

Return to the auxiliary ground state preserves matter parity, including virtual excursions outside the code. This does not conserve particle number or protect the hopping coefficients against finite-coupling corrections.

7.6 Calibrating existing controls at finite coupling

The settings above cancel unwanted terms only at second order. At the finite coupling actually used, higher-order terms shift the balance. We therefore re-tune the Hamiltonian's own adjustable parameters so that the exact single-link behaviour has the desired form. The finite-coupling construction keeps Eq. (57) and solves for its existing controls. It does not delete effective operators by hand. The two operating points $\lambda = 0.02$, $h = 40$ and 80 were fixed before the new calculation. At each point six link equations determine

$$x = (w, c_{rr}, c_{ru}, c_{ur}, c_{uu}, h_0). \tag{63}$$

The equations use link spectra and operators only; neither masses, mixing angles nor star-evolution fidelities enter their solution. In other words, nothing is fitted to the final performance score or to observed particle data.

To define the link operator unambiguously, let E_C embed the bare block code with all auxiliaries in their ground states on a cluster C. Let U_C contain the orthonormal code-connected eigenvectors of its physical Hamiltonian, with dimensionless eigenvalues ε_j in units G. For $A_C = E_C^\dagger U_C$, let Q_C be its unitary polar factor (the unitary matrix closest to A_C). The effective operator in bare code coordinates is

$$K_C = \lambda^{-2} Q_C \, \mathrm{diag}(\varepsilon_j) \, Q_C^\dagger. \tag{64}$$

This is an exact effective Hamiltonian for the cluster: it has the true low-lying energies, expressed in the basis of the original ("bare") code states. The calculations select the eigenvectors with the largest code overlaps among the computed low states and check residuals and overlap singular values. For the link the explicit weight-sector solve has $560 \times 4 = 2240$ states. SU(3) invariance reconstructs the full 36-dimensional two-site code operator from this sector.

Write its occupation blocks as $\langle nm|K_{01}|kl\rangle = a_{nm,kl} I + b_{nm,kl} S$ and set $A_{nm} = a_{nm,nm}$. The six conditions are

$$\begin{aligned} a_{01,10} &= 0, \qquad & b_{nm,nm} &= 0 \quad (n, m = 0, 1), \\ A_{11} + 2A_{10} - 2A_{01} - A_{00} + 4h_0 &= 0. \end{aligned} \tag{65}$$

The first removes hopping that leaves colour behind; the next four cancel diagonal colour swaps. The last equalizes the scalar energies of the centre-occupied and centre-empty two-particle configurations after the duplicated centre fields are removed. These six scalar constraints do not impose number conservation.

The maximum equation residual is below 7×10^{-12}. A separate one-link implementation reproduces the parameters to 7.8×10^{-12} and the effective link matrices to 2.5×10^{-11} in maximum-entry error. This is a reproducible finite-coupling calibration, not a parameter-free law or a symmetry-protected cancellation.

Control	$h = 40$	$h = 80$
w	0.604383902703	0.612853584455
c_{rr}	0.002301827373	0.001838802012
c_{ru}	0.002281295100	0.004622001684
c_{ur}	0.002281099731	0.004621598857
c_{uu}	0.597826268170	0.605397602935
h_0	40.003762716560	80.007590524561

Table 5: Finite-link calibration at $\lambda = 0.02$. All six values enter the physical Hamiltonian before diagonalization. Rounded values are printed; machine-precision inputs are archived.

7.7 The effective star and its tested time window

With the controls fixed, we assemble an effective Hamiltonian for the whole four-site star from the single-link and two-link results, and compare its time evolution with ideal fermion hopping. Every cluster contains the full field h_0 on its centre and h on each of its leaves. The connected two-link correction and assembled star are

$$\begin{aligned} O_{0uv} &= K_{0uv} - K_{0u} - K_{0v} + h_0 Z_0, \\ K_\star^{(\le 2)} &= \sum_{v=1}^{3} K_{0v} - 2h_0 Z_0 + \sum_{1 \le u < v \le 3} O_{0uv}, \\ Z_0 &= \mathrm{diag}(1,1,1,-1,-1,-1)_0. \end{aligned} \tag{66}$$

Here O_{0uv} is the part of the two-link result that is not already explained by the two single links. The subtractions count each onsite field once. At each calibrated point the two-link physical problem was recomputed on 554,400 states, including all its returns. Its eigenpair residual is below 1.1×10^{-13} and its reconstruction error on the code below 3.3×10^{-10}. The independent calibration/dynamics verification imports these saved two-link matrices; it does not independently repeat that large physical solve.

Equation (66) retains pair creation and colour-dependent returns. It omits connected three-link corrections and microscopic bare-state dressing (the small admixture of excited states in the true low-energy states). Evolution uses scaled time $\tau = G\lambda^2 t$ and the finite-link colour-carrying coefficient $b = b_{01,10}$, giving

$$\tau_{\mathrm{hop}} = \frac{\pi}{2|b|}, \qquad b = -0.239242543801 \ (h = 40), \quad -0.241588068343 \ (h = 80). \tag{67}$$

One "hop time" τ_{hop} is the time an ideal hop takes to move a particle fully from one site to the next. For a normalized initial state in the 54-state singlet-background code, the comparison measures

$$F(\tau) = |\langle \psi_{\mathrm{ideal}}(\tau) | \psi_{\mathrm{eff}}(\tau) \rangle|^2, \qquad \mathcal{L}(\tau) = 1 - \|\mathcal{V}^\dagger \psi_{\mathrm{eff}}(\tau)\|^2, \tag{68}$$

where $\mathcal{V}$ is the encoding in Eq. (47) and the ideal dynamics uses its colour-carrying fermion hops. The fidelity F is the probability that the simulated state matches the ideal one; the leakage $\mathcal{L}$ is the probability of having left the code. The boundary colour is included in both evolutions.

The fixed test uses all 54 code basis states and 48 new complex superpositions, eight in each of six conserved colour-count sectors, with seed 2026110402. Only the superpositions are new; the basis states had been tested before. A time grid resolves the fastest spectral phase increment to at most 0.125 radians. The criterion is sampled $F \ge 0.99$ and $\mathcal{L} \le 0.01$ throughout the stated window.

h	Hop times	Minimum F	Maximum $\mathcal{L}$	Criterion
40	4	0.998663453	0.059090%	Pass
40	8	0.994821818	0.220187%	Pass
40	16	0.979588079	0.855914%	Fail
80	4	0.999673836	0.014379%	Pass
80	8	0.998733156	0.053057%	Pass
80	16	0.994990366	0.206802%	Pass

Table 6: Calibrated effective-star results on 102 starts. Extrema are over the sampled inputs and times, not certified bounds for all states or continuous time. A separate implementation reproduces these metrics within 6×10^{-11}.

The improvement is substantial but has a finite scope. An independent prospective test (one whose pass criterion was fixed in advance) of the *uncalibrated* second-order settings at $h = 40$, $\lambda = 0.02$ gave $F_{\min} = 0.960767$ and $\mathcal{L}_{\max} = 1.06101\%$ over eight hops on 78 starts. That failure stands; Table 6 tests different physical coefficients, fixed by Eq. (65), and a new random seed. The earlier four-hop comparisons remain valid within their shorter window. By itself, the successful calibrated $h = 80$ point establishes sampled effective dynamics. Section 7.10 supplies the separate bound connecting it to microscopic fidelity.

7.8 Why virtual pairs still limit transport

This subsection identifies the main source of the remaining error. Even after calibration the link creates and annihilates pairs through

$$V_{\text{pair}} = a_p I + b_p S, \qquad (a_p, b_p) \simeq \begin{cases} (0.087825, -0.383950), & h = 40, \\ (0.086409, -0.383568), & h = 80. \end{cases} \tag{69}$$

Large detuning makes an actual change of particle number unlikely, but a short virtual visit to the zero- or four-particle sector changes the two-particle evolution. Eliminating those sectors at an energy E gives the energy-dependent correction

$$\Sigma_2(E) = -K_{2Q}(K_{QQ} - EI)^{-1}K_{Q2}, \qquad Q = \{N = 0, 4\}. \tag{70}$$

This is the Schur complement; replacing it by a constant evaluated at the mean code energy is a diagnostic approximation, not an exact propagator.

The leading signs have a simple explanation. Let $\Delta_0 > 0$ be the gap for a centre-occupied pair to visit $N = 0$, and $\Delta_4 < 0$ the corresponding gap for a leaf-occupied pair to visit $N = 4$. Neglecting smaller scalar and colour terms in these denominators gives

$$\begin{aligned} \delta K_{\{0,v\}} &\simeq -\frac{(a_p I + b_p S_{0v})^2}{\Delta_0} = -\frac{a_p^2 + b_p^2}{\Delta_0} I - \frac{2a_p b_p}{\Delta_0} S_{0v}, \\ \delta K_{\{u,v\}} &\simeq -\frac{(a_p - b_p)^2}{\Delta_4} I. \end{aligned} \tag{71}$$

The second expression uses the antisymmetric empty pair, on which $S = -1$. For the onsite fields alone the gaps are $\pm 2(h_0 + h)$; at leading order $h_0 \simeq h$ they are $\pm 4h$. The calibrated scalar returns modify the actual denominators. Thus these are leading large-detuning formulas, not exact finite-coupling identities with gaps $4h_0$ and $-4h$.

Since $a_p b_p < 0$, the $N = 0$ return adds a positive colour-swap term as well as a negative scalar shift. The $N = 4$ return shifts the leaf-occupied configurations upwards. Returns through

different links also connect different two-particle configurations and different background states; diagonal counterterms cannot remove all of them. The error amplitude therefore has a natural $1/h$ scale and can accumulate coherently with time. It is not a constant leakage rate.

At the mean code energy, the projected Schur-correction norm divided by $|b|$ is 0.00969 for $h = 40$ and 0.00476 for $h = 80$. Doubling h reduces the sampled sixteen-hop infidelity by 4.08 and leakage by 4.14, consistent with a leading amplitude of order $1/h$ over this window (an amplitude error of order $1/h$ gives probability errors of order $1/h^2$, hence the factor of about four). Hypothetically removing all number-changing terms raises both sixteen-hop fidelities above 0.99999; that deletion isolates the mechanism but is not an allowed physical repair.

The calibrated two-link correction has norm 6.17×10^{-5} and 6.40×10^{-5} in scaled energy units at the two points. Removing it changes the sixteen-hop minimum fidelity by only 1.49×10^{-4} and 7.51×10^{-5}, respectively. It is retained in Table 6. These small changes identify its contribution in the chosen full-field convention; they do not bound omitted three-link effects.

7.9 Bounding the dressing of the full star

The simulations above use the effective model, not the full microscopic Hamiltonian. The remaining subsections ask how far the two can differ, and derive bounds on the difference; the bounds are evaluated in floating point rather than certified with interval arithmetic. The independently reproduced dynamics concerns $K_\star^{(\le 2)}$. Transferring its accuracy to the full physical star requires two bounds: one on the dressing of the complete star code, and one on the difference between its exact band Hamiltonian and Eq. (66). Here the *dressing* is the extent to which the true low-energy states differ from the bare code states, and the *band* is the set of true low-energy states that correspond to the code. We first bound the dressing directly, then use a residual calculation to control both errors together.

Throughout these estimates, matrix norms are spectral norms (the largest factor by which a matrix can stretch a vector); a superoperator norm is induced by that matrix norm. A useful conditional statement makes the distinction precise. Let E be the bare star-code isometry and W its polar-aligned isometry into an exact invariant spectral band, so that $D = E^\dagger W$ is positive. Set $s = \lambda_{\min}(D)$ and let K_{exact} be that band's Hamiltonian in the same scaled units. If $\delta = \|K_{\text{exact}} - K_\star^{(\le 2)}\|_2$, then

$$\left\| E^\dagger e^{-iH_\star t} E - e^{-iK_\star^{(\le 2)}\tau} \right\|_2 \le \epsilon_{\text{dress}} + |\tau|\delta, \qquad \epsilon_{\text{dress}} = 2(1-s) + (1-s^2). \tag{72}$$

In words, the total error is a fixed amount due to dressing, ϵ_{dress}, plus an amount that grows linearly in time due to the difference δ between the exact and effective Hamiltonians. Here $\tau = G\lambda^2 t$. To see this, split $E = WD + E_\perp$. The exact compressed evolution is $De^{-iK_{\text{exact}}\tau}D + E_\perp^\dagger e^{-iH_\star t}E_\perp$; the two errors are bounded by $2\|I - D\|$ and $\|E_\perp\|^2 = 1 - s^2$. Duhamel's identity adds the term $|\tau|\delta$. For any normalized input, a probability or fidelity difference is at most twice the resulting amplitude bound.

The measured one- and two-link overlaps, approximately 0.999942 and 0.999894, do not determine s for three links, and summing their overlap defects is an estimate, not a lower bound. A lower bound follows from the equation for an invariant subspace. Write $H_\star/G = H_D + B$, where H_D contains the selectors, auxiliary energies and onsite fields, and B contains the auxiliary vertices and pair counterterms. Let P project on the bare star code and $Q = 1 - P$. The selector and auxiliary energies vanish on P and are at least one on Q; the fields commute with them and have norm at most $f = \lambda^2(h_0 + 3h)$. Hence $S(X) = H_{D,Q}X - XH_{D,P}$ satisfies $\|S^{-1}\|_2 \le 1/\sigma$ with $\sigma = 1 - 2f > 0$. The idea is to describe the true low-energy subspace as a

tilted copy of the code space, $\{(\phi, X\phi)\}$, where the operator X gives the small excited-state admixture. The subspace $\{(\phi, X\phi)\}$ is invariant exactly when

$$S(X) = -B_{QP} - B_{QQ}X + XB_{PP} + XB_{PQ}X. \tag{73}$$

Set $X_0 = -S^{-1}B_{QP}$, $\beta = (\|B_{QQ}\|_2 + \|B_{PP}\|_2)/\sigma$ and $\gamma = \|B_{PQ}\|_2/\sigma$. If $\beta < 1$ and $(1-\beta)^2 > 4\gamma\|X_0\|_2$, Eq. (73) defines a contraction on the ball of radius

$$x = \frac{1 - \beta - \sqrt{(1-\beta)^2 - 4\gamma\|X_0\|_2}}{2\gamma}, \tag{74}$$

so it has a solution with $\|X\|_2 \le x$. When $\gamma = 0$, use the continuous limit $x = \|X_0\|_2/(1-\beta)$. The contraction constant is at most $\beta + 2\gamma x < 1$. Upper bounds may replace the norms throughout. The normalized isometry onto this subspace has $E^\dagger W = (1 + X^\dagger X)^{-1/2}$, which is already positive, so $1 - s^2 \le x^2/(1+x^2)$. To identify it as the lowest band, one also checks $p + \|B_{PQ}\|_2 x < 1 - \|H_\star/G - H_{\text{sel+aux}}\|_2$, where p bounds $\|PH_\star P/G\|_2$ and $H_{\text{sel+aux}}$ is the sum of the selectors and auxiliary penalties. The left side bounds the graph-subspace energies; Weyl's inequality bounds the next eigenvalue from below by the right side. This inequality holds at both operating points.

The vertex contribution to $\|X_0\|_2$ can be evaluated without a star-sized diagonalization; the counterterm contribution is bounded separately. A vertex on one link leaves the other blocks in the code, so their fields cancel in S and each denominator is the link's own. Vertices exciting different auxiliaries give orthogonal vectors. Thus $\|X_0\|_2^2$ is the norm of a sum of link matrices embedded in the 6^4-dimensional code, plus a small counterterm contribution bounded separately. The other norms use the triangle inequality over links. Table 7 gives the results.

Quantity	$h = 40$	$h = 80$
σ	0.8720	0.7440
β	0.3514	0.4131
$\|X_0\|_2$ at most	0.01644	0.01651
x	0.02544	0.02829
Three links: $1 - s^2$ at most	6.47×10^{-4}	8.00×10^{-4}
Three links: s at least	0.999676	0.999600
Three links: ϵ_{dress} at most	1.294×10^{-3}	1.600×10^{-3}
Three links: leading-order loss, not a bound	2.70×10^{-4}	2.73×10^{-4}
One link: exact loss, bound (10^{-4})	1.155, 1.459	1.159, 1.491
Two links: exact loss, bound (10^{-4})	2.108, 3.533	2.125, 3.833

Table 7: Bound on the full-star overlap, for the calibrated controls of Table 5. The last two rows compare the same construction with exactly solved smaller clusters. Values are in floating point, not interval arithmetic.

This leaves a definite requirement on δ. The evolutions $e^{-iK\tau}$ are unitary, so angles between states add. Let $a_0 = \arccos\sqrt{F}$ and $l_0 = \arcsin\sqrt{\mathcal{L}}$ be the worst sampled values of Eq. (68) in a window, and let $\eta = 2\arcsin(|\tau|\delta/2)$. With bare-state preparation and readout,

$$F_{\text{micro}} \ge [\cos(a_0 + \eta) - \epsilon_{\text{dress}}]_+^2\,, \qquad \mathcal{L}_{\text{micro}} \le 1 - [\cos(l_0 + \eta) - \epsilon_{\text{dress}}]_+^2\,, \tag{75}$$

where $[z]_+ = \max(z, 0)$. These nontrivial bounds require $|\tau|\delta \le 2$ and $a_0 + \eta, l_0 + \eta \le \pi/2$; otherwise use the trivial bounds zero and one. All reported cases satisfy these conditions. The conclusions apply to the sampled inputs and times. At $h = 80$ over sixteen hops, $\delta = 0$ would give $F_{\text{micro}} \ge 0.9918$ and $\mathcal{L}_{\text{micro}} \le 0.53\%$, and the one-percent criterion requires $\delta \le 1.13 \times 10^{-4}$.

Over eight hops the requirement is $\delta \le 9.0 \times 10^{-4}$ at $h = 80$ and $\delta \le 2.7 \times 10^{-4}$ at $h = 40$. For scale, the stored sum of connected two-link terms has norm 1.08×10^{-4} at $h = 80$.

Auxiliary-flip counting puts a connected three-link term first at physical order λ^6 at fixed controls, or λ^4 after dividing by λ^2: an auxiliary must be restored, so each link enters through two vertices or one counterterm. This does not bound its coefficient or the higher-order remainder. In particular, a small net two-link correction can contain cancellations and does not itself bound the next cluster.

An all-time leakage bound of $2(1 - s^2)$ would also be too small in general: a bare vector with overlaps s and $\sqrt{1 - s^2}$ on two eigenstates acquires leakage $4s^2(1 - s^2)$ when their phases differ by π. Equation (72) avoids that error. The residual method below supplies the missing control of the band dynamics, in a slightly different choice of basis.

7.10 Controlling the omitted returns with a residual

A small effective Hamiltonian is useful only if the states it describes remain close to an invariant subspace of the full Hamiltonian. We can test that directly without diagonalizing the full star. The approach resembles checking an approximate eigenvector: apply the full Hamiltonian to a trial set of states and measure how much of the result falls outside that set. If this "residual" is small compared with the energy gap to other states, the trial states must be close to true ones. Construct a trial isometry T from the bare code plus its first- and second-order virtual excursions, normalize its columns, and align it so that $E^\dagger T$ is positive. No calibrated coefficient is changed. In this subsection energies are initially measured in units of G: write $\bar{H} = H_\star/G$ and define

$$\widetilde{K} = T^\dagger \bar{H} T, \qquad R = \bar{H}T - T\widetilde{K}, \qquad r = \|R\|_2. \tag{76}$$

The residual R is the part of the exact Hamiltonian action that leaves the trial subspace. It is evaluated with all terms of $H_\star$, including those absent from the trial itself. Thus this is an a-posteriori error estimate (one computed from the actual approximate solution), not an assumption that the perturbation series has terminated.

Let e_c be a lower bound on the first eigenvalue above the desired band, e_+ the largest eigenvalue of $\widetilde{K}$, and κ half its spectral width. Set $\gamma = e_c - e_+ > 0$, $\nu = r/\gamma < 1$ and $\widetilde{s} = \sigma_{\min}(E^\dagger T)$. Appendix E proves that the exact band, aligned to T, has a Hamiltonian K' satisfying

$$\begin{aligned} &\|K' - \widetilde{K}\|_2 \le \Delta, && \Delta = \frac{3\nu(\nu\kappa + r)}{1 - \nu^2}, \\ &s \ge s_* := \cos\!\big(\arccos \widetilde{s} + \arcsin \nu\big), && g_* = 1 - \widetilde{s} + \nu^2 + \sqrt{1 - \widetilde{s}^2}\,\nu. \end{aligned} \tag{77}$$

The angle sum is less than $\pi/2$ here. The number g_* bounds the mismatch between the bare code and this exact-band basis. Keeping this term is essential: a good spectral approximation alone need not reproduce interference in the bare-state readout.

The usable amplitude allowances are

$$\epsilon' = 2g_* + 1 - s_*^2, \qquad \delta' = \left\| \frac{\widetilde{K}}{\lambda^2} - K_\star^{(\le 2)} \right\|_2 + \frac{\Delta}{\lambda^2}. \tag{78}$$

Then Eq. (72) holds with ϵ', δ' in place of $\epsilon_{\text{dress}}, \delta$. This controls the entire omitted-band contribution; it does not separately identify a three-link coefficient. In particular, the first norm in Eq. (78) also contains the error in the trial isometry.

The full-star calculation gives $r = 1.957 \times 10^{-5}$ and 1.982×10^{-5} at $h = 40, 80$. The respective separating gaps are 0.5685 and 0.4396. The resulting δ' values are 6.070×10^{-6}

and 9.499×10^{-6} in scaled energy units, while ϵ' is at most 5.426×10^{-4} and 5.483×10^{-4}. Substitution into Eq. (75), with its stated domain restrictions, gives Table 8.

h	Hop window	F_{micro} at least	$\mathcal{L}_{\text{micro}}$ at most	δ'
40	8	0.99369	0.332%	6.070×10^{-6}
80	8	0.99760	0.165%	9.499×10^{-6}
80	16	0.99375	0.326%	9.499×10^{-6}

Table 8: Numerically evaluated microscopic error bounds for bare-state preparation and readout, using the same 102 inputs and sampled times as Table 6. Fidelity is rounded down and leakage up. These are bounds obtained from residuals, not a direct simulation of the complete microscopic time evolution.

So, for the stated sampled test and subject to the numerical qualification below, the bounds place the full microscopic model within the one-percent criterion. The theorem controls operator norms on the whole retained code. The reported fidelity and leakage bounds nevertheless inherit the *sampled* input and time scope of the effective-law comparison. They are not a worst-case fidelity calculation over all code states or a continuous-time window. The residuals, norms and eigensolvers use floating-point arithmetic, with numerical checks and margins rather than interval certification. Both star calculations have been replayed from preserved inputs, and the residual estimate has been independently derived and tested on complex Hermitian models. These distinctions are detailed in Appendix F.

This closes the omitted-return question for the stated finite-model test, subject to the numerical qualification. It does not establish robustness. The earlier one-parameter finite differences held the two-link matrix fixed. Even in that approximation, recomputing the one-link law at the proposed tolerance boundaries produces failures; a fully recomputed parameter-perturbed star has not been tested. They should therefore be read as sensitivity estimates, not safe operating intervals. The supplied three-register vertices, exact calibrated controls and prepared background remain assumptions. Neither scalable preparation nor integration with the earlier clock and coherent spectator blocks is supplied by this calculation.

8 Conclusions

Pair exchange on the selected four-qutrit block generates two symmetry-selected virtual hopping channels with fixed colour ratios. A real mixture can cancel the unwanted hopping coefficient on one isolated link. The complete cancellation family still has large unavoidable shared-link returns, so this mechanism alone does not supply a junction law.

A prepared antisymmetric background does support local colour-carrying fermion hops on the small junction. This applies the established colour–orbital symmetry correspondence. It requires shared entanglement and at most three empties; the hopping Hamiltonian does not uniquely select the background. Separate particle and background colour rotations explain which number-conserving diagonal colour term threatens the code.

Neutral auxiliaries eliminate the second-order shared-link returns and preserve a combined parity. Calibrating six existing controls against finite-link operator equations then improves the dynamics: at $h = 80$, $\lambda = 0.02$, the independently reproduced effective law passes the stated one-percent test through sixteen hops on 102 inputs. The $h = 40$ point fails at sixteen hops. No observed masses, mixing angles or star-dynamics scores were used to choose the six calibrated coefficients.

The remaining error is consistent with virtual pair creation, whose leading return amplitude decreases with detuning. This is useful finite-model progress, not an all-order protection

mechanism. A full-star residual calculation now controls the dressing and omitted returns: at $h = 80$ through the sampled sixteen-hop experiment it gives microscopic fidelity at least 0.99375 and leakage at most 0.326%, with floating-point rather than interval certification. This is a finite-model accuracy result, not parameter-error tolerance or an all-input guarantee. Deriving the three-register controlled vertices, preparing a scalable background, and establishing robustness are the next requirements. Integration with the earlier clock/spectator model also remains open. No result here establishes relativistic fermions, local SU(3) gauge dynamics or observed particle masses.

Data, code and authorship

The supporting code, saved numerical inputs, audit scripts and review records are deposited in a versioned Zenodo archive[24]. Its general DOI identifies the support series; the reference also pins version 1.0.0 used for this paper. Appendix F describes the evidence and its limits. The archive provides a fresh-copy runner for the principal matrix, error-bound and finite-calibration checks. Some historical exploratory scripts still require local inputs; the release documentation identifies that limitation. The earlier clock–colour archive[20] is a separate work and does not certify these new results. The cancellation algebra and sampled calibrated dynamics have received separate computational checks, but human scientific review remains pending. The supplied independent calibration check shares the saved two-link matrices, as stated above.

AI tools were used to assist with drafting, code construction, algebraic checking and editorial revision. Their outputs are not independent authorship or peer review. The author is responsible for reviewing and approving the final scientific claims and manuscript.

A Spectrum and boundary algebra of the block

This appendix derives the selector spectrum quoted in Eq. (12) and the compression formulas of Section 3. The class sum Σ_4 is central in the permutation algebra, meaning that it commutes with every permutation, so it takes a single value on each symmetry type of state. These symmetry types are labelled by Young diagrams (arrangements of four boxes in left-justified rows). Its eigenvalue on a Young diagram λ is the sum of box contents, $\sum_{(i,j)\in\lambda}(j-i)$, where box (i,j) sits in row i and column j. For $(4),(3,1),(2,2),(2,1,1)$ these values are respectively $6,2,0,-2$. Since

$$M = 4I + \Sigma_4 + S_{01} + S_{23}, \qquad L = 3I + \frac{3}{2}\Sigma_4 - \frac{1}{2}(S_{01} + S_{23}), \tag{79}$$

the disjoint pair swaps give a simultaneous diagonalization: Σ_4, S_{01} and S_{23} all commute, so M and L can be read off from their joint eigenvalues. Table 9 supplies the complete 81-state accounting.

Equation (3) specifies every code coefficient, so the boundary law in Eq. (13) can be obtained by ordinary contraction of Levi-Civita tensors. Its form also follows from covariance (the requirement that it transform correctly under colour rotations): an operator transforming as a colour matrix unit has a traceless colour part and a scalar part, with matrices on the neutral multiplicity factor. Normalization fixes their combination to the stated form once A_i is known. Summing over $a = b$ gives I_6 at each physical site, as required.

To derive Eq. (14), substitute Eq. (13) twice into $\sum_{a,b} e^{(xi)}_{ab} e^{(yj)}_{ba}$. The matrix-unit product contributes $A_i \otimes A_j \otimes S_c$. The remaining scalar terms combine to $(I_4 - A_i \otimes A_j) \otimes I_9/3$. No omitted colour operator is allowed: $\mathbf{3} \otimes \mathbf{3}$ has two inequivalent irreducible summands, each

Young diagram	Colour representation	Σ_4	M	L	Multiplicity
$(2,1,1)$	$\mathbf{3}$	-2	0	1	3
$(2,1,1)$	$\mathbf{3}$	-2	2	0	6
$(2,2)$	$\bar{\mathbf{6}}$	0	2	4	6
$(2,2)$	$\bar{\mathbf{6}}$	0	6	2	6
$(3,1)$	$R_{(2,1)}$	2	6	6	30
$(3,1)$	$R_{(2,1)}$	2	8	5	15
(4)	$R_{(4,0)}$	6	12	11	15

Table 9: Joint spectrum, derived from the permutation representation and checked independently with physical-qutrit matrices. The unwanted old $M = 2$ states are the six states in the third row.

with multiplicity one, so the identity and the swap are the only colour-invariant operators on two triplets.

For the first-order bound of Eq. (18), internal swaps contribute only operators on one occupation factor, tensored with colour identity. Their trace with $\tau_y^x \tau_y^y \otimes S_c$ vanishes. Cross swaps contribute products of I and τ_z by Eq. (14), whose trace with τ_y also vanishes. Finally, $\mathrm{Tr}(S_c^\dagger S_c) = 9$ and $\mathrm{Tr}(\tau_y^2) = 2$, giving $\|J\|_F^2 = 36$ and $\mathrm{Tr}(J^\dagger T) = 18$. This proves Eq. (18) without relying on the numerical span rank.

B Exact second-order cancellation certificate

This appendix shows exactly why the coefficient matrix of Eq. (28) removes every unwanted single-flip process, and then finds all real coupling matrices that do the same. The single-flip polynomial supplies both a certificate for Eq. (28) and a classification of the real cancellation family. A *certificate* here is a short exact calculation that anyone can check by hand or with computer algebra. Coefficients retain the normalization $g_0 = -4/15$.

Consider the single-flip blocks on block x, and abbreviate

$$A = c_{++}c_{-+}, \quad B = c_{+-}c_{--}, \quad U = c_{+r}c_{-r}, \quad W = c_{+u}c_{-u}, \quad V = c_{+r}c_{-u} + c_{+u}c_{-r}. \tag{80}$$

The two independent colour coefficients for each value of the unchanged occupation on y vanish precisely when the following four numerators vanish:

$$\begin{aligned} 68453A + 141526B + 321264U + 166320V + 110880W &= 0, \\ 2079A - 4158B + 7624U + 3696V + 2464W &= 0, \\ 141526A + 68453B + 321264U - 166320V + 110880W &= 0, \\ -4158A + 2079B + 7624U - 3696V + 2464W &= 0. \end{aligned} \tag{81}$$

The denominators, 2494800 and 92400, are nonzero and irrelevant to these zero conditions. The corresponding equations for a flip on y follow by transposing c. This covers all eight directed single-flip blocks after Hermitian conjugation.

For Eq. (28),

$$A = B = \rho, \qquad U = \alpha\rho, \qquad W = -\beta\rho, \qquad V = 0. \tag{82}$$

The differences of the first and third equations and of the second and fourth vanish immediately. The sums reduce to

$$419958 + 642528\alpha - 221760\beta = 0, \qquad -4158 + 15248\alpha - 4928\beta = 0, \tag{83}$$

which hold for the rational values in Eq. (27). Symmetry of c proves the same result for y. Equation (26), with $c_{++}c_{--} = 1$ and $c_{+-}c_{-+} = \rho^2 = 56560/39771$, then gives Eq. (29) exactly.

Every nonzero term in this coefficient matrix has at least one e_+ or e_- port. Its first-order projection is therefore zero by Eq. (13). The coefficient matrix contains both parity characters, so this cancellation is not an exact symmetry statement about the full Hamiltonian. Setting c_{-u} to zero breaks the certificate and gives nonzero single-flip entries, providing a direct control.

The complete real first-order-silent family

The same polynomial system describes every real cancellation with a nonzero pure-swap hop. Use zero-based port indices $(0, 1, 2, 3) = (e_+, e_-, r, u)$ and write $A_0 = c_{00}$, $B_0 = c_{01}$, $C_0 = c_{10}$ and $D_0 = c_{11}$. First-order vanishing requires $c_{22} = c_{23} = c_{32} = c_{33} = 0$. Eliminating the single-flip polynomials and their transposes gives

$$\begin{aligned} A_0C_0 &= B_0D_0, & A_0B_0 &= C_0D_0, \\ c_{02}c_{12} &= \alpha A_0C_0, & c_{03}c_{13} &= -\beta A_0C_0, \\ c_{02}c_{13} + c_{03}c_{12} &= 0, \end{aligned} \tag{84}$$

with the analogous three equations after transposition. Nonzero hopping and $a = 0$ imply $B_0C_0 = \rho^2 A_0D_0 \neq 0$. The first two equations therefore force $D_0 = sA_0$, $C_0 = sB_0$ and $B_0 = t\rho A_0$, with $s, t = \pm 1$. The full family is

$$c = \begin{pmatrix} A_0 & t\rho A_0 & z & \eta dz \\ st\rho A_0 & sA_0 & \alpha X/z & -\eta d\alpha X/z \\ z' & \alpha Y/z' & 0 & 0 \\ \eta' dz' & -\eta' d\alpha Y/z' & 0 & 0 \end{pmatrix}, \quad \begin{aligned} X &= st\rho A_0^2, & Y &= t\rho A_0^2, \\ A_0, z, z' &\neq 0, & \eta, \eta' &= \pm 1. \end{aligned} \tag{85}$$

Indeed the product constraints give $c_{03}/c_{02} = \eta\sqrt{\beta/\alpha}$, and similarly for the transposed pair. Conversely, substitution into the eight single-flip polynomials and the hopping law verifies every member. Apart from overall scale and four discrete signs, the two free ratios are z/A_0 and z'/A_0. Their inverse rescalings preserve these cancellation equations, but need not preserve other terms in the Hamiltonian.

Finally we compute the shared-link return of Section 4.3 for every member of the family. A centre-only shared-link matrix element has polynomial

$$\frac{297}{1000}c_{02}c_{12} - \frac{3}{20}c_{02}c_{13} + \frac{3}{20}c_{03}c_{12} - \frac{33}{250}c_{03}c_{13}. \tag{86}$$

For $\eta = +1$ this gives the maximum magnitude $|A_0|^2\rho[297\alpha/1000 + 33\beta/250 + 3\sqrt{\alpha\beta}/10]$. For the other orientation a different leaf-occupation matrix element supplies the same magnitude. The sixteen discrete sign choices are included in the independent rational review. Since the hop has magnitude $|A_0|^2(39064/331425)$, the ratio in Eq. (32) follows and is independent of z, z'. This is a restriction on this real selector/port family, not on all possible qutrit networks.

C Permutation classification of the junction

This appendix proves Proposition 6.1: it checks, for each of the 24 ways of rearranging four registers, whether the two junction histories can agree up to that rearrangement and a scalar factor. Let a permutation p act by moving the content of register i to register $p(i)$. Operator

products act right to left. Cycle notation such as $(0\,1\,2\,3)$ means $0 \to 1 \to 2 \to 3 \to 0$. With pure swaps, Eq. (38) gives

$$\pi_L = (0\,1\,2\,3), \qquad \pi_R = (0\,3\,2\,1), \qquad \pi_L \pi_R^{-1} = (0\,2)(1\,3). \tag{87}$$

These explicit cycles also fix the background assignments in Eq. (41).

For the general law $aI + bS$, expand the two three-factor products in $\mathbb{C}[S_4]$, the group algebra (formal linear combinations of permutations). On $(\mathbb{C}^3)^{\otimes 4}$ the 24 permutation matrices span a 23-dimensional space. The sole missing ideal is generated by the totally antisymmetric sum, since $\Lambda^4(\mathbb{C}^3) = 0$: four registers with only three colours cannot be totally antisymmetric, so that one combination of permutations acts as zero. Thus if a difference has coefficients d_p, it is the zero physical operator if and only if

$$d_p = \lambda \, \mathrm{sgn}(p) \qquad \text{for every } p \in S_4 \tag{88}$$

for some common scalar λ.

This converts $L_c - \vartheta P R_c = 0$ into a finite polynomial system. For $b \neq 0$, homogeneity allows $b = 1$. An exact rational Gröbner calculation (a standard computer-algebra method for deciding whether polynomial equations have a common solution) for each of the 24 choices of P gives the unit ideal, meaning no solution, for 23 choices; for $P = (0\,2)(1\,3)$ its reduced basis is $\{\vartheta - 1, a\}$. For $b = 0$, the nonzero products are $a^3 I$, so $P = I$ and $\vartheta = 1$. The computational certificate is reproduced by the independent manuscript audit, using its own permutation multiplication and the physical antisymmetrizer relation. Incorporating the supplied ribbon sign changes the total exchange multiplier to $-\vartheta$.

For independent nonzero coefficient pairs (a_i, b_i) on the three links, the cases ("charts") are classified by which b_i vanish. If all vanish, only $P = I$ occurs. With one nonzero b_i, the histories coincide; the link swap supplies an additional relation only for $a_i/b_i = \pm 1$, a symmetric or antisymmetric projector. With two nonzero b_i, both active links must be pure swaps, giving their oriented three-cycle. With three, all must be pure swaps, giving (02)(13). A zero link makes both histories vanish and is excluded. These eight charts explain the exceptional projector cases without interpreting them as generic transport statistics.

D Link phases, rotation classes and their scope

This appendix collects calculations about rotations and spin that were part of the original draft. The spin-related calculations in the original draft answer a different question from colour transport. They are useful diagnostics, but do not establish a relativistic spin–statistics theorem for the network.

D.1 Planar charge–flux composites

In the abelian quantum double of $\mathbb{Z}_N$ (the family of lattice models that includes Kitaev's toric code for $N = 2$), a charge q and flux m have topological twist

$$\theta_{q,m} = \exp(2\pi i q m / N). \tag{89}$$

This is the phase of rotating a framed topological excitation by a full turn, not by itself an ordinary three-dimensional spin representation[11]. The $N = 2$ composite $(q, m) = (1, 1)$ has twist -1 and fermionic self-exchange. Pure charge and pure flux have twist $+1$. For $N = 3$, the possible twists are cube roots of unity and none is -1. Reversing a braiding orientation conjugates the phase; comparisons must use consistent conventions.

The finite planar checks reproduce these results. A private internal twist can change an internal rotation phase without entering another particle's hopping path. It therefore supplies no general implication from that rotation phase to the external exchange sign. The successful shared-link construction assumes that both paths see the same attached flux and the specified ribbon routing.

D.2 Closed-node rotation tests

Two finite graphs were used for the rotation calculation. Graph V has the six vertices and twelve edges of an octahedron, with eight triangular plaquettes. Graph F is its face-adjacency graph, with eight vertices, twelve edges and six square plaquettes. The implemented Graph F has one degree of freedom per face; it is *not* a 24-vertex system retaining three independent vertices on every triangular face.

For a link-phase Hamiltonian (hopping with a fixed phase on each link), a spatial graph permutation may be a symmetry only after a compensating diagonal gauge transformation, that is, after also changing the phase of each site. Let U_x, U_y be such lifts of two commuting half-turns in $D_2 \cong \mathbb{Z}_2 \times \mathbb{Z}_2$. The invariant commutator sign is

$$U_x U_y = s\, U_y U_x, \qquad s = \pm 1. \tag{90}$$

The sign distinguishes the two $U(1)$ projective classes relevant here. Informally, $s = -1$ means that the lifts of the two half-turns anticommute rather than commute, the same algebraic behaviour as half-turns of a spin-one-half particle; the caveats that follow explain why this alone does not establish spin. The squares of the individual lifts depend on their phase normalization. In particular, $s = -1$ is not by itself a statement that a time-reversal operator squares to -1, so any resulting doublet should not automatically be called Kramers-degenerate. Nor does this finite subgroup alone determine a unique continuum angular momentum j.

The supplied enumeration contains 316 declared symmetry-compatible patterns, of which 144 have $s = -1$ and 172 have $s = +1$. It includes finite clock families and selected continuous-phase examples; it is not an exhaustive enumeration of all $U(1)$ configurations. For outward plaquette phases and a symmetry-compatible choice of representatives, it finds

$$s = (-1)^n, \qquad n = \frac{1}{2\pi} \sum_p \Phi_p. \tag{91}$$

The integer n is consistent with the product of plaquette holonomies being one on the closed surface. Branch choices at $\Phi_p = \pi$ must be handled by whole symmetry orbits when comparing its parity.

Declared family	Graph V: projective patterns	Graph F: projective patterns
$\mathbb{Z}_2$	0/4	4/8
$\mathbb{Z}_3$	0/3	0/9

Table 10: Finite enumeration results. Counts refer to the specified symmetry families, not all possible qutrit or qubit Hamiltonians.

A uniform $\pi/4$ triangular flux on V gives $n = 1$ and energies $-2\sqrt{2}$ (multiplicity two) and $+\sqrt{2}$ (multiplicity four) in the source normalization. These multiplicities and further point-group characters are compatible with restricted spinor representations. They are not a calculation of a relativistic dispersion or of two-node exchange in three dimensions.

Goldhaber's charge–monopole analysis supplies a continuum spin–statistics connection for its own physical assumptions[25]. A finite graph with prescribed link phases has not

established those assumptions: it has neither a dynamical Coulomb field nor a computed three-dimensional exchange process. The closed-node result is therefore retained as a finite rotation-class diagnostic. The original draft's separate ring-beat formula is omitted because it does not match its stated Hamiltonian normalization and is not needed for any transport result.

E Proof and evaluation of the residual bound

This appendix proves the residual bound of Section 7.10 and explains how its ingredients were computed. The idea is the matrix version of a familiar fact: if a trial vector is almost an eigenvector (small residual) and the nearest unwanted eigenvalue is well separated (large gap), then the trial vector is close to a true eigenvector. This argument applies to a finite Hermitian Hamiltonian; it does not rely on a perturbative truncation of the exact evolution. Write $\bar{H} = H_\star/G$, let E and T be isometries with m columns, and assume $E^\dagger T \geq \widetilde{s} I > 0$. Let $\widetilde{K} = T^\dagger \bar{H} T$ and $R = \bar{H}T - T\widetilde{K}$. In particular, $T^\dagger R = 0$. Suppose a bound $e_c \leq \lambda_{m+1}(\bar{H})$ exceeds $e_+ = \lambda_{\max}(\widetilde{K})$. The min–max principle gives $\lambda_m(\bar{H}) \leq e_+ < e_c$, identifying a separated lowest band of dimension m.

Let P_{low} project onto that band. On its complement the spectrum starts at e_c. Projecting the residual equation onto this complement gives a Sylvester equation (a linear matrix equation of the form $AX - XB = C$) whose two spectra are separated by $\gamma = e_c - e_+$. Its integral solution has operator norm at most $1/\gamma$ times the norm of its right-hand side. Consequently

$$\|(I - P_{\mathrm{low}})T\|_2 \leq r/\gamma = \nu. \tag{92}$$

For $\nu < 1$, choose an isometry W' onto the exact band with $D = T^\dagger W' \geq 0$. Then $D \geq \sqrt{1-\nu^2}\, I$ and $W' = TD + N$, where $T^\dagger N = 0$ and $\|N\|_2 \leq \nu$. Writing $K' = W'^\dagger \bar{H} W'$ gives the exact identity

$$K'D - D\widetilde{K} = N^\dagger R. \tag{93}$$

Subtract the midpoint of the spectrum from $\widetilde{K}$, so that its norm becomes κ. With $d = \sqrt{1-\nu^2}$, the commutator with D is at most $2\kappa(1-d)$. Hence

$$\|K' - \widetilde{K}\|_2 \leq \frac{2\kappa(1-d) + \nu r}{d} \leq \frac{3\nu(\nu\kappa + r)}{1-\nu^2} = \Delta. \tag{94}$$

The second, looser bound is the one used in the saved calculation. This proves the secular, or time-growing, part of the estimate: an error of size Δ in the Hamiltonian produces a phase error that grows in proportion to time.

The readout also depends on the basis, because states are prepared and measured in the bare code rather than in the exact band. From $W' = TD + N$ and the positivity of $E^\dagger T$,

$$\begin{aligned} \|E^\dagger W' - I\|_2 &\leq (1 - \widetilde{s}) + (1 - d) + \sqrt{1 - \widetilde{s}^2}\, \nu \\ &\leq (1 - \widetilde{s}) + \nu^2 + \sqrt{1 - \widetilde{s}^2}\, \nu = g_*. \end{aligned} \tag{95}$$

The triangle inequality for the largest principal angle between subspaces gives $\sigma_{\min}(E^\dagger W') \geq s_*$ of Eq. (77), provided the stated angle sum is below $\pi/2$. Split $E = W'W'^\dagger E + E_\perp$. Invariance removes the cross terms in the compressed evolution. The two factors of $E^\dagger W'$ cost at most $2g_*$, while the complementary term has norm at most $1 - s_*^2$. Duhamel's identity then gives

$$\left\| E^\dagger e^{-iH_\star t} E - e^{-iK_\star^{(\leq 2)}\tau} \right\|_2 \leq \epsilon' + |\tau|\delta', \qquad \tau = G\lambda^2 t. \tag{96}$$

For the sharper probability bounds, compare the two unitary band evolutions first, using their state-angle distance, then subtract the bounded readout error. This gives Eq. (75) with ϵ', δ' and its positive-part and angle conditions.

What the numerical calculation evaluates

The trial before normalization is $\Psi = E + X_1 + X_2$, where X_1 contains one auxiliary excitation and X_2 contains the zero- and two-excitation returns, including shared-link terms. These terms are orthogonal to E, so $E^\dagger\Psi = I$ and $T = \Psi(\Psi^\dagger\Psi)^{-1/2}$ is polar-aligned. The trial has at most two excited auxiliaries; applying the full Hamiltonian can produce three, and all such components are retained in R. This finite support makes its norm computable without a full time-evolution calculation.

The code is $(\mathbb{C}^2 \otimes \mathbb{C}^3)^{\otimes 4}$, of dimension 1296. All trial and residual Gram operators commute with global SU(3). In the colour factor $\mathbf{3}^{\otimes 4}$, the diagrams $(4), (3,1), (2,2), (2,1,1)$ all contain weight $(2,1,1)$, with weight multiplicities $1, 2, 1, 1$. Restricting to that weight therefore preserves the norm of every invariant code operator: each irreducible multiplicity block is still present. (A *weight* is a fixed count of how many registers hold each colour; every colour representation that occurs contains states of this weight, so nothing is lost by working only with them.) The restricted colour space has dimension 12, and including the 16 occupations gives the 192 columns used in the calculation. In physical-qutrit coordinates this is weight $(6,5,5)$. This argument, rather than an assumption that one arbitrary weight is representative, licenses the reduction. The boundary colour register is inert under the star Hamiltonian; tensoring all operators with its identity preserves these norms, so the bound also applies to states entangled with that boundary.

A Weyl bound on the next eigenvalue is obtained from the selector gap one minus the norm of the entire perturbation, including the onsite fields and pair counterterms. The calculation evaluates the three Gram matrices $\Psi^\dagger\Psi$, $\Psi^\dagger\bar{H}\Psi$ and $(\bar{H}\Psi)^\dagger\bar{H}\Psi$. Direct residual-vector and smaller exact-band checks test the subtraction used to extract r^2. Floating-point norm margins and replay agreement provide numerical evidence; they are not a proof of directed-rounding enclosures.

F Verification and reproducibility

This appendix records which calculations support which claims, and how independently each was checked. The paper combines results from several small studies. Table 11 distinguishes code replay from scientific review. A count of successful checks is a description of a harness, not a substitute for a proof or a measure of independent evidence.

The source bundles are named

`qutrit_occupation_pair_transport_2026-11-03`,
`qutrit_flux_attached_junction_2026-11-03`,
`qutrit_colour_swap_hop_search_2026-11-04`,
`qutrit_qubit_fermion_network_2026-11-03`,
`qutrit_qubit_integration_review_2026-11-03`,
`one_twist_spin_statistics_2026-11-03`,
`qutrit_relational_transport_completion_2026-11-04`,
`fermions_colour_independent_check_2026-11-04`,
`qutrit_local_dressed_junction_2026-11-04`,
`fermions_colour_novelty_check_2026-11-04`,
`fermions_colour_review_2026-11-04`,
`qutrit_aux_finite_calibration_2026-11-04`.

These identifiers locate the historical studies within the public supporting archive[24]. Its `verification/` directory contains input manifests, source snapshots, saved numerical matrices, the exact port replay, independent audits and the resulting logs. The release's `reproduce.py`

Package identifier	Checks reported	Scope and status
QQF1	29/29	Finite exchange algebra; independent integration review adds 31/31 and 9/9 with scope corrections.
OTS1	57/57; 8/8	Planar ribbons and finite rotation enumeration; constructor-checked, not independently reviewed.
Occupation transport	48/48, 25/25, 10/10	Pair compression, virtual returns and exact rational audit; independently replayed, not an independent derivation review.
OA1	22/22; 2/2	Planar dressing and colour comparison; reviewed with corrections, including vacuous rows and the background permutation.
A0S1	18/18	Specified-family search; its exact ratio certificates are stronger than its incomplete candidate search.
Selection checks	35/35; 11/11	Full two-channel formula and follow-up arithmetic; the old universal parity conclusion is superseded by Section 4.2.
Manuscript audit	20/20	Independent block, basis, permutation and cancellation checks. Selector fault: 18/20, exit 1.
Relational extension	16/16	Auxiliary grading and fixed-number colour routing. Four faults give 11/16, 12/16, 15/16, 14/16, all exit 1.
Independent check	60/60	Separate builder and a non-perturbative route for the cancellation, parity, family and permutation claims. Selector fault: exit 1. Independence refers to the calculation, not to human peer review.
Local dressed junction	50/50	Exact code identities and short-window effective dynamics; the longer uncalibrated test below fails. A symmetrized-background control fails.
Novelty maps	19/19	Evaluated maps to known statistics constructions, separate colour rotations and background-sector energies. Wrong-sign control fails.
Colour-symmetry audit	9/9	Independent exhaustive two-site generators and all-colour code invariance. Removing colour carriage gives 7/9, exit 1.
External manuscript review	Reproduced	Independent algebra and 78-start eight-hop test. The uncalibrated point fails; separate intake replay 16/16.
Finite calibration	Two fixed points	Six physical controls; 102 inputs, 4/8/16 hops. Independent link and dynamics reproduction imports the saved two-link matrices.
Error-bound audit	9/9	Conditional polar estimate and explicit counterexample to the proposed leakage bound; fault 8/9, exit 1.

Table 11: Evidence ledger. Several rows are constructor consistency checks. In particular, equality of propagators follows from equality of their generators; it is not an additional derivation of the encoding.

runner executes the principal checks in a temporary copy. The original research files were not altered by the editorial audit.

The independent matrix audit constructs the four-qutrit embedding directly from ϵ_{abc} and builds its own permutations. The exact port polynomials are separately exported from the unchanged source calculation. Their cancellation is checked symbolically, then corroborated by a separately constructed floating 560-state inverse. The 24-permutation classification is independently rebuilt with exact polynomial arithmetic. The pre-existing exact rational transport audit also replays 10/10 in the copied tree.

The new cancellation is an editorial-audit result, not an outcome registered before the original search. The relational extension has its own prospective specification and does not reuse the original search's outcome labels. The subsequent external computational review re-derived the cancellation family and junction obstruction; it did not certify every historical component or the microscopic dynamics. The first manuscript-audit attempt retained two incorrect expected permutation tuples and encountered a JSON integer-serialization error; the failed record is preserved, and the corrected convention is explicitly tested.

The local-junction development initially omitted a same-order return through two different links and reported a false pass. The corrected calculation retained that term and rejected the cancelling star law. The auxiliary study failed at $h = 10$; its post-hoc four-hop $h = 40$ continuation was followed by a failed prospective eight-hop test. Finite-coupling calibration is a new law with the distinct fixed test in Section 7.7; it does not relabel the earlier failure. All development records are retained. One reported limiting-value control compares two zeros and is non-discriminating; it is not evidence for convergence. The exact low-band comparisons provide the relevant order check. The novelty revision replayed the 50-check producer and the 19-check mapping script, with no coefficient changes. It also checked the colour generators independently on all 36 two-site states and code invariance on the complete 54-state singlet code.

The novelty assessment uses primary literature and explicit maps where the models differ. It identifies the antisymmetric-background sign and the three-colour limit as known. The literature search was not exhaustive and did not trace forward citations systematically. Finding no identical formulation of the colour-carrying composite is therefore not evidence of priority. The paper claims the stated finite-model calculations, rather than a new statistics mechanism. The centre-charge proposition and class-sum bounds are standard representation-theory consequences.

All physical simulations in the source studies retain the declared initialized symmetry sector, not necessarily the entire tensor-product Hilbert space. Some historical clock-evolution inputs remain outside this release, and certain archived exploratory scripts retain machine-specific paths. The release README distinguishes those records from the portable reproduction commands. This limitation does not affect the included principal matrix checks, sampled finite-calibration replay or reported residual-bound arithmetic.

Scope of the finite-calibration verification

The separate verification script constructs its own 2240-state link solver and star/background encoding. It reproduces the calibrated parameters, link operators and sampled 102-start results, while importing the original connected two-link matrices. The editorial replay repeats those calculations without changing coefficients. The ideal evolution control passes; a 2% wrong-clock control fails through the same dynamics pipeline.

The verification also proposes full-star error and tolerance bounds. Those proposals are not accepted as numerical certificates: its three-link overlap is extrapolated from smaller

clusters, and its proposed sixth-order norm bound uses the net two-link norm without a resolvent/remainder proof. The stated premise $M_{\mathrm{ct}}E = 0$ is false for the defined pair counterterm: on the explicit link sector its norm is about 0.959 at $h = 40$ and 0.972 at $h = 80$, with nonzero off-code parts as well. This does not invalidate the calibrated link solve. The additional leakage claim is refuted by the two-level example in Section 7.9. The nine-row audit verifies that example and the valid conditional inequality; replacing the corrected bound with the rejected one fails its control.

A further analytical counterterm adjustment improves the supplied effective-model fidelity by about 40% of the remaining infidelity. It re-solves the six link controls but keeps the old two-link matrix. It is a post-test diagnostic, not a consistently recomputed new physical candidate, and is excluded from Table 6. Likewise, tolerance figures from two finite-difference evaluations per control are local sensitivity estimates, not established safe intervals. The paper retains the verified calibration result. The separate residual calculation in Section 7.10 supplies a new bound on the omitted dynamics; it does not rehabilitate the earlier invalid estimates or establish tolerance intervals.

Scope of the full-star overlap bound

The bound in Table 7 is not covered by the separate verification or its review. The final manuscript audit checked its derivation and added the missing positivity and contraction hypotheses. Its numerical evidence remains distinct from the earlier calibration review. The script is preserved under `verification/final_math_review_2026-11-04/snapshot/technical_notes/fermions_colour_bounds_check_2026-11-04/` in the support archive. It uses its own operators on the whole eight-qutrit space. The bound holds against the exact one- and two-link overlaps, and substituting a wrong auxiliary energy makes that check fail. The theorem itself had no violation on 3349 random finite models with exact lowest bands, while dropping its Neumann and quadratic terms gives 1727 violations.

Final residual and reference audit

The further package `qutrit_star_band_residual_2026-11-04` constructs the trial and evaluates the full-star residual. The final manuscript audit independently derived Appendix E, including the stronger intermediate inequality (94), and checked it on 500 complex Hermitian models with 2000 finite-time comparisons. Removing the dressing allowance produces a failing control. An independent tableau enumeration verifies that the chosen weight contains every colour irrep. Both star calculations were replayed from preserved copies without changing coefficients; the original producer's 1500-model theorem test, its deliberate fault, and the one-link exact-band comparison were also replayed. Reusing a producer is a reproducibility check, not an independent implementation of its full-star tensors.

The audit record is `verification/final_math_review_2026-11-04/`. It preserves the previous paper, input hashes, checks and a reference-by-reference ledger. All 24 references were checked against DOI registration metadata, publisher records or arXiv; Ball's title was corrected to *Fermions without Fermion Fields*. The original electronic metadata dates of the Fulton–Harris book and the online-first date of Allerstorfer et al. differ from their cited print/volume years; these are documented rather than treated as nonexistent references. No human peer review is implied by these computational checks.

References

[1] Richard Cleve, Daniel Gottesman, and Hoi-Kwong Lo. How to share a quantum secret. *Physical Review Letters*, 83:648–651, 1999. doi: 10.1103/PhysRevLett.83.648.

[2] John Kogut and Leonard Susskind. Hamiltonian formulation of Wilson's lattice gauge theories. *Physical Review D*, 11:395–408, 1975. doi: 10.1103/PhysRevD.11.395.

[3] D. Horn. Finite matrix models with continuous local gauge invariance. *Physics Letters B*, 100:149–151, 1981. doi: 10.1016/0370-2693(81)90763-2.

[4] David P. DiVincenzo, Dave Bacon, Julia Kempe, Guido Burkard, and K. Birgitta Whaley. Universal quantum computation with the exchange interaction. *Nature*, 408:339–342, 2000. doi: 10.1038/35042541.

[5] P. Jordan and E. Wigner. Über das Paulische äquivalenzverbot. *Zeitschrift für Physik*, 47:631–651, 1928. doi: 10.1007/BF01331938.

[6] Sergey B. Bravyi and Alexei Yu. Kitaev. Fermionic quantum computation. *Annals of Physics*, 298:210–226, 2002. doi: 10.1006/aphy.2002.6254.

[7] Frank Verstraete and J. Ignacio Cirac. Mapping local Hamiltonians of fermions to local Hamiltonians of spins. *Journal of Statistical Mechanics: Theory and Experiment*, 2005: P09012, 2005. doi: 10.1088/1742-5468/2005/09/P09012.

[8] R. C. Ball. Fermions without Fermion fields. *Physical Review Letters*, 95:176407, 2005. doi: 10.1103/PhysRevLett.95.176407.

[9] Kanav Setia, Sergey Bravyi, Antonio Mezzacapo, and James D. Whitfield. Superfast encodings for fermionic quantum simulation. *Physical Review Research*, 1:033033, 2019. doi: 10.1103/PhysRevResearch.1.033033.

[10] Michael Levin and Xiao-Gang Wen. Fermions, strings, and gauge fields in lattice spin models. *Physical Review B*, 67:245316, 2003. doi: 10.1103/PhysRevB.67.245316.

[11] Alexei Yu. Kitaev. Fault-tolerant quantum computation by anyons. *Annals of Physics*, 303:2–30, 2003. doi: 10.1016/S0003-4916(02)00018-0.

[12] Michael A. Levin and Xiao-Gang Wen. String-net condensation: A physical mechanism for topological phases. *Physical Review B*, 71:045110, 2005. doi: 10.1103/PhysRevB.71.045110.

[13] Yu-An Chen, Anton Kapustin, and Đorđe Radičević. Exact bosonization in two spatial dimensions and a new class of lattice gauge theories. *Annals of Physics*, 393:234–253, 2018. doi: 10.1016/j.aop.2018.03.024.

[14] Loïc Herviou. Exact SU(N) symmetry resolution of bosonic Hubbard models: general boson–fermion mappings and flat-band singlet-cluster states. arXiv:2610.14183, 2026. URL https://arxiv.org/abs/2610.14183. Preprint, version 1, 9 October 2026.

[15] Linda Sansoni, Fabio Sciarrino, Giuseppe Vallone, Paolo Mataloni, Andrea Crespi, Roberta Ramponi, and Roberto Osellame. Two-particle bosonic-fermionic quantum walk via integrated photonics. *Physical Review Letters*, 108:010502, 2012. doi: 10.1103/PhysRevLett.108.010502.

[16] Jonathan C. F. Matthews, Konstantinos Poulios, Jasmin D. A. Meinecke, Alberto Politi, Alberto Peruzzo, Nur Ismail, Kerstin Wörhoff, Mark G. Thompson, and Jeremy L. O'Brien. Observing fermionic statistics with photons in arbitrary processes. *Scientific Reports*, 3: 1539, 2013. doi: 10.1038/srep01539.

[17] William Fulton and Joe Harris. *Representation Theory: A First Course*, volume 129 of *Graduate Texts in Mathematics*. Springer, New York, 1991. doi: 10.1007/978-1-4612-0979-9.

[18] Sergey Bravyi, David P. DiVincenzo, and Daniel Loss. Schrieffer–Wolff transformation for quantum many-body systems. *Annals of Physics*, 326:2793–2826, 2011. doi: 10.1016/j.aop.2011.06.004.

[19] Herman Feshbach. Unified theory of nuclear reactions. *Annals of Physics*, 5:357–390, 1958. doi: 10.1016/0003-4916(58)90007-1.

[20] David Elliman. Qutrit clocks and colour: Pairing algebra and finite-gap transport. Zenodo, https://doi.org/10.5281/zenodo.23900249, 2026. Research preprint, version 2.0. Version DOI: 10.5281/zenodo.23912746.

[21] Rene Allerstorfer, Matthias Christandl, Dmitry Grinko, Ion Nechita, Maris Ozols, Denis Rochette, and Philip Verduyn Lunel. Monogamy of highly symmetric states. *Communications in Mathematical Physics*, 407:21, 2026. doi: 10.1007/s00220-025-05483-3. Published online 8 December 2025; author manuscript arXiv:2309.16655v2.

[22] Fatemeh Moradi Kalarde, Sadra Boreiri, Xiangling Xu, Lucas Tendick, Salman Beigi, Paolo Perinotti, Tommaso Guaita, and Marc-Olivier Renou. Fermions are fundamentally more nonlocal than bosons. arXiv:2606.12363, 2026. URL https://arxiv.org/abs/2606.12363. Preprint, version 1, 10 June 2026.

[23] Kyung-Su Kim and Hosho Katsura. Exact hole-induced SU(N) flavor singlets in certain $U = \infty$ SU(N) Hubbard models. *Physical Review Research*, 6:013307, 2024. doi: 10.1103/PhysRevResearch.6.013307.

[24] David Graham Elliman. Supporting code and evidence for Fermions and Colour: Transport and its limits in qutrit–qubit networks. Zenodo, https://doi.org/10.5281/zenodo.23974017, 2026. Version 1.0.0. Version DOI: 10.5281/zenodo.23974018.

[25] Alfred S. Goldhaber. Connection of spin and statistics for charge-monopole composites. *Physical Review Letters*, 36:1122–1125, 1976. doi: 10.1103/PhysRevLett.36.1122.