Two-Query Chirality in Tetrahedral Quantum Echoes

Exact Parallel Readout, a Nine-Dimensional Query Defect, and a Certified Adaptive Advantage

Lluis Eriksson

August 2026

Abstract

Four balanced Pauli kicks along the vertices of a regular tetrahedron, followed by their reverse echo, generate 24 order-labelled qubit unitary channels. All 24 words have the same trace and split into two 12-element orbits of the proper tetrahedral group. One query sees only the common trace: a Bell tester is globally optimal and gives the same performance in the two orbit sectors. We show that two queries recover information that is invisible at one query. The quadratic output decomposes as $\mathbf{1} \oplus \mathbf{2} \oplus \mathbf{3} \oplus \mathbf{3}$, and two exact chiral energies $E_\pm$ enter the optimum through a scalar clamp. This yields closed formulas for the Bell-square optimum and for the globally optimized parallel strategy. The optimal parallel input is generically not two Bell pairs.

At the algebraic pulse point with common half trace $q = 1/2$, the parallel value is

$$P_{\mathrm{par}} = \frac{5+\sqrt{15}}{24}.$$

Every causal two-query protocol has success at most 9/24: fixed trace and fixed Pauli-vector norm remove one direction from the universal ten-dimensional quadratic query space. We give an explicit rational two-comb normalization and an exact weighted-frame identity attaining this cap. Hence

$$P_{\mathrm{causal}} = \frac{3}{8} > P_{\mathrm{par}}, \qquad P_{\mathrm{causal}} - P_{\mathrm{par}} = \frac{4-\sqrt{15}}{24}.$$

This is an analytic adaptive advantage for a non-group unitary ensemble, not a numerical SDP observation. A quotient-ring certificate verifies all 24 echo words, the chiral invariants, the parallel formula, and the causal projector identity. A uniform query-norm estimate further proves that the same fixed adaptive tester beats every parallel tester on an explicit open pulse interval around the algebraic point. We make no claim about indefinite causal order or noisy hardware.

1 Introduction

Minimum-error discrimination of unitary channels changes qualitatively when the unknown gate may be queried more than once. For a uniformly distributed projective representation of a finite group, covariance often makes parallel strategies optimal. Outside the group case, sequential and more general architectures can be strictly better [1]. Explicit analytic families exhibiting the mechanism remain comparatively scarce: general tester semidefinite programs certify an optimum, but usually hide why a particular inter-query memory helps.

This paper gives a closed example derived from a concrete pulse architecture. Let four control Hamiltonians point to the vertices of a regular tetrahedron. Executing the four kicks in an order $\sigma \in S_4$ and closing the path by a reverse echo gives a qubit unitary W_σ. The 24 labels do not form a unitary group. They form two compound geometrically uniform orbits under A_4, and every word has the same trace. That common trace makes one-query readout blind to the difference

Table 1: Scope of the exact results. "Causal" includes parallel, sequential, and adaptive two-query circuits with delayed final measurement.

Task	Result here	Limitation
One query, 12 or 24 labels	Exact global tester optimum	Recalled as baseline
Two queries, Bell-square	Exact for every pulse	Bell-square is not globally parallel-optimal
Two queries, parallel	Exact for every pulse	Formula concerns noiseless channels
Two queries, causal	Exact at $q = 1/2$; strict advantage on an explicit interval	Rank cap, not full causal optimum, away from $q = 1/2$
Collision classification	Exact on the fundamental interval	Labels and channels distinguished

between the two orbit geometries. At two queries, third and fourth moments survive and split the sectors.

The main contributions are as follows.

1. We derive the complete two-copy representation reduction. Two architecture-specific invariants E_+ and E_- are explicit rational functions of the pulse parameter. They determine the exact Bell-square optimum by a one-dimensional clamp.
2. We optimize over every parallel input, including arbitrary ancillas. Symmetry reduces the input marginal to a singlet weight a, after which the remaining optimization is elementary. At $q = 1/2$ the answer is $(5 + \sqrt{15})/24$, attained at $a_* = (5 - \sqrt{15})/10$; two Bell pairs are strictly suboptimal.
3. We prove an architecture-specific rank cap. Any purified causal two-query protocol is a fixed linear image of $W_\sigma^{\otimes 2}$, whose span has dimension at most nine rather than the universal ten.
4. We saturate that cap at $q = 1/2$. An explicit rational deterministic comb Ξ_* and two positive orbit weights obey $KCK = K$. This converts the 24 effective output states into a scalable tight frame and proves the exact causal optimum 3/8.
5. The separation is not confined to one isolated pulse. A uniform Lipschitz estimate gives an explicit interval on which the same adaptive tester remains better than the globally optimal parallel strategy.

The ingredients surrounding the result are classical: group-covariant state discrimination [2, 3, 4], quantum testers [5, 6], and unitary-design dimension bounds [7, 1]. The new point is their exact synthesis for this fixed-trace, two-orbit echo family, including a rational adaptive certificate and a strict closed-form parallel-versus-causal separation. The tetrahedral word itself is rederived below, so the argument is logically self-contained. The appendices give the full block reduction, a six-Kraus realization of the inter-query map, and the collision-elimination and certificate details that are easy to obscure in a short presentation.

2 The tetrahedral echo family

Let

$$r_1 = (1, 1, 1), \quad r_2 = (1, -1, -1), \quad r_3 = (-1, 1, -1), \quad r_4 = (-1, -1, 1).$$

Then $r_a \cdot r_a = 3$, $r_a \cdot r_b = -1$ for $a \neq b$, and $\sum_a r_a = 0$. With the Pauli vector $\boldsymbol{\sigma}$, define

$$K_a(y) = \exp[-iyr_a \cdot \boldsymbol{\sigma}] = cI - isr_a \cdot \boldsymbol{\sigma}, \quad c = \cos(\sqrt{3}y), \quad s = \frac{\sin(\sqrt{3}y)}{\sqrt{3}},$$

so $c^2 + 3s^2 = 1$. The first entry of an order acts first:

$$U_\sigma = K_{\sigma_4} K_{\sigma_3} K_{\sigma_2} K_{\sigma_1}, \qquad W_\sigma = U_\sigma U^*_{\sigma^R}.$$

Equivalently,

$$W_\sigma = K_{\sigma_4}(y)\cdots K_{\sigma_1}(y)K_{\sigma_4}(-y)\cdots K_{\sigma_1}(-y).$$

Write

$$W_\sigma = q_\sigma I - i v_\sigma \cdot \boldsymbol{\sigma}, \qquad q_\sigma^2 + \|v_\sigma\|^2 = 1.$$

Proposition 2.1 (Common half trace)**.** *For every $\sigma \in S_4$, $q_\sigma = q$, where, with $z = s^2$,*

$$q(z) = 1 - 32z^2 + 128z^3 - 256z^4, \qquad 0 \le z \le \frac{1}{3}. \tag{1}$$

The even and odd permutations form two free 12-label orbits $\mathcal{O}_+$ and $\mathcal{O}_-$ under left multiplication by A_4.

Proof. If a partial product is $aI - ib\cdot\boldsymbol{\sigma}$, left multiplication by $cI - isr\cdot\boldsymbol{\sigma}$ sends

$$(a,b) \longmapsto (ca - sr\cdot b,\ cb + sar + sr\times b).$$

Eight iterations for one representative of either parity give the common scalar

$$(c^2 - s^2)(c^6 + 13c^4s^2 + 35c^2s^4 + 79s^6).$$

Substituting $c^2 = 1 - 3s^2$ gives Equation (1). Proper tetrahedral rotations act as A_4 on labels and conjugate the corresponding words, so one representative per parity suffices. The verifier independently expands all 24 words. $\square$

Set $r^2 = 1 - q^2$. Irreducibility of the real tetrahedral representation gives for each parity

$$\sum_{\sigma\in\mathcal{O}_\pm} v_\sigma = 0, \qquad \sum_{\sigma\in\mathcal{O}_\pm} v_\sigma v_\sigma^{\mathsf{T}} = 4r^2 I_3. \tag{2}$$

These second moments coincide. The two orbits first separate in higher moments.

It is useful to replace z by

$$t = \frac{c}{s} = \sqrt{\frac{1-3z}{z}}.$$

The endpoint $z = 0$ is understood as $t \to \infty$. Define the off-diagonal quadratic energies

$$E_+(t) = \frac{32768(t-1)^8(t+1)^2(t^5 - t^4 + 8t^3 - 8t^2 + 15t + 1)^2}{(t^2+3)^{16}}, \tag{3}$$

$$E_-(t) = \frac{32768(t-1)^2(t+1)^8(t^5 + t^4 + 8t^3 + 8t^2 + 15t - 1)^2}{(t^2+3)^{16}}. \tag{4}$$

For a seed vector in the indicated orbit,

$$E_\pm = 2\sum_{a<b} v_a^2 v_b^2, \qquad D_\pm = \frac{2r^4}{3} - E_\pm.$$

Although Equation (2) makes one-query performance identical, the pair $(D_\pm, E_\pm)$ remembers the orbit.

3 One query as a symmetry baseline

Let $|\Omega\rangle = (|00\rangle + |11\rangle)/\sqrt{2}$ and $|\psi_\sigma\rangle = (W_\sigma \otimes I)|\Omega\rangle$. Vectorization gives

$$\langle\psi_\tau|\psi_\sigma\rangle = \frac{1}{2}\operatorname{Tr}(W_\tau^* W_\sigma) = q^2 + v_\tau\cdot v_\sigma.$$

For either 12-state orbit, or for their 24-state union with label multiplicity, the nonzero spectrum of the average state is

$$q^2, \qquad \frac{r^2}{3}, \quad \frac{r^2}{3}, \quad \frac{r^2}{3}. \tag{5}$$

Every one-query tester can be symmetrized by A_4. Its input marginal then commutes with the irreducible two-dimensional spin lift and is therefore $I_2/2$. Hence the Bell input loses no performance. The Yuen–Kennedy–Lax conditions [8] in $\mathbf{1} \oplus \mathbf{3}$ give

$$P_1^{(N)}(q) = \frac{\left(|q| + \sqrt{3(1-q^2)}\right)^2}{N}, \qquad N \in \{12, 24\}. \tag{6}$$

At $q = \pm 1/2$ the four eigenvalues in Equation (5) are equal. Thus a 12-label orbit is a tight frame in dimension four and $P_1^{(12)} = 1/3$; the 24-label value is $1/6$. This tightness is a one-query statement and does not imply a unitary two-design.

4 Two-copy representation and exact Bell-square readout

For two Bell probes, the normalized output vector is $|W_\sigma\rangle\!\rangle^{\otimes 2}/2$. The relevant quadratic feature space is contained in

$$\mathrm{Sym}^2(\mathbf{1} \oplus \mathbf{3}) = \mathbf{1}^{\oplus 2} \oplus \mathbf{1}' \oplus \mathbf{1}'' \oplus \mathbf{3}^{\oplus 2}.$$

The two invariant coordinates are proportional for every label, leaving one zero direction. The two copies of $\mathbf{3}$ decouple: in coordinates this is the identity $v \cdot Q(v) = 3v_1v_2v_3 = 0$ on either echo orbit.

Proposition 4.1 (Bell-square frame)**.** *The average state of the 24 Bell-square outputs has spectrum*

$$0, \quad m, \quad \lambda_D^{\times 2}, \quad \lambda_V^{\times 3}, \quad \lambda_E^{\times 3},$$

where

$$m = q^4 + \frac{r^4}{3}, \qquad \lambda_D = \frac{D_+ + D_-}{4}, \qquad \lambda_V = \frac{2q^2r^2}{3}, \qquad \lambda_E = \frac{E_+ + E_-}{6}.$$

The generic rank is nine. It is never a tight frame on that support.

Proof. The block eigenvalues follow from Schur averaging and the definitions of $D_\pm, E_\pm$. Tightness would in particular require $m = \lambda_V$, or

$$6(q^2)^2 - 4q^2 + 1 = 0,$$

whose discriminant is negative. □

The ensemble is compound geometrically uniform, not a single unitary group. Consequently the ordinary pretty-good measurement need not be optimal. The state-discrimination dual can nevertheless be averaged over A_4. Let

$$T = \frac{2r^4}{3}, \qquad E_c = \mathrm{clamp}\left(\frac{2r^4}{5}, [\min(E_+, E_-), \max(E_+, E_-)]\right),$$

where $\mathrm{clamp}(x, [u, v]) = \min(v, \max(u, x))$, and put

$$D_c = T - E_c, \qquad H = \sqrt{2D_c} + \sqrt{3E_c}.$$

Theorem 4.2 (Exact Bell-square optimum)**.** *For all $z \in [0, 1/3]$,*

$$P_{\mathrm{Bell}^{\otimes 2}} = \frac{1}{24}\left(\sqrt{q^4 + r^4/3} + \sqrt{6}\,|q|r + H\right)^2. \tag{7}$$

Proof. Give the $+$ seed constraint dual weight $w \in [0,1]$ and the $-$ seed weight $1-w$. Then $E(w) = wE_+ + (1-w)E_-$ and $D(w) = T - E(w)$. Minimizing the positive dual blocks and using the Schur multiplicities reduces the optimum to

$$\frac{1}{24}\max_{0\le w\le 1}\left[\sqrt{m}+\sqrt{2D(w)}+\sqrt{6q^2r^2}+\sqrt{3E(w)}\right]^2.$$

The concave function $\sqrt{2(T-E)}+\sqrt{3E}$ has its unconstrained maximum at $E = 3T/5 = 2r^4/5$. Projection to the available interval gives E_c. The corresponding covariant seed operators satisfy the primal completeness conditions, so the bound is attained. Appendix A gives the complete isotypic block table, the reduced dual program, and an explicit attaining primal construction. $\square$

5 Optimization over every parallel input

Consider an arbitrary parallel two-query strategy with any retained reference. Averaging it over A_4 preserves success. The marginal on the two queried qubits commutes with the diagonal spin action and therefore has the form

$$\rho_a = aP_{\mathrm{sing}} + \frac{1-a}{3}P_{\mathrm{trip}}, \qquad 0 \le a \le 1. \tag{8}$$

Every such marginal can be purified. Define

$$h = \frac{(4q^2-1)^2}{9}, \qquad K = \sqrt{8}\,|q|r + \frac{2H}{\sqrt{3}}.$$

Theorem 5.1 (Global parallel optimum)**.** *If $q^2 = 1$, all labels are the same channel and $P_{\mathrm{par}} = 1/24$. Otherwise, for fixed a,*

$$P_{\mathrm{par}}(a) = \frac{1}{24}\left[\sqrt{h+(1-h)a}+K\sqrt{1-a}\right]^2.$$

The optimum over all parallel inputs is

$$P_{\mathrm{par}}^{\mathrm{opt}} = \begin{cases} \dfrac{1}{24}\left(1+\dfrac{K^2}{1-h}\right), & hK^2 \le (1-h)^2,\\ \dfrac{1}{24}(\sqrt{h}+K)^2, & hK^2 > (1-h)^2. \end{cases} \tag{9}$$

In the first branch an attaining singlet weight is

$$a_* = 1 - \frac{K^2}{(1-h)(K^2+1-h)};$$

in the second branch $a_* = 0$.

Proof. In the purification of Equation (8), the scalar block has squared norm $h + (1-h)a$. Every non-scalar block in Theorem 4.2 scales by $1-a$. The same reduced primal-dual calculation therefore gives the displayed $P_{\mathrm{par}}(a)$. Put $u = \sqrt{(1-h)(1-a)}$. The bracket becomes

$$\sqrt{1-u^2}+\frac{K}{\sqrt{1-h}}u, \qquad 0 \le u \le \sqrt{1-h}.$$

Maximizing this planar inner product gives the two branches and a_*. The symmetrization, purification, and block rescaling are derived explicitly in Appendix B. $\square$

Let z_* be the unique root in $(0,1/3)$ of $q(z)=1/2$:

$$z_* = 0.167964937759754\ldots, \qquad y_* = \frac{1}{\sqrt{3}}\arcsin\sqrt{3z_*} = 0.455698535295322\ldots.$$

At this point the clamp is interior and Equation (9) simplifies.

Corollary 5.2 (Non-Bell optimum at the tight one-query point)**.** *At $q=1/2$,*

$$P_{\mathrm{par}}^{\mathrm{opt}} = \frac{5+\sqrt{15}}{24} = 0.369707639425309\ldots, \qquad a_* = \frac{5-\sqrt{15}}{10}. \tag{10}$$

Thus two Bell pairs are not the optimal parallel input, despite a Bell input being globally optimal for one query.

6 The nine-dimensional causal cap

The universal two-query dimension bound for qubit unitaries uses $\dim \mathrm{Sym}^2(\mathbb{C}^4)=10$ [1]. The echo family lies on a fixed conjugacy class and loses one dimension.

Theorem 6.1 (Fixed-trace query defect)**.** *For any purified causal protocol using the same unknown W_σ twice, the 24 final pure states span a space of dimension at most nine. Therefore*

$$P_{\mathrm{causal}}^{(2)} \le \frac{9}{24} = \frac{3}{8}. \tag{11}$$

Proof. Defer all intermediate measurements and purify every fixed operation. The final vector is then a fixed linear image of $W_\sigma \otimes W_\sigma$. Since $W_\sigma = qI - iv_\sigma\cdot\boldsymbol{\sigma}$ with common q and common $\|v_\sigma\|^2=r^2$, its quadratic coordinates consist of one scalar, three coordinates qv_a, and six symmetric products v_av_b. The trace relation $\sum_a v_a^2 = r^2$ removes one of these ten directions. Thus the final span has dimension at most $1+3+5=9$, independently of the intermediate memory. For N equiprobable pure hypotheses supported in dimension d, the dual operator I_d/N gives $P \le d/N$. $\square$

The bound is stronger than the universal 10/24 cap, but a dimension bound need not be attainable. We now exhibit exact attainment.

7 An exact adaptive certificate at $q=1/2$

Use the Choi ordering $B_2A_2B_1A_1$, and let the deterministic two-comb normalization act on $A_2B_1A_1$. In the computational basis define

$$\Xi_* = \frac{1}{54}\begin{pmatrix} 12 & 0 & 0 & 0 & 0 & 0 & -8 & 0\\ 0 & 7 & 0 & 0 & 5 & 0 & 0 & -8\\ 0 & 0 & 20 & 0 & 0 & 0 & 0 & 0\\ 0 & 0 & 0 & 15 & 0 & 0 & 5 & 0\\ 0 & 5 & 0 & 0 & 15 & 0 & 0 & 0\\ 0 & 0 & 0 & 0 & 0 & 20 & 0 & 0\\ -8 & 0 & 0 & 5 & 0 & 0 & 7 & 0\\ 0 & -8 & 0 & 0 & 0 & 0 & 0 & 12 \end{pmatrix}. \tag{12}$$

Its spectrum and normalization are

$$\mathrm{spec}(\Xi_*) = \left\{0^{\times 2}, (7/27)^{\times 2}, (10/27)^{\times 4}\right\}, \qquad \mathrm{Tr}_{A_2}\Xi_* = I_{B_1A_1}/2. \tag{13}$$

Thus Ξ_* is a valid deterministic causal normalization with maximally mixed first input. The realization theorem for quantum combs then identifies these positive normalization constraints with a causally ordered two-slot circuit, with the final measurement postponed to the end [9]. Appendix C supplies a six-Kraus channel and its Stinespring implementation, making this comb realization constructive.

Let $t_* > 0$ be the unique positive root

$$f(t) = t^8 + 12t^6 - 10t^4 - 20t^2 - 239 = 0. \tag{14}$$

This is exactly the $q = 1/2$ condition. Put

$$|J_\sigma^{(2)}\rangle = |W_\sigma\rangle\!\rangle^{\otimes 2}, \qquad K_{\sigma\tau} = \langle J_\sigma^{(2)}|(I_{B_2} \otimes \Xi_*)|J_\tau^{(2)}\rangle.$$

All diagonal entries of K equal one. Define the field element

$$\delta(t) = -\frac{t(27686401t^6 + 85868681t^4 - 2778560421t^2 + 1179263131)}{34125376000}.$$

At $t = t_*$ set

$$c_+ = \frac{3}{8} - \delta(t_*), \qquad c_- = \frac{3}{8} + \delta(t_*). \tag{15}$$

Numerically,

$$c_+ = 0.09473473348\ldots, \qquad c_- = 0.65526526651\ldots.$$

In particular both weights are strictly positive. Let $C = \mathrm{diag}(c_+ I_{12}, c_- I_{12})$, with labels ordered by parity.

Lemma 7.1 (Exact scalable-frame identity)**.** *At the algebraic point Equation (14),*

$$KCK = K, \qquad 12(c_+ + c_-) = 9. \tag{16}$$

Proof. Write every pulse as $(tI - ir_a \cdot \boldsymbol{\sigma})/\sqrt{t^2+3}$. Each echo entry is then a rational function with denominator $(t^2+3)^4$. Hence every entry of $KCK - K$ belongs to $\mathbb{Q}(t)$ with denominator nonzero at t_*. Reducing the 576 numerators modulo $f(t)$ gives zero. By A_4 covariance it is enough to check one row representative from each parity orbit. Positivity of $c_\pm$ follows by rational interval evaluation on the isolating interval $1.71860 < t_* < 1.71862$. The accompanying verifier performs these quotient-ring operations directly; no floating-point SDP is used. □

Theorem 7.2 (Exact adaptive advantage)**.** *At $q = 1/2$, the optimal success probability among all causally ordered two-query protocols is*

$$P^{(2)}_{\mathrm{causal}} = \frac{3}{8}. \tag{17}$$

It strictly exceeds the global parallel optimum:

$$P^{(2)}_{\mathrm{causal}} - P^{\mathrm{opt}}_{\mathrm{par}} = \frac{4-\sqrt{15}}{24} > 0. \tag{18}$$

Proof. Let $|\phi_\sigma\rangle$ be the normalized effective state obtained by applying $I \otimes \Xi_*^{1/2}$ to $|J_\sigma^{(2)}\rangle$. Its Gram matrix is K. Define

$$M_\sigma = c_\pm |\phi_\sigma\rangle\langle\phi_\sigma|, \qquad \sigma \in \mathcal{O}_\pm.$$

The identity $KCK = K$ says that $\sum_\sigma M_\sigma$ is the orthogonal projector onto the span of the effective states. Complete it arbitrarily on the unused orthogonal complement. This gives a valid final POVM for the deterministic comb Ξ_*. Its success is

$$\frac{1}{24}\sum_\sigma c_\pm |\langle\phi_\sigma|\phi_\sigma\rangle|^2 = \frac{12(c_+ + c_-)}{24} = \frac{3}{8}.$$

The upper bound is Theorem 6.1; hence the strategy is optimal. Subtract Equation (10) to obtain Equation (18). □

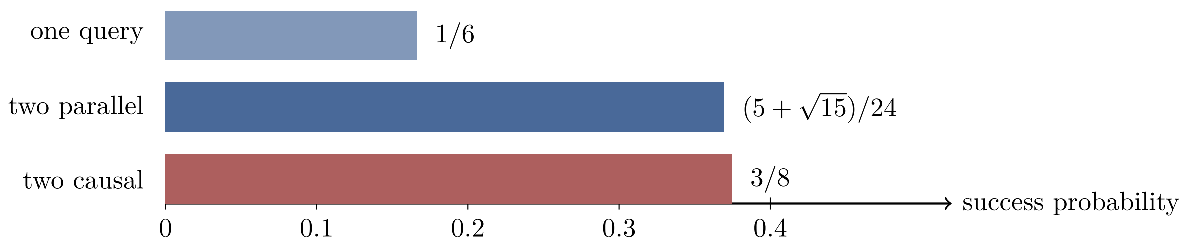

Figure 1: Exact values for 24 labelled orders at $q = 1/2$. The second use more than doubles the one-query success, and causal memory gives the strict gain $(4 - \sqrt{15})/24$ over every parallel input.

The exact point also forces an open interval of adaptive advantage. This is useful because it separates the phenomenon from algebraic fine tuning.

Corollary 7.3 (Certified open pulse interval)**.** *Let $y_* = 0.455698535295322\ldots$ be the pulse in Equation (10). Keep the comb and final POVM of Theorem 7.2 fixed while varying the 24 echo words with y. Then this causal strategy strictly outperforms every parallel two-query strategy whenever*

$$|y - y_*| < \frac{4 - \sqrt{15}}{1536\sqrt{3}} = 4.7742903163\ldots \times 10^{-5}. \tag{19}$$

Consequently the optimal causal value is strictly larger than the parallel optimum throughout that interval.

Proof. Each echo is a product of eight factors whose Hermitian generators have norm $\sqrt{3}$, hence $\|W'_\sigma(y)\| \leq 8\sqrt{3}$. After two queries, the derivative of any normalized purified output has norm at most $16\sqrt{3}$. Therefore the success probability of any fixed tester is $32\sqrt{3}$-Lipschitz. Taking the supremum preserves the same constant, so $P^{\mathrm{opt}}_{\mathrm{par}}(y)$ is also $32\sqrt{3}$-Lipschitz. The difference between the fixed causal strategy and the parallel optimum can thus decrease from its value $(4 - \sqrt{15})/24$ at speed at most $64\sqrt{3}$. The radius in Equation (19) follows. $\square$

This separation does not follow from the standard parallelization results. Those results cover binary unitary sequences or covariant estimation when the queried unitaries form a group representation [5]; the 24 echo words here form two conjugacy orbits and are not closed under multiplication. Non-group ensembles can exhibit a sequential advantage in general [1], but that fact supplies neither the value nor a strategy for this family. Equations (12)–(16) provide the missing architecture-specific certificate.

8 Channel collisions and interpretation

The discrimination problem is label-valued. At exceptional pulses, distinct labels can implement the same $SU(2)$ word, and W and $-W$ define the same physical channel. On the fundamental branch $0 \leq z \leq 1/3$, all exact word collisions occur at four points:

Pulse point	$SU(2)$ word classes	Physical channel classes
$z = 0$	one class of size 24	one class of size 24
$z = 1/4$	six classes of size 4	three classes of size 8
$64z^3 - 4z - 1 = 0$	twelve classes of size 2	same twelve classes
$256z^5 - 64z^4 - 16z^3 + 12z^2 - 1 = 0$	six pairs and twelve singles	same 18 classes
generic z	24 singles	24 singles

The point z_* used in Theorem 7.2 is generic, so its 24 physical channels are distinct. Appendix D proves completeness of the table by resultants, Sturm isolation, and direct projective filtering.

The operational phenomenon can now be stated precisely. A single query depends only on the second moments Equation (2), and therefore only on the common scalar q. Two queries expose $E_+ - E_-$, a higher-order orbit invariant. The causal comb Equation (12) uses the first output coherently when preparing the second input; it rescales the two orbit frames just enough to make the nine-dimensional query cap attainable. This is not causal nonseparability. It is an advantage inside ordinary causally ordered quantum circuits.

9 Reproducibility and limitations

The release contains three independent lightweight checks with deliberately different certification levels.

1. A word audit expands all 24 eight-pulse echoes symbolically over $\mathbb{Q}[c,s]/(c^2+3s^2-1)$, verifies the common trace, both tetrahedral moments, and the collision resultants; numerical root isolation and tolerant grouping are used as diagnostics for the physical collision classes.
2. A representation audit checks the two-copy block spectrum, scalar clamp, and parallel fixture at $q=1/2$ using NumPy with an explicit tolerance. The all-pulse parallel theorem is established by the analytic block reduction, not by this numerical replay.
3. A causal audit works in $\mathbb{Q}[t]/(f)$ at $q=1/2$. It verifies Equation (13), reconstructs the 24 effective states, proves positivity of $c_\pm$, and reduces every entry of $KCK-K$ to zero. It then checks the exact 3/8 success and the strict algebraic gap.

Thus the central causal certificate is exact and uses no numerical SDP. The other two scripts are transparent corroborating audits, while the universal parallel reduction and collision completeness remain mathematical proofs in the text. The dependency file records both SymPy and NumPy.

Several limitations are deliberate. We do not optimize indefinite-causal- order processes. We do not model pulse noise, finite measurement statistics, or laboratory compilation cost. The 24 hypotheses are known implementations with an unknown label, not an unknown continuous unitary. Finally, the dimension-nine obstruction is architecture-specific: it comes from common trace and norm, and is not asserted for arbitrary qubit ensembles.

10 Conclusion

Tetrahedral echo order has a genuinely query-dependent memory. One query compresses the two parity orbits to the same $\mathbf{1}\oplus\mathbf{3}$ spectrum. Two queries recover their different quadratic energies, make the Bell-square input suboptimal, and admit a complete parallel solution. At the algebraic point $q=1/2$, a rational causal comb saturates a sharpened nine-dimensional query bound and strictly outperforms all parallel architectures. The resulting separation is exact, analytic, and certified over all 24 pulse orders. The explicit neighborhood in Corollary 7.3 shows that the advantage is robust as a mathematical query phenomenon, even though no laboratory noise model is asserted.

A The full two-copy block reduction

This appendix expands the representation steps used in Theorems 4.1, 4.2 and 5.1. For $x=(q,-iv)\in\mathbf{1}\oplus\mathbf{3}$, the symmetric square decomposes under A_4 as

$$\mathrm{Sym}^2(\mathbf{1}\oplus\mathbf{3})=\mathbf{1}^{\oplus 2}\oplus\mathbf{1}'\oplus\mathbf{1}''\oplus\mathbf{3}_V\oplus\mathbf{3}_E.$$

The two invariant coordinates are q^2 and $-\|v\|^2/\sqrt{3}$. They are proportional for all labels, so only one invariant direction is occupied. The first copy of $\mathbf{3}$ is represented by qv. The second is represented by

$$Q(v)=\sqrt{2}\,(v_2v_3,v_3v_1,v_1v_2).$$

A direct recurrence for either parity seed gives $v_1 v_2 v_3 = 0$; the proper tetrahedral rotations are signed coordinate permutations, so the identity holds on each entire orbit. Hence

$$\langle v, Q(v)\rangle = 3\sqrt{2}\, v_1 v_2 v_3 = 0,$$

and Schur averaging cannot couple $\mathbf{3}_V$ to $\mathbf{3}_E$. This is the claimed decoupling of the two equivalent three-dimensional summands.

For a seed in orbit $\varepsilon \in \{+, -\}$, the squared block norms are

block	occupied $\mathbf{1}$	$\mathbf{1}' \oplus \mathbf{1}''$	$\mathbf{3}_V$	$\mathbf{3}_E$
dimension	1	2	3	3
seed norm squared	m	D_ε	$2q^2r^2$	E_ε

where $m = q^4 + r^4/3$ and $D_\varepsilon + E_\varepsilon = 2r^4/3$. This table immediately yields the spectrum in Theorem 4.1.

For completeness, the reduced minimum-error SDP can also be solved without suppressing the scalar step. Let $w \in [0, 1]$ be the multiplier of the $+$ seed constraint and write

$$E(w) = wE_+ + (1 - w)E_-, \qquad D(w) = \frac{2r^4}{3} - E(w).$$

On a multiplicity-free block of dimension d, minimizing a positive scalar dual block that dominates a rank-one component of squared norm e contributes $\sqrt{de}$ to the trace square root. The Schur-complement condition for all four occupied blocks therefore gives

$$\frac{1}{24} \max_{0 \le w \le 1} \left[\sqrt{m} + \sqrt{2D(w)} + \sqrt{6}\,|q|r + \sqrt{3E(w)}\right]^2. \tag{20}$$

The primal has strictly feasible seeds after adding an arbitrarily small multiple of the identity, so finite-dimensional SDP strong duality applies. Thus Equation (20) is not merely an upper bound. The only variable term is $\sqrt{2(T - E)} + \sqrt{3E}$ with $T = 2r^4/3$. Its stationary point is $E = 3T/5 = 2r^4/5$; projecting it onto the interval with endpoints E_+, E_- gives exactly the clamp in Theorem 4.2.

B Why one singlet parameter solves every parallel input

After A_4 symmetrization, an arbitrary two-qubit query marginal is

$$\rho_a = aP_{\text{sing}} + \frac{1 - a}{3} P_{\text{trip}}, \qquad 0 \le a \le 1.$$

This is exhaustive because the diagonal spin representation is $\mathbf{1} \oplus \mathbf{3}$, with both summands irreducible. Every ρ_a has a purification, so optimizing a includes arbitrary retained references.

The singlet is fixed by $W_\sigma \otimes W_\sigma$. On the triplet, the normalized invariant overlap is the spin-one character divided by three,

$$\eta(q) = \frac{\chi_1(W_\sigma)}{3} = \frac{4q^2 - 1}{3}, \qquad h = \eta(q)^2.$$

Consequently the occupied scalar block has squared norm $h + (1 - h)a$, while every non-scalar block in Equation (20) is multiplied by $1 - a$. Substitution gives, without an ansatz on the reference system,

$$P_{\text{par}}(a) = \frac{1}{24}\left[\sqrt{h + (1 - h)a} + K\sqrt{1 - a}\right]^2.$$

Put $u = \sqrt{(1-h)(1-a)}$. The bracket becomes

$$\sqrt{1-u^2} + \frac{K}{\sqrt{1-h}}u, \qquad 0 \le u \le \sqrt{1-h}.$$

It is the inner product of $(\sqrt{1-u^2}, u)$ with a fixed vector. The unconstrained maximizer lies inside the interval exactly when $hK^2 \le (1-h)^2$; otherwise the endpoint $a = 0$ is optimal. This proves both branches and the stated a_* in Theorem 5.1. At $q = 1/2$, $h = 0$ and the displayed expressions reduce algebraically to Equation (10).

C A concrete six-Kraus realization of the causal comb

The positivity and partial-trace equations in Equation (13) already imply physical realizability by the comb theorem. Here is an explicit dilation. Identify the stored reference with a copy of A_1, and use the output-major vectorization $|L\rangle\rangle \in A_2 \otimes (B_1 A_1)$. Then $J(\Lambda) = 2\Xi_*$ is the Choi matrix of a channel $\Lambda : B_1 A_1 \to A_2$. One exact Kraus family is

$$L_1 = \frac{1}{9}\begin{pmatrix} -4 & 0 & 0 & -5 \\ 0 & 0 & 1 & 0 \end{pmatrix}, \qquad L_2 = \frac{1}{9}\begin{pmatrix} 0 & -1 & 0 & 0 \\ 5 & 0 & 0 & 4 \end{pmatrix},$$

$$L_3 = \frac{2\sqrt{15}}{9}\begin{pmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \qquad L_4 = \frac{2\sqrt{15}}{9}\begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{pmatrix},$$

$$L_5 = \frac{2\sqrt{5}}{9}\begin{pmatrix} -1 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}, \qquad L_6 = \frac{2\sqrt{5}}{9}\begin{pmatrix} 0 & -1 & 0 & 0 \\ -1 & 0 & 0 & 1 \end{pmatrix}.$$

Direct multiplication gives

$$\sum_{\mu=1}^{6} L_\mu^* L_\mu = I_4, \qquad \sum_{\mu=1}^{6} |L_\mu\rangle\rangle\langle\langle L_\mu| = 2\Xi_*.$$

Thus a circuit can prepare a Bell purification of the first maximally mixed input, store its reference, and after receiving B_1 apply the Stinespring isometry

$$V = \sum_{\mu=1}^{6} L_\mu \otimes |\mu\rangle_M.$$

The A_2 output is sent through the second unknown channel and M is retained for the final measurement. The orbit-weighted rank-one POVM in the proof of Theorem 7.2 completes this explicit sequential circuit. Hence the attainment is not only an abstract feasible comb matrix.

D Collision completeness and replay details

For each of the 276 unordered label pairs, the word recurrence first computes $\|v_\sigma - v_\tau\|^2$ in $\mathbb{Q}[c,s]/(c^2 + 3s^2 - 1)$. After setting $z = s^2$, reduction by $c^2 = 1 - 3z$ leaves an expression $A(z) + cB(z)$. If $B = 0$, its roots are tested directly. Otherwise every physical root must divide

$$A(z)^2 - (1-3z)B(z)^2.$$

Factoring these resultants gives 17 distance types. Sturm isolation on $[0, 1/3]$, followed by substitution of the physical branch $c = \sqrt{1-3z}$, removes squared-resultant artifacts and leaves exactly the four rows of the collision table. Finally, $W_\sigma = -W_\tau$ is possible only when the common scalar q vanishes, which adds only the projective identifications at $z = 1/4$. This proves completeness, rather than merely discovering candidate collision points numerically.

The supplementary archive contains the manuscript source, manifest, and three fail-closed scripts:

```
python verification/verify_coherent_order_echo.py
python verification/verify_two_use_a4.py
python verification/verify_two_use_causal_qhalf.py
```

The first performs symbolic 24-word and 276-pair elimination, followed by numerical root screening and tolerant class grouping. The second uses NumPy to check the two-copy spectrum, clamp, and parallel fixture at $q = 1/2$. The third works in the exact field $\mathbb{Q}[t]/(f)$, verifies the Kraus/comb normalization, isolates the physical root, and reduces every numerator of $KCK - K$ to zero. Accordingly, only the third is advertised as a fully exact machine certificate; the first two are reproducible diagnostics supporting the analytic arguments above. None invokes an SDP solver or random sample.

AI assistance disclosure

AI tools assisted with symbolic exploration, literature triage, code drafting, and language editing. The author selected the claims and assumptions. The central causal identity is replayed exactly; the remaining scripts provide symbolic and numerical corroboration at the scopes stated above. AI output is not treated as independent evidence.

References

[1] Jessica Bavaresco, Mio Murao, and Marco Túlio Quintino. Unitary channel discrimination beyond group structures: Advantages of sequential and indefinite-causal-order strategies. *Journal of Mathematical Physics*, 63:042203, 2022.

[2] Yonina C. Eldar and G. David Forney, Jr. On quantum detection and the square-root measurement. *IEEE Transactions on Information Theory*, 47(3):858–872, 2001.

[3] Yonina C. Eldar, Alexandre Megretski, and George C. Verghese. Optimal detection of symmetric mixed quantum states. *IEEE Transactions on Information Theory*, 50(6):1198–1207, 2004.

[4] Hari Krovi, Saikat Guha, Zachary Dutton, and Marcus P. da Silva. Optimal measurements for symmetric quantum states with applications to optical communication. *Physical Review A*, 92:062333, 2015.

[5] Giulio Chiribella, Giacomo M. D'Ariano, and Paolo Perinotti. Memory effects in quantum channel discrimination. *Physical Review Letters*, 101:180501, 2008.

[6] Anna Jenčová and Martin Plávala. Conditions for optimal input states for discrimination of quantum channels. *Journal of Mathematical Physics*, 57:122203, 2016.

[7] David Gross, Koenraad Audenaert, and Jens Eisert. Evenly distributed unitaries: On the structure of unitary designs. *Journal of Mathematical Physics*, 48:052104, 2007.

[8] Horace P. Yuen, Robert S. Kennedy, and Melvin Lax. Optimum testing of multiple hypotheses in quantum detection theory. *IEEE Transactions on Information Theory*, 21(2):125–134, 1975.

[9] Giulio Chiribella, Giacomo M. D'Ariano, and Paolo Perinotti. Theoretical framework for quantum networks. *Physical Review A*, 80:022339, 2009.