

Partial Progress on Multi-slot Virtual Channel Conjugation

An AI-generated research note in Yuxuan (Vincent) Zhang's Agentic Research log

14 September 2026 · Version 1 · Partial progress; human review pending

Abstract

We give an exact signed simulation of the complex conjugate of an unknown memoryless quantum channel using repeated queries. For fixed input and output dimensions with $b > 1$, its quasiprobability cost approaches one at the optimal order $1/n$. We also give an exact finite-partition formula for the state-preparation subcase, a dimension-uniform two-query optimum for qubit output, and rational certificates showing strict adaptive advantage for three and four qubit-channel queries. These are partial resolutions of QIQCOP problem op_06e9f0c7b3b62f3b [1], not a general finite-query solution. Proofs and certificate verifiers accompany the note. Independent AI audits are disclosed; neither human peer review nor priority is claimed.

1. Model and main statements

Let $a = \dim A$, $b = \dim B$, and let $\mathcal{N}$ be an unknown CPTP map from A to B . A slot is one call to the same channel. Each physical branch is a deterministic, fixed-order n -slot quantum comb with arbitrary finite quantum memory. The target is complex conjugation in fixed bases, not the adjoint:

$$\mathcal{N}^*(X) = \overline{\mathcal{N}(\overline{X})}, \quad J_{\mathcal{N}^*} = J_{\mathcal{N}}^T.$$

The quantity $g_n(a, b)$ is the infimum of the sum of absolute real coefficients in a universal exact signed combination of such branches. Positive and negative coefficients can be grouped into weights $p, q \geq 0$, where $p - q = 1$ and the cost is $1 + 2q$. Parallel branches form a subclass. Unused calls may be discarded, so $g_{n+1} \leq g_n$. For $b = 1$, the unique channel is the trace map and the cost is one.

General bound. For $b > 1$, set $D = ab$ and $k = \min(n, D - 1)$. Then

$$1 + \frac{2(b-1)}{n+1} \leq g_n(a, b) \leq \min\left\{1 + a(b-1), 1 + \frac{2(D-1)k}{n+k}\right\}.$$

Consequently, at every fixed nontrivial pair of dimensions,

$$\inf_n g_n = \lim_{n \rightarrow \infty} g_n = 1, \quad g_n - 1 = \Theta(1/n).$$

Finite-query results. Two queries obey $g_2(a, 2) = a + 2/3$, and adaptivity does not help at two queries in any dimensions. For qubit channels, the certified intervals are

$2.25387912 < g_3 < 2.25389587$ and $2.00714780 < g_4 < 2.00715161$; both have strict advantage over parallel access. Sections 3-4 give stronger mixed-state lower bounds, and Section 5 specifies the exact rational endpoints.

2. A physical parallel-comb construction

Use unnormalized Choi matrices, $\text{Tr}_B J_{\mathcal{N}} = I_A$. Order registers as $P, I_1, O_1, \dots, I_n, O_n, F$, with a -dimensional input-type registers and b -dimensional output-type registers. Linking contracts the comb against the transposed Choi matrix in each slot. A positive operator C describes a deterministic parallel comb if

$$\text{Tr}_F C = R_{PI^n} \otimes I_{O^n}, \quad R \geq 0, \quad \text{Tr}_{I^n} R = I_P.$$

These conditions implement CPTP preprocessing and postprocessing with memory [6]. Define, with spectator identities understood,

$$V = \frac{1}{na^{n-1}} \sum_{i=1}^n F_{PI_i} F_{O_i F}, \quad W = \frac{I}{a^n b}.$$

Here F_{XY} is the swap. Every active swap pair transposes the inserted Choi matrix, while each inactive slot contributes $\text{Tr} J_{\mathcal{N}} = a$. Thus V implements the exact target. Its output partial trace equals $R_V \otimes I_{O^n}$, where R_V is the input swap sum divided by na^{n-1} and has input partial trace I_P . These are linear normalization conditions; V need not be positive. The operator W is the physical depolarizing parallel comb.

Group (P, F) and each (I_i, O_i) into D -dimensional registers. The numerator S of V becomes a star-swap sum. The Jucys-Murphy content bound [3, Appendix D, Lemma 4] gives

$$-kI \leq S \leq nI, \quad -\frac{Dk}{n}W \leq V \leq DW.$$

For completeness, Schur-Weyl duality permits at most D rows in a Young diagram. The last box has content at least $-(D-1)$ and at most n ; each of the n swaps is also bounded below by $-I$. This yields the stated lower endpoint.

Define two normalized operators

$$C_+ = \frac{nV + DkW}{n + Dk}, \quad C_- = \frac{DW - V}{D - 1}.$$

Both are positive by the spectral bounds. Both satisfy the linear parallel-comb normalization because their coefficients of V, W sum to one. Their preprocessing operators are positive by partial trace, so they are physical deterministic branches, not merely positive maps on Choi states. Finally,

$$q = \frac{(D-1)k}{n+k}, \quad p = 1 + q, \quad pC_+ - qC_- = V.$$

This proves the upper bound with cost $p + q$. The other upper bound discards all but one call and uses the known optimum $g_1 = 1 + a(b-1)$ [2, Theorem 4]. At finite n the target is exact; it is the signed cost, not channel error, that has a limit. For $n \geq D-1$, the constructed cost is

$1 + 2(D-1)^2/(n+D-1)$. The spectral endpoints and state-transposition construction are prior work [3]; the step checked here is their physical parallel-comb lift with channel normalization.

3. Lower bounds from replacement channels

Restrict the unknown channel to $\mathcal{N}_\rho(X) = \text{Tr}(X)\rho$ and fix the external input. Every adaptive branch becomes a CPTP map on $\rho^{\otimes n}$: each query discards its input and supplies a fresh copy, so all known circuit operations can be absorbed into a single map. Replacement-channel arguments already appear in [2]; here they bound signed multi-query cost.

For pure $\rho = \psi$, let $d_n = \binom{n+b-1}{n}$. The Haar-average fidelity of any physical branch with the conjugate pure state is bounded by $d_n/d_{n+1} = (n+1)/(n+b)$ [3, Theorem 1]. Indeed, after restricting its input to the symmetric subspace, its Choi matrix has trace d_n , while the Haar moment is the symmetric projector divided by d_{n+1} . Hence exact signed fidelity one implies

$$1 = pf_+ - qf_- \leq (1+q)\frac{n+1}{n+b}, \quad q \geq \frac{b-1}{n+1}.$$

This includes arbitrary adaptive memory. Combined with Section 2 it proves the fixed-dimension asymptotic statements, including strict cost improvement over one query for sufficiently large n .

Exact finite-partition formula for arbitrary mixed states

In the state-preparation subcase $a = 1$, the following formula evaluates the optimal cost using only integer partitions and rational arithmetic. Pad a partition λ of n with zeros to b entries. Define the Weyl dimension

$$D_b(\lambda) = \prod_{1 \leq i < j \leq b} \frac{\lambda_i - \lambda_j + j - i}{j - i}.$$

Row i is addable if $i = 1$ or $\lambda_{i-1} > \lambda_i$. With $(x)_+ = \max(x, 0)$, set

$$q_\lambda = \frac{1}{n} \sum_{i \text{ addable}} (i - 1 - \lambda_i) + \frac{D_b(\lambda + e_i)}{D_b(\lambda)}.$$

Then

$$g_n(1, b) = 1 + 2 \max_{\lambda \vdash n, \ell(\lambda) \leq b} q_\lambda, \quad g_n(a, b) \geq g_n(1, b).$$

The formula is exact but need not simplify to a single elementary expression. It optimizes over freely chosen CPTP branches, rather than a fixed white-noise line. That distinction and the signed-channel/diamond-norm framework are established in [5, Theorem 3 and Section 7.1]; symmetrized multicopy extensions are established in [4, Theorem 2]. The next section supplies the sector-resolved evaluation. Novelty of this specialization is unestablished.

4. Proof of the state formula and its consequences

Let register 0 be the output and let V be the average of the output-input swaps. Its Choi map sends $\rho^{\otimes n}$ to ρ^T and satisfies $\text{Tr}_0 V = I$. Group any signed implementation into positive matrices P, M with difference realizing transposition and partial traces pI, qI .

Average both branches over input permutations. Polarization makes their difference uniquely V : powers $X^{\otimes n}$ span the permutation-invariant operator space. Equality on density matrices extends by homogeneity to positive X , where the target is $(\text{Tr} X)^{n-1} X^T$, then by polynomial identity to all Hermitian X .

Next twirl by the common $U(b)$ action on all Choi registers. This preserves V and the branch constraints. The input sector is $W_\lambda \otimes [\lambda]$. The multiplicity-free Pieri rule decomposes the output factor times W_λ into sectors $W_{\lambda+e_i}$, and V acts on block (λ, i) by the scalar $(\lambda_i + 1 - i)/n$ [3]. If V_+, V_- are its spectral parts, positivity therefore forces $M > V_-$ **after this joint twirl**. This inequality is not asserted for arbitrary noncommuting positive decompositions.

The block projector obeys

$$\text{Tr}_0 \Pi_{\lambda, i} = \frac{D_b(\lambda + e_i)}{D_b(\lambda)} \Pi_\lambda, \quad A := \text{Tr}_0 V_- = \sum_\lambda q_\lambda \Pi_\lambda.$$

The coefficient follows from symmetry and the ratio of full traces; the symmetric-group multiplicity cancels. Thus $qI \geq A$, proving $q \geq q_* := \max q_\lambda$. This lower bound is attained by the common positive filler

$$H = \frac{I_0}{b} \otimes (q_* I - A), \quad P = V_+ + H, \quad M = V_- + H.$$

These matrices have difference V and partial traces $(1 + q_*)I, q_* I$, proving the formula. The supplement gives elementary reductions $g_n(1, 2) = 1 + 2/(n+1)$ and $g_n(1, 3) = 1 + 4/(n+1)$. The qubit formula is already a corollary of [3] and the fidelity bound; it is not a novelty claim.

For $b \geq 2$, a rectangular partition $\lambda = (m^r, 0^{b-r})$ with $n = mr$ and $1 \leq r < b$ gives $q_\lambda = r(b-r)/(n+r^2)$ by telescoping Weyl's product. For other n , pass to the next multiple using monotonicity. This gives

$$g_n(a, b) \geq 1 + 2 \max_{1 \leq r < b} \frac{r(b-r)}{r|n/r| + r^2}.$$

In particular,

$$\liminf_{n \rightarrow \infty} n(g_n - 1) \geq 2\lfloor b^2/4 \rfloor.$$

For example, the finite formula gives $g_{12}(1, 4) = 3/2$, witnessed by $(6, 6, 0, 0)$; the pure-state bound is only $19/13$. A state protocol cannot automatically be lifted to a channel comb without checking causality, so this calculation does not solve the general channel problem.

5. Two-query exactness and adaptive separations

Two-slot structure. For all a, b , the optimal two-slot adaptive and parallel costs coincide. To see this, twirl the three output-type registers of each branch by the same $U \in U(b)$. Moving these unitaries through the Choi link cancels the corresponding unitary in the conjugated target, so universality is preserved. Schur-Weyl duality gives

$$\text{Tr}_F C = X_{PI_1 I_2} \otimes I_{O_1 O_2} + Y_{PI_1 I_2} \otimes F_{O_1 O_2}.$$

Last-slot causality requires identity on O_2 . Its traceless projection yields $Y \otimes (F - I/b) = 0$, hence $Y = 0$ for $b > 1$. The remaining trace equations and positivity are precisely the parallel-comb conditions. The case $b = 1$ is trivial.

Qubit output. The supplied dimension-symbolic primal/dual certificate proves

$$g_2(a, 2) = g_2^{\text{parallel}}(a, 2) = a + \frac{2}{3} \quad (a \geq 1).$$

For $a \geq 3$ the full real twirled commutant has 21 dimensions: five symmetric and one skew input-type directions, and four symmetric and one skew output-type directions. Polynomial expansion through degree two gives exact operator constraints. The witnesses in `general_two_slot_proof.md` are rational functions of a , checked identically rather than by interpolation. Both primal branches are parallel. The dual uses actual universality and fixed-order causal equations; block trace pairings retain the inverse Gram factors of the nonorthonormal representation basis. Each block is a constant rational positive semidefinite matrix times a positive rational function for $a \geq 3$. The low-dimensional cases follow separately from the state formula and the exact $a = b = 2$ certificate.

Three and four qubit-channel slots. Rational primal/dual witnesses give

$$\frac{154885394229}{68719476736} \leq g_3(2, 2) \leq \frac{4840204505}{2147483648},$$

$$g_3^{\text{parallel}}(2, 2) \geq \frac{314540901607}{137438953472} > 2.28858626,$$

$$\frac{1077579273}{536870912} \leq g_4(2, 2) \leq \frac{538790657}{268435456},$$

$$g_4^{\text{parallel}}(2, 2) \geq \frac{2169915}{1048576} > 2.06939220.$$

The upper endpoints are below the corresponding parallel lower bounds, proving strict adaptive advantage. Also $g_4(2, 2) < g_3(2, 2)$. Together with the two-slot theorem, three is the first possible slot count for adaptive advantage in this problem. Neither interval is an exact closed form for its optimum.

6. What the computational certificates check

For qubit channels, twirling input-type and output-type registers independently, then averaging with entrywise conjugation, preserves the full fixed-order problem. On $m = n + 1$ qubits the permutation commutant has Catalan dimension. For $n = 4$ its real symmetric and skew parts have dimensions 26 and 16, giving $26^2 + 16^2 = 932$ real symmetric joint directions. Skew-times-skew directions are retained.

Recursive partial-trace equations impose fixed-order causality on each positive branch. Write

$J = I/2 + X$, with $\text{Tr}_O X = 0$. Universality means the constant and linear coefficients give J^T and every higher coefficient vanishes. Polarization turns this into complete operator identities. Slot permutations are symmetrized only in these identities for the difference; the physical branches are not forced to be slot-symmetric, which would remove adaptive strategies.

The four-slot system has 3,330 adaptive or 4,506 parallel equations in 1,864 branch coefficients. Saved matrix-entry frames are proved faithful by exact determinants, correct dimensions, and symmetric/skew ordering checks. Exact highest-weight reductions give nine positivity blocks of sizes 1, 4, 5, 4, 16, 20, 5, 20, 25. Rational elimination checks positivity. A dual witness proves a lower bound through exact stationarity and its pairing with positive branches. Integer products use explicit overflow bounds or arbitrary-size arithmetic. Verification uses no optimizer, numerical eigenvalue threshold, or saved floating-point seed.

Independent AI review and repairs before release

One AI reviewer independently rederived the general construction and replacement lower bound before reading the proof, then checked the symbolic two-slot proof. A second independently checked the state formula and its prior-art context. An adversarial reviewer tested the comb model and certificate implementation, including independent direct channel-polynomial contractions for four slots. Full reports and executable checks are in the source archive. These processes are not human peer review or formal proof-assistant verification.

The adversarial review found two pre-release verifier defects. The three-slot code converted saved rational dual equations to floating point when checking provenance: a tiny rational right-hand-side perturbation could pass that check and forge a lower bound. This was replaced by exact dyadic-lattice membership and augmented integer-row matching, with a full-basis independence check. The four-slot verifier accepted misordered symmetric/skew metadata; ordering and expected dimensions are now enforced. The original certificate rows passed independent exact provenance checks, so no endpoint above changed. The released checks include regressions rejecting these malformed inputs.

The four-slot polynomial cross-check exhaustively certifies the supplied primal's universal action through all 1,820 channel-monomial groups, including complex directions, independently of the main constraint generator; it does not sample channels. Full-model coverage and the all-dimension and all-query theorems additionally rest on the analytic arguments. The release records file hashes and executable test results. Historical local bundles should not be substituted for its repaired verifiers.

7. Scope, attribution, and reproducibility

The broad question remains open: no general formula for the optimum has been established for arbitrary finite n, a, b , and there is no proof that the leading coefficient $n(g_n - 1)$ has a limit for general channels. Even for $a = b = 2$, the proved asymptotic coefficient range is only

$2 \leq \liminf \leq \limsup \leq 18$. The results concern fixed-order combs, not indefinite causal order, and do not establish efficient circuit synthesis. Sampling variance scales with squared signed cost, but each run uses n calls; a lower g_n alone does not prove a reduction in total channel queries, whose standard scaling includes ng_n^2 .

The one-slot optimum is [2]. State-transposition fidelity, white-noise endpoints and the swap spectrum are [3]; multicopy symmetrization is [4]; the free signed-CPTP framework is [5]. The QIQOP page already records asymptotic state-preparation progress. The additional claims offered here for scrutiny are the general-channel parallel lift, the sector-resolved state evaluation, the two-slot statements, and the finite adaptive certificates. The literature search was limited; absence of a matching source is not evidence of priority.

The downloadable source archive contains this editable manuscript, its PDF builder, detailed proof supplements, rational certificates, verification programs, and independent AI audit reports. From its root, install the listed verification dependencies and run `verify_release.py` without Python optimization. No network, optimizer or numerical seed is required after dependencies are installed. Numerical discovery helpers are supplied for context but are not run by certificate verification. The README distinguishes exact certificate checks, exact finite enumeration, and floating-point sanity checks.

Provenance. This note and its proofs were generated by AI agents and checked by separate AI agents as described above. It is published as a claim with its working attached in the Agentic Research log, not as a human-refereed paper. External mathematical review and any catalog inclusion remain pending. The QIQOP problem should remain Unsolved.

References

- [1] QIQOP. Multi-slot overhead of virtual channel conjugation. Problem `op_06e9f0c7b3b62f3b`. Accessed 14 September 2026. https://qiqc-op.com/problem/op_06e9f0c7b3b62f3b/
- [2] C. Zhu, Z. Tang, G. Zhen, Y. Li, G. Bai and X. Wang. Simulation of Adjoints and Petz Recovery Maps for Unknown Quantum Channels. arXiv:2602.05828v1 (2026), Theorems 1 and 4. <https://arxiv.org/abs/2602.05828v1>
- [3] V. Brzić, D. Grinko, M. Studziński and M. T. Quintino. Optimal pure state cloning and transposition are complementary channels. arXiv:2603.23628v2 (2026), Theorems 1 and 3, Appendix D. <https://arxiv.org/abs/2603.23628v2>
- [4] Q. Dong, M. T. Quintino, A. Soeda and M. Murao. Implementing positive maps with multiple copies of an input state. Physical Review A 99, 052352 (2019), Theorem 2. <https://arxiv.org/abs/1808.05788v2>
- [5] B. Regula, R. Takagi and M. Gu. Operational applications of the diamond norm and related measures in quantifying the non-physicality of quantum maps. Quantum 5, 522 (2021), Theorem 3 and Section 7.1. <https://doi.org/10.22331/q-2021-08-09-522>
- [6] G. Chiribella, G. M. D'Ariano and P. Perinotti. Theoretical framework for quantum networks. Physical Review A 80, 022339 (2009). <https://doi.org/10.1103/PhysRevA.80.022339>