A permutation-orbit criterion for the advantage of limited entanglement assistance

Yuxuan Zhang[1,2]
[1]École Polytechnique Fédérale de Lausanne
[2]Princeton University

27 November 2026 · Version 1.0

Abstract

When does a quantum channel transmit more classical information with a limited supply of entanglement than by time sharing between the unassisted and fully assisted endpoints? We give a necessary and sufficient condition involving a single input state on finitely many channel uses and the entropy of its permutation-averaged output. The threshold uses regularized channel invariants. At the true threshold any strict finite-block witness is sound, with no finite-size correction. A uniform comparison with the ensemble support function costs only $O(\log n)$ bits at block length n and fixed output dimension.

1 The criterion

Let $\mathcal{N} : A \to B$ be a finite-dimensional memoryless quantum channel and $\mathcal{E}$ any Stinespring complement. Let $C(e)$ denote the classical capacity under a budget of e preshared ebits per use. The resource is initially maximally entangled, its log Schmidt rank counts towards the budget, and the decoding error tends to zero. We use the closed asymptotic rate region. The time-sharing benchmark joins $(0, C(0))$ to the *least-cost* point attaining the fully assisted capacity, and stays constant thereafter. Rates and entropies are in bits.

Write $\mathcal{T}_n(y) = \frac{1}{n!}\sum_{\pi\in S_n} P_\pi y P_\pi^\dagger$ for the average over permutations of the n output factors. For one arbitrary input state ρ on $A^{\otimes n}$ put

$$f_n(\rho,\lambda) = S(\mathcal{T}_n(\mathcal{N}^{\otimes n}\rho)) + (1-\lambda)S(\rho) - S(\mathcal{E}^{\otimes n}\rho), \tag{1}$$

$$L(\lambda) = \lim_{n\to\infty}\frac{1}{n}\max_\rho f_n(\rho,\lambda), \qquad 0 \le \lambda \le 1. \tag{2}$$

The complement entropy is unchanged by an isometric change of complement. Hence these quantities can be specified entirely by the channel. Define

$$C_0 = L(1), \qquad C_1 = L(0).$$

When $C_1 > C_0$, also define

$$h_* = \lim_{\lambda\downarrow 0}\frac{C_1 - L(\lambda)}{\lambda}, \qquad \lambda_* = \frac{C_1 - C_0}{h_*}. \tag{3}$$

Theorem 1. *The limits in (2) exist. When $C_1 > C_0$, the limit in (3) also exists. C_0 is the unassisted capacity, C_1 the fully assisted capacity, and h_* the minimum entanglement cost attaining C_1 in the closed region. If $C_1 = C_0$, time sharing is optimal at every cost. Otherwise $0 < \lambda_* \leq 1$, and strict improvement over endpoint time sharing at some cost occurs if and only if there exist a finite n and one density matrix ρ satisfying*

$$f_n(\rho, \lambda_*) > nC_0. \tag{4}$$

This supplies a regularized intrinsic criterion for the question posed by Wilde [3] in the formulation of the QIQCOP Zoo [4]. Testing (4) at specified channel constants uses entropies of explicit matrices. The threshold itself may require arbitrarily large block lengths. Deciding equality or certifying the numerical values of the regularized constants in finite time remains open.

2 Replacing the ensemble by the orbit of a single input

For an ensemble $\{p_i, \rho_i\}$ of n-use inputs write $\bar{\rho} = \sum_i p_i \rho_i$. The limited-entanglement coding theorem of Shor, with regularization over channel uses [1, 2], gives the support function $L_{\mathrm{op}}(\lambda) = \max_e [C(e) - \lambda e]$ as $\sup_n F_n(\lambda)/n$, where

$$F_n(\lambda) = \max_{\{p_i, \rho_i\}} \left[S(\mathcal{N}^{\otimes n} \bar{\rho}) + \sum_i p_i \big((1-\lambda) S(\rho_i) - S(\mathcal{E}^{\otimes n} \rho_i) \big) \right].$$

Indeed the ensemble rate in Shor's theorem has the displayed expression at $\lambda = 0$ and its cost is $\sum_i p_i S(\rho_i)$. Time sharing and closure preserve the supremum of a linear functional. Product ensembles give superadditivity of F_n, so the supremum is its normalized limit. Budgets may be restricted to $[0, \log_2 \dim A]$, since every ensemble point lies in that interval and full assistance is attainable there.

Set $O_n(\lambda) = \max_\rho f_n(\rho, \lambda)$. For fixed ρ, use the uniform ensemble of all permutations of this one input. The entropy of the mean output is $S(\mathcal{T}_n(\mathcal{N}^{\otimes n} \rho))$, and both individual entropies are invariant under permutation. This proves

$$O_n(\lambda) \leq F_n(\lambda). \tag{5}$$

It is the direction that makes a finite strict witness sound.

Here is a uniform upper bound in the other direction. In the Schur–Weyl decomposition $B^{\otimes n} = \bigoplus_j U_j \otimes V_j$, let U_j denote the irreducible unitary-group space and V_j the symmetric-group multiplicity factor for that action. Write $u_j = \dim U_j$, $v_j = \dim V_j$, a_n for the number of summands, and set

$$\Omega_n = \frac{1}{a_n} \bigoplus_j \frac{I_{U_j \otimes V_j}}{u_j v_j}, \qquad K_n = a_n \max_j u_j.$$

A permutation-invariant state has the form $\bigoplus_j p_j \sigma_j \otimes I_{V_j}/v_j$ and hence is bounded above by $K_n \Omega_n$. This is a universal symmetric state as in [6]. For fixed $d = \dim B$, $a_n \leq (n+1)^d$ and the Weyl dimension formula gives $u_j \leq (n+d)^{d(d-1)/2}$. Thus $\log K_n = O_d(\log(n+1))$.

For any state y the cross entropy against Ω_n is unchanged by the twirl. The domination just proved and operator monotonicity of the logarithm give

$$-\operatorname{Tr} y \log_2 \Omega_n = S(\mathcal{T}_n y) + D(\mathcal{T}_n y \| \Omega_n) \leq S(\mathcal{T}_n y) + \log_2 K_n.$$

For singular $\mathcal{T}_n y$ one can add a positive multiple of the identity on both sides before taking the logarithm and pass to the limit on its support. For an arbitrary ensemble in F_n bound the entropy of its average output by its cross entropy against Ω_n, use linearity, and apply this inequality to each member. Consequently

$$O_n(\lambda) \;\le\; F_n(\lambda) \;\le\; O_n(\lambda) + \log_2 K_n. \tag{6}$$

This proves (2) and $L = L_{\mathrm{op}}$. The difference in (6) vanishes per use, uniformly in the slope.

There is also a direct proof of the limiting assertion with no reference state. Fix a finite ensemble of n-use inputs of rational type p with a labels. In l blocks make one tensor input containing exactly lp_i copies of ρ_i. The uniform permutation of the l whole blocks produces a type-averaged output ω, whose single-block marginals equal $\bar{y} = \sum_i p_i \mathcal{N}^{\otimes n}\rho_i$. If P is the probability of the type under independent sampling with law p, then $\bar{y}^{\otimes l} \ge P\omega$. Since $P \ge (l+1)^{-a}$, the relative entropy identity gives

$$S(\omega) \ge lS(\bar{y}) - a\log_2(l+1).$$

The full twirl of nl uses can only increase this entropy. The individual input and complement entropies of the tensor input add exactly. Rational types approximate any finite ensemble by continuity, and O_n is superadditive because twirling a product over all slots includes twirling its two blocks separately. This yields the same equality of regularized supports.

3 Identifying the relevant slope

At $\lambda = 1$, refine each input of an ensemble into pure states. Concavity of complement entropy shows that the objective defining $F_n(1)$ can only increase. On a pure input the channel and complement output entropies coincide. Thus $F_n(1)$ equals the maximum unassisted Holevo information for the n-use channel, giving $C_0 = C(0)$ on regularization. It also bounds an entanglement-assisted rate by $C_0 + e$, so assistance of vanishing rate does not shift the zero-budget endpoint. The other endpoint $L(0)$ is the maximum achievable rate.

Suppose $C_1 > C_0$ and let h be the smallest cost at capacity in the compact closed rate region on the bounded cost interval. For a positive slope λ let (e_λ, R_λ) be an optimizer. Comparison with (h, C_1) gives

$$e_\lambda \le \frac{C_1 - L(\lambda)}{\lambda} \le h, \qquad C_1 - \lambda h \le R_\lambda \le C_1.$$

Every limit point as $\lambda \downarrow 0$ attains full capacity and costs at most h. It therefore has cost h. This proves (3) with $h_* = h$, and the inequality $C_1 \le C_0 + h$ implies $0 < \lambda_* \le 1$.

On $[0, h]$ the endpoint chord is exactly the line $C_0 + \lambda_* e$. Above h, $C(e)$ already equals its ceiling. Thus the operational curve lies strictly above the benchmark somewhere precisely when $L(\lambda_*) > C_0$. If (4) holds, (5) and $F_n/n \le L$ give this strict inequality immediately. Conversely, if $L(\lambda_*) > C_0$, convergence in (2) yields a finite block and an input attaining (4). If both endpoints coincide, monotonicity of $C(e)$ gives a constant curve. This proves Theorem 1.

4 Two controls

For the noiseless d-dimensional channel the support is $(2 - \lambda)\log_2 d$. Indeed the output entropy and the input entropy cost are each bounded by $\log_2 d$ per use, and the maximally mixed input attains these bounds. Thus $h_* = \log_2 d$ for $d > 1$ and $\lambda_* = 1$. The witness reduces to an output entropy larger than $n \log_2 d$, which is impossible.

For a positive control take the qubit channel $\mathcal{N}(\rho) = \frac{3}{4}\rho + \frac{1}{4}Z\rho Z$ and let $\mu = 1/10$. The exact dephasing trade-off of Brádler et al. [5, Thm. 4] gives $C_0 = 1$, $C_1 = 2 - h_2(1/4)$ and $h_* = 1$. For even n take one product input with half its factors $\mathrm{diag}(\mu, 1-\mu)$ and the rest reversed. Its input entropy per use is $h_2(\mu)$. The complement entropy per use is $h_2((1+\sqrt{1-3\mu(1-\mu)})/2)$. The output twirl is diagonal and uniform within each Hamming-weight class. If K is the sum of independent binomials with parameters $(n/2, \mu)$ and $(n/2, 1-\mu)$, its entropy is $H(K) + \mathbb{E}\log_2 \binom{n}{K}$. Substitution into (4) gives the following excess per use, rounded to six decimal places.

n	16	64	128	256
$f_n/n - C_0$	−0.013819	−0.000141	0.002039	0.003122

This illustrates the behaviour of one chosen family of witnesses. The dephasing advantage is already known.

Relation to prior work. Shor and Hsieh–Wilde supply the coding theorem; Brádler et al. compute explicit examples. Zhu, Zhuang and Shor exhibit superadditivity with limited assistance [7]. Permutations also occur in Shor's coding proof and in Wang and Hayashi's study of coding with prescribed shared states and encoding restrictions [8]. Here the orbit is an auxiliary ensemble on one arbitrary block input for the initial-ebit trade-off. The contribution is the single-input finite strict test at the regularized slope and the uniform estimate (6). Internal proof review and computational checks are documented in the supplement. Independent specialist review and priority remain open.

References

[1] P. W. Shor, *The classical capacity achievable by a quantum channel assisted by a limited entanglement*, Quantum Inf. Comput. **4** (2004), 537–545. Paper.

[2] M.-H. Hsieh and M. M. Wilde, *Trading classical communication, quantum communication, and entanglement in quantum Shannon theory*, IEEE Trans. Inf. Theory **56** (2010), 4705–4730. arXiv:0901.3038.

[3] M. M. Wilde, *Quantum Information Theory*, 2nd ed., Cambridge (2017), Sec. 22.5. Book.

[4] QIQCOP Zoo, *Structural criterion for classical–entanglement trade-off advantage*. Problem statement.

[5] K. Brádler, P. Hayden, D. Touchette and M. M. Wilde, *Trade-off capacities of the quantum Hadamard channels*, Phys. Rev. A **81** (2010), 062312. arXiv:1001.1732.

[6] M. Mosonyi and T. Ogawa, *Strong converse exponent for classical–quantum channel coding*, Commun. Math. Phys. **355** (2017), 373–426. arXiv:1409.3562.

[7] E. Y. Zhu, Q. Zhuang and P. W. Shor, *Superadditivity of the classical capacity with limited entanglement assistance*, Phys. Rev. Lett. **119** (2017), 040503. arXiv:1704.06955.

[8] K. Wang and M. Hayashi, *Permutation enhances classical communication assisted by entangled states*, IEEE Trans. Inf. Theory (2021). arXiv:2001.01845.