

Exact Soundness of the Multipartite Product Test

July 13, 2026

Abstract

The product test is the optimal two-copy test with perfect completeness for multipartite pure product states. Let $\text{PT}_n(\omega)$ be its largest acceptance probability over n -partite pure states whose squared overlap with the closest product state is ω , with finite local dimensions allowed to vary. Soleimanifar and Wright determined $\text{PT}_n(\omega)$ for $\omega \geq 1/2$ and asked whether $\text{PT}_n(\omega) \rightarrow 1/2$ as $\omega \rightarrow 0$. Lovitz and Lowe recently determined the bipartite curve and asked whether it remains extremal in the multipartite low-overlap regime. We answer both questions affirmatively. If $k = \lfloor 1/\omega \rfloor$, then for every $n \geq 2$ and $0 < \omega \leq 1$,

$$\text{PT}_n(\omega) = \frac{1}{2} (1 + k\omega^2 + (1 - k\omega)^2).$$

Consequently, the soundness tends to $1/2$ as ω tends to zero, and an extremal state can always be chosen as a bipartite Schmidt state tensored with product factors. Thus the formula also gives the exact soundness curve of the best two-copy, perfect-completeness test. The proof is a one-party Schmidt-decomposition induction, inspired by the SVD analysis of the swap norm by Bhangale, Khot, Liu, and Minzer. A sharp bound on off-diagonal Schmidt terms reduces the induction to an elementary capped-purity optimization, whose branches combine by concatenating probability vectors.

Note. This paper was fully AI generated by GPT-5 and the Codex harness, prompted by Yang Liu. The proof blueprint supplied to Codex was prepared by GPT 5.6 Sol.

1 Introduction

Let

$$|\psi\rangle \in \mathcal{H}_1 \otimes \cdots \otimes \mathcal{H}_n$$

be a multipartite pure state. The product test takes two copies of $|\psi\rangle$, performs a swap test between the two copies of each local system, and accepts if every local test accepts. If F_i swaps the two copies of $\mathcal{H}_i$, its acceptance projector is

$$Q_n = \bigotimes_{i=1}^n \frac{1 + F_i}{2}. \quad (1)$$

Every product state is accepted with probability one. The parameter that governs soundness is the squared closest-product overlap

$$\omega(\psi) = \max_{\substack{|\phi_i\rangle \in \mathcal{H}_i \\ \|\phi_i\|=1}} |\langle \phi_1 \otimes \cdots \otimes \phi_n | \psi \rangle|^2. \quad (2)$$

Its negative logarithm is one standard convention for the geometric measure of entanglement [WG03].

The two-register swap test already appears in quantum fingerprinting [BCWdW01], and closely related products of local symmetrizers occur in observable constructions for multipartite concurrence [MKB05]. Harrow and Montanaro [HM13] established dimension- and party-independent soundness bounds for the product test and proved a pointwise optimality statement: among all tests on two copies that accept every pure product state with certainty, the product test rejects every input as often as possible. Their estimates show that the acceptance probability is $1 - \Theta(1 - \omega)$, but the constants in the regime of small ω were left open. Soleimanifar and Wright [SW22] formalized the acceptance probability as an extremal function of the overlap. We make the dimension convention in their dimension-suppressed notation explicit by defining

$$\text{PT}_n(\omega) = \sup_{\substack{d_1, \dots, d_n \in \mathbb{N}, \\ |\psi\rangle \in \bigotimes_{i=1}^n \mathbb{C}^{d_i}, \\ \|\psi\|=1, \omega(\psi)=\omega}} \langle \psi |^{\otimes 2} Q_n | \psi \rangle^{\otimes 2}, \quad (3)$$

where the supremum ranges over all finite local dimensions. They proved

$$\text{PT}_n(\omega) = 1 - \omega + \omega^2 \quad (n \geq 2, \omega \geq 1/2),$$

using a recursive Schmidt-decomposition analysis. For $\omega < 1/2$ they obtained the upper bound $2/3 + \omega^2/3$, left the exact value open, and asked in particular whether $\text{PT}_n(\omega)$ tends to $1/2$ as $\omega \rightarrow 0$.

Lovitz and Lowe [LL26, Corollary 6.3] subsequently computed the exact bipartite curve. In the notation below, their value is

$$\frac{1}{2} (1 + k\omega^2 + (1 - k\omega)^2), \quad k = \left\lfloor \frac{1}{\omega} \right\rfloor. \quad (4)$$

Padding their bipartite optimizer with product factors gives the same expression as a lower bound for every $n \geq 2$. They asked explicitly whether this bound is tight for $n > 2$ and $\omega < 1/2$ [LL26, Remark 6.4]. Our main result proves the matching multipartite upper bound.

Theorem 1.1. *For every $n \geq 2$ and $0 < \omega \leq 1$, let*

$$k = \left\lfloor \frac{1}{\omega} \right\rfloor, \quad r = 1 - k\omega.$$

Then

$$\text{PT}_n(\omega) = \frac{1}{2} (1 + k\omega^2 + r^2). \quad (5)$$

More strongly, for every fixed finite choice of local Hilbert spaces and every state $|\psi\rangle$ on them,

$$\langle \psi |^{\otimes 2} Q_n | \psi \rangle^{\otimes 2} \leq \frac{1}{2} (1 + k\omega(\psi)^2 + (1 - k\omega(\psi))^2), \quad k = \left\lfloor \frac{1}{\omega(\psi)} \right\rfloor. \quad (6)$$

If two local dimensions are at least $\lceil 1/\omega \rceil$, equality is attained by a state entangled only across those two systems.

The proof sharpens the existing Schmidt recursion in two ways. Decompose $|\psi\rangle$ across one party versus the rest, and retain every Schmidt branch. The product-test acceptance has an exact recurrence in which an off-diagonal pair of orthogonal Schmidt vectors contributes at most $1/2$. The diagonal branches are controlled inductively. Their product overlaps need not be bounded individually by ω ; instead, if λ_a is a Schmidt weight and ω_a is the product overlap inside its branch, then $\lambda_a \omega_a \leq \omega$. Each branch produces a probability vector capped by ω_a . Scaling it by λ_a and concatenating all the vectors gives a single probability distribution capped by ω . Maximizing the purity of such a distribution yields exactly the piecewise-quadratic expression in (5).

The organization of this proof was suggested by the swap-norm viewpoint of Bhangale, Khot, Liu, and Minzer [BKLM24]. Their coordinatewise swap form is the same quartic expression as the product-test acceptance after a function is normalized as a quantum amplitude tensor, and their J -SVD is the corresponding Schmidt decomposition. Their high-swap argument peels off one SVD coordinate at a time, while their general inverse theorem uses random restrictions and common refinements of several SVDs. We use none of that deeper machinery: the positive projector in the quantum test gives the sharp off-diagonal constant $1/2$, and capped-purity concatenation closes the recursion directly. Recent work has also related multi-copy product tests to convergent hierarchies for the geometric measure of entanglement [WRG⁺26]; the present result concerns the exact extremal behavior of the original two-copy test.

Section 2 records the swap-form identity and the scalar capped-purity lemma. Section 3 proves the exact upper and lower bounds. Section 4 gives the limiting, structural, and testing consequences.

2 Preliminaries

2.1 The product test and the dimension convention

All Hilbert spaces in this paper are finite-dimensional. On two copies of $\mathcal{H}_1 \otimes \cdots \otimes \mathcal{H}_n$, let F_i swap the two copies of the i th local system. The operators F_i commute, so Q_n in (1) is an orthogonal projector. We write

$$p_n(\psi) = \langle \psi |^{\otimes 2} Q_n | \psi \rangle^{\otimes 2}.$$

The maximum in (2) is attained by compactness. A nonzero finite-dimensional pure state always has positive product overlap: some nonzero coefficient in any product basis supplies a product vector with nonzero overlap. Thus exact statements use $\omega > 0$; the value at $\omega = 0$ is only a limit.

We follow the convention of Soleimanifar and Wright in (3): local dimensions are not parameters of PT_n . The upper bound in Theorem 1.1 is statewise and hence holds for fixed dimensions. The matching construction may require larger spaces, specifically two local dimensions at least $\lceil 1/\omega \rceil$.

2.2 The classical swap form

We spell out the normalization behind the connection to [BKLM24]. Let Σ be a finite set, let $f : \Sigma^n \rightarrow \mathbb{C}$ be nonzero, and use the uniform norm

$$\|f\|_{2,\text{unif}}^2 = \mathbb{E}_{x \in \Sigma^n} |f(x)|^2.$$

Associate to f the normalized amplitude state

$$|\psi_f\rangle = \frac{1}{|\Sigma|^{n/2} \|f\|_{2,\text{unif}}} \sum_{x \in \Sigma^n} f(x) |x_1\rangle \cdots |x_n\rangle. \quad (7)$$

For independent uniform $x, y \in \Sigma^n$, independently preserve or swap the pair (x_i, y_i) in each coordinate, and denote the resulting pair by (x', y') . The BKLM swap form is

$$\text{swap}(f) = \mathbb{E}_{x, y, (x', y')} \left[f(x) f(y) \overline{f(x')} \overline{f(y')} \right]. \quad (8)$$

Expanding $Q_n = 2^{-n} \sum_{S \subseteq [n]} F_S$, where F_S swaps exactly the coordinates in S , gives

$$p_n(\psi_f) = \frac{\text{swap}(f)}{\|f\|_{2,\text{unif}}^4}. \quad (9)$$

Indeed, the summand for S is the normalized average of $\overline{f(x)f(y)} f(x') f(y')$; the coordinate-swap involution exchanges this expression with the one in (8). This also makes clear that complex conjugations and the uniform normalization are essential.

Under (7), a product function becomes a product-state amplitude tensor after its factors are normalized. Similarly, a matrix SVD of f across coordinates J and $[n] \setminus J$ is, up to the harmless uniform-versus-counting normalization and the usual conjugation of the right singular vector, a Schmidt decomposition across the same bipartition. Identity (9) is motivational; the proof below is stated entirely in quantum language.

2.3 Capped purity

The scalar function controlling the soundness curve is

$$h(t) = \max \left\{ \sum_j p_j^2 : p_j \geq 0, \sum_j p_j = 1, \max_j p_j \leq t \right\}, \quad 0 < t \leq 1, \quad (10)$$

where arbitrary finite probability vectors and trailing zeroes are allowed.

Lemma 2.1 (Capped purity). *If $k = \lfloor 1/t \rfloor$ and $s = 1 - kt$, then*

$$h(t) = kt^2 + s^2. \quad (11)$$

An optimizer has k entries equal to t , one entry equal to s when $s > 0$, and all remaining entries zero.

Proof. Suppose two positive coordinates $p_i \geq p_j$ are both strictly below t . Put $\delta = \min\{t - p_i, p_j\}$ and replace them by $p_i + \delta$ and $p_j - \delta$. The cap and total mass are preserved, while the change in purity is

$$(p_i + \delta)^2 + (p_j - \delta)^2 - p_i^2 - p_j^2 = 2\delta(p_i - p_j) + 2\delta^2 > 0.$$

Repeating this transfer leaves at most one positive coordinate below the cap. The unit total mass then forces exactly k capped entries and the residual $s = 1 - kt$. If $t = 1/m$, then $k = m$ and $s = 0$; this is the reciprocal-integer boundary. The resulting vector is feasible and has the purity in (11). $\square$

3 Exact soundness

We first isolate the exact recurrence and the off-diagonal estimate.

Lemma 3.1 (One-party Schmidt recurrence). *Let $n \geq 2$ and Schmidt-decompose the first party from the rest as*

$$|\psi\rangle = \sum_a \sqrt{\lambda_a} |a\rangle |v_a\rangle, \quad (12)$$

where the positive λ_a sum to one and the $|v_a\rangle$ are orthonormal. Define

$$Q_R = \bigotimes_{i=2}^n \frac{\mathbf{I} + F_i}{2}, \quad q_{ab} = \langle v_a | \langle v_b | Q_R | v_a \rangle | v_b \rangle.$$

Then

$$p_n(\psi) = \sum_{a,b} \lambda_a \lambda_b q_{ab}. \quad (13)$$

Moreover,

$$q_{aa} = p_{n-1}(v_a), \quad q_{ab} \leq \frac{1}{2} \quad (a \neq b). \quad (14)$$

Proof. Let $F_R = \prod_{i=2}^n F_i$ be the global swap of the two copies of all remaining parties. Since Q_R projects onto the simultaneous +1 eigenspaces of the local swaps,

$$F_R Q_R = Q_R, \quad Q_R \preceq \frac{\mathbf{I} + F_R}{2}. \quad (15)$$

The second relation is the corresponding containment of projector ranges.

Expanding the first local symmetric projector in $Q_n = ((\mathbf{I} + F_1)/2)Q_R$, the identity term contributes

$$\frac{1}{2} \sum_{a,b} \lambda_a \lambda_b q_{ab}.$$

The first-party swap term contributes

$$\frac{1}{2} \sum_{a,b} \lambda_a \lambda_b \langle v_b | \langle v_a | Q_R | v_a \rangle | v_b \rangle.$$

But $\langle v_b | \langle v_a | = \langle v_a | \langle v_b | F_R$, so the first relation in (15) makes this cross matrix element equal to q_{ab} . This proves (13).

The diagonal identity in (14) is the definition of the product test on the remaining parties. For $a \neq b$, the second relation in (15) and orthogonality give

$$\begin{aligned} q_{ab} &\leq \langle v_a | \langle v_b | \frac{1 + F_R}{2} | v_a \rangle | v_b \rangle \\ &= \frac{1}{2} (1 + |\langle v_a | v_b \rangle|^2) = \frac{1}{2}. \end{aligned}$$

□

Proposition 3.2 (Statewise upper bound). *For every $n \geq 1$ and every n -partite pure state $|\psi\rangle$,*

$$p_n(\psi) \leq \frac{1 + h(\omega(\psi))}{2}. \quad (16)$$

Proof. We use induction on n . For $n = 1$, every pure state is a product state, so $p_1(\psi) = \omega(\psi) = h(1) = 1$.

Let $n \geq 2$, use (12), and for each branch write

$$\omega_a = \max_{|\pi\rangle \text{ product on parties } 2, \dots, n} |\langle \pi | v_a \rangle|^2.$$

Choosing $|a\rangle$ on the first party and a closest product vector to $|v_a\rangle$ on the other parties shows that

$$\lambda_a \omega_a \leq \omega(\psi) =: \omega \quad \text{for every } a. \quad (17)$$

By Lemma 3.1, its off-diagonal bound, and the inductive hypothesis,

$$\begin{aligned} p_n(\psi) &\leq \sum_a \lambda_a^2 \frac{1 + h(\omega_a)}{2} + \frac{1}{2} \sum_{a \neq b} \lambda_a \lambda_b \\ &= \frac{1}{2} + \frac{1}{2} \sum_a \lambda_a^2 h(\omega_a). \end{aligned} \quad (18)$$

It remains to compare the last sum with $h(\omega)$. For every a , let $q^{(a)}$ be a finite probability vector attaining $h(\omega_a)$ in Lemma 2.1, and concatenate the scaled vectors by setting

$$p_{a,j} = \lambda_a q_j^{(a)}. \quad (19)$$

The $p_{a,j}$ form a probability distribution. By (17), every atom satisfies

$$p_{a,j} \leq \lambda_a \omega_a \leq \omega.$$

Its purity is exactly

$$\sum_{a,j} p_{a,j}^2 = \sum_a \lambda_a^2 \sum_j (q_j^{(a)})^2 = \sum_a \lambda_a^2 h(\omega_a).$$

The definition of $h(\omega)$ therefore bounds this quantity by $h(\omega)$. Substitution in (18) proves (16). □

The concatenation in (19) is the closure step that is lost if one keeps only the leading Schmidt branch. It uses the full family of constraints $\lambda_a \omega_a \leq \omega$ without requiring either λ_a or ω_a separately to be at most ω .

Proposition 3.3 (Bipartite extremizer). *Let $n \geq 2$ and $0 < \omega \leq 1$. Put $k = \lfloor 1/\omega \rfloor$ and $r = 1 - k\omega$. If two local dimensions are at least $\lceil 1/\omega \rceil$, there is a state $|\Psi_\omega\rangle$ with*

$$\omega(\Psi_\omega) = \omega, \quad p_n(\Psi_\omega) = \frac{1}{2}(1 + k\omega^2 + r^2). \quad (20)$$

Proof. Take Schmidt weights

$$p_1 = \cdots = p_k = \omega, \quad p_{k+1} = r \quad \text{if } r > 0,$$

and on two parties define

$$|\Psi_\omega^{(2)}\rangle = \sum_j \sqrt{p_j} |j\rangle |j\rangle.$$

The required Schmidt rank is k when $r = 0$ and $k + 1$ otherwise, namely $\lceil 1/\omega \rceil$. The closest-product overlap of a bipartite pure state is its largest Schmidt weight, so it is exactly ω .

For completeness, if ρ_A and ρ_B are the two marginals, expansion of the four terms in $Q_2 = ((1 + F_A)/2)((1 + F_B)/2)$ gives

$$p_2(\Psi_\omega^{(2)}) = \frac{1}{4} (1 + \text{Tr } \rho_A^2 + \text{Tr } \rho_B^2 + 1) = \frac{1}{2} \left(1 + \sum_j p_j^2 \right).$$

The final sum is $k\omega^2 + r^2$. Tensoring arbitrary pure product states on the other $n - 2$ parties preserves both the closest-product overlap and the acceptance probability. $\square$

Proposition 3.2, Lemma 2.1, and Proposition 3.3 prove Theorem 1.1.

4 Consequences and discussion

The high-overlap regime and reciprocal breakpoints. For $\omega > 1/2$, one has $k = 1$, and (5) becomes

$$\text{PT}_n(\omega) = 1 - \omega + \omega^2.$$

At $\omega = 1/2$, the floor convention gives $k = 2$, but the value is still $3/4$, so the same formula holds on the closed interval $[1/2, 1]$. More generally, the curve is continuous and quadratic between the breakpoints $\omega = 1/m$. At such a breakpoint, the residual vanishes; equivalently, it may be viewed as one additional capped atom, so the two neighboring quadratic descriptions agree.

The low-overlap limit. Since $k\omega^2 \leq \omega$ and $0 \leq r < \omega$, Lemma 2.1 gives

$$0 \leq h(\omega) \leq \omega + \omega^2.$$

Consequently,

$$\lim_{\omega \rightarrow 0} \text{PT}_n(\omega) = \frac{1}{2},$$

answering the question of Soleimanifar–Wright and the multipartite tightness question of Lovitz–Lowe. The formula is independent of the number of parties.

Bipartite extremality. For every prescribed ω , an optimizer in the dimension-free problem can be chosen to have entanglement on only two parties. Thus genuinely multipartite entanglement provides no strict advantage in maximizing false acceptance at fixed closest-product overlap. This is an existence statement about an optimizer; the proof does not classify every equality case. A natural next question is a stability version: whether every near-extremizer must be close, in an appropriate sense, to a state with a similarly capped Schmidt spectrum.

Optimality among two-copy tests. Let M be the accepting POVM element of any two-copy test having perfect completeness for all pure product states. The vectors

$$\{(|\phi_1\rangle \otimes \cdots \otimes |\phi_n\rangle)^{\otimes 2}\}$$

span $\bigotimes_i \text{Sym}^2(\mathcal{H}_i)$, whose projector is exactly Q_n . Perfect completeness and $0 \preceq M \preceq I$ imply $M \succeq Q_n$; this is the argument of Harrow–Montanaro [HM13, Section 5]. Hence every such test accepts every input at least as often as the product test. Theorem 1.1 therefore gives the exact best possible two-copy, perfect-completeness soundness curve, not only the performance of one particular measurement.

Higher-copy perfect-completeness tests project each local collection of copies onto its symmetric subspace [HM13], and recent hierarchies study their relation to closest-product overlap [WRG+26]. Determining comparably sharp finite-copy curves beyond two copies remains an interesting direction.

References

- [BCWdW01] Harry Buhrman, Richard Cleve, John Watrous, and Ronald de Wolf. Quantum fingerprinting. *Physical Review Letters*, 87(16):167902, 2001.
- [BKLM24] Amey Bhangale, Subhash Khot, Yang P. Liu, and Dor Minzer. On approximability of satisfiable k -CSPs: VI, 2024.
- [HM13] Aram W. Harrow and Ashley Montanaro. Testing product states, quantum Merlin–Arthur games and tensor optimization. *Journal of the ACM*, 60(1):1–43, 2013.
- [LL26] Benjamin Lovitz and Angus Lowe. Nearly tight bounds for testing tree tensor network states. *IEEE Transactions on Information Theory*, 72(5):3074–3097, May 2026.

- [MKB05] Florian Mintert, Marek Kuś, and Andreas Buchleitner. Concurrence of mixed multipartite quantum states. *Physical Review Letters*, 95(26):260502, 2005.
- [SW22] Mehdi Soleimanifar and John Wright. Testing matrix product states. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1679–1701. Society for Industrial and Applied Mathematics, 2022.
- [WG03] Tzu-Chieh Wei and Paul M. Goldbart. Geometric measure of entanglement and applications to bipartite and multipartite quantum states. *Physical Review A*, 68(4):042307, 2003.
- [WRG⁺26] Lisa T. Weinbrenner, Albert Rico, Kenneth Goodenough, Xiao-Dong Yu, and Otfried Gühne. Complete hierarchies for the geometric measure of entanglement, 2026. arXiv:2601.23243v2.