The One-Step Subspace in the Quantum Wielandt Inequality

2026-08-22

Verdict: scope-restriction (open).

1 The assertion

The quantum Wielandt inequality of Sanz–Pérez-García–Wolf–Cirac is stated in [SPGWC10, Theorem 1].[1] Let $\mathcal{E}_A$ be a quantum channel on $M_D(\mathbb{C})$ with Kraus operators $A_1, \ldots, A_a$. For each n, set

$$S_n(A) = \operatorname{span}\{A_{i_1} \cdots A_{i_n} : 1 \leq i_1, \ldots, i_n \leq a\}.$$

Thus $S_n(A)$ is the subspace generated by words of length exactly n. Let

$$i(A) = \min\{n \geq 1 : S_n(A) = M_D(\mathbb{C})\},$$

when this minimum exists. Likewise, the paper's positive word-length convention gives the primitivity index

$$q(\mathcal{E}_A) = \min\{q \geq 1 : H_q(A, \varphi) = \mathbb{C}^D \text{ for every } \varphi \neq 0\}.$$

The paper writes d for the number of linearly independent Kraus operators. Equivalently, after replacing the displayed Kraus list by a minimal spanning list, one has

$$d = \dim S_1(A).$$

With this convention, Theorem 1 asserts that a primitive channel satisfies

$$i(A) \leq (D^2 - d + 1)D^2.$$

The same theorem also gives two better estimates under additional hypotheses on the one-step subspace. If $S_1(A)$ contains an invertible matrix, then

$$i(A) \leq D^2 - d + 1.$$

[1] For traceability, this note corresponds to issue #1049. The local source passage is `Papers/0909.5347/main.tex:360--365`, `Papers/0909.5347/main.tex:572--590`, and `Papers/0909.5347/main.tex:645--684`.

If $S_1(A)$ contains a non-invertible matrix with a nonzero eigenvalue, then

$$i(A) \leq D^2.$$

These are hypotheses about arbitrary matrices in $S_1(A)$. The paper does not require the special matrix to be one of the Kraus operators in a chosen Kraus representation.

2 What has been formalized

The formal statements use the same subspaces $S_n(A)$ and define the Kraus rank by

$$\mathrm{kr}(A) = \dim S_1(A).$$

This is the paper's number d written invariantly. Replacing d by $\mathrm{kr}(A)$ is therefore not a change in the mathematical bound. [2]

The general inequality is formalized for the same index $i(A)$:

$$i(A) \leq (D^2 - \mathrm{kr}(A) + 1)D^2.$$

The proof passes through the usual eventual full-rank formulation and a blocking argument, but the conclusion is the equality $S_n(A) = M_D(\mathbb{C})$ for one fixed length n. It is not merely a statement about the increasing sum $S_0(A) + \cdots + S_n(A)$.[3]

The two improved estimates now use the paper's one-step subspace hypotheses. For the invertible case, the formal statement assumes a matrix $X \in S_1(A)$ and X invertible; under that assumption it proves

$$S_{D^2 - \mathrm{kr}(A)+1}(A) = M_D(\mathbb{C}).$$

For the non-invertible case, the formal statement assumes a matrix $X \in S_1(A)$, X non-invertible, and a nonzero eigenvalue of X; under that assumption it proves

$$S_{D^2}(A) = M_D(\mathbb{C}).$$

These are the hypotheses in [SPGWC10, Theorem 1]. The older chosen-Kraus statements are still present, but only as corollaries obtained by taking $X = A_{i_0}$.[4]

[2] The formal definitions are `MPSTensor.wordSpan` in `TNLean/Wielandt/SpanGrowth/CumulativeSpan.lean:57--64` and `MPSTensor.krausRank` in `TNLean/Wielandt/Primitivity/Definitions.lean:84--90`. The local source identifies d as the number of linearly independent Kraus operators in `Papers/0909.5347/main.tex:360--365`; the proof of Lemma 1 uses $\dim T_1(A) = d$ in `Papers/0909.5347/main.tex:579--583`.

[3] The general theorem is `MPSTensor.iIndex_le_general_of_isPrimitivePaper` in `TNLean/Wielandt/Inequality/Bounds.lean:389--468`. The formal equivalence between the primitive condition used by Sanz–Pérez-García–Wolf–Cirac and eventual full Kraus rank under normalization is in `TNLean/Wielandt/Primitivity/Equivalence.lean:169--191`.

[4] The relevant declarations are `MPSTensor.wordSpan_eq_top_of_isPrimitivePaper_of_mem_wordSpan_one_of_isUnit` and `MPSTensor.iIndex_le_of_mem_wordSpan_one_of_isUnit` for the invertible case, and `MPSTensor.wordSpan_eq_top_of_isPrimitivePaper_of_mem_wordSpan_one_of_noninvertible_eigenvector` and `MPSTensor.iIndex_le_sq_of_mem_wordSpan_one_of_noninvertible_eigenvector` for the non-invertible case, all in `TNLean/Wielandt/Inequality/Bounds.lean`. The shorter chosen-generator statements in the same file are corollaries.

3 Resolved one-step-subspace issue

The former distinction was small but genuine. Suppose first that the paper gives an invertible matrix $X \in S_1(A)$. Since X is a linear combination of the A_i, multiplication by X sends $S_n(A)$ into $S_{n+1}(A)$. Since X is invertible, this multiplication is injective on $M_D(\mathbb{C})$. This is the linear-algebra reason why the dimensions of the spaces $S_n(A)$ cannot stagnate for too long; it is the dimension-growth step used in the proof of the sharper bound.

The formal proof now uses this one-step-subspace version. It adjoins the chosen $X \in S_1(A)$ as a redundant generator and proves that the exact word spans and the Kraus rank are unchanged. The span equality is `MPSTensor.wordSpan_oneStepAugment_eq`, and the Kraus-rank equality is `MPSTensor.krausRank_oneStepAugment`, both in `TNLean/Wielandt/SpanGrowth/InvertibleWordSpan.lean`. This reduces the paper's hypothesis to the single-generator argument without changing the channel invariant being bounded.

The same issue appears in the non-invertible case. The source hypothesis gives a matrix $X \in S_1(A)$ which is singular but has a nonzero eigenvalue. The formal proof again uses the redundant-generator construction, so the theorem is stated for an arbitrary such X, not only for a distinguished Kraus matrix.

4 Primitive-to-normal theorem with a chosen positive fixed point

The theorem `MPSTensor.isNormal_of_isPrimitiveMPS_with_posDef`[5] proves normality from a chosen fixed point ρ satisfying the local complementary transfer-map gap predicate `IsPrimitiveMPS` and the additional hypothesis $\rho > 0$.

This is the formal argument for Proposition 3(c)⇒(b) of [SPGWC10]: the proof first views the hypotheses as strong irreducibility with a positive-definite fixed point, then applies the eventual-full-Kraus-rank theorem. The positive definiteness is used in the trace-pairing argument $B \mapsto \operatorname{Re}\operatorname{tr}(B^\dagger \rho B)$, where it gives a nondegenerate quadratic form and, by compactness of the unit sphere, a uniform positive lower bound. With only $\rho \geq 0$, this bound can vanish on the kernel of ρ.

Consequently, the explicit $\rho > 0$ hypothesis is a genuine hypothesis of this auxiliary route, not an accidental artefact. This theorem should not be cited as the source theorem itself. The formal statements following the paper remain the primitive-channel statements in `TNLean/Wielandt/Inequality/Bounds.lean` and the primitive/eventual-full-rank equivalence in `TNLean/Wielandt/Primitivity/Equivalence.lean`. CPSV16 Section 2.3 should be cited through those primitive-channel statements rather than through the auxiliary positive-definite-fixed-point theorem.

[5] `TNLean/Wielandt/Primitivity/StronglyIrreducibleToFullRank.lean`.

5 What is not a discrepancy

The replacement of d by $\mathrm{kr}(A) = \dim S_1(A)$ is not a paper error and not a formal weakening. It records explicitly the linearly independent Kraus count used by the source proof.

Nor is the use of the increasing spaces

$$T_n(A) = S_0(A) + S_1(A) + \cdots + S_n(A)$$

by itself a theorem-level deviation. Some intermediate formal results prove $T_{D^2}(A) = M_D(\mathbb{C})$ under normality. These results are weaker than the full quantum Wielandt theorem, but the theorem in `TNLean/Wielandt/Inequality/Bounds.lean` proves the stated bound for $i(A)$, where $i(A)$ is defined using the spaces $S_n(A)$ themselves. [6]

Finally, the lower-level use of normality is not itself a mismatch. In these formal statements, normality is eventual full Kraus rank. Under the paper's normalization, the formal statements prove the equivalence between the paper's primitive condition and eventual full Kraus rank, so using normality as an intermediate hypothesis is legitimate.

6 Scope of the index comparison

For a nonzero vector $\varphi \in \mathbb{C}^D$, set

$$H_n(A, \varphi) = S_n(A)\varphi,$$

so that $q(\mathcal{E}_A)$ defined above is the least positive n such that $H_n(A, \varphi) = \mathbb{C}^D$ for every nonzero φ. As for $i(A)$, the value is defined to be zero when no such positive integer exists. Proposition 1 of [SPGWC10] states

$$q(\mathcal{E}_A) \leq i(A)$$

for every quantum channel.

The available comparison assumes trace-preserving normalization and uniform vector-spreading primitivity. It proves the same inequality under these hypotheses, and is therefore a specialization of Proposition 1 rather than the source theorem with its full hypothesis set. A related comparison assumes instead that $i(A)$ is finite.[7]

The missing cases are mathematically degenerate but must still be represented explicitly. If the channel is not primitive, then the admissible set defining

[6] The cumulative-span theorem is `MPSTensor.cumulativeSpan_eq_top` in `TNLean/Wielandt/SpanGrowth/NonzeroTraceProduct.lean:164--197`; the theorem for the spaces $S_n(A)$ is the general theorem cited above. The blueprint makes this distinction in `blueprint/src/chapter/ch08_wielandt.tex:461--490` and states the general bound in `blueprint/src/chapter/ch08_wielandt.tex:492--548`.

[7] The two declarations are `MPSTensor.qIndex_le_iIndex_of_isPrimitivePaper` and `MPSTensor.qIndex_le_iIndex`. The definitions are in `TNLean/Wielandt/Primitivity/Definitions.lean`, and the restricted theorem is in `TNLean/Wielandt/Inequality/Bounds.lean`.

$q(\mathcal{E}_A)$ is empty and $q(\mathcal{E}_A) = 0$. If it does not have eventual full Kraus rank, the admissible set for $i(A)$ is empty and $i(A) = 0$. A source-faithful theorem can therefore be obtained by proving these two zero-index characterizations and splitting on primitivity, using the existing primitive-channel inequality in the nontrivial case.

7 Scope of the strong-irreducibility equivalence

Proposition 3(c) of [SPGWC10] requires that 1 be the only peripheral eigenvalue and that its eigenspace be one-dimensional, generated by a positive-definite fixed point. The one-dimensionality is essential: the identity channel has peripheral spectrum $\{1\}$ and the positive-definite fixed point I, but its fixed space is not one-dimensional. The present comparison strengthens the source condition by also requiring irreducibility of the transfer map. Consequently, the proved equivalence is a strengthened version of Proposition 3 rather than its literal third clause.[8]

To remove the restriction, define the literal source predicate using the positive-definite fixed point, the one-dimensional fixed eigenspace, and the absence of other peripheral eigenvalues. Then prove that a normalized completely positive map with these properties has no nontrivial invariant projection: such a projection would produce an additional positive fixed direction, contradicting fixed-space simplicity. This implication supplies the extra irreducibility conjunct of the existing predicate, after which the established equivalences yield the literal Proposition 3. Alternatively, the three equivalence directions may be reproved directly for the source predicate.

8 Conclusion

The general quantum Wielandt bound has been formalized with the right index and with the paper's linearly independent Kraus count made explicit. The two improved estimates have also been upgraded to the paper's hypotheses: the distinguished matrix is any element of $S_1(A)$ with the required spectral property. The chosen-Kraus statements remain useful corollaries, but they should not be cited as the theorem itself.

Two local scope restrictions remain: the index comparison assumes primitivity (or, in the companion form, finiteness of $i(A)$), and the strong-irreducibility predicate carries an additional irreducibility conjunct. The two preceding sections record the zero-index and source-predicate bridges needed to remove these restrictions.

The remaining caution is theorem-surface hygiene. Cumulative-span declarations and bundled analysis chains are useful support lemmas for the proof,

[8] The strengthened predicate is `MPSTensor.IsStronglyIrreduciblePaper`; the equivalences are `MPSTensor.primitivePaper_iff_stronglyIrreducible` and `MPSTensor.hasEventuallyFullKrausRank_iff_stronglyIrreducible`. See `TNLean/Wielandt/Primitivity/Definitions.lean` and `docs/glossary.md`.

while the Wielandt results are the exact-length statements about $S_n(A)$ and the index $i(A)$.

References

[SPGWC10] M. Sanz, D. Pérez-García, M. M. Wolf, and J. I. Cirac. A quantum version of Wielandt's inequality. *IEEE Transactions on Information Theory*, 56(9):4668–4673, 2010.