

Sharp Finite-Length Block Separation for Canonical-Form Tensors

2026-07-15

Verdict: scope-restriction (resolved).

1 The assertion

The sharp block-injectivity proposition of Cirac, Pérez-García, Schuch, and Verstraete is the point at issue [CPGSV16, lines 340–345].¹

Let A be a tensor in canonical form with bond-dimension parameter D , and let $A_1, \dots, A_g$ be a basis of normal tensors for A . The paper characterizes such a basis by minimality: no two chosen normal tensors may be related by an invertible similarity together with a phase. It then defines block-injective canonical form as the statement that, after passing to these blocks, simultaneous word evaluations span the whole product algebra

$$\bigoplus_{j=1}^g M_{D_j}(\mathbb{C}).$$

Equivalently, for every block tuple $X = (X_j)_j$ there should be a linear combination of common physical words whose evaluation on block j gives X_j for every j .

The source first invokes the quantum Wielandt theorem: after at most D^4 blockings, each normal tensor becomes injective. It then invokes the cross-block argument of Pérez-García–Verstraete–Wolf–Cirac [PGVWC07]: if L is a blocking length after which every normal block is injective, then after $3(L + 1)(D - 1)$ blockings the tensor is block injective. The proposition records the resulting bound in the following form: after blocking at most $3D^5$ spins, every tensor in canonical form is in block-injective canonical form. The statement is at lines 340–345 of the local source, immediately after the definition of block-injective canonical form at lines 318–322.

¹For traceability, this note corresponds to issue #1043. Related formal work is recorded in issue #587 and issue #822; earlier proof attempts are PR #682 and PR #813.

2 The older per-copy formulation

The older auxiliary horizontal block hypotheses make the finite-length conclusion explicit on every flattened copy. They record the following trace-separation assumption: there is a length N such that, if matrices $\Delta_j \in M_{D_j}(\mathbb{C})$ satisfy

$$\sum_{j=1}^g \text{tr}(\Delta_j A_j^w) = 0 \quad \text{for every word } w \text{ of length } N,$$

then $\Delta_j = 0$ for every j . By nondegeneracy of the trace pairing on the product algebra, this is the dual form of the finite-length spanning property above.

Several exact sufficient hypotheses have been proved. The first is the finite-length product-algebra span itself: the tuples

$$w \mapsto (A_1^w, \dots, A_g^w)$$

span $\bigoplus_j M_{D_j}(\mathbb{C})$ for one common length. The second is an entrywise linear-independence criterion: the scalar functions $w \mapsto (A_j^w)_{\alpha\beta}$, indexed by the block and by a matrix entry in that block, are linearly independent at one common length. This implies the product-algebra span, hence the trace-separation property. The third is a selector-word formulation of the sharp block-injectivity proposition: after a common blocking length each block is injective, and after a further finite length there are word polynomials which are the identity on one chosen block and vanish on all the others. Concatenating an injective prefix with such a selector suffix gives the full product-algebra span.

In the formal files, these three finite-length criteria suffice to produce the older per-copy trace-separation conclusion.² These statements remain correct restricted results. The sharp proposition is now proved separately for a representative-indexed BNT sector decomposition, so repeated copies are grouped before the separation argument.

3 The counterexample boundary

A tempting shortcut is false. One cannot recover the finite-length trace-separation witness from blockwise injectivity, left-canonicity, and nonzero weights alone. Take two one-dimensional scalar blocks over a one-letter alphabet, with both block tensors equal to 1 and with weights 1 and 2. Each block is injective and left-canonical, and the weights are nonzero and distinct. However every word evaluation is the same on the two blocks. The tuple span is the

²The criteria are named `WordTupleSpanTop`, `wordEntryFamily`, and `PropBlockInjective`. The trace-separation assumption is `HorizontalCFData.biCF`. The constructors from these hypotheses to `HorizontalCFData` structures are in the modules under `TNLean/MPS/MPD0/BiCFDerivation/`. Relevant implications include `hasBiCF_of_wordTupleSpanTop`, `hasBiCF_of_wordEntryFamily_linearIndependent`, `hasBiCF_of_propBlockInjective`, and corresponding `HorizontalCFData` constructors such as `horizontalCFData_of_propBlockInjective`. The blueprint records these restricted criteria in its finite-witness remark.

diagonal line in $\mathbb{C} \oplus \mathbb{C}$, not the whole product algebra. Equivalently, choosing $\Delta_1 = 1$ and $\Delta_2 = -1$ makes the trace pairing vanish on every word while the tuple (Δ_1, Δ_2) is nonzero.

This example is formalized as `duplicateScalarBlocks` and its lemmas in `TNLean/MPS/MPDO/BiCFDerivation/Counterexample.lean`, and is recorded in `docs/counterexamples.md`, lines 29–35. It shows that the additional finite-length witness in the formal hypotheses is genuinely stronger than the other horizontal canonical-form assumptions. It does not refute the source’s sharp block-injectivity proposition. The duplicate scalar blocks are related by similarity and phase, and therefore do not form a minimal basis of normal tensors in the sense of the source characterization.

4 Consequence for Lemma L

The same counterexample marks the restriction in the older theorem `blockwise_insert_eq_of_mpv_agree`, which assumes `biCF` on the full family of copies. The source’s Lemma L (Appendix C.3, lines 1835–1858) instead writes the canonical-form tensor as $\bigoplus_{j,q} \mu_{j,q} X_{j,q} A_j X_{j,q}^{-1}$, groups all copies over the same minimal representative A_j , and obtains the coefficient $c_j^{(N)} = \sum_q \mu_{j,q}^N$. The nonvanishing power-sum lemma then supplies a length at which $c_j^{(N)} \neq 0$.

This representative-grouped argument is now formalized. At a fixed length, simultaneous word-tuple separation of the representatives forces $c_j^{(N)}(YA_j - ZA_j) = 0$ for every j ; the bounded nonvanishing power-sum argument then gives $YA_j = ZA_j$. A second theorem proves the required eventual word-tuple separation when the BNT representatives have already been blocked to become injective. Thus multiplicities and unequal nonzero copy weights are no longer an obstruction to the formal Lemma L argument.

The first-site action has also been transported through physical blocking. For a block $I = (i_0, \dots, i_L)$, the induced matrix is

$$Y_{I,J}^{[L+1]} = Y_{i_0,j_0} \prod_{k=1}^L \delta_{i_k,j_k},$$

and its insertion satisfies

$$(Y^{[L+1]} A^{[L+1]})^I = (YA)^{i_0} A^{i_1} \dots A^{i_L}.$$

Consequently, an all-length first-site equality before blocking gives the corresponding equality after blocking. Conversely, if the length- L words of A span the full matrix algebra, equality of the two blocked insertions implies $YA = ZA$. These auxiliary statements provide the transport needed to combine the source’s block-injectivity argument with Lemma L; they introduce no additional hypothesis beyond block injectivity for the converse.

5 Classification

The original counterexample should not be read as a source erratum. A small counterexample shows only that per-copy blockwise injectivity, left-canonicity, and nonzero weights do not imply simultaneous separation. It does not satisfy the source’s minimal-representative condition.

The representative-level source theorem is now composed with the canonical-form hypotheses. Positive-length block injectivity propagates from length L to length $L + 1$ without a normalization hypothesis: the identity $S_{L+1}(A) = \text{span}\{A_i\}_i S_L(A)$, together with the same identity at length L , gives the result. Physical blocking of a sector decomposition has also been defined explicitly: it preserves the representative and copy indices, replaces A_j by $A_j^{[p]}$, replaces $\mu_{j,q}$ by $\mu_{j,q}^p$, and gives the coefficient identity

$$c_{P^{[p]},j}^{(N)} = c_{P,j}^{(Np)}.$$

The tensor obtained from this blocked decomposition has the same MPV family as the physical blocking of the original tensor.

The representative-separation step has now been obtained without asserting preservation of the full canonical-form predicate. If $p > 0$ and every $A_j^{[p]}$ is injective, then pair separation for the original representatives at length $6p$ transports through the bijection of blocked and unblocked words to pair separation for the blocked representatives at length 6. It follows that $P^{[p]}$ has simultaneous representative word-tuple span at every blocked length at least $6(g - 1) + 1$.

The first-site transport was first proved under an explicit common block-injectivity hypothesis. Given a common positive L for which every representative is L -block injective, use $p = L + 1$; positive-length propagation makes every $A_j^{[L+1]}$ injective. Blocked representative separation and the representative-grouped Lemma L give equality of the induced insertions on every $A_j^{[L+1]}$. The original length- L span then recovers equality of the unblocked insertions. This length shift is essential: recovery from a physical block of length $L + 1$ uses the span of its length- L tail.

The first task is now closed without adding a hypothesis. For an irreducible left-canonical representative, quantum Perron–Frobenius theory gives a finite peripheral period m . Its self-overlap along multiples of m tends to m , whereas the normalized BNT hypothesis makes the same subsequence tend to 1. Thus $m = 1$, the transfer map is primitive, and the representative is normal. The finitely many normal representatives have a common positive block-injectivity length by taking a common multiple of their individual lengths. Consequently the representative-indexed Lemma L now follows from the BNT canonical-form predicate itself.

The numerical Wielandt part has since been sharpened to the exact common length used in the source estimate. If the representative dimensions are D_j and the flattened canonical-form tensor has bond dimension D , then $D_j \leq D$. Quantum Wielandt gives an injective length $L_j \leq D_j^4$ for the j th representative,

and positive-length propagation therefore makes every representative injective at the single exact length D^4 . This exact common length has been formalized.³

The direct-sum interpolation is now complete. Put g for the number of minimal representatives and set $L_0 = D^4$. The two-block argument gives the full pair span at length $S = 3(L_0 + 1)$. For a chosen block k , choose an anchor block $j \neq k$ and use this pair span to realize an arbitrary pair $(X_k, 0)$. Multiply that word polynomial by one identity-zero selector for each of the remaining $g-2$ blocks. The resulting simultaneous tuple equals X_k on block k and vanishes on every other block, at length

$$(g-1)S = 3(g-1)(D^4 + 1).$$

Summing these block-supported tuples over k realizes every element of the full product algebra. No separate injective prefix is used. The length-zero identity occurs only as the neutral base of the finite product when there are no remaining selectors; it is not a padding hypothesis or a separate case in the theorem.

If $g = 1$, ordinary block injectivity gives the required span at length D^4 . If $g \geq 2$, positivity of the representative dimensions and multiplicities gives $g \leq D$. The remaining arithmetic is

$$3(D^4 + 1)(D - 1) = 3(D^5 - D^4 + D - 1) \leq 3D^5,$$

because $D - 1 \leq D^4$.

The formal proof is recorded by the following declarations:

- `MPSTensor.exists_mem_span_wordTuple_eq_and_eq_zero_of_pairWordTupleSpanTop` realizes the arbitrary anchor-pair value;
- `MPSTensor.wordTupleSpanTop_mul_pred_of_forall_pairWordTupleSpanTop` and `MPSTensor.wordTupleSpanTop_mul_pred_of_forall_pairTraceSeparatingAt` prove the direct-sum interpolation at length $(g-1)S$;
- `MPSTensor.wordTupleSpanTop_threeBlock_mul_pred_of_blocksNotGaugePhaseEquiv_c1` specializes the result to the source length $3(g-1)(L_0 + 1)$;
- `MPSTensor.SectorDecomposition.basisCount_le_totalDim` proves $g \leq D$;
- `MPSTensor.IsBNTCanonicalForm.exists_basis_wordTupleSpanTop_le_three_totalDim_pow_five` gives a positive simultaneous-span length at most $3D^5$.

The public horizontal canonical-form predicate has now been replaced by the representative-indexed formulation. It quantifies a BNT sector decomposition whose total bond dimension equals the original bond dimension and which is

³Formal declarations: `MPSTensor.SectorDecomposition.basisDim_le_totalDim`, `MPSTensor.isNBlkInjective_pow_four_of_isNormal_leftCanonical`, and `MPSTensor.IsBNTCanonicalForm.basis_isNBlkInjective_totalDim_pow_four`.

literally related to the doubled-index tensor by one global gauge. Complete MPV equality transports the first-site identities to the BNT representatives. The representative-grouped Lemma L concludes equality of the opposite-corner insertions on every minimal representative. Linearity extends this equality to the full weighted direct sum, and gauge conjugation transports the result to the original MPO tensor. The older per-copy `HorizontalCFData` theorems remain available only as explicitly restricted auxiliary statements.

6 The Appendix C.2 sector entropy argument

The one-letter simultaneous-span assumption in the formal sector entropy theorem is not an additional hypothesis absent from the cited source. At the start of Case II, the source fixes a simple tensor K in block-injective canonical form [CPGSV16, line 1628]. By Definition 2.6, block-injective canonical form says exactly that the one-letter tuples of the minimal normal representatives span

$$\bigoplus_{j=1}^g M_{D_j}(\mathbb{C}).$$

The inverse tensor used immediately afterwards, and hence the one-site projections in the proof of Proposition 4.13, depend on this hypothesis. The formal assumption `WordTupleSpanTop S.basis 1` is therefore the direct algebraic form of the source’s biCF assumption.

It would be a strictly stronger assertion to begin with an unblocked canonical form and remove this hypothesis by choosing a finite blocking length. Although SAL passes from a tensor to any positive physical blocking, the converse does not follow from the periodic block identity: an L -blocked chain records only the original chain lengths NL and only cuts whose lengths are multiples of L . In particular, SAL of $K^{[L]}$ does not by itself recover the equalities $I_m(K) = I_{m+1}(K)$ at arbitrary chain lengths and cuts. The proposed reverse transport in issue #4119 is therefore unavailable without a new mathematical argument not present in Appendix C.2, and it is not needed for faithfulness to the printed proposition.

Consequently, the sector entropy theorem should retain the simultaneous one-letter span and identify it as the source biCF hypothesis. The finite blocking theorem remains the justification that an arbitrary canonical-form tensor can be placed in biCF when the surrounding result is itself stated only up to physical blocking; it does not turn the Appendix C.2 conclusion into an unblocked SAL statement.

This closes the multiplicity and block-separation gap in Lemma L. It does not by itself complete Proposition 4.13: constructing the letterwise vertical decomposition from the resulting reducing projections, and proving the subsequent positivity and isometric normalization of its sectors, are separate parts of that proposition.

The previously open gap is now resolved: the formal theorem matches the source’s hypothesis set and conclusion, through a different proof of the

simultaneous-interpolation step. BNT canonical form supplies normality, left-canonicity, normalized self-overlap, and pairwise gauge-phase inequivalence; no finite separation witness is added as a hypothesis. The source-level Lemma L needs only existence of a suitable finite length, while the formal theorem also establishes the stated bound $3D^5$.

References

- [CPGSV16] J. Ignacio Cirac, David Pérez-García, Norbert Schuch, and Frank Verstraete. Matrix product density operators: Renormalization fixed points and boundary theories. Preprint; published version: *Annals of Physics* 378 (2017), 100–149, 2016.
- [PGVWC07] David Pérez-García, Frank Verstraete, Michael M. Wolf, and J. Ignacio Cirac. Matrix product state representations. *Quantum Information and Computation*, 7(5–6):401–430, 2007.