Theorem Surface Audit for Canonical Form, BNT, and the Non-periodic MPS Fundamental Theorem

2026-05-05

Verdict: clarification (open).

1 Purpose

This note compares the theorem surface in the formal development with the statements in the source papers. It is not an issue index. The goal is to separate four kinds of results:

1. named statements appearing directly in the papers;
2. cited mathematical results used by the papers;
3. auxiliary formal statements needed to make an implicit proof step exact;
4. stronger or alternative statements that should not be presented as part of the paper-level proof unless they become necessary.

2 Named source statements

The 2016 source has a small number of named statements in the canonical-form and BNT part:

- normal tensor and canonical form definitions [CPGSV16, lines 233–255];
- basis of normal tensors definition and characterization [CPGSV16, lines 271–280];
- injective and block-injective canonical form definitions, together with the quantum-Wielandt and block-injective-after-blocking assertions [CPGSV16, lines 317–345];
- the proportional Fundamental Theorem and the equal-MPV corollary [CPGSV16, lines 347–356].

The 2021 review gives the same normal tensor, BNT, proportional theorem, and equal-MPV corollary in Section IV [CPGSV21, lines 1828–1899].

These are the paper-level targets. They should remain the prominent theorem statements in the blueprint.

3 Why the formalization has more declarations

The paper proof suppresses finite-dimensional and finite-index operations that must be named in a proof assistant. In this part of the development the extra formal statements usually fall into one of the following categories.

Reduction statements. The canonical-form proof splits the paper's "put the tensor in canonical form" step into invariant-subspace extraction, block-diagonal decomposition, trace-preserving normalization, period-removing blocking, and primitive-block normality. These are not new mathematical claims beyond the source reduction, but they need separate hypotheses and conclusions so later theorems can apply them.

Common blocking statements. The paper says that periodicity is removed by blocking a common multiple of the periods. The formalization must record positivity of the chosen block length, persistence of injectivity under further blocking, and the simultaneous application to a finite family of representative sectors. These statements are auxiliary, but they are mathematically faithful to the period-removal step.

Wielandt proof-chain statements. The source proof of the quantum-Wielandt bound uses Lemma 1, eigenvalue extraction, vector spreading, and the fixed-length rank-one step as an argument, not as separate MPS theorem statements [SPGWC10, Theorem 1 proof]. Formal conjunctions that bundle these inputs, such as the *Wielandt analysis* and *Wielandt chain* declarations, should therefore be lemmas. The visible theorem is the cumulative or fixed-length Wielandt conclusion, with the stronger fixed-length claim distinguished from the cumulative span bound.

BNT comparison statements. The source BNT characterization selects one representative from each phase-gauge equivalence class and compares the power sums of weights. The formalization separates this into overlap decay, eventual linear independence, permutation rigidity, sector-weight comparison, and phase absorption. The rectangular block-matching theorem, allowing the two sector decompositions to have different bases before matching, is a genuine formal conclusion for the block-matching part, while the final global weighted gauge equivalence remains the paper-level equal-MPV target. The overlap-to-linear-independence and change-of-basis declarations are linear-algebra support lemmas for the permutation-rigidity theorem, not additional BNT theorems from the source. The repeated-copy part should be written around the displayed coefficients

$$c_{S,j}^{(N)} = \sum_{q=1}^{r_j^S} (\mu_{j,q}^S)^N$$

and the resulting multiset equality

$$\{\{\mu_{j,q}^S : 1 \le q \le r_j^S\}\} = \{\{\mu_{j,q}^T : 1 \le q \le r_j^T\}\}.$$

Here the double braces denote multisets, so repeated weights are counted with their multiplicities. In this notation, BNT linear independence separates the coefficient arrays, and Newton–Girard recovers the weights. The formal declarations that only package these two algebra steps in sector-data notation are support lemmas, even when their names mention the Fundamental Theorem. One previously used formal route recovered the multiplicities r_j^S, r_j^T by first extrapolating eventual equality of the positive-power sums to the exponent-zero identity

$$\sum_{q=1}^{r_j^S} (\mu_{j,q}^S)^0 = \sum_{q=1}^{r_j^T} (\mu_{j,q}^T)^0,$$

and then applying the same-cardinality Newton–Girard theorem. This is not a gap when all weights are nonzero, because the geometric-sequence extrapolation lemma proves the $N = 0$ identity from eventual equality at positive lengths. It is nevertheless not the route displayed in [CPGSV16, Lemma `Lem:app_simple`, lines 1155–1163]; the source applies the unequal-cardinality power-sum lemma after the BNT representatives have been matched. Without the nonzero-weight hypothesis, positive powers would only determine the nonzero submultiset, so any use that counts zero weights must record an additional $N = 0$ hypothesis or a nonzero coefficient hypothesis. This classification follows the protocol in `docs/paper-gaps/policy.tex`; the detailed comparison is recorded in `docs/paper-gaps/cpsv16_power_sum_alternative_route.tex`.

4 External lemmas need explanations

When a result is needed but the MPS source only cites it, the formal development should give a pedagogical explanation at the point where the result enters the proof. A citation alone is not enough in this situation. The explanation should include:

1. the exact mathematical statement used;
2. the hypotheses in the MPS notation, after blocking or reindexing if necessary;
3. a short proof idea, or a precise indication that the proof is imported from the cited source;

4. the formal declaration that supplies the result, if it has already been formalized.

This applies in particular to the following families of external input.

- The quantum Wielandt theorem supplies finite-length matrix-algebra span from normality or primitivity [SPGWC10]. The paper cites it in the canonical-form reduction; the formal text should say whether it is using cumulative span, exact-length span, injectivity of the map Γ_L, or block-injectivity of a direct sum.
- The block-injective canonical-form bound uses the older MPS injectivity and block-injectivity argument of [PGVWC07]. The formal statement should identify the common blocked physical alphabet and the direct-sum matrix algebra that is being spanned. The direct-sum proof input in the older source is recorded in `docs/paper-gaps/pgvwc07_direct_sum_input.tex`.
- Newton–Girard identities are elementary, but they are not proved in the MPS source. When they are used to recover coefficients, multiplicities, or weight multisets, the blueprint should explain the finite power-sum recursion. Declarations that only restate these finite-sequence facts in BNT notation should be lemmas, even when they are applied inside a paper-level theorem.
- Perron–Frobenius and peripheral-spectrum facts for quantum channels are standard inputs from finite-dimensional channel theory, but they are not interchangeable. A formal theorem should specify whether it uses irreducibility, primitivity, trace preservation, a faithful fixed point, or a rank-one peripheral projection.

This rule also gives a retirement test. If an external-looking lemma has no source citation, no proof explanation, and no current proof path needing it, it should not remain as a headline theorem in the paper-level blueprint.

5 Dependency-chain classification

The non-periodic Fundamental Theorem route should be read with the following separation of proof roles.

5.1 Final target trace-back

For the non-periodic route, the source target is the Fundamental Theorem for canonical-form MPS and its equal-MPV corollary, after the tensor has been put in normal form and the periodicity has been removed by blocking. In the notation of [CPGSV16, lines 347–356], the proportional statement compares two canonical-form tensors with the same family of MPVs up to a scalar sequence and concludes a block permutation, phases, and blockwise gauges. The equal-MPV corollary is the special case in which the scalar sequence is identically 1, but the proof still passes through the BNT coefficient comparison.

The formal route currently has three distinct layers that should not be identified with one another.

1. The single-block injective theorem `MPSTensor.fundamentalTheorem_singleBlock` is the injective-block gauge theorem. It is a source-level theorem only for one injective block, not for the multi-block BNT comparison.

2. The SectorBNT equal-MPV declarations `MPSTensor.ft_sector_bnt_equal_sector_data` and `MPSTensor.ft_sector_bnt_equal_global_gauge` are the current formal statements closest to the BNT part of the source proof. They work with raw sector weights and derive the sector matching, copy-count equalities, and matched copy-weight identities on the SectorBNT surface. The corrected active, nonzero SectorBNT equal-MPV theorem is obtained after the canonical-form supplier has put the input tensors on this surface with the required unit-modulus-copy normalization. The literal ambient source corollary retains the separate zero-weight-coordinate obstruction.

3. The sector-decomposition theorem `MPSTensor.fundamentalTheorem_equalMPV_sectorDecomposition_hetero_of_sectorMatching` is now an auxiliary rectangular comparison statement: from a supplied sector-basis matching, linear independence, and equality of total MPV families, it proves the corresponding equality of sector-weight multisets, up to the phases in the matching. It should not be presented as the paper-level equal-MPV corollary.

The current proportional theorem[1] is complete on the SectorBNT surface. Every representative there has a positive copy count, every stored copy weight is nonzero, and the BNT data includes the converse coverage needed to exclude extraneous representatives. From the proportional MPV hypothesis it obtains a bijection of sectors, matched bond dimensions, unit phases, and invertible block gauges. Separate global-gauge declarations strengthen this sectorwise conclusion and should not be conflated with the source theorem.

Literal source obstruction. The BNT definition in arXiv:1606.00608, lines 271–274, requires normality, spanning of the total MPV, and eventual linear independence, but it does not require every candidate to occur with a nonzero coefficient. Consequently an unrelated normal tensor may be adjoined with coefficient zero at every length. The same tensor then has a one-element and a two-element BNT, so the equality of cardinalities asserted by Theorem II.1 is false under the printed hypotheses.[2]

[1] Formal declaration: `MPSTensor.fundamentalTheorem_proportional_canonicalForm`.

[2] Formal declaration: `MPSTensor.exists_isCPSVBasisOfNormalTensors_unequal_card`; see `docs/paper-gaps/cpsv16_bnt_uniqueness_zero_coefficient.tex`.

Thus the exact missing hypothesis surface is converse coverage of the active canonical-form blocks, not a coefficient identity inside the proved SectorBNT theorem. With this active, nonzero restriction, the proportional sector theorem is complete. Without it, the printed theorem requires a source correction rather than another matching argument. In CPSV16, the equal-MPV ambient conjugacy corollary has the additional zero-weight-coordinate and ambient-dimension obstruction recorded separately.

The formal declaration `MPSTensor.CPSVCanonicalFormData.fundamentalTheorem_equal_ambient_canonicalForm` now covers the nonzero equal-ambient sub-case. It assumes that both tensors have the same ambient bond dimension and that their canonical-form data satisfy the separate weight normalization. Exact active reconstructions and a coisometry-alignment argument then give ambient gauge equivalence. This does not remove the heterogeneous ambient-dimension obstruction or the zero-family boundary of the printed corollary.

CPSV21 uses the same weak BNT definition, so the extra zero-coefficient representative still refutes the chosen-BNT cardinality conclusion of Theorem 4.4. It does not refute Corollary 4.5. The latter concerns the ambient canonical tensors, whose actual direct-sum coefficients are positive in the standing construction. The current equal SectorBNT theorem is an analogue of that corollary; identifying it with the source statement still requires a bridge from CPSV21 canonical form to the SectorBNT hypotheses.

Conditional after-blocking statements. The after-blocking comparison theorems are useful only after the source canonical-form, BNT selection, block-injective span, and BNT matching inputs have been supplied. The common-phase-cover and BNT-cover formulations are conditional statements using those inputs. They should stay visibly conditional until the hypotheses are produced from the paper's assumptions. In particular, the finite-length span or phase-cover assumptions are not extra hypotheses in the source theorem; they are the formal names for proof inputs that still have to be obtained from the source argument.

Data structures and transport lemmas. Objects such as `SectorBasisOverlapSpanHypotheses` are formal comparison data for the proof. They should not be cited as paper theorems. Lemmas converting one data structure into another are needed in the formal proof chain, but they are auxiliary unless they state one of the source mathematical claims directly.

External finite-dimensional inputs. The formal route also contains named statements for linear algebra that the paper uses without proof: Newton–Girard recovery, finite-dimensional Perron–Frobenius structure, Burnside–Jacobson algebra generation, and quantum Wielandt bounds. These may be important mathematical theorems in their own right, but in this proof chain they enter as cited inputs. Their statements should be explained, and their formal names should normally be lemma-level when they are merely restated in MPS notation.

6 Current demotions

The following declarations are auxiliary and should not be read as independent paper-level theorems.

- The explicit-coefficient equal-MPV result is named

  `fundamentalTheorem_equalMPV_of_explicit_coefficients`.

  It is now lemma-level. It is not the full equal-MPV corollary from the papers, since it assumes the coefficient arrays and their nonzero limits.

- The former proportional result on the old strict-representative surface has been deleted from the active formal route. It used explicit coefficient arrays, their limits, and non-vanishing assumptions, while the source theorem derives the relevant BNT coefficient comparison from the proportional MPV hypothesis. This was a wrong theorem surface, not merely a theorem-versus-lemma presentation choice.

- The unused fixed-size BNT Vandermonde MPV declarations and the unused scalar Vandermonde separation lemma have been removed. The checked block-separation proof uses mixed-transfer peeling, and the equal-case BNT weight recovery uses the finite power-sum identity.

- The equal-span permutation rigidity statement with asymptotically orthonormal overlaps is a lemma. It is a useful span route, but the proportional BNT permutation theorem is the theorem used by the main proportional Fundamental Theorem in Chapter 11.

- The older theorem from coefficient-array comparison data to sector comparison without leftover-zero-block bookkeeping has been removed. The active after-blocking route keeps this bookkeeping visible because the formal equality relation includes the empty word. This is not terminology from [CPGSV16]; the source only says that canonical forms may have zero blocks through the inequality $\sum_k D_k \le D$ [CPGSV16, lines 217–219].

- The power-sum multiset recovery statement

  `power_sum_eq_implies_multiset_eq`

  is a lemma-level restatement of the elementary Newton–Girard input used in the equal-MPV proof. The former declaration

  `perBlock_sameMPV_of_equalMPV_CFBNT`

extracting the blockwise equality conclusion from the common-block equal-MPV lemma has likewise been removed from the active Lean tree. These old names are retained here only to identify the retired one-copy route; neither should be counted as an additional source-level Fundamental Theorem.

- The implication

  $$\texttt{sameMPV}_2\texttt{\_implies\_proportionalMPV}_2$$

  from equal MPV families to proportional MPV families with proportionality constant 1 is also lemma-level. The displayed subscript 2 denotes the Unicode subscript 2 used in the Lean declaration name. It is a logical conversion between hypotheses, not a separate source theorem.

- The BNT sector-weight recovery declarations `weight_multiset_eq_of_sameMPV_bnt` and `exists_copies_eq_and_weight_multiset_eq_of_sameMPV_bnt` are lemma-level. They formalize the coefficient comparison $c_{S,j}^{(N)} = c_{T,j}^{(N)}$ and the elementary recovery of copy counts and weight multisets; they are not additional BNT theorems beyond the source's basis-normal-tensor comparison.

- The phase-matched sector-comparison adapters `fundamentalTheorem_equalMPV_sectorDecomposition_hetero_of_phaseMatch` and `fundamentalTheorem_equalMPV_sectorDecomposition_hetero_of_matched_basis` are also lemma-level. The names are legacy names for rectangular comparison adapters; they transport the same coefficient and multiset formulas across a supplied phase or matched-basis identification. The theorem-level surface begins only when the proof supplies the actual matching data from the BNT hypotheses.

7 Statements to keep visibly conditional

Several declarations should stay in the development, but their statements must remain visibly conditional because they assert stronger conclusions or start from more explicit data than the source theorem.

- The conditional after-blocking Fundamental Theorem assumes the post-blocking sector hypotheses. It should not be cited as the unconditional canonical-form theorem.

- The explicit-coefficient equal case assumes coefficient arrays rather than deriving the BNT power-sum data from bare equality of MPVs.

- The conditional proportional global-gauge theorem assumes scalar-free matched-sector coefficient identities. Mere projective proportionality instead carries a common length-dependent scalar c_N, which need not be geometric. Thus the global-gauge conclusion is an intentionally stronger result, not unfinished work for the sectorwise proportional theorem.

8 Retirement criterion

An auxiliary declaration should be removed or moved out of the main blueprint route when all three conditions hold:

1. it is not a statement in the cited paper and is not a cited outside result;
2. no current proof of a paper-level target uses it;
3. it states a stronger finite-length, span, or synchronization hypothesis than the source proof requires.

If a declaration is still used but only as a formal intermediate result, it should be a lemma or a clearly conditional theorem, not a headline replacement for the paper theorem.

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.

[CPGSV21] J. Ignacio Cirac, David Pérez-García, Norbert Schuch, and Frank Verstraete. Matrix product states and projected entangled pair states: Concepts, symmetries, theorems. *Reviews of Modern Physics*, 93(4):045003, 2021.

[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.

[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.