Basis-of-Representatives Scope Restriction in the Fundamental Theorem Formalization

2026-05-10

Verdict: scope-restriction (resolved).

1 The source argument

Cirac–Pérez-García–Schuch–Verstraete 2016 states the general proportional Fundamental Theorem as Corollary `II_cor2` of [CPGSV16]. The theorem applies to two tensors A and B whose matrix-product vectors are proportional, each presented in its basis-of-normal-tensor (BNT) canonical form. The canonical form for A is a direct-sum decomposition

$$A^i = \bigoplus_{j=1}^{g_A} I_{r_{A,j}} \otimes \mu_{A,j}(A^i_{j,1} \oplus \cdots \oplus A^i_{j,r_{A,j}}),$$

where each modulus class j contains $r_{A,j} \geq 1$ representative blocks $A_{j,1}, \ldots, A_{j,r_{A,j}}$ all sharing the same weight modulus $|\mu_{A,j}|$, and the moduli are strictly ordered across classes: $|\mu_{A,1}| > |\mu_{A,2}| > \cdots > |\mu_{A,g_A}|$. The same structure applies to B with multiplicity counts $r_{B,j}$.

The proof of Corollary `II_cor2` in [CPGSV16] proceeds in several steps. Step 3 (lines 1156–1163 of the arXiv version) derives the blockwise power-sum identity

$$\sum_{q=1}^{r_{A,j}} \mu^N_{A,j,q} = \sum_{q=1}^{r_{B,j}} (\nu_{B,j,q}\, e^{i\phi_j})^N, \qquad \text{for all } j \text{ and all } N \geq 1, \tag{1}$$

from the BNT linear-independence input together with the equal-MPV hypothesis (so the equality of matrix-product vectors forces the equality of the sums of Nth powers). Step 4 invokes Lemma `Lem:app_simple` of the same paper: if two finite sequences $(\lambda_{a,k})_{k=1}^{x_a}$ and $(\lambda_{b,k})_{k=1}^{x_b}$ satisfy

$$\sum_{k=1}^{x_a} \lambda^N_{a,k} = \sum_{k=1}^{x_b} \lambda^N_{b,k} \qquad \text{for } N = 1, \ldots, \max(x_a, x_b),$$

then $x_a = x_b$ and the two sequences are equal as multisets. Applied to (1), this recovers $r_{A,j} = r_{B,j} =: r_j$ and the multiset equality $\{\mu_{A,j,q}\} = \{e^{i\phi_j}\nu_{B,j,q}\}$ for each modulus class j. Step 6 then assembles the global gauge isomorphism as $Y = \bigoplus_j I_{r_j} \otimes Y_j$, using the per-class dimension matching $r_{A,j} = r_{B,j}$ from Step 4.

The mathematical structure of this argument is that the multiplicities r_j and the weight multisets are recovered from a finite family of polynomial identities in N, without any passage to a limit.

2 The basis-of-representatives restriction

The surviving formal normal-CF-BNT structure[1] no longer imposes strictly decreasing moduli. Its weight condition is only non-increasing by modulus, and all retained weights are nonzero. Thus equal-modulus blocks are allowed. The additional BNT separation field says instead that distinct blocks are not gauge-phase equivalent.

The remaining restriction is therefore not strict modulus ordering. It is that `IsNormalCanonicalFormBNT` is a basis-of-representatives surface: it keeps one representative for each gauge-phase class already selected in the block family. It does not retain, as separate summands of the same structure, the repeated copies of a single source BNT tensor together with their individual weights $\mu_{j,q}$, coefficient contribution

$$c_N^{(j)} = \sum_q \mu_{j,q}^N,$$

and copy gauges. That raw multiplicity data is carried in the formalization by `SectorDecomposition` and the SectorBNT comparison theorems, not by `IsNormalCanonicalFormBNT`.

In the one-copy special case the source decomposition reduces to $A^i = \bigoplus_{j=1}^g \mu_j A_j^i$ with a single copy of each basis tensor. Then the identity (1) becomes the single-term equality $\mu_{A,j}^N = (e^{i\phi_j}\nu_{B,j})^N$, Steps 3–4 of the source proof become vacuous, and the global gauge in Step 6 takes the simpler form $Y = \bigoplus_j Y_j$ (no identity tensor factor). The present normal-CF-BNT surface is a representative-family surface of this kind; the full copy multiplicity comparison is provided elsewhere.

3 Missing components for the general case

Four items are currently absent from the normal-CF-BNT representative surface; each is required to recover the source-level multiplicity data with $r_j \geq 1$.

Lemma `Lem:app_simple` (finite Newton–Girard uniqueness). The source paper states, in its appendix, that two finite sequences whose power

[1] `TNLean/MPS/BNT/Construction.lean`: `IsNormalCanonicalFormBNT`. The earlier non-normal wrapper was retired after its downstream consumers were replaced by explicit separated-block hypotheses.

sums agree for $N = 1, \ldots, \max(x_a, x_b)$ must be equal as multisets. This is a finite algebraic fact, provable by Newton–Girard identities: the power sums determine the elementary symmetric polynomials, which determine the multiset of roots. A module named `ScalarPowerSumIdentity`[2] was added at an early stage of the project, but it does not specialize to this particular fixed-length statement in the form needed by Step 4 of the source proof. The lemma needs to be formalized as a standalone result, separate from any BNT context.

BNT power-sum identity. The identity (1) requires a proof: starting from the BNT structure on A and B and the equal-MPV hypothesis, one must derive, for each modulus class j, the equality of the two power sums in (1). The formal development no longer packages coefficient arrays for the discarded proportional branch as standalone hypotheses. The source derivation of (1) from the BNT linear-independence data and the equal-MPV condition is still not carried out. This derivation is logically prior to applying `Lem:app_simple`.

Multi-copy sector structure. The current representative surface `IsNormalCanonicalFormBNT` allows equal moduli, but it does not carry repeated gauge-phase-equivalent copies with their individual weights. The full theorem therefore needs the `SectorDecomposition`/SectorBNT multiplicity structure, where each sector admits an arbitrary finite list of copy weights and gauges. Without this explicit multiplicity data, the general-multiplicity versions of Corollary `II_cor2` and of the gauge assembly step cannot be stated on the normal-CF-BNT surface alone.

Global gauge assembly. With $r_j \geq 1$ copies per sector, the global gauge isomorphism takes the form $Y = \bigoplus_j I_{r_j} \otimes Y_j$, where I_{r_j} is the identity on the multiplicity space and Y_j is the per-block gauge on the bond dimension. Constructing and verifying this more general gauge requires the dimension match $r_{A,j} = r_{B,j}$ from Step 4, which in turn requires `Lem:app_simple` and the power-sum identity. The simpler $Y = \bigoplus_j Y_j$ gauge used in the one-copy case is already present.

4 What the current formalization does prove

Within the basis-of-representatives restriction, the formalization establishes the equal-MPV branch of the Fundamental Theorem argument. The BNT separation hypotheses yield cross-block overlap decay, and the equal-MPV hypothesis at a single witness length is used to match the per-block gauge. The former strict-modulus proportional branch has been deleted: the surviving block-matching boundary uses the explicit conclusion `BlockPermutationGaugePhaseConclusion` only where the BNT comparison has already been supplied. Thus no current theorem is presented as deriving the source proportional comparison from the discarded coefficient hypotheses.

The old dominant-normalization field and the old strict-ordering field are no longer part of `IsNormalCanonicalFormBNT`. When a CPSV argument needs a unit-modulus witness for a sector, that datum is carried as an explicit

[2] Introduced in commit `dd54b79e`.

SectorBNT hypothesis, for example by the corresponding field of `SectorDecomposition`.

5 Downstream audit status

The downstream CPSV16 surfaces that currently consume the restricted normal-CF-BNT predicate have been audited against the basis-of-representatives boundary.[3] Each of these public surfaces either assumes or constructs an already chosen representative family, non-increasingly ordered by weight modulus. Equal moduli are allowed, but repeated gauge-phase-equivalent copies and their individual copy weights are not reconstructed by this surface. These statements do not introduce a source-faithful canonical-form existence theorem and do not recover the repeated-copy sector data of the source BNT decomposition. Their docstrings and the corresponding Wielandt blueprint corollary therefore carry the same basis-of-representatives marker.

This audit does not apply that marker to the sector-BNT development based on `IsBNTCanonicalForm`. That predicate is the paper-faithful sector surface used for the CPSV16 coefficient and weight-multiplicity arguments, where repeated copies inside one sector remain explicit.

6 Conclusion

The present formalization proves the proportional Fundamental Theorem on the already selected basis-of-representatives surface. The general case of Corollary `II_cor2` in [CPGSV16], allowing repeated copies inside a sector, has since been supplied on the SectorBNT surface. The standalone Newton-Girard uniqueness input is formalized as `Matrix.sum_pow_eq_implies_card_eq_and_multiset_eq_of_le_max_card`; the BNT coefficient identity from the equal-MPV hypothesis is derived by `MPSTensor.coeff_identity_via_global_gauge`, using the phase-valued auxiliary result `MPSTensor.coeff_identity_via_matched_mpv_phase`; the resulting power-sum expression and weight recovery are supplied by `MPSTensor.SectorDecomposition.coeff_eq_sum_weight_pow`, `MPSTensor.matched_sector_weight_multiset_eq`, and `MPSTensor.matched_sector_weight_equiv`; and the global gauge assembly with the multiplicity identity factors is provided by `MPSTensor.ft_sector_bnt_equal_global_gauge`. The corresponding tracked items issue #1560, issue #1561, and issue #1562 are closed.

Thus this note remains a record of why the older one-copy surface is restricted. The paper-faithful multi-copy route is the `MPSTensor.IsBNTCanonicalForm` / SectorBNT declaration path, and it does not carry the one-copy

[3] The relevant Lean surfaces are `IsNormalCanonicalFormBNT`, `IsNormalCanonicalFormBNT.ofSeparatedData`, `IsNormalCanonicalFormBNT.cross_overlap_tendsto_zero`, `IsNormalCanonicalFormBNT.isBNT`, `IsNormalCanonicalFormBNT.exists_common_blockTensor_isInjective`, and `exists_common_blockTensor_isInjective_two_of_isNormalCanonicalFormBNT`.

scope-restriction marker. The local scope-restriction marker on the one-copy BNT comparison in the blueprint remains correct.

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.