Scope of the Formal Canonical-Form Reductions Relative to PGVWC07

2026-05-12

Verdict: scope-restriction (resolved).

1 The source theorem

The translation-invariant canonical-form theorem of Pérez-García, Verstraete, Wolf, and Cirac starts from an arbitrary translation-invariant MPS representation on a finite ring [PGVWC07, Theorem `Th:TIcanonical`, lines 742–763]. It states that the matrices can always be decomposed as

$$A_i = \begin{pmatrix} \lambda_1 A_i^1 & 0 & 0 \\ 0 & \lambda_2 A_i^2 & 0 \\ 0 & 0 & \cdots \end{pmatrix},$$

where $1 \geq \lambda_j > 0$ and each block satisfies:

$$\sum_i A_i^j A_i^{j\dagger} = \mathbf{1}, \qquad \sum_i A_i^{j\dagger} \Lambda^j A_i^j = \Lambda^j,$$

for a diagonal positive full-rank matrix Λ^j, and $\mathbf{1}$ is the only fixed point of

$$\mathcal{E}_j(X) = \sum_i A_i^j X A_i^{j\dagger}.$$

If the initial representation has bond dimension D, the canonical representation has bond dimension at most D.

The statement does not assume injectivity, a condition C1, blockwise primitivity, a prior block decomposition, strict separation of the weights, or that only one block occurs in each weight class. These properties enter elsewhere in the paper as later hypotheses, consequences, or genericity conditions. For example, the injectivity statement appears only after the canonical form has been obtained [PGVWC07, Proposition `prop-inj`, lines 911–950], and the generic uniqueness discussion is a separate theorem [PGVWC07, Theorem `thm-uniq`, lines 1002–1015].

2 Faithful proof order

The proof of Theorem `Th:TIcanonical` is short but highly ordered, and the formalization should preserve that order rather than replace it by later canonical-form theorems. The source proof proceeds as follows.

1. Lines 765–770 normalize the spectral radius and, when the positive fixed point is full rank, gauge the tensor by $X^{-1/2}$ and $X^{1/2}$ to obtain the unital condition $\sum_i B_i B_i^\dagger = \mathbf{1}$.

2. Lines 771–783 treat a non-full-rank positive fixed point. The support projection P_R is not an independent hypothesis: it is derived from the spectral decomposition of the singular positive fixed point, using the positivity contradiction in lines 775–783. This is the point at which the invariant-subspace hypothesis is proved; it is not an assumption imported from a prepared block decomposition.

3. Lines 785–815 use the derived relation $A_i P_R = P_R A_i P_R$ to split the trace over the finite ring into the R and $R^\perp$ restrictions. The cyclic trace calculation is a finite-ring argument: the P_R-term passes through the product and cancels against $P_{R^\perp}$. This is the source of the recursive block decomposition and of the bond-dimension non-increase.

4. Lines 816–826 iterate the same splitting on each block and then use a non-scalar fixed point $X \neq \mathbf{1}$ to produce the positive non-full-rank fixed point $\mathbf{1} - \lambda_1^{-1} X$. This is the step that enforces uniqueness of the identity fixed point in each final block.

5. Lines 827–832 apply the analogous argument to the dual map and use a unitary conjugation to diagonalize the resulting full-rank positive fixed point, producing the matrices Λ^j in the theorem.

Thus the source theorem is not a primitive-block theorem, an injective-block theorem, or a prepared-family theorem. Nor is the invariant-projection lemma by itself enough: the projection must be obtained from the positive fixed-point argument in the cited lines, and the finite-ring trace split must then be carried out before the recursive decomposition is asserted. A formal proof of Theorem `Th:TIcanonical` should reconstruct these steps from an arbitrary translation-invariant representation and use later notions only as lemmas after their hypotheses have been obtained in this proof order.

3 Normalization orientation

There is also a convention difference in the direction of the canonical normalization. In the source theorem, the block transfer map is written as

$$\mathcal{E}_j(X) = \sum_i A_i^j X A_i^{j\dagger}$$

and imposes the unital condition

$$\sum_i A_i^j A_i^{j\dagger} = \mathbf{1}.$$

The local canonical-form chapter usually works in the dual trace-preserving, or left-canonical, orientation

$$\sum_i A_i^{j\dagger} A_i^j = \mathbf{1}.$$

This is a convention difference when it is accompanied by the corresponding adjoint transfer map or conjugate-transposed Kraus family. It is nevertheless important to state the convention explicitly: a theorem written only in the left-canonical orientation should not be read literally as PGVWC07 Theorem `Th:TIcanonical` unless the dualization or gauge conversion has also been supplied.

4 The formal reductions

The currently checked statements in the canonical-form reduction chapter prove conditional reductions from prior block hypotheses. In particular, the primitive fixed-point reduction assumes that the blocks are injective, trace-preserving in the chosen gauge, strictly ordered by weight modulus, nonzero, positive-dimensional, and primitive. The peripheral-spectrum variant assumes the same injectivity, normalization, strict weight ordering, nonzero weights, and positive block dimensions, together with the statement that each block transfer map has peripheral spectrum $\{1\}$. The normal-canonical-form reduction starts from a weighted block decomposition already equivalent to the original state and assumes irreducibility, trace-preserving normalization, primitivity, pairwise distinct weight moduli, nonzero weights, and positive bond dimensions.

These are correct conditional consequences, but they are not formalizations of the full PGVWC07 translation-invariant canonical-form theorem. Their hypotheses are stronger and more structured than the source theorem's hypotheses. Treating them as the full source theorem would import extra assumptions into any later argument that cites canonical-form existence.

5 Current proof-path audit

The earlier existence path now has the following formal pieces.

1. `MPSTensor.spectralUnitalGauge_isUnital_of_transferMap_eigenvector` formalizes the spectral-radius normalization part of Theorem `Th:TIcanonical`, lines 765–769. Given a positive definite eigenvector ρ with $\mathcal{E}_A(\rho) = r\rho$ and $r > 0$, it proves that $B_i = r^{-1/2}\rho^{-1/2}A_i\rho^{1/2}$ satisfies $\sum_i B_i B_i^\dagger = \mathbf{1}$.

2. `MPSTensor.unitalGauge_isUnital_of_transferMap_fixedPoint` formalizes the full-rank fixed-point gauge in Theorem `Th:TIcanonical`, lines 767–769. Given a positive definite fixed point ρ of $\mathcal{E}_A$, it constructs the right-canonical gauge $B_i = \rho^{-1/2} A_i \rho^{1/2}$ and proves $\sum_i B_i B_i^\dagger = \mathbf{1}$. The companion theorem `MPSTensor.sameMPV_unitalGauge` records preservation of the finite-ring coefficients under this similarity.

3. `MPSTensor.exists_singular_posSemidef_fixedPoint_of_unital_nonScalar_fixedPoint` formalizes the fixed-point shift in Theorem `Th:TIcanonical`, lines 819–826. For a unital block, a Hermitian fixed point which is not a scalar multiple of the identity gives the singular positive fixed point $\lambda_{\max}\mathbf{1} - X$. The companion theorem `MPSTensor.exists_twoBlock_decomp_of_unital_nonScalar_fixedPoint` composes this shift with the support-projection split, giving a strict two-block MPV-equivalent decomposition. The formal hypothesis uses "non-scalar" rather than merely $X \neq \mathbf{1}$, since every scalar multiple of the identity is fixed by a unital linear map and cannot force a nontrivial split.

4. `MPSTensor.fixed_eq_scalar_of_isIrreducibleTensor_unital` formalizes the fixed-point endpoint in Theorem `Th:TIcanonical`, lines 816–826. Once the block is irreducible and unital, every fixed point of its transfer map is a scalar multiple of $\mathbf{1}$. The proof uses the MPS irreducible-tensor/irreducible-map bridge and the already formalized channel theorem `Kraus.fixed_eq_scalar_of_irreducible_unital` (Wolf Theorem 6.6). Thus the identity is the unique fixed point after the usual scalar normalization of the fixed-point line.

5. `MPSTensor.exists_unitary_diag_posDef_adjointFixedPoint_of_unital_of_isIrreducibleTensor` formalizes the dual-map diagonalization in Theorem `Th:TIcanonical`, lines 827–832, conditional on having already reached an irreducible unital block. The proof applies the existing trace-preserving diagonalization theorem to the conjugate-transposed Kraus family $i \mapsto A_i^\dagger$. This is not a duplicate spectral proof: unitality of A is trace preservation for $A^\dagger$, and translating back gives a unitary gauge $B_i = U^\dagger A_i U$ with

$$\sum_i B_i B_i^\dagger = \mathbf{1}, \qquad \sum_i B_i^\dagger \Lambda B_i = \Lambda,$$

for diagonal positive-definite Λ.

6. `MPSTensor.exists_unital_data_of_irreducible` formalizes the Perron–Frobenius unital-gauge step in Theorem `Th:TIcanonical`, lines 765–770, for a single irreducible nonzero block. It obtains a positive-definite eigenvector ρ for the transfer map by applying the adjoint Perron–Frobenius theorem to the conjugate-transposed Kraus family, and then forms $B_i = r^{-1/2}\rho^{-1/2} A_i \rho^{1/2}$ with $\sum_i B_i B_i^\dagger = \mathbf{1}$. This is a scoped single-block input, not the full recursive canonical-form theorem.

7. `MPSTensor.exists_pgvwc07_unital_dualDiag_data_of_irreducible` composes the single-block source steps in Theorem `Th:TIcanonical`, lines 765–770 and 816–832. Starting from one irreducible nonzero block, it constructs the Perron–Frobenius unital representative $B_i = r^{-1/2}\rho^{-1/2}A_i\rho^{1/2}$, proves that the fixed points of $\mathcal{E}_B$ are scalar multiples of the identity, and then applies the dual fixed-point diagonalization to a unitary conjugate of B. This removes a previous composition gap for a single block. It is still not the full source theorem, because the construction has not yet been threaded through the recursive finite-ring block decomposition, the positive weights, and the total bond-dimension bound.

8. `MPSTensor.exists_pgvwc07_unital_dualDiag_blockwise` applies the preceding single-block theorem to a prepared family of nonzero irreducible summands with unit weights. It constructs the positive spectral-radius weights, unital block representatives, scalar transfer-map fixed points, and diagonal positive-definite dual fixed points block by block, while preserving the weighted MPV family. This closes the blockwise assembly step after the irreducible summands have already been obtained. It is still not the full source theorem, because the statement assumes the prepared nonzero irreducible decomposition.

9. `MPSTensor.exists_pgvwc07_unital_dualDiag_from_arbitrary` composes the positive-length nonzero-block decomposition with the preceding blockwise PGVWC07 theorem. Starting from an arbitrary tensor, it discards all-zero blocks, which vanish on nonempty rings, and puts the retained blocks into the unital orientation, with positive spectral-radius weights, scalar transfer-map fixed points, and diagonal positive-definite dual fixed points. The structural dimension identity for the original direct sum gives $\sum_k D_k \le D$; no empty-word trace calculation is used.

10. `MPSTensor.PGVWC07PositiveLengthWitness` and `MPSTensor.exists_pgvwc07_positiveLengthWitness` record the preceding positive-length theorem as a single package of data. The fields contain the nonzero unital blocks, positive real weights, diagonal positive-definite dual fixed points, scalar transfer-map fixed-point conclusions, positive block dimensions, positive-length MPV equality, and the total bond-dimension bound. Nonvanishing of the weights is a direct consequence of their positive-real form, so nonvanishing is not an additional field of the witness. This is not a new proof route; it is an equivalent formulation of the proved positive-length theorem, intended to prevent later arguments from repeating the long existential conclusion.

11. `MPSTensor.PGVWC07PositiveLengthWitness.exists_weight_normalization` formalizes the finite-family scalar normalization corresponding to Theorem `Th:TIcanonical`, lines 765–766, for a witness with at least one nonzero block. Dividing the positive real weights by their maximum

gives normalized positive weights with $|\nu_k| \leq 1$ and equality for one block. The theorem also records the resulting global factor s^N in length-N MPV coefficients.

12. `MPSTensor.PGVWC07PositiveLengthWitness.block_count_pos_of_exists_ne_zero_mpv` proves that a positive-length witness is nonempty whenever the original tensor has at least one nonzero positive-length MPV coefficient. If the block family were empty, the block-diagonal MPV expansion would be an empty sum at every positive length, contradicting the witness equality with the chosen nonzero coefficient.

13. `MPSTensor.exists_pgvwc07_normalized_exact_form_after_rescaling_of_exists_ne_zero_mpv` composes the arbitrary-input witness theorem with the preceding nonemptiness lemma and the exact weight-normalization theorem. Under the nonzero positive-length MPV hypothesis, it produces a positive scalar s and normalized weights $|\nu_k| \leq 1$, with equality for one block, such that the globally rescaled tensor $s^{-1}A$ has exactly the same positive-length MPV coefficients as the normalized weighted block tensor. This formalizes the exact finite-ring effect of the "without loss of generality" spectral-radius normalization in Theorem `Th:TIcanonical`, lines 765–766.

14. `MPSTensor.exists_pgvwc07_normalized_exact_form_after_rescaling_or_forall_pos_mpv_eq_zero` records the arbitrary-input zero/nonzero dichotomy for the exact rescaled formulation. Either all positive-length MPV coefficients of the original tensor vanish, or the preceding exact rescaled normalized PGVWC07 form exists. The next theorem absorbs the zero branch into an empty retained block family.

15. `MPSTensor.exists_pgvwc07_normalized_exact_form_after_rescaling_allow_empty` removes the disjunction in the preceding theorem by adopting the positive-length convention that the zero branch has an empty nonzero-block family. Thus the positive-weight, unital, scalar-fixed-point, and dual-diagonal block conditions are vacuous in the branch where every positive-length MPV coefficient vanishes. In the nonzero branch the family is nonempty and still has a normalized weight of norm one. This is an exact positive-length formulation: after a positive global rescaling of the original tensor, the normalized weighted block tensor has the same coefficients on every nonempty ring, and $\sum_k D_k \leq D$.

16. `MPSTensor.exists_irreducible_blockDecomp` formalizes the recursive invariant-projection splitting used in the proof of Theorem `Th:TIcanonical`, lines 765–833. It starts from an arbitrary tensor and gives a finite direct sum of irreducible blocks with the same matrix-product-vector coefficients and the structural identity $\sum_k D_k = D$. This is a faithful partial proof step. It is not the full source theorem, because it does not construct the weights λ_j, the unital block orientation, the

diagonal full-rank matrices Λ^j, the uniqueness of the identity fixed point, or the final bond-dimension bound.

17. `MPSTensor.exists_irreducible_blockDecomp_nonzeroBlocks` discards the all-zero irreducible summands at positive lengths and retains only blocks with a nonzero Kraus operator. It proves positive-length MPV equality, positive block dimensions, and $\sum_k D_k \leq D$ directly from the preceding structural identity.

18. `MPSTensor.exists_tp_gauge_from_arbitrary` applies Perron–Frobenius normalization block by block after the all-zero summands have been discarded. It produces nonzero irreducible blocks in the local left-canonical orientation $\sum_i (B_k^i)^\dagger B_k^i = \mathbf{1}$ and preserves the bond-dimension bound. This is the dual orientation of the unital condition in Theorem `Th:TIcanonical`; the conversion is a convention issue only when the adjoint transfer map is carried along explicitly.

19. `MPSTensor.exists_normalCanonicalForm_of_primitive_blockDecomp` is a prepared-block theorem. It assumes a weighted block decomposition, irreducibility, trace-preserving normalization, primitivity, pairwise distinct weight moduli, nonzero weights, and positive bond dimensions. It should remain a conditional corollary, not a citation target for PGVWC07 Theorem `Th:TIcanonical`.

Thus the primitive, peripheral-primitive, and normal-canonical-form reductions listed above should not be read as the PGVWC07 translation-invariant canonical-form theorem. The blueprint instead uses the checked positive-length statement as the theorem statement for Theorem `Th:TIcanonical`. The finite-family part of the paper's convention $1 \geq \lambda_j > 0$ is formalized by rescaling all positive weights by their maximum. In exact coefficients this introduces the global factor s^N at length N, so the primary exact statement rescales the original tensor by s^{-1}.

6 Formal boundary

The currently checked primitive, peripheral-primitive, and normal-canonical-form reductions should therefore be read as conditional canonical-form theorems rather than as formalizations of the arbitrary-input source theorem.[1] They may be used after their hypotheses have been proved from earlier source-faithful reductions, but they should not be cited as giving canonical form directly from an arbitrary translation-invariant MPS representation.

[1]The formal declarations are listed in Section 5.

7 Current conclusion

The formalization now has a theorem cited for PGVWC07 Theorem `Th:TIcanonical`, in the positive-length convention used by the blueprint: `MPSTensor.exists_pgvwc07_normalized_exact_form_after_rescaling_allow_empty`. It starts from an arbitrary translation-invariant MPS representation. After a positive global rescaling, it constructs a finite family of positive weights, unital blocks, diagonal full-rank dual fixed points, scalar transfer-map fixed points, and the bond-dimension bound. The nonzero-block family is allowed to be empty exactly in the case where every positive-length MPV coefficient vanishes.

The preceding theorem `MPSTensor.exists_pgvwc07_unital_dualDiag_from_arbitrary` supplies the unnormalized positive-length block family and its bond-dimension bound directly, without a separate empty-word convention.

The remaining conditional primitive, peripheral-primitive, and normal-canonical-form reductions are later consequences used only after their extra hypotheses have been established. They are not alternate citations for Theorem `Th:TIcanonical`.

References

[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.