The Commuting-Form-to-SAL Implication for Simple MPDOs

2026-07-31

Verdict: scope-restriction (resolved).

1 Source statement

The proposition labelled *4to2* in [CPGSV16, Appendix C.2, lines 1597–1619] considers a normal MPDO block $\mathcal{K}$. Its finite-chain operators are positive scalar multiples of products of translates of one positive two-site bond,

$$\sigma^{(N)}(\mathcal{K}) = c_N \prod_{i=0}^{N-1} B_{i,i+1}, \qquad c_N > 0, \qquad [B_{i,i+1}, B_{j,j+1}] = 0.$$

Together with zero correlation length, this hypothesis is asserted to imply saturation of the area law.

The phrase "of the form" refers back to the single bond in equation *expression* at source lines 1571–1578. It is therefore a statement about one translation-invariant bond family, not a separate existence statement at each chain length.

2 The theorem cited through Beigi

The reference denoted by *Beigi* is Salman Beigi, *Classification of the phases of 1D spin chains with commuting Hamiltonians*, arXiv:1105.1019v2, J. Phys. A **45** (2012) 025306, doi:10.1088/1751-8113/45/2/025306.

The relevant assertion is Lemma 2.1, labelled *lem:comm* in the arXiv source and credited there to Bravyi and Vyalyi. If Hermitian operators X_{AB} and Y_{BC} commute after their natural embeddings, then

$$\mathcal{H}_B \cong \bigoplus_{\beta \in V} \mathcal{H}_{\beta_l} \otimes \mathcal{H}_{\beta_r},$$

such that X_{AB} acts trivially on every $\mathcal{H}_{\beta_r}$ and Y_{BC} acts trivially on every $\mathcal{H}_{\beta_l}$. Beigi applies this lemma to neighboring translates of one local interaction. Translation invariance permits the same one-site decomposition at every site; the local decomposition is independent of the system size.

The proof uses the representation theory of finite-dimensional C^*-algebras. It first forms the commutant of the operator-Schmidt factors of X_{AB} and then identifies that algebra, after a unitary change of basis, with

$$\bigoplus_{\beta\in V} \mathbf{1}_{\beta_l} \otimes \mathbf{L}(\mathcal{H}_{\beta_r}).$$

The commutant has the complementary action on the left tensor factors.

3 Present formal distinction

The predicate `MPOTensor.HasCommutingForm` is chainwise: for every $N \geq 2$ it allows a new two-site bond B_N. It is weaker than the source hypothesis and does not support the comparison of the N, $N+1$, and $N+2$ chains used by the source trace-factorization argument.

The formal model fixes one bond whose commuting translates realize every finite chain.[1]

The source ZCL hypothesis is `MPOTensor.IsSourceZCL`, the scale-invariant idempotence condition for the physical-trace transfer. The predicate `MPOTensor.IsZCL` concerns the doubled-index completely positive transfer map and is not a faithful substitute for the proposition labelled *4to2*.

4 Resolution of the implication

The finite-dimensional C^*-algebra part of the Bravyi–Vyalyi argument is now available. For a star-subalgebra $S \subseteq \mathbf{L}(H_B)$, one orthonormal identification

$$H_B \cong \bigoplus_{\beta} H_{\beta_l} \otimes H_{\beta_r}$$

simultaneously puts S in $\bigoplus_\beta \mathbf{1}_{\beta_l} \otimes \mathbf{L}(H_{\beta_r})$ and its commutant in $\bigoplus_\beta \mathbf{L}(H_{\beta_l}) \otimes \mathbf{1}_{\beta_r}$. This complementary action is `StarSubalgebra.exists_complementary_action_orthonormalBasis`.

The coefficient passage in Beigi's proof is now also available. Fixing the outer indices produces middle-site matrices $X_{aa'}$ and $Y_{cc'}$; commutation of the two overlapping operators implies $[X_{aa'}, Y_{cc'}] = 0$ for all outer indices. Applying the complementary-action theorem to the star-commutant of the $Y_{cc'}$ gives one orthonormal decomposition on which all $X_{aa'}$ and $Y_{cc'}$ act on complementary factors. This is `Matrix.exists_middle_spatial_decomposition_of_overlappingLifts_commute`.

The conjugated coordinate form of Beigi's Lemma 2.1 is `Matrix.exists_unitary_blockActions_of_overlappingLifts_commute`. It constructs one

[1] The bond family is `MPOTensor.TranslationInvariantBondData`; the realization equalities are fields of `MPOTensor.EtaLocalStructureData`.

unitary identification of the middle space and combines the coefficient matrices into Hermitian operators $R_{A\beta_l}$ and $S_{\beta_r C}$. Applying the inverse coordinate identifications and cancelling the unitary factors gives the two equalities for the original operators in Beigi's statement. This is `Matrix.exists_unitary_blockSum_equalities_of_overlappingLifts_commute`. Its hypotheses are exactly the source hypotheses: Hermiticity of both overlapping operators and commutation of their natural three-factor lifts. The block operators are Hermitian, as in the source; no unitarity property is asserted for them.

The cyclic three-site coordinate identification is now available. The equivalence `MPOTensor.threeSiteOverlappingEquiv` sends a three-site configuration to the ordered left-associated triple. Under this one common equivalence, `MPOTensor.EtaLocalStructureData.reindex_bondAt_zero_eq_leftOverlappingLift` and `MPOTensor.EtaLocalStructureData.reindex_bondAt_one_eq_rightOverlappingLift` identify the first two adjacent translates with the left and right overlapping lifts of the same pair-indexed bond. This bond is Hermitian by positivity, and `MPOTensor.EtaLocalStructureData.overlappingLifts_pairBond_comm` transports their length-three commutation to the abstract lifts. Hence the formal hypotheses of Beigi's lemma are obtained without an additional assumption.

The application of the finite-matrix theorem to this pair-indexed bond is `MPOTensor.EtaLocalStructureData.exists_unitary_blockActions_of_pairBond`. It supplies one unitary middle-space identification and the two Hermitian block actions required by Beigi's lemma.[2]

The local neighboring operators have now been extracted without assuming a Hayashi decomposition. The theorem `MPOTensor.EtaLocalStructureData.exists_positive_eta_pairBond_decomposition` uses the same one-site unitary at both ends of the fixed bond and proves

$$(U^* \otimes U^*)B(U \otimes U) = \bigoplus_{k,h} \mathbb{1}_{H_{k,l}} \otimes \eta_{k,h} \otimes \mathbb{1}_{H_{h,r}}, \qquad \eta_{k,h} \succeq 0.$$

Since the bond and unitary are chain-independent, this is one neighboring family for every translate.

The bond-local cyclic product is formalized for every $N \geq 2$. The theorem `MPOTensor.reindex_product_embedLocalOperator_of_etaPair_decomposition` constructs the dependent cyclic regrouping

$$H^{\otimes N} \simeq \bigoplus_{k_0,\ldots,k_{N-1}} \bigotimes_{n=0}^{N-1} (H_{k_n,r} \otimes H_{k_{n+1},l})$$

and proves that the product of the translated, locally decomposed bond is $\bigoplus_k \bigotimes_n \eta_{k_n,k_{n+1}}$, with $k_N = k_0$. The length-two specialization is included: its two windows have the opposite cyclic orders $(0,1)$ and $(1,0)$.

[2]The application theorem is in `TNLean/MPS/MPDO/CommutingFormSpatialBridge`; its coordinate and commutation hypotheses are supplied by `TNLean/MPS/MPDO/CommutingOverlappingCoordinates`.

The tensor-power conjugation and positive realization scalar are combined with this cyclic product. The theorem `MPOTensor.EtaLocalStructureData.exists_positive_cyclic_eta_block_decomposition` chooses the unitary, sector dimensions, and positive neighboring family once, independently of the chain length, and proves for every $N \geq 2$ that

$$E_N^*(U^*)^{\otimes N}\sigma^{(N)}(\mathcal{K})U^{\otimes N}E_N = c_N \bigoplus_{k_0,\ldots,k_{N-1}} \bigotimes_{n=0}^{N-1} \eta_{k_n,k_{n+1}}, \qquad c_N > 0, \qquad k_N = k_0.$$

Local fix (proportionality scalar). The converse argument at source lines 1603–1605 starts from the relation $\sigma^{(N)}(\mathcal{K}) \propto \prod_n B_{n,n+1}$ but invokes the coefficient-free equation *sigmaNK2* without retaining the corresponding constant. The formal theorem proves the proportionality-preserving identity displayed above. It does not prove the coefficient-free equation printed at source lines 1581–1588.

This omission cannot be repaired from the proportional commuting-bond form alone. Take physical dimension $d = 1$ and bond dimension $D = 2$, with sole matrix-product letter $M^{00} = I_2$. Then

$$\sigma^{(N)} = \mathrm{tr}(I_2^N) = 2.$$

The scalar two-site bond $B = 1$ realizes this family with $c_N = 2$ for every N. Since the physical space is one-dimensional, any positive neighboring decomposition has a single one-dimensional sector and a single scalar $a \geq 0$. A coefficient-free identity for all $N \geq 2$ would therefore give $a^2 = 2$ and $a^4 = 2$, which is impossible. Thus a fixed local rescaling does not follow from the displayed proportionality data.

The example is not a normal injective block. It therefore does not contradict the source setting, but shows why a scalar-rescaling route would have needed to use normal injectivity. Pursuing that route would require a comparison of the matrix-product family with the fixed bond-product family proving that its proportionality constants have the form

$$c_N = a^N$$

for one $a > 0$. Only then can a be absorbed once into the bond, or equivalently into the neighboring operators, independently of N. The completed proof does not use this route: it compares the normalized marginals of the selected fixed-product tensor with those of the source tensor directly.

The finite-state part of this comparison is now isolated. The tensor `MPOTensor.cyclicEdgeWeightTensor` has virtual dimension d^2 and generates exactly one prescribed cyclic nearest-neighbor scalar weight at every positive length. The datum `MPOTensor.TranslationInvariantBondData.FixedProductTensorData` records one positive-dimensional tensor whose operator family is exactly the product of the translates of the fixed bond for every $N \geq 2$. Given such a datum, the theorem `MPOTensor.EtaLocalStructureData.eventuallyNonzeroProportionalMPV_fixedProductTensor` derives eventual nonzero proportionality to the source tensor without adding a length-one hypothesis.

The fixed tensor is now constructed directly from an arbitrary MPOTensor.TranslationInvariantBondData. The theorem MPOTensor.TranslationInvariantBondData.fixedProductTensorData assembles the chain-independent Beigi sector coordinates and neighboring operators into a cyclic scalar weight in the transformed one-site coordinates, then conjugates the resulting tensor back through the same one-site unitary. It uses only positivity and commutation of the periodic translates.

The selected fixed tensor now carries the same positive physical-sector factorization used in its construction. More precisely, MPOTensor.TranslationInvariantBondData.exists_positive_physicalSectorFactorization_fixedProductTensorData gives a factorization of the exact tensor in MPOTensor.TranslationInvariantBondData.fixedProductTensorData; its neighboring operators are positive semidefinite and its reconstructed physical bond is the original commuting bond. Thus no equality of closed matrix-product operators is used to replace one local tensor by another. This completes the fixed-representative construction only. It neither factorizes the original source tensor nor proves a normality or blocking comparison between that tensor and the selected fixed tensor.

The coordinate qualification is necessary. In physical dimension $d = 2$, let X be the Pauli flip and $B = \mathbb{1} + X \otimes X$. Then $B \succeq 0$, all periodic translates commute, and

$$\langle 0^N, (\prod_n B_{n,n+1}) 0^N \rangle = 2$$

for every $N \geq 2$. At $N = 2$ the repository's periodic product is $B^2 = 2B$; for $N \geq 3$ the empty edge set and the full cycle give the two terms with zero boundary. A scalar edge weight in the original coordinates would make this entry a^N, so the cases $N = 2$ and $N = 3$ would require both $a^2 = 2$ and $a^3 = 2$. Thus an arbitrary fixed bond does not carry MPOTensor.TranslationInvariantBondData.CyclicEdgeWeightForm in the original coordinates. Equation *sigmaNK2* gives that scalar form only after the one-site unitary coordinate change.

An arbitrary tensor representing the products need not become normal after one positive rescaling. Indeed, adjoining a zero virtual direct-sum block changes neither the closed operator at any positive length nor the product identity, but no nonzero rescaling removes that invariant summand. Proposition *3to4* does not assert normality or uniqueness for such an arbitrary representation.

For the bond selected in the SAL proof, no canonical-form reduction is needed. The source physical-sector factorization gives the coefficient-one identity directly, so the original source tensor itself represents the bond product. This is MPOTensor.PhysicalSectorFactorization.fixedProductTensorData. Under the normal single-block hypothesis, its normality and the exact all-length identity are recorded by MPOTensor.exists_normal_fixedProductTensorData_of_isSAL. Thus the source-selected presentation has rescaling $s = 1$.

For a different proportional commuting-bond presentation, the conditional theorem `MPOTensor.EtaLocalStructureData.mpo_eq_positive_power_product_of_normal_smul` shows that if sC is normal for one $s > 0$, then

$$\sigma^{(N)}(\mathcal{K}) = s^N \prod_n B_{n,n+1} \qquad (N \geq 2).$$

First, the normal-tensor fundamental theorem gives a possible unit phase ζ^N. The self-overlap of the normal source tends to one, so the fixed product is nonzero at two sufficiently large consecutive lengths. The positive realization constants at those lengths force $\zeta = 1$. This is a valid strengthening for a presentation already known to have a normal rescaling. The existence of such a rescaling for every proportional commuting-bond presentation is not claimed in Appendix C.2 and is not needed for the source-selected SAL-to-commuting-form implication.

The scaled identity uses only the fixed commuting bond and its positive realization scalar. It does not assert zero correlation length, rank-one trace factorization, a quantum Markov decomposition, saturation of the area law, or the proposition labelled *4to2*.

The fourth-region calculation after the cyclic identity requires a coordinate clarification. Tracing one physical site of the length-N cyclic eta product leaves a coupled sum of the right partial trace of $\eta_{k_{N-1},q}$ and the left partial trace of η_{q,k_1}:

$$\sum_q \operatorname{tr}_{H_{q,l}}\left(\eta_{k_{N-1},q}\right) \otimes \operatorname{tr}_{H_{q,r}}\left(\eta_{q,k_1}\right).$$

The same sector q occurs in both factors, so rank-one factorization does not separate this into two independent sums. Source line 1613 instead invokes ZCL to replace this marginal by a trace of two boundary sites on the length-$(N+1)$ chain. In edge coordinates the middle operator $\eta_{q,h}$ is then fully traced, while the two adjacent operators retain their right and left partial traces. The displayed formula at source line 1616 suppresses these partial-trace symbols.

The conditional algebraic identity after this replacement is formalized by `Matrix.eta_fourth_region_trace_formula`. It assumes $\operatorname{tr}(\eta_{q,h}) = a_q b_h$ and separates the two boundary sums. It does not prove the ZCL replacement of the one-site marginal, derive the rank-one factorization, construct a quantum Markov decomposition, or prove SAL.

The ZCL replacement itself is now formalized. If E is the physical-trace transfer and $E^2 = \lambda E$ with $\lambda > 0$, then for every $N \geq 0$ the normalized N-site marginal obtained by tracing two sites from the length-$(N+2)$ chain equals the one obtained by tracing one site from the length-$(N+1)$ chain. Entrywise, the two expressions are

$$\frac{\operatorname{tr}(M^{u,v}E^2)}{\operatorname{tr}(E^{N+2})} \quad \text{and} \quad \frac{\operatorname{tr}(M^{u,v}E)}{\operatorname{tr}(E^{N+1})},$$

and the numerator and denominator of the first are both multiplied by λ. This is `MPOTensor.reducedBlockState_succ_succ_eq_succ_of_isSourceZCL`.

It proves the first sentence of the source's line 1613 argument, but it does not separate the two boundary sector sums.

There is a separate obstruction at source line 1613. The phrase "arguing as in Lemmas *propSN* and *SALZCL*" does not show that source ZCL alone gives $\mathrm{tr}(\eta_{k,h}) = a_k b_h$. Lemma *propSN* obtains primitivity from SAL, which is the conclusion being proved in Proposition *4to2*; Lemma *SALZCL* then uses the invalid trace-power-to-rank-one inference recorded in `docs/paper-gaps/cpgs v17_pf_rank_one.tex`. Formally, source ZCL yields an idempotent normalized pairing operator, but for the opposite product it yields only a square–cube identity and constant positive-power traces. The four-sector example `MPOTensor. ActiveSectorSpanningCounterexample.virtual_spanning_does_not_for ce_rectangular_remainder_zero` satisfies the recorded positivity, primitivity, trace-one, and spanning conditions while its trace matrix is not idempotent. The capstone theorem `MPOTensor.ActiveSectorSpanningCounterexample.ten sor_has_sal_sourceZCL_and_non_rankOne_activeTraceMatrix` adds SAL, TNLean's up-to-scalar source-ZCL predicate, and the source-selected inverse-map factorization, and proves that its active trace matrix is not rank one. This failure belongs to the selected four-sector description: the same tensor has a one-sector physical factorization with positive trace-one neighboring operator, formalized in `MPOTensor.ActiveSectorSpanningCounterexample.exists_ oneSector_neighboringTraceFactorization`. Thus the four-sector trace matrix is not an existential obstruction to a different factorization. The raw tensor is also not the normal tensor assumed in Appendix C.2 Case I, and its normal representative fails the paper's literal ZCL diagram. The example therefore isolates two proof-path boundaries rather than refuting Lemma C.5 under its complete standing hypotheses. The claimed rank-one conclusion from source ZCL alone remains unavailable and is not used in the completed proof, which instead obtains a positive rank-one factorization of the two-step coefficient.

Primitivity cannot be assumed for an arbitrary selected Beigi factorization. The theorem `MPOTensor.CommutingBondTraceMatrixObstruction.norm al_sourceZCL_fixed_commutingBond_does_not_force_traceMatrix_pr imitive` gives an injective normal bond-dimension-one product MPDO with source ZCL and one exact fixed positive commuting-bond presentation. The chosen two-sector Beigi-style factorization has positive neighboring operators and trace matrix

$$T = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}.$$

The second physical summand is inactive, but it is retained by this nonminimal middle-space decomposition, so its full matrix is not primitive. The same product tensor also has a one-sector factorization with the primitive 1×1 trace matrix (1). Thus this example does not exclude the existence of a primitive Beigi factorization. It shows instead that Beigi's Lemma 2.1, which supplies only the direct-sum tensor-factor decomposition, must be supplemented by a source-faithful minimal or visible-sector selection before primitivity can be used.

Such a visible-sector selection is now available without SAL. A sector is called cyclic-active when it has an outgoing nonzero neighboring operator followed by a possibly empty return path. The explicit first edge excludes the empty reflexive walk. This is `MPOTensor.PhysicalSectorFactorization.IsCyclicActiveSector`. Every nonzero finite cyclic neighboring product uses only cyclic-active sectors, and therefore deleting all other sectors leaves every finite cyclic bond product unchanged; see `MPOTensor.PhysicalSectorFactorization.cyclicNeighboringProduct_eq_zero_of_not_isCyclicActiveSector`.

For an injective tensor with positive semidefinite neighboring operators, the cyclic-active sectors are nonempty and their trace matrix is primitive. The proof uses the nondegeneracy of the matrix trace pairing twice. First, every nonzero virtual sector lies on a directed two-cycle, so all virtual matrices outside the cyclic-active set vanish and the retained virtual matrices still span the full matrix algebra. The directed-cut argument then gives one strongly connected component. Second, every retained sector has a return of length two, and every retained edge closes through a third retained sector; the resulting coprime return lengths give aperiodicity. This is `MPOTensor.PhysicalSectorFactorization.cyclicActiveSectorTraceMatrix_isPrimitive`. No inverse-map sector weights, SAL, or Hayashi decomposition occur in this argument.

Source ZCL gives a positive normalization $\widehat{T} = \lambda^{-1}T$ of this restricted matrix satisfying

$$\widehat{T}^2 = \widehat{T}^3, \qquad \operatorname{rank}(\widehat{T}^2) = 1, \qquad \operatorname{tr}(\widehat{T}^m) = \operatorname{tr}(\widehat{T}) \quad (m \geq 1).$$

The combined result is `MPOTensor.PhysicalSectorFactorization.exists_normalized_cyclicActiveSectorTraceMatrix_rank_pow_two_eq_one_of_isSourceZCL`. It concerns the restricted matrix only. In particular, it neither makes the unreduced trace matrix primitive nor asserts that $\widehat{T}$ itself has rank one.

This last distinction determines how the new result can enter Proposition *4to2*. The two-boundary expression at source lines 1613–1616 contains the one-step coefficient $T_{q,h} = \operatorname{tr}(\eta_{q,h})$. It does not contain T^2. Applying the ZCL marginal replacement once further, and hence tracing three boundary sites of a chain longer by two sites, produces two successive fully traced middle edges. Their coefficient is

$$(T^2)_{q,h} = \sum_{r \in C} \operatorname{tr}(\eta_{q,r}) \operatorname{tr}(\eta_{r,h}).$$

The second iteration of the normalized marginal identity is now `MPOTensor.reducedBlockState_add_three_eq_succ_of_isSourceZCL`. The coefficient calculation in cyclic-active physical-sector coordinates is `MPOTensor.PhysicalSectorFactorization.sum_cyclicActive_partialTraceRight_threeSiteSectorClosure`. Its normalized form, with each traced edge scaled by λ^{-1}, is `MPOTensor.PhysicalSectorFactorization.sum_cyclicActive_normalized_trace_mul_trace_eq_pow_two` and yields precisely the coefficient $(\widehat{T}^2)_{q,h}$. Also, `MPOTensor.PhysicalSectorFactorization.map_cyclicNeighboringProduct_eq_zero_of_not_forall_isCyclicActiveSector` shows that an arbitrary linear coordinate contraction preserves the deletion of every vanishing cyclic summand. Thus the proved rank-one statement supplies the separating boundary coefficient in this three-boundary form.

The normalized marginal equality has now been transported through the complete physical-sector direct sum while retaining the two outer partial traces. The result is `MPOTensor.PhysicalSectorFactorization.exists_cyclicActiveFourthRegion_formula_of_isSourceZCL`. Separating the retained chain at an arbitrary distinguished site, normalizing the positive left and right path factors, and collecting their trace products gives `MPOTensor.PhysicalSectorFactorization.exists_hayashiMarkovDecomposition_cyclicActiveCut_of_isSourceZCL`. Thus the Markov decomposition for arbitrary left and right lengths is available in the physical-sector coordinates. Substituting the admissible entropy-cut lengths, applying the all-cut Markov criterion, and transporting through the physical isometry gives `MPOTensor.PhysicalSectorFactorization.isSAL_of_isSourceZCL`. The cut calculation and basis transport are therefore complete under an explicit positive physical-sector factorization. The older conditional theorem `Matrix.eta_fourth_region_trace_formula` remains a coordinate model for a one-step rank-one coefficient; it is not used to identify T_C with T_C^2.

The explicit-factorization theorem `MPOTensor.PhysicalSectorFactorization.isSAL_of_isSourceZCL` remains a conditional scope-restricted result: by itself it does not prove the source implication

$$\text{one translation-invariant commuting bond} + \text{source ZCL} \Longrightarrow \text{SAL}$$

because it applies to the tensor carrying the factorization. For the source theorem, now formalized by `MPOTensor.EtaLocalStructureData.isSAL_of_isSourceZCL`, the required positive physical-sector factorization is instead constructed on the selected fixed tensor from the bond in equation *sigmaNK2* at source lines 1581–1605. The following comparison transfers the resulting normalized marginals, rather than assuming injectivity, normality, or zero correlation length for the selected tensor.

The recurrent-sector comparison no longer requires a normality or blocking comparison with that selected representative. Express the original injective tensor in the physical coordinates of the selected factorization. The length-two closed-operator identity removes all cross-sector matrices and all matrices outside the positive-length cyclic support. The remaining matrices still span the full auxiliary matrix algebra. Positivity gives a nonzero closed coefficient through every cyclic-active sector, while the length-three identity sends every zero neighboring operator to a zero ordered product of the corresponding two sector spaces. A proper reachability cut would therefore split the full matrix algebra into two nonzero subspaces whose products in one order vanish, which is impossible. This is `MPOTensor.EtaLocalStructureData.selectedFixedProduct_cyclicActive_reflTransGen_neighboringOperator_ne_zero`. It proves that the cyclic-active sectors of the selected factorization form one recurrent component without asserting injectivity, normality, or a virtual gauge for the selected fixed tensor. This closes the selected-sector visibility step tracked in issue #5128.

The coefficient comparison at source lines 1606–1617 is now formalized. Expanding the normalized reduced state after tracing one suffix site and after tracing two suffix sites gives surviving boundary contractions with coefficients T_C^2 and T_C^3, respectively. Recurrence supplies a nonzero retained path between every pair of cyclic-active sectors, so equality of the two reduced states first shows that these two powers have the same positive support. Irreducibility then makes every entry of T_C^2 strictly positive. A positive two-edge summand gives a retained word of the common length three for every ordered pair; comparison at that fixed length yields

$$\widehat{T}^2 = \widehat{T}^3.$$

Here $\widehat{T} = \lambda^{-1} T_C$ and the common normalization is the ratio of the selected length-five and length-four traces. This is `MPOTensor.EtaLocalStructureData.exists_normalized_selectedFixedProduct_cyclicActiveSectorTraceMatrix_pow_two_eq_pow_three`. It uses only source ZCL and injectivity of the original tensor, not of the selected fixed-product tensor. This closes the adjacent coefficient-extraction step tracked in issue #5129.

This square–cube identity has now been assembled with the strictly positive two-step coefficient. Nonemptiness of the cyclic-active sector set follows from the spanning family of original one-site matrices: if the sector set were empty, that family would have zero span rather than the full auxiliary matrix algebra. The normalized trace matrix is therefore primitive, with its second power as a strictly positive power, and `Matrix.IsPrimitive.exists_pos_pow_two_eq_vecMulVec_of_pow_two_eq_pow_three` gives strictly positive vectors a, b satisfying

$$(\lambda^{-1} T_C)^2 = ab^{\mathsf{T}}.$$

The one-step and iterated marginal identities, and the positive selected traces, are transported separately by `MPOTensor.EtaLocalStructureData.selectedFixedProduct_reducedBlockState_succ_succ_eq_succ_of_isSourceZCL`, `MPOTensor.EtaLocalStructureData.selectedFixedProduct_reducedBlockState_add_three_eq_succ_of_isSourceZCL`, and `MPOTensor.EtaLocalStructureData.exists_pos_trace_mpo_selectedFixedProduct_of_isSourceZCL`. Substituting the positive factorization into the three-suffix formula and regrouping at an arbitrary distinguished site produces the Hayashi direct sum. Positive proportionality then identifies the selected normalized marginal with that of the original source tensor in the same one-site physical coordinates. This is `MPOTensor.EtaLocalStructureData.exists_hayashiMarkovDecomposition_selectedPhysicalCut_of_isSourceZCL`. It assumes injectivity of the source tensor, an eta-local structure, and source ZCL. It assumes no MPDO positivity at this intermediate stage and no injectivity, normality, or zero correlation length for the selected fixed-product tensor. MPDO positivity enters only in the final SAL capstone, where it supplies positivity of the source operators in the selected physical coordinates.

Applying the all-cut Markov criterion to these physical-coordinate marginals and transporting saturation through the one-site isometry gives `MPOTensor.EtaLocalStructureData.isSAL_of_isSourceZCL`. Positivity of the physical-coordinate periodic operators follows from the MPDO hypothesis by isometric conjugation, and their traces agree with the original traces, which are nonzero at every positive length by source ZCL. The theorem assumes exactly the source tensor, its MPDO and injectivity hypotheses, the fixed eta-local structure, and source ZCL. It assumes no injectivity, normality, ZCL, or SAL property of the selected fixed-product tensor.

Scope restriction (sectorwise source ZCL). The proposition labelled *prop4to2* at source lines 1801–1808 states that the GSNNCH form alone implies SAL. Its proof applies Proposition *4to2* to each summand, although Proposition *4to2* assumes ZCL at source lines 1597–1600. The surrounding main theorem also includes ZCL in condition (iii), at source lines 851–893. Thus the printed statement omits a hypothesis used by its proof.

The final sentence of the proof is nevertheless a separate mathematical claim: a finite orthogonal direct sum on fixed one-site subspaces satisfies SAL whenever every summand satisfies SAL. The present scope-restricted theorem therefore assumes SAL for the summands, a positive proportional local orthogonal-sum identity at every positive chain length, supplied one-site projections supporting every nonempty sector marginal, and positive natural multiplicities. It must not be presented as a proof that GSNNCH alone implies SAL. The structure `MPOTensor.ProportionalOrthogonalCommutingSectorFamily` records fixed one-site projections and commuting sector products whose periodic products realize the sector operators up to positive scalars. Projection idempotence, bond support, and positive proportional realization imply that the normalized sector chain and every nonempty sector marginal on a chain of length at least two are supported on the matching projection at the first site; these are `MPOTensor.ProportionalOrthogonalCommutingSectorFamily.firstSiteMatrix_mul_normalizedMPO` and `MPOTensor.ProportionalOrthogonalCommutingSectorFamily.firstSiteMatrix_mul_reducedBlockState`. The scalar factors through the first-site action and then cancels under trace normalization. The earlier `MPOTensor.OrthogonalCommutingSectorFamily` is the scalar-one specialization. The chain-length restriction is sufficient here, since the mutual-information comparison in SAL only uses chains of length at least four.

The sector-family structure does not contain the outer tensor or its multiplicities, and therefore cannot by itself supply the proportional outer sum at every positive length. The theorem `MPOTensor.isSAL_of_proportionalLocalOrthogonalSum` assumes numbers $c_N > 0$ such that

$$\sigma^{(N)}(M) = c_N \sum_s n_s \sigma^{(N)}(K_s)$$

and cancels c_N from the normalized chain and all its marginals. It assumes sectorwise SAL and proves the non-circular final local direct-sum implication; it neither derives the sectorwise SAL premise nor uses ZCL. The earlier exact theorem `MPOTensor.isSAL_of_localOrthogonalSum` is its specialization with $c_N = 1$.

In the canonical BNT specialization, ambient source ZCL does supply the missing sectorwise source-ZCL premise. For a simple tensor already in the chosen canonical form, `MPOTensor.IsSimpleCanonicalForm.exists_commonWeightAbsorbedBasisMPOTensor_isSourceZCL` chooses the canonical BNT data, proves that the weights over each normal representative are independent of the copy index, and proves source ZCL for every absorbed representative. Its final step uses the absorbed-transfer identity `MPOTensor.commonWeightAbsorbedBasisMPOTensor_isSourceZCL`. Hence every canonical absorbed simple-BNT sector has source ZCL under ambient source ZCL. This is precisely the sectorwise hypothesis used when Proposition *4to2*, at source lines 1597–1619, is applied sector by sector in Proposition *prop4to2*, at source lines 1801–1808.

This conclusion belongs to the canonical absorbed presentation, not to an arbitrary GSNNCH presentation of the same finite-chain operators. Source ZCL is a property of the tensor's physical-trace transfer, whereas GSNNCH records only the represented periodic operators. For example, in physical dimension one the one-dimensional tensor (1) and the virtual direct sum

$$(1) \oplus \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}$$

generate the same scalar operator at every positive length. The first tensor has source ZCL, while the second transfer has a nonzero nilpotent summand and satisfies no equation $E^2 = \lambda E$ with $\lambda > 0$. Thus representation- level source ZCL cannot be inferred from GSNNCH alone without selecting the canonical absorbed sectors.

The downstream theorem `MPOTensor.isSAL_of_proportionalOrthogonalCommutingSectorFamily_of_sectorwise_isSourceZCL` requires only that each sector generate MPDOs, be one-site injective, and have source ZCL. It does not require normality. Source ZCL itself forces a nonzero bond dimension by `MPOTensor.IsSourceZCL.bondDim_ne_zero`, so no normality assumption is needed to install the nonzero-dimensional bond-space hypothesis in the all-cut Markov argument.

For the canonical absorbed sectors, the standing biCF one-letter span and copy-independent weights give one-site injectivity directly through `MPOTensor.commonWeightAbsorbedBasisMPOTensor_isInjective`. Ambient source ZCL descends to each absorbed sector by `MPOTensor.commonWeightAbsorbedBasisMPOTensor_isSourceZCL`. Thus the only remaining positivity question is the chain length not covered by the commuting-bond product formula.

The proportional orthogonal commuting sector family records a positive commuting bond whose periodic product is proportional to the sector operator by a positive scalar for every $N \geq 2$. Hence `MPOTensor.ProportionalOrthogonalCommutingSectorFamily.mpo_posSemidef_of_two_le` proves positivity at every such length. The theorem `MPOTensor.ProportionalOrthogonalCommutingSectorFamily.isMPDO_of_mpo_one_pos` shows that the sharp additional interface is only

$$\rho^{(1)}(K_s) \geq 0.$$

Together these facts give full sector MPDO positivity. No positivity at length $N \geq 2$ and no spectral normality of an absorbed sector is assumed separately.

The corrected canonical statement is `MPOTensor.isSAL_of_commonWeightAbsorbedBasisMPOTensor_of_proportionalSectors_of_isSourceZCL`. It assumes the chosen canonical decomposition, the reconstructing block-diagonal gauge and positive-length MPV equality, copy-independent weights, nonnilpotence of the representative physical-trace transfers, the standing biCF one-letter product-algebra span, the proportional orthogonal commuting sector family, one-site positivity of every absorbed sector, and ambient source ZCL. The exact canonical BNT outer sum, one-site injectivity, sectorwise source ZCL, and the preceding positivity result then supply every hypothesis of the general orthogonal-sector SAL theorem. The conclusion is SAL at all original cuts.

The length-one premise cannot be removed merely by reading the source's commuting-product display: TNLean's unambiguous periodic bond realization begins at $N = 2$. The standing Case II biCF hypothesis already gives the one-letter simultaneous span needed for injectivity at every original cut.

Verdict. Proposition *4to2* is formalized by `MPOTensor.EtaLocalStructureData.isSAL_of_isSourceZCL`. The positive fixed-product representative, its cyclic-active recurrent component, the normalized square–cube identity, and the positive rank-one factorization of the two-step coefficient are obtained from the source tensor. The resulting Hayashi decomposition holds at every one-site cut, and isometric invariance returns SAL to the original physical coordinates. Normality is not a hypothesis of this theorem.

For Proposition *prop4to2*, the positive proportional outer-sum theorem reduces the conclusion to sectorwise SAL. In the canonical Case II application, ambient source ZCL supplies sectorwise source ZCL, the biCF one-letter span supplies sector injectivity, and one-site positivity completes sector positivity because the commuting bonds already give positivity for all $N \geq 2$. The corrected canonical theorem is therefore complete at the ambient-source-ZCL and one-site-positivity boundary, with no independent absorbed-sector normality or full sector-MPDO premise.

The literal proposition *prop4to2* is still not proved as printed. Its statement assumes only GSNNCH, while Proposition *4to2*, invoked in its proof, also assumes ZCL. The surrounding main theorem supplies ambient ZCL in condition (iii), and the corrected ZCL-conditioned canonical-BNT application is covered by the preceding argument. An arbitrary GSNNCH presentation does not inherit tensor-level source ZCL, as the nilpotent-summand example above shows. Thus ambient source ZCL remains a separate scope restriction on the available implication.

5 Completed proof

1. Apply the all-cut Markov criterion to the decompositions supplied by `MPOTensor.EtaLocalStructureData.exists_hayashiMarkovDecomposition_selectedPhysicalCut_of_isSourceZCL`. Positivity and nonvanishing

of the physical-coordinate periodic operators follow from the MPDO and source-ZCL hypotheses.

2. Transport SAL back through the selected one-site physical isometry. These two steps are `MPOTensor.EtaLocalStructureData.isSAL_of_isSourceZCL`, whose blueprint entry now carries a `\leanok` marker.

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.