The Perron–Frobenius Rank-One Step in Cirac–Perez-Garcia–Schuch–Verstraete 2017 Appendix C.2

2026-04-30

Verdict: false-source (open).

1 The assertion

Appendix C.2 of Cirac–Pérez-García–Schuch–Verstraete relates the strong area law and zero correlation length for a simple matrix product density operator to an outer-product factorization of a finite trace matrix [CPGSV16, Appendix C.2, Lemmas C.3–C.4].[1]

The paper constructs positive semidefinite neighboring operators $\eta_{k,h}$ and defines

$$T_{k,h} = \mathrm{tr}(\eta_{k,h}).$$

Thus T is entrywise non-negative. Under the strong area law the source argues that T is primitive, in the sense that some power of T has all entries strictly positive. Under zero correlation length it obtains

$$\mathrm{tr}(T^N) = \mathrm{tr}(T)$$

for every positive integer N, with $\mathrm{tr}(T) = 1$ after normalization. The proof then invokes fixed-point theory for primitive non-negative matrices and concludes that T has one eigenvalue equal to 1 and all other eigenvalues equal to 0. From this it concludes the desired factorization

$$T_{k,h} = a_k b_h.$$

The last implication is the point at issue. Primitivity and the equalities of positive-power traces control the nonzero eigenvalues of T, but they do not

[1] For traceability, this note corresponds to issue #1041. The underlying question was isolated in issue #832; the finite counterexample was added in pull request #849; and the positive-semidefinite replacement theorem was added in pull request #917. The inverse-map provenance question is recorded in issue #4270; the preceding audit and obstruction verdict are recorded in issue #3593. The source passage is `Papers/1606.00608/MPDO-22-12-17-2.tex:1394--1499`.

exclude nilpotent Jordan structure on the generalized eigenspace for the eigenvalue 0. Such nilpotent structure is invisible to traces of positive powers and can prevent T from having rank one.

2 The counterexample

Consider the real 3×3 matrix

$$T = \begin{pmatrix} \frac{1}{2} & \frac{1}{2} & 0 \\ \frac{1}{6} & \frac{1}{6} & \frac{2}{3} \\ \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \end{pmatrix}.$$

All entries are non-negative, and

$$T^2 = P := \frac{1}{3}\begin{pmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{pmatrix}.$$

Hence T is primitive, since T^2 has all entries strictly positive. Also $P^2 = P$ and $PT = TP = P$, so $T^N = P$ for every $N \geq 2$. Since $\operatorname{tr}(T) = \operatorname{tr}(P) = 1$, it follows that

$$\operatorname{tr}(T^N) = \operatorname{tr}(T) = 1$$

for every positive integer N.

Nevertheless T is not an outer product. If $T = ab^\top$, then the entry $T_{1,1} = 1/2$ implies $a_1 \neq 0$. The entry $T_{1,3} = 0$ then forces $b_3 = 0$, but this contradicts $T_{2,3} = 2/3$. Thus the universal matrix implication

$$\text{primitive} + \text{constant positive-power traces} \Longrightarrow \text{rank one}$$

is false for finite entrywise non-negative real matrices. [2]

This example has the form $T = P + \frac{1}{6}xy^\top$, where $x = (1, -1, 0)^\top$ and $y = (1, 1, -2)^\top$. The perturbation satisfies $(xy^\top)^2 = 0$ and is annihilated by P on both sides. It therefore leaves all traces of positive powers unchanged from the Perron projector while raising the rank of T.

The same matrix has a genuinely rectangular realization of the pairing in lines 1490–1497 of Appendix C.2. Set

$$L = \begin{pmatrix} \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \\ \frac{1}{6} & \frac{1}{6} & -\frac{1}{3} \end{pmatrix}, \qquad R = \begin{pmatrix} 1 & 1 \\ 1 & -1 \\ 1 & 0 \end{pmatrix}.$$

Direct multiplication gives

$$RL = T, \qquad LR = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}, \qquad (LR)^2 = LR.$$

[2] The formal counterexample is `TNLean/Archive/PerronFrobeniusRankOneCounterexample.lean`, especially `counterexample` and `counterexample_with_rectangular_pairing`. It is also summarized in `docs/counterexamples.md:19--27`.

Thus the counterexample satisfies not only the scalar trace identities but also the full rectangular idempotence underlying the displayed zero-correlation-length equation at source lines 1490–1493. Nevertheless T is not an outer product, so this equation does not justify the rank-one assertion at lines 1498–1499.

3 What the operator-valued ZCL identity implies

The source uses more than the trace equalities before passing to Perron–Frobenius theory. With $Le_k = l_k$ and $(Rv)_k = (r_k \mid v)$, one has $T = RL$, while the displayed ZCL identity in lines 1489–1493 is

$$LR = L(RL)R = (LR)^2.$$

This identity makes the operator LR idempotent. Multiplication on the left by R and on the right by L gives only

$$T^2 = T^3,$$

not $T = T^2$. The rectangular matrices displayed above satisfy the full hypothesis $T = RL$ and $(LR)^2 = LR$, while $T^2 = T^3 = P$. Thus retaining the displayed operator equation, without a further property of the inverse-map families l_k and r_k, does not remove the nilpotent part of T.

There is nevertheless a direct rectangular proof of the trace-power display at source lines 1494–1497. Cyclicity of trace gives

$$\operatorname{tr}(T) = \operatorname{tr}(RL) = \operatorname{tr}(LR) = \operatorname{tr}((LR)^2) = \operatorname{tr}(T^2).$$

Together with $T^2 = T^3$, this implies $T^N = T^2$ for every $N \geq 2$, and hence

$$\operatorname{tr}(T^N) = \operatorname{tr}(T)$$

for every positive integer N. [3]

The exact sufficient conclusion is transparent. If the inverse-map construction were shown to give $T^2 = T$, then T itself would be idempotent. An idempotent real matrix has trace equal to the dimension of its range; hence trace one makes that range one-dimensional and gives $T = ab^\top$. This argument is formalized as `Matrix.hasRankOneFactorization_of_mul_self_eq_self`. The remaining source-faithful task is therefore to derive $T^2 = T$, or another condition excluding the generalized zero-eigenspace, from a property of the concrete inverse-map tensors not currently recorded in `MPOTensor.ExplicitEtaOperators`.

[3] Formal declaration: `Matrix.trace_pow_eq_trace_of_rectangular_idempotent` in `TNLean/Algebra/PerronFrobenius/RankOne`.

3.1 The pairing-idempotence route

The formal development derives $T^2 = T$ from the displayed operator identity under one additional hypothesis on the inverse-map tensors. Suppose the family $(l_k)_k$ is linearly independent, so that L is injective. The operator identity $(LR)^2 = LR$ gives

$$L\,T^2 = (LR)^2 L = (LR)L = L\,T,$$

and cancelling L on the left yields $T^2 = T$. If instead the functionals $(r_k)_k$ are linearly independent, the same argument in the dual space gives the idempotence of $T^\top$, hence of T. Together with the trace normalization $\mathrm{tr}(T) = 1$ and the idempotent trace-one criterion above, this recovers the outer-product factorization asserted by Lemma C.5.[4]

The linear-independence hypothesis is a scope restriction relative to the source. Lemma C.5 neither assumes nor derives it: the construction preceding the lemma uses the injectivity of $\mathcal{K}$ to prove that the sector decomposition of the trace matrix is trivial, that is, that T is primitive [CPGSV16, Appendix C.2, lines 1457–1470]. Primitivity constrains only the support pattern of the nonnegative matrix T; a primitive T can arise from linearly dependent l_k, so the cited passage does not make the family $(l_k)_k$ or $(r_k)_k$ linearly independent. The proof of Lemma C.5 then uses primitivity in the Perron–Frobenius trace-power step examined above, not any independence of the families. The formal lemmas therefore state the independence as an explicit hypothesis; whether the inverse-map construction supplies it — for instance through the injectivity of $\mathcal{K}$ — remains to be established at the matrix product density operator call site.

3.2 The sector pairing at the matrix product density operator side

The formal development now records the paper's pairing on the matrix product density operator side. In an abstract real vector space, let $|l_k)$ be the closed sector tensors and $(r_k|$ the corresponding functionals. It assumes the two displayed identities of the source,

$$T_{k,h} = (r_k|l_h), \qquad \sum_k |l_k)(r_k| = \sum_{k,h} T_{k,h}|l_k)(r_h|,$$

from lines 1478–1493 of Appendix C.2. These equations make the pairing operator $M = \sum_k |l_k)(r_k|$ idempotent, without any independence hypothesis on $(l_k)_k$ or $(r_k)_k$. Pairing the idempotence with the functional map R on the left and the coefficient map L on the right, as in Section 3, gives $T^2 = T^3$ unconditionally; restricting the pairing operator to the finite-dimensional span of $(l_k)_k$ then identifies $\mathrm{tr}(T)$ with $\mathrm{tr}(T^2)$, so every positive power of T has the same trace

[4] The formal statements are `Matrix.mul_self_eq_self_of_pairing_idempotent`, `Matrix.mul_self_eq_self_of_pairing_idempotent_dual`, and `Matrix.hasRankOneFactorization_of_pairing_idempotent` in `TNLean/Algebra/PerronFrobenius/RankOne`, added in PR #3685.

as T. This recovers the trace identity displayed at lines 1494–1497 of Appendix C.2 unconditionally, at the matrix product density operator side. [5]

If in addition the vectors $(l_k)_k$ are linearly independent, then the coefficient map L is injective and

$$LT^2 = M^2L = ML = LT,$$

so $T^2 = T$, the strictly stronger idempotence used for the rank-one factorization below. Trace one and idempotence together then give the factorization directly. Primitivity separately makes each $|l_k)$ nonzero, so an independent family of subspaces carrying the vectors is sufficient to derive their linear independence. [6]

The inverse-map construction has now been connected to this argument without introducing additional hypotheses. The closed tensors $(l_k)_k$ and $(r_k)_k$ are constructed from an injective simple tensor, the complex matrix is defined by

$$T_{k,h} = (r_k|l_h) = \mathrm{tr}(\eta_{k,h}),$$

and its pairing operator is identified with the physical-trace transfer. Source zero correlation length therefore supplies a number $\lambda > 0$ for which the normalized pairing operator is idempotent. The rectangular factorization $S = LR$ and $T = RL$ then proves

$$(\lambda^{-1}T)^2 = (\lambda^{-1}T)^3, \qquad \mathrm{tr}((\lambda^{-1}T)^N) = \mathrm{tr}(\lambda^{-1}T)$$

for every positive integer N. [7]

The normalized identities can also be placed on the same real trace matrix for which the strong-area-law argument proves primitivity. After discarding only the zero-weight indices and choosing the coherent sector phases, set

$$T^*_{k,h} = \mathrm{Re}\,\mathrm{tr}(\eta_{k,h}) \qquad (p_kp_h \neq 0).$$

The left tensor of every zero-weight sector vanishes. Consequently the physical-trace transfer still has the active rectangular factorization LQ, while positivity of the neighboring operators identifies QL with the complexification of T^*. The same scalar $\lambda > 0$ therefore gives

$$T^* \text{ primitive}, \qquad (\lambda^{-1}T^*)^2 = (\lambda^{-1}T^*)^3, \qquad \mathrm{tr}((\lambda^{-1}T^*)^N) = \mathrm{tr}(\lambda^{-1}T^*)$$

[5] The formal declarations are `MPOTensor.SectorPairingData.pairing_sq_eq_pairing_cube` and `MPOTensor.SectorPairingData.tracePowersConstant` in `TNLean/MPS/MPDO/SimpleLocalStructure`, resting on the general matrix statements `Matrix.pow_two_eq_pow_three_of_pairing_idempotent`, `Matrix.trace_eq_trace_sq_of_pairing_idempotent`, `Matrix.pow_eq_pow_two_of_pow_two_eq_pow_three`, and `Matrix.tracePowersConstant_of_pairing_idempotent` in `TNLean/Algebra/PerronFrobenius/RankOne`, added for GitHub issue #3714.

[6] The formal declarations are `MPOTensor.SectorPairingData.pairing_operator_idempotent`, `MPOTensor.SectorPairingData.mul_self_eq_self`, `MPOTensor.SectorPairingData.l_ne_zero`, `MPOTensor.SectorPairingData.linearIndependent_l`, `MPOTensor.sal_zcl_implies_rank_one_T_of_pairing_idempotent`, and `MPOTensor.sal_zcl_implies_rank_one_T_of_sector_supports` in `TNLean/MPS/MPDO/SimpleLocalStructure`.

[7] Formal declarations: `MPOTensor.closedSectorTraceMatrix_apply`, and `MPOTensor.closedSectorTraceMatrix_normalized_relations` in `TNLean/MPS/MPDO/SectorPairingTransfer`.

for every positive integer N. [8]

Write $\mathcal{B}_{\mathcal{K}}$ for the physical-trace transfer of $\mathcal{K}$. If SAL is accompanied by the paper's literal identity $\mathcal{B}_{\mathcal{K}}^2 = \mathcal{B}_{\mathcal{K}}$, the one-site normalization clause in SAL first gives $\mathcal{B}_{\mathcal{K}} \neq 0$. Thus literal ZCL implies source ZCL, and the same primitive active trace matrix and normalized relations follow without any additional hypothesis. This specialization is formalized by `MPOTensor.exists_activeTraceMatrix_relations_of_isSAL_of_literal_ZCL` in `TNLean/MPS/MPDO/InverseMapActiveSectorZCL`.

This conclusion is the scale-invariant consequence of the displayed operator equation; it does not imply that the normalized trace matrix is idempotent or semisimple, nor that it has rank one. The remaining pairing route would have to prove linear independence of one of the two closed tensor families, or supply a different argument excluding the generalized zero-eigenspace. The direct sum in the source's sector splitting at lines 1436–1448 concerns the physical indices; after those indices are contracted, the closed tensors live in a common bond space. Their membership in corresponding members of an independent family of subspaces — or their independence by another argument — therefore does not follow from the displayed equations alone.

3.3 Virtual spanning does not eliminate the nilpotent part

Injectivity gives more than primitivity in the present formal development: the sector virtual matrices span the full virtual matrix algebra. This additional conclusion still does not imply the idempotence of the opposite product. Consider

$$L = \begin{pmatrix} \frac{1}{6} & \frac{1}{6} & \frac{1}{3} & \frac{1}{3} \\ 0 & \frac{1}{6} & \frac{1}{6} & -\frac{1}{3} \end{pmatrix}, \qquad Q = \begin{pmatrix} 1 & 1 \\ 1 & 1 \\ 1 & -1 \\ 1 & 0 \end{pmatrix}.$$

Writing l_k for the kth column of L and r_k for the kth row of Q, direct multiplication gives

$$LQ = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}, \qquad T = QL = \begin{pmatrix} \frac{1}{6} & \frac{1}{3} & \frac{1}{2} & 0 \\ \frac{1}{6} & \frac{1}{3} & \frac{1}{2} & 0 \\ \frac{1}{6} & 0 & \frac{1}{6} & \frac{2}{3} \\ \frac{1}{6} & \frac{1}{6} & \frac{1}{3} & \frac{1}{3} \end{pmatrix}.$$

The square of T is

$$T^2 = \begin{pmatrix} \frac{1}{6} & \frac{1}{6} & \frac{1}{3} & \frac{1}{3} \\ \frac{1}{6} & \frac{1}{6} & \frac{1}{3} & \frac{1}{3} \\ \frac{1}{6} & \frac{1}{6} & \frac{1}{3} & \frac{1}{3} \\ \frac{1}{6} & \frac{1}{6} & \frac{1}{3} & \frac{1}{3} \end{pmatrix}.$$

[8] Formal declarations: `MPOTensor.PhysicalSectorFactorization.activeSectorTraceMatrix_normalized_relations_of_isSourceZCL`, `MPOTensor.exists_rephased_inverseMap_activeSectorTraceMatrix_normalized_relations`, and `MPOTensor.exists_activeTraceMatrix_relations_of_isSAL_of_isSourceZCL` in `TNLean/MPS/MPDO/InverseMapActiveSectorZCL`.

Thus T is primitive, $\mathrm{tr}(T) = 1$, $T^2 = T^3$, and $T \neq T^2$. Equivalently, the exact missing expression is nonzero:

$$Q(1 - LQ)L = T - T^2 \neq 0.$$

Moreover, the four matrices $l_k r_k$ span $M_2(\mathbb{R})$. For example, after vectorizing them in the order $(00, 01, 10, 11)$, the determinant of the resulting four-by-four matrix is $1/108$. Consequently full virtual spanning does not exclude the generalized zero-eigenspace.

This example is realized at the level of physical slices by the diagonal matrix product operator tensor

$$K^{ij} = \delta_{ij}\, l_i r_i.$$

Its doubled-index tensor is injective because the four nonzero slices span $M_2(\mathbb{C})$, and its physical-trace transfer is $\sum_i K^{ii} = LQ$, so it satisfies source zero correlation length. These assertions, together with primitivity, trace normalization, the spanning identity, and the nonzero nilpotent remainder, are proved in `TNLean/Archive/PerronFrobeniusVirtualSpanningCounterexample.lean`. [9]

The construction does not yet identify L and Q with the particular closed tensors selected from a Hayashi decomposition by `MPOTensor.zeroWeightReparameterizedInverseMapPhysicalSectorFactorization`. This distinction is essential. The theorem above supplies a directly chosen rectangular realization; it does not supply the Hayashi structure, normalized tail entry, inverse-map indices, and rephasing data from which the source factorization is constructed. Thus it is not a counterexample to the full statement of Lemma C.5.

The obstruction instead proves that every presently exported algebraic consequence—source ZCL, injectivity, primitivity, trace normalization, rectangular pairing idempotence, constant positive-power traces, and full virtual spanning—is insufficient. A source-faithful completion must derive an additional identity from SAL and from the provenance of the inverse-map tensors. In the notation above, the needed assertion is precisely

$$Q(1 - LQ)L = 0,$$

or any equivalent statement excluding the nilpotent generalized zero-eigenspace of QL.

3.4 Positive semidefinite neighboring operators do not force the rectangular vanishing

The preceding obstruction can also be realized inside the physical-sector factorization with every neighboring operator positive semidefinite. This can be

[9] Combined formal declaration: `TNLean.Archive.PerronFrobeniusVirtualSpanningCounterexample.injective_sourceZCL_tensor_with_nilpotent_sector_pairing`.

seen from a second, independent witness with bond dimension two and four one-dimensional physical sectors. Let

$$L = \begin{pmatrix} \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} \\ 1 & -1 & 1 & -1 \end{pmatrix}, \qquad Q = \begin{pmatrix} 1 & \frac{1}{8} \\ 1 & \frac{1}{8} \\ 1 & -\frac{1}{8} \\ 1 & -\frac{1}{8} \end{pmatrix}.$$

For each sector k, take the virtual matrix to be the outer product of the kth column of L with the kth row of Q. These four matrices span $M_2(\mathbb{C})$: in the coordinate order $(00, 01, 10, 11)$, the matrix having these four outer products as columns has determinant $-1/64$. Placing them on the four diagonal physical entries therefore gives an injective tensor with a scalar physical-sector factorization.

The two rectangular products are

$$LQ = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}, \qquad QL = \begin{pmatrix} \frac{3}{8} & \frac{1}{8} & \frac{3}{8} & \frac{1}{8} \\ \frac{3}{8} & \frac{1}{8} & \frac{3}{8} & \frac{1}{8} \\ \frac{1}{8} & \frac{3}{8} & \frac{1}{8} & \frac{3}{8} \\ \frac{1}{8} & \frac{3}{8} & \frac{1}{8} & \frac{3}{8} \end{pmatrix}.$$

Hence LQ is a nonzero idempotent physical-trace transfer, while QL is strictly positive, primitive, and trace one. Its entries are the scalar neighboring operators and are therefore positive semidefinite. Nevertheless $(QL)^2 \neq QL$; for example, the $(0,1)$ entry changes from $1/8$ to $1/4$. Equivalently,

$$Q(1 - LQ)L = QL - (QL)^2 \neq 0.$$

Thus injectivity, full virtual-matrix spanning, positive semidefinite neighboring operators, primitivity, trace normalization, and the source zero-correlation-length rectangular identity can all hold without the missing vanishing. [10]

The induced finite-chain state is the classical cyclic law

$$p_N(i_1, \ldots, i_N) = \prod_{n=1}^{N} T_{i_n, i_{n+1}}, \qquad i_{N+1} = i_1,$$

where $T = QL$ is doubly stochastic and T^2 is the matrix all of whose entries are $1/4$. Let h be the common Shannon entropy, in bits, of a row of T; the four rows are permutations of one another. For every proper nonempty block of length L in a chain relevant to SAL,

$$H_L = 2 + (L-1)h, \qquad H_N = Nh.$$

Consequently $I_L = H_L + H_{N-L} - H_N = 4 - 2h$, independently of L. Formally, the same conclusion is obtained without logarithmic calculations: after one site

[10] The formal tensor and the combined counterexample theorem are in `TNLean/MPS/MPDO/ActiveSectorSpanningCounterexample`; see `MPOTensor.ActiveSectorSpanningCounterexample.virtual_spanning_does_not_force_rectangular_remainder_zero`.

is traced out, the diagonal state is an ordinary Markov path, and conditioning on its middle label gives a four-sector Hayashi decomposition at every admissible cut. The theorem `MPOTensor.ActiveSectorSpanningCounterexample.tensor_isSAL` therefore proves the strong area law for this tensor.

The inverse-map provenance is now explicit. The refined Hayashi decomposition has four one-dimensional sectors of weight $1/4$, with diagonal left and right density matrices determined by QL. Its normalized four-site tail is LQ. The four nonzero physical slices form a basis of the virtual matrix algebra; computing the dual basis shows that the source inverse map recovers the displayed closed tensors L and Q exactly. Consequently its neighboring operators agree with those of `MPOTensor.ActiveSectorSpanningCounterexample.factorization`, and the same source-selected construction satisfies

$$Q(1 - LQ)L \neq 0.$$

These statements are proved in `MPOTensor.ActiveSectorSpanningCounterexample.inverseMap_provenance_preserves_rectangular_remainder`. Together with SAL, injectivity, and source ZCL, these facts are assembled in `MPOTensor.ActiveSectorSpanningCounterexample.tensor_has_sal_sourceZCL_and_non_rankOne_activeTraceMatrix`. The active trace matrix has no factorization $T_{k,h} = a_k b_h$, while $Q(1 - LQ)L \neq 0$. Hence the example proves the full raw-tensor conjunction, including inverse-map provenance. It does not refute Lemma C.5 under its complete standing hypotheses: the lemma is inside Case I at source lines 1374–1381, where $\mathcal{K}$ is assumed to be an injective normal tensor. The raw witness is not transfer-normalized, while its normal representative loses the literal ZCL identity. This normalization qualification corrects the earlier verdict tracked in GitHub issue #4270.

4 Renormalization maps and the global normalization boundary

The standing hypotheses of Theorem 4.9 require a simple tensor in biCF with a basis of normal tensors; they are imposed at source lines 849–893 and restated at line 1628. They also inherit the global canonical-form convention at source line 246: every BNT coefficient has modulus at most one, and at least one has modulus one. The distinction between this single global unit witness and a stronger sectorwise convention is recorded in `docs/paper-gaps/cpsv16_global_vs_persector_unit_witness.tex`.

Write $\mathcal{K}$ for the four-sector witness and regard its doubled physical index as an MPS index. Its four nonzero physical matrices span $M_2(\mathbb{C})$, so the doubled tensor is injective at one site. Its raw transfer map is

$$\mathcal{E}_{\mathcal{K}}(X) = \left(2X_{00} + \frac{1}{32}X_{11}\right)\begin{pmatrix} \frac{1}{8} & 0 \\ 0 & 2 \end{pmatrix}, \qquad \mathcal{E}_{\mathcal{K}}^2 = \frac{5}{16}\mathcal{E}_{\mathcal{K}}.$$

Thus the raw doubled tensor has spectral radius 5/16. Set

$$A = \frac{4}{\sqrt{5}}\mathcal{K}, \qquad \mu = \frac{\sqrt{5}}{4}.$$

Then $\mathcal{K} = \mu A$, and $\mathcal{E}_A$ has spectrum $\{1,0,0,0\}$ and is idempotent. The singleton family $\{A\}$ is the normal basis for $\mathcal{K}$, but its only coefficient is $\mu < 1$. Consequently the raw tensor $\mathcal{K}$ does not meet the source's global unit-weight convention. The normalized representative A does: it is injective, hence normal; its singleton BNT has coefficient one; and the same one-site spanning identity gives biCF with blocking length one. Moreover,

$$\mathcal{B}_A = \frac{4}{\sqrt{5}}LQ = \frac{4}{\sqrt{5}}\begin{pmatrix}1 & 0\\ 0 & 0\end{pmatrix}$$

is not nilpotent, so the sole BNT element is simple in the sense used in Theorem 4.9.

These two representatives lie on opposite sides of the source's two literal normalizations. The raw tensor has $\mathcal{B}_{\mathcal{K}} = LQ = P$, so it satisfies the literal diagram $\mathcal{B}_{\mathcal{K}}^2 = \mathcal{B}_{\mathcal{K}}$, but it lacks a global unit-modulus BNT coefficient. The representative A has the required unit BNT coefficient, but

$$\mathcal{B}_A^2 = \frac{4}{\sqrt{5}}\mathcal{B}_A \neq \mathcal{B}_A,$$

so it fails the paper's literal zero-correlation-length diagram. The scalar absorption at source line 1660 occurs only after that literal ZCL identity has already been assumed; it does not transport either literal ZCL or the renormalization maps across a scalar rescaling.

The raw tensor nevertheless gives a useful positive calculation for the trace-preserving completely positive map condition. Let $\tau = QL$ and define

$$f(i,j) = (i \bmod 2) + 2\left\lfloor \frac{j}{2} \right\rfloor.$$

The sector matrices obey

$$S_iS_j = \tau_{ij}S_{f(i,j)}, \qquad \sum_{(i,j):f(i,j)=k} \tau_{ij} = 1$$

for each $k \in \{0,1,2,3\}$. Consequently the refinement and coarse-graining Kraus operators

$$R_{ij} = \sqrt{\tau_{ij}}|i,j\rangle\langle f(i,j)|, \qquad C_{ij} = |f(i,j)\rangle\langle i,j|$$

resolve the respective input identities. Their maps are trace-preserving and completely positive, and direct substitution gives

$$\mathsf{C}[\mathcal{K}_2(X)] = \mathcal{K}_1(X), \qquad \mathsf{R}[\mathcal{K}_1(X)] = \mathcal{K}_2(X)$$

for every virtual matrix X. Applying each map twice and passing to the canonical blocked coordinates gives the same two identities for $\mathcal{K}^{[2]}$. These statements are proved by `MPOTensor.ActiveSectorSpanningCounterexample.tensor_isRFPViaTS`, `MPOTensor.ActiveSectorSpanningCounterexample.blockTwo_tensor_isRFPViaTS`, and the general blocking theorem `MPOTensor.IsRFPViaTS.blockTwo`.

4.1 Pair-preserving two-to-four-site maps

There is a sharp trace obstruction to applying the twice-iterated maps while retaining every fixed pair of outer sector labels. For an arbitrary physical-sector factorization, put

$$C_{k,h} = \operatorname{tr}(\eta_{k,h}), \qquad \Theta_{k,h} = \bigoplus_{l,m} \eta_{k,l} \otimes \eta_{l,m} \otimes \eta_{m,h}.$$

Additivity of trace on direct sums and multiplicativity on tensor products give

$$\operatorname{tr}(\Theta_{k,h}) = \sum_{l,m} C_{k,l} C_{l,m} C_{m,h} = (C^3)_{k,h}.$$

It follows that a trace-preserving family which sends every $\eta_{k,h}$ to the corresponding $\Theta_{k,h}$, or every $\Theta_{k,h}$ to the corresponding $\eta_{k,h}$, forces

$$C = C^3.$$

This project-derived result is a *conditional helper*: it identifies the necessary trace equality for a pair-preserving construction, but it neither asserts that the source maps preserve such pairs nor constructs a label-mixing alternative. The trace identity and its two directions are proved in `MPOTensor.PhysicalSectorFactorization.trace_fourSiteNeighboringOperator`, `MPOTensor.PhysicalSectorFactorization.neighboringTraceMatrix_eq_pow_three_of_pairwise_coarsening`, and `MPOTensor.PhysicalSectorFactorization.neighboringTraceMatrix_eq_pow_three_of_pairwise_refinement`.

For the selected four-sector factorization of the active-spanning witness, $C = QL$. Rectangular idempotence of LQ gives $C^2 = C^3$, but the explicit entries satisfy $C_{0,1} = 1/8$ and $(C^2)_{0,1} = 1/4$. Hence

$$C^2 = C^3, \qquad C \neq C^3,$$

and neither direction of a pair-preserving trace-preserving map can exist for these four selected sectors. This is proved in `MPOTensor.ActiveSectorSpanningCounterexample.not_exists_pairwise_fourSite_coarsening` and `MPOTensor.ActiveSectorSpanningCounterexample.not_exists_pairwise_fourSite_refinement`.

This exclusion is *proof-path drift*, not an obstruction to renormalization. The trace-preserving completely positive maps above use the many-to-one rule $f(i, j)$ and therefore coarsen the selected labels; after iteration they still prove the renormalization equations for $\mathcal{K}^{[2]}$. They are not among the fixed-(k,h) families excluded by the trace calculation. Likewise, the one-sector factorization has no such four-sector obstruction. Thus this calculation rules out the selected pair-preserving route but does not construct the general coarsened or label-mixing maps sought in GitHub issue #6793; that issue remains open.

This positive result is not invariant under the normalization above. Put $P = \mathrm{diag}(1, 0)$. The physical-trace transfer of the normalized representative and that of its two-site blocking are

$$\mathcal{B}_A = \frac{4}{\sqrt{5}} P, \qquad \mathcal{B}_{A^{[2]}} = \mathcal{B}_A^2 = \frac{16}{5} P.$$

Definition 4.1 would force $\mathcal{B}_{A^{[2]}}$ to be literally idempotent, because trace preservation identifies the traces of its one-site and two-site closures for every virtual boundary matrix. But $(16P/5)^2 \neq 16P/5$. Hence $A^{[2]}$ is not a renormalization fixed point. This is proved by `MPOTensor.ActiveSectorSpanningCounterexample.blockTwo_normalizedTensor_not_isRFPViaTS`. [11]

There are two different zero-correlation-length statements here. TNLean's `MPOTensor.IsSourceZCL` interpretation is scale-invariant: it asks for $\mathcal{B}_A^2 = \lambda \mathcal{B}_A$ with some $\lambda > 0$. The paper's diagrammatic identity is literal and selects a fixed scalar representative. Both the strong area law and TNLean's up-to-scalar predicate are unchanged by multiplying the tensor by $4/\sqrt{5}$, but the paper's diagram is not. Hence A is a counterexample only to the broadened implication obtained by replacing the literal diagram with `MPOTensor.IsSourceZCL`; it is not an unqualified counterexample to Theorem 4.9. This statement concerns the four-sector witness A only; the scalar repeated-copy counterexample below is an exact counterexample to the printed (iv)$\Rightarrow$(v) implication.

The scalar identities, the raw renormalization maps, and the failure of Definition 4.1 for $A^{[2]}$ are kernel-checked. This example leaves the source (ii)$\Rightarrow$(v) boundary unsettled: $\mathcal{K}$ satisfies literal ZCL but not the global unit-weight convention, while A satisfies the global convention but not literal ZCL. Packaging the normality, singleton BNT, biCF, simplicity, and SAL facts for A remains useful, but cannot turn its up-to-scalar ZCL relation into the paper's literal diagram. The distinct raw copy example below settles (iv)$\Rightarrow$(v) without asserting condition (ii).

5 A positive-semidefinite route

The formal development proves a corrected sufficient matrix theorem. Let T be a finite real square matrix. If T is positive semidefinite, $\mathrm{tr}(T) = 1$, and $\mathrm{tr}(T^N) = \mathrm{tr}(T)$ for every positive integer N, then T has an outer-product factorization $T = ab^\top$. [12] Primitivity is not needed for this corrected matrix

[11] The source-hypothesis and normalization-sensitive obstruction are tracked in GitHub issue #5388.

[12] The formal statements are `Matrix.PosSemidef.trace_powers_constant_implies_rank_one` and `Matrix.primitive_trace_powers_constant_implies_rank_one_of_posSemidef` in `TNLean/Algebra/PerronFrobenius/RankOne.lean`.

theorem once positive semidefiniteness and trace normalization are assumed.

The proof is spectral. Positive semidefiniteness makes T Hermitian and diagonalizable with non-negative real eigenvalues λ_i. The trace normalization gives $\sum_i \lambda_i = 1$. The trace-power hypothesis at $N = 2$ gives $\sum_i \lambda_i^2 = 1$. For non-negative numbers with sum one, this forces exactly one λ_i to be 1 and all the others to be 0. Therefore the diagonal form has rank one, and unitary conjugation gives an outer-product factorization of T.

6 Audit of the positive-semidefinite route

The neighboring-operator construction displayed in the source does not by itself exhibit T as a Gram matrix. Immediately before Lemma C.5, the paper writes

$$T_{k,h} = (r_k \mid l_h),$$

with two different families $(r_k)_k$ and $(l_h)_h$ [CPGSV16, Appendix C.2, lines 1473–1482]. No equality or adjoint relation between these two families is stated. Consequently this is a cross-pairing, rather than a Gram representation, and it gives neither symmetry nor positive semidefiniteness of T.

The present type of explicit neighboring operators records only $\eta_{k,h} \geq 0$ for each ordered pair (k, h). This hypothesis implies $T_{k,h} \geq 0$ entrywise, but contains no relation between $\eta_{k,h}$ and $\eta_{h,k}$. Thus matrix-level positive semidefiniteness cannot be deduced from that type alone.

There is one special positive-semidefinite case in the formal development. The sector-reduced Hayashi–Markov family is defined by

$$\eta_{k,h} = \operatorname{tr}_C(\rho^{(k)}_{b_2 C}) \otimes \operatorname{tr}_A(\rho^{(h)}_{A b_1}).$$

Since each reduced density matrix has trace one, its trace matrix is the all-ones matrix. The theorem `MPOTensor.ExplicitEtaOperators.traceMatrixRe_ofHayashiMarkov_posSemidef` proves its positive semidefiniteness. This does not discharge Lemma C.5: the sector-reduced family has not been identified with the inverse-map family constructed in lines 1413–1455, and the trace of its all-ones matrix is the number of sectors, not one unless there is a single sector.

The positive-semidefinite route must therefore use more of the matrix product density operator construction than the separate inequalities $\eta_{k,h} \geq 0$. One possible completion is to derive an adjoint or Gram relation from the inverse-map tensors. Positive semidefiniteness of the source trace matrix is not presently established by the cited passage.

This corrected theorem does not refute the intended matrix product density operator result; it supplies one sufficient matrix-level hypothesis. For this route, the open question is whether the trace matrix $T_{k,h} = \operatorname{tr}(\eta_{k,h})$ is Hermitian or positive semidefinite. The pairing-idempotence route of Section 3.1 avoids that hypothesis, but instead requires the concrete sector pairing, the zero-correlation-length operator identity, and linear independence. Positivity of each $\eta_{k,h}$ gives only $T_{k,h} \geq 0$ entry by entry, which implies neither matrix-level positive semidefiniteness nor Hermitian symmetry. [13]

7 Audit of the inverse-map provenance

The remaining possibility was that the particular inverse-map tensors used in Appendix C.2 might satisfy an additional identity which is not visible in the rectangular pairing equation. An audit of the complete source passage does not reveal such an identity.

In the proof of Lemma C.4, at lines 1406–1471, the strong area law is used to obtain the physical-sector factorization, positive neighboring operators, and primitivity of the trace matrix. The last point uses injectivity to rule out more than one primitive block. This argument does not prove an adjoint relation between the closed left and right tensors, their linear independence, or the vanishing

$$Q(1 - LQ)L = 0.$$

The physical direct sum separates the uncontracted physical sector spaces. After the physical legs are closed at lines 1473–1482, the vectors l_h and functionals r_k lie in common bond spaces, and the source gives only $T_{k,h} = (r_k \mid l_h)$.

The proof of Lemma C.5, at lines 1489–1499, then uses zero correlation length to obtain the pairing-operator identity and its trace-power consequence. No further property of the inverse map occurs between this identity and the rank-one conclusion. At the common-matrix level this leaves the gap described above. The source's separate Case I assumption that $\mathcal{K}$ is an injective normal tensor must nevertheless be retained when judging the full lemma.

The cyclic finite-chain form at lines 1581–1593 does not repair this step. It is obtained later from the same positive neighboring family and states a direct sum of cyclic products. Its formal, scalar-preserving version is `MPOTensor.EtaLocalStructureData.exists_positive_cyclic_eta_block_decomposition`. Neither the source formula nor this theorem asserts independence, a Gram relation, kernel stabilization, or rectangular vanishing for the closed tensors. Taking traces of cyclic products returns the already established positive-power trace identities.

Consequently the strongest explicit common-matrix conclusion remains

$$T^2 = T^3, \qquad \mathrm{tr}(T^N) = \mathrm{tr}(T) \quad (N \geq 1),$$

together with entrywise nonnegativity and primitivity. The formal counterexample shows that these conclusions, even supplemented by injectivity, full virtual spanning, and positivity of every neighboring operator, do not imply the

[13] The entrywise nonnegativity result is `MPOTensor.ExplicitEtaOperators.traceMatrixRe_nonneg`. The positive-semidefinite reformulation for the rank-one conclusion is `MPOTensor.sal_zcl_implies_rank_one_T_of_posSemidef`. The corresponding blueprint entries have labels `thm:mpdo_explicit_eta_trace_matrix_nonneg` and `thm:psd_trace_powers_rank_one_t` in `blueprint/src/chapter/ch21_mpdo_rfp_simple_local_primitivity.tex`.

missing vanishing. The explicit Hayashi decomposition and inverse-map calculation now show that the same nonzero nilpotent remainder occurs for the source-selected factorization. The all-cut Markov decomposition proves the strong area law for the same raw tensor. This tensor satisfies the displayed SAL and literal ZCL relations, but it is not the injective normal tensor assumed in Case I. Its normal representative loses literal ZCL. Thus the example isolates the normalization-sensitive step needed for $T_{k,h} = a_k b_h$; it does not refute Lemma C.5 under all standing hypotheses.

This obstruction does not survive the full normal Case-I hypotheses: the completed route below excludes it and formalizes the structural corollary at lines 1503–1506. It remains relevant to the later Case-II passage. The sectorwise analytic part of that passage is now complete. For every absorbed BNT representative $\mathcal{K}_s = \mu_s A_s$, the formalization proves one-site injectivity, positivity on every nonempty periodic chain, saturation of the area law, and the literal identity

$$\mathcal{T}_{\mathcal{K}_s}^2 = \mathcal{T}_{\mathcal{K}_s}.$$

The remaining obstruction is normality after coefficient absorption. If A_s is normal, then the transfer map of $\mu_s A_s$ has spectral radius $|\mu_s|^2$. The global canonical-form convention supplies at least one unit-modulus coefficient, but it does not imply $|\mu_s| = 1$ for every sector. Consequently the normal Case-I structural theorem cannot be applied to all absorbed BNT representatives. The constructions from an explicitly supplied rank-one neighboring trace factorization remain valid restricted results, but they do not prove that such a factorization exists for every absorbed representative.

The explicit example discussed in the conclusion realizes this normalization obstruction under every condition used in the simple Case-II passage, at the normalized fixed-representative boundary. It refutes the absorption inference, not the source assertion that some suitable factorization exists.

8 The unit-closure-trace necessary condition

The existential assertion of condition (iv) quantifies over all possible physical-sector factorizations of each absorbed representative, so a counterexample must exclude every factorization and a proof may choose any. The formalization now records a constraint that every qualifying factorization imposes on the tensor itself, independently of the choice.

Let $\mathcal{K}$ be any MPO tensor admitting a physical-sector factorization with positive semidefinite neighboring operators $\eta_{k,h}$ and real families satisfying

$$\operatorname{tr}(\eta_{k,h}) = a_k b_h, \qquad \sum_k a_k b_k = 1.$$

Closing the physical legs of the factorized slices and using unitary invariance of the trace writes the physical-trace transfer as a rectangular product $\mathcal{T}_{\mathcal{K}} = LQ$

with $L_{\beta,k} = \mathrm{tr}((l_k)_\beta)$ and $Q_{k,\alpha} = \mathrm{tr}((r_k)_\alpha)$, and the opposite product is the rank-one coefficient matrix

$$(QL)_{k,h} = \mathrm{tr}(\eta_{k,h}) = a_k b_h.$$

Cyclic invariance of the trace gives $\mathrm{tr}((LQ)^N) = \mathrm{tr}((QL)^N)$ for every $N \geq 1$, and all positive powers of the normalized rank-one matrix equal the matrix itself, of trace one. Write $\rho^{(N)}(\mathcal{K})$ for the periodic operator denoted $\sigma^{(N)}(\mathcal{K})$ in the source. Hence

$$\mathrm{tr}\big(\mathcal{T}_{\mathcal{K}}^N\big) = 1, \qquad \mathrm{tr}\,\rho^{(N)}(\mathcal{K}) = 1, \qquad N \geq 1.$$

The formal statements are `MPOTensor.PhysicalSectorFactorization.physTraceTransfer_eq_leftTraceMatrix_mul_rightTraceMatrix`, `MPOTensor.PhysicalSectorFactorization.NeighboringTraceFactorization.trace_physTraceTransfer_pow_eq_one`, and `MPOTensor.PhysicalSectorFactorization.NeighboringTraceFactorization.trace_mpo_eq_one` in `TNLean/MPS/MPDO/NeighboringTraceObstruction.lean`.

Three consequences follow. First, the normalization $\mathrm{tr}(T) = 1$ invoked at line 1498 is a constraint on the tensor rather than a convention: a tensor with a closed chain of non-unit trace admits no factorization of the required form. Second, the per-representative assertion of condition (iv) carries the implicit normalization $\mathrm{tr}(\mathcal{T}_{\mathcal{K}_s}) = 1$ for each absorbed representative, whereas the ambient normalization at lines 1755–1759 constrains only the weighted sum over representatives. Whether the standing hypotheses of condition (ii) imply this per-representative unit trace is part of the open content of the all-sector structure statement; if some admissible tensor violates it, the violation immediately refutes the printed per-representative assertion for that tensor. Third, under literal zero correlation length the condition says that the physical-trace transfer of each absorbed representative is an idempotent of unit trace; the four-sector active-trace witness is consistent with this, since its physical-trace transfer is a rank-one coordinate projection.

9 Conclusion

The proof in Appendix C.2 contains an unjustified matrix step as written: a primitive entrywise non-negative matrix with constant traces of positive powers need not have rank one. The obstruction is a nilpotent Jordan component at the zero eigenvalue, which is not detected by the trace identities.

The formal development proves two corrected sufficient matrix theorems. The first derives the rank-one factorization from positive semidefiniteness, trace normalization, and the trace-power identities. The second derives $T^2 = T$ from the sector-pairing identity and linear independence of the closed sector tensors; trace one then gives the same rank-one factorization directly. These criteria remain available as independent matrix results. The unrelated aggregate bundle that paired a three-site Markov witness with an independently supplied trace matrix has been retired: it imposed no relation between the state, the matrix, and an MPO tensor, and therefore added no consequence to either criterion.

The completed normal Case-I route supplies a third, tensor-attached result at the source hypothesis boundary. The active-sector conclusion is formalized in `MPOTensor.exists_activeSectorTraceMatrix_rank_one_coefficients_of_isSAL_of_literal_ZCL`, and the all-sector coefficient statement is formalized in `MPOTensor.exists_neighboringOperator_trace_rank_one_coefficients_of_isSAL_of_literal_ZCL`. Normality and literal ZCL, together with the physical-sector positivity and primitivity data, force the active sector to be unique and hence $T^2 = T$. For the project factors $T = QL$, this gives

$$Q(1 - LQ)L = QL - (QL)^2 = 0.$$

Thus the nilpotent-zero obstruction is excluded on this normal Case-I surface; idempotence of LQ alone would still yield only $T^2 = T^3$.

For the all-sector statement, retain the active coefficient functions on $I_* = \{k : p_k \neq 0\}$ and set

$$a_k = b_k = 0 \qquad (p_k = 0).$$

In the zero-weight-reparameterized and rephased factorization selected by the construction, $p_k = 0$ makes both every incoming and every outgoing neighboring operator vanish. On active pairs, the all-sector theorem first gives $\mathrm{Re}\,\mathrm{tr}(\eta_{k,h}) = a_k b_h$. Every neighboring operator is positive semidefinite, hence Hermitian with real trace. Therefore

$$\mathrm{tr}(\eta_{k,h}) = a_k b_h$$

for every pair of sector labels, and filtering the diagonal sum to I_* gives $\sum_k a_k b_k = 1$. This is a project-derived realization of the printed all-index coefficient equation in Lemma C.5: CPSV16 quantifies sector labels without defining TNLean's active–inactive split or its ambient filler sectors.

The Case-II sectorwise analytic slice is formalized separately. Literal physical-trace idempotence first descends to every weighted canonical block by `MPOTensor.weighted_basis_physTraceTransfer_sq_of_literal_ZCL`, and then to every absorbed BNT representative by `MPOTensor.commonWeightAbsorbedBasisMPOTensor_physTraceTransfer_sq_of_literal_ZCL`. Together with the projector, positivity, and entropy arguments, `MPOTensor.commonWeightAbsorbedBasisMPOTensor_caseII_properties_of_literal_ZCL` gives, for every s,

$$\mathcal{K}_s \text{ injective}, \qquad \sigma^{(N)}(\mathcal{K}_s) \geq 0 \ (N \geq 1), \qquad \mathcal{K}_s \text{ SAL}, \qquad \mathcal{T}_{\mathcal{K}_s}^2 = \mathcal{T}_{\mathcal{K}_s}.$$

This literal identity is distinct from the scale-invariant predicate `MPOTensor.IsSourceZCL`. The proof temporarily derives that predicate with scalar 1 only to invoke the existing entropy theorem; it retains the literal identity as a separate conclusion.

The absorbed tensor $\mathcal{K}_s = \mu_s A_s$ is an absorbed BNT representative, not automatically a normal representative. Normality of A_s fixes the spectral radius of its transfer map to one, whereas coefficient absorption changes it to $|\mu_s|^2$. The source's global normalization gives only one unit-modulus coefficient across the whole canonical form and does not supply $|\mu_s| = 1$ for every s.

The printed absorbed-normality step is now formally refuted under all conditions used in Appendix C.2 by the two-sector, bond-one theorem `MPOTensor.CaseIIAbsorptionCounterexample.ambient_isSAL_isSimpleCanonicalForm_and_firstAbsorbed_not_isNormalTensor`. For every $N > 0$, its diagonal periodic operator is positive semidefinite and has trace two. Explicit trace-preserving completely positive maps witness the Definition 4.1 closure equations, hence SAL, and its physical-trace transfer is the identity, giving literal ZCL. Its normal representatives have disjoint doubled-physical support, global weights $(1/\sqrt{2}, 1)$, normalized fixed-representative simple canonical form, and biCF. Nevertheless, the first absorbed representative has transfer spectral radius $1/2$ and is not normal. This repairs the missing full-condition statement and decides the printed normality inference used in Proposition `prop2to3`; it does not refute the asserted all-sector conclusion or the literal implication (iv)⇒(v).

That implication is separately refuted by the one-representative, two-copy construction in `TNLean/MPS/MPDO/Theorem49RepeatedCopyCounterexample.lean`, `MPOTensor.CaseIIAbsorptionCounterexample.printed_theorem49_iv_to_v_is_false`. Its raw copy weights are 1 and $1/2$. The theorem packages the exact BNT canonical assembly and global unit-weight condition, MPDO positivity, nonnilpotence of the sole BNT representative, simultaneous one-letter span, representative MPDO positivity, the distinct-layer equation, the physical-sector factorization, and the positive rank-one neighboring-trace data. Thus every printed standing hypothesis and all of condition (iv) have explicit component witnesses. The two-site block also generates MPDOs. Nevertheless it fails Definition 4.1: transfer idempotence would equate the traces $1 + 2^{-4}$ and $1 + 2^{-2}$. The literal implication is therefore an *unfaithful statement*; the attempt to regard outer-BNT assembly as its only unfinished step is *proof-path drift*.

Proposition `prop2to5` remains a viable missing statement because it starts from condition (ii). The earlier Case-II argument at source lines 1646–1665 uses literal ZCL to make repeated-copy weights independent of the copy label and absorb their common value before the short proof of `prop2to5` invokes Proposition C.7. For the rescaling-stable example, the separate explicit vertical identification is complete in `MPOTensor.RescalingStableLengthDependentRFP.R_oneLabelBNTAlgebraTensorClause`; it does not supply sectorwise normality and is not used in this descent argument.

The structural corollary at lines 1503–1506 is now formalized by `MPOTensor.exists_physicalSectorFactorization_rank_one_coefficients_of_isSAL_of_literal_ZCL`. It retains the selected physical-sector factorization itself: an isometry U with TNLean's orientation

$$U\kappa_{\beta,\alpha}U^\dagger = \bigoplus_k (l_k)_\beta \otimes (r_k)_\alpha,$$

its left and right factor matrices, and ambient positive semidefinite neighboring operators. Together with the exact complex trace identity above and the normalization, this completes the normal Case-I corollary on the literal-ZCL surface. It asserts no additional closure property of the sector tensors. The orientation is the convention of `MPOTensor.PhysicalSectorFactorization`, rather than an identification with a separately drawn graphical orientation in the source.

This one-factorization result is a *conditional helper*. It assumes a neighboring trace factorization of the whole tensor to which the channels are applied. Condition (iv) instead supplies a separate physical factorization for each BNT representative and leaves its raw repeated-copy weights unconstrained. The scalar two-copy example above shows that no direct-sum construction can turn these hypotheses alone into conclusion (v).

The projector-controlled construction remains relevant to Proposition `prop2to5`, hence to the viable source implication (ii)⇒(v). After ZCL has forced common copy weights and those values have been absorbed, Proposition C.7 constructs a channel pair within each outer BNT factorization. A direct concatenation of the inner sector labels is not valid: cross-BNT neighboring operators vanish by support orthogonality, while a single global rank-one equation $\mathrm{tr}(\eta_{k,h}) = a_k b_h$ would generally make cross terms nonzero. The missing construction must instead use the outer support projectors to control a direct sum of channel pairs. This all-sector missing statement is tracked in GitHub issue #6632, retargeted from the refuted raw (iv)⇒(v) reading.

The arbitrary-boundary part of this construction is now formalized by `MPSTensor.trace_evalWord_gauge_toTensorFromBlocks_mul` and `MPOTensor.physCloseN_eq_sum_commonWeightAbsorbedBasis_of_gauge_toTensor`. Let G be the invertible virtual matrix implementing the gauge equivalence between the raw tensor and its sector decomposition, and let X be an arbitrary virtual boundary. The boundary entering a fixed representative is the sum of the diagonal copy blocks of $G^{-1}XG$ over that representative. The same transformation is valid at every chain length. Copy-weight equality is used only among copies of one representative, so this calculation imposes no relation between distinct BNT coefficients. The projector-controlled combination is now formalized by `MPOTensor.blockTwo_isRFPViaTS_of_commonWeightAbsorbedBNTChannels`. For a complete outer projective resolution, the measured projection on one blocked site is the reindexed matrix $P_s \otimes I$. Compression selects the matching absorbed representative at both required closure lengths; its supplied channel is then applied.

The completeness hypothesis is removed by `MPOTensor.blockTwo_isRFPViaTS_of_commonWeightAbsorbedBNTChannels_of_orthogonalSupports`. For pairwise orthogonal two-sided supports, $Q = \sum_s P_s$ is an orthogonal projection and both relevant closures lie in its first-site corner. The projector-controlled active maps are completely positive and preserve the trace selected by Q. Measuring $I - Q$ and preparing a fixed density matrix therefore completes each active map to a channel, while the complementary term vanishes on the closures. This resolves GitHub issue #6807 without asserting that the active Markov projection is the ambient identity. Thus two routes to the source implication remain: the source-context factorization in GitHub issue #6775, or the direct construction of one blocked channel pair for every absorbed representative in GitHub issue #6793.

The composition from supplied factorization data is complete end to end. Given, for every absorbed representative, a physical-sector factorization with positive semidefinite neighboring operators and normalized rank-one trace coefficients, the theorem `MPOTensor.blockTwo_isRFPViaTS_of_bntLayerOrthogonal_of_physicalSectorFactorization_of_commonWeight` constructs the pairwise orthogonal outer supports from layer orthogonality and one-site spanning, produces each representative channel pair by Proposition C.7, and assembles them into the blocked channel pair of the complete tensor. The coarsest merging is packaged generically by `MPOTensor.PhysicalSectorFactorization.NeighboringTraceFactorization.ofOneSector`: a single-sector factorization whose sole neighboring operator is positive semidefinite of unit trace suffices. The sole remaining obligation of the source implication (ii)⇒(v) is therefore the existence of these per-representative factorizations.

The all-cut Markov proof and the source-selected inverse-map calculation show that the raw four-sector tensor satisfies the displayed SAL, literal ZCL, and inverse-map relations while the trace matrix of those four fixed sectors is not rank one. This is not an existential obstruction. Reindexing a physical label as two binary labels gives a one-sector factorization with

$$l_0 = \tfrac{1}{4}I, \quad l_1 = Z, \qquad r_0 = I, \quad r_1 = \tfrac{1}{8}Z.$$

Its sole neighboring operator is

$$I \otimes \tfrac{1}{4}I + \tfrac{1}{8}Z \otimes Z,$$

which is positive semidefinite and has trace one. Thus the selected inverse-map sectors are nonminimal. This factorization and its normalized trace coefficients are formalized in `MPOTensor.ActiveSectorSpanningCounterexample.exists_oneSector_neighboringTraceFactorization`. Their non-rank-one trace matrix refutes only the proof path tied to those four selected sectors.

The corresponding low-level implication is refuted by a different four-letter tensor. Let

$$L = \begin{pmatrix} 1 & 1 & 1 & 1 \\ 1 & 2 & -7 & 4 \end{pmatrix}, \qquad R = \begin{pmatrix} 1/4 & -3/100 \\ 1/4 & -1/100 \\ 1/4 & 1/100 \\ 1/4 & 3/100 \end{pmatrix},$$

and take the four diagonal physical slices to be the outer products of the corresponding columns and rows. Then $LR = \operatorname{diag}(1, 0)$, so the physical trace transfer is literally idempotent, while RL has rank two. The four slices span $M_2(\mathbb{C})$, and the displayed scalar sector factorization has positive neighboring operators. Hence the tensor is injective and satisfies SAL. Perron normalization expresses it as a nonzero scalar multiple of a normal tensor.

The obstruction applies to every physical-sector factorization, not merely to the displayed one. The scalar slice and the two slices with simple spectra

$1, 2, -7, 4$ and $-3, -1, 1, 3$ force every left and right sector factor to be one-dimensional. The physical isometry then assigns each physical label to a unique sector. The neighboring trace matrix consequently contains a nonzero two-by-two minor, whereas a factorization $\mathrm{tr}(\eta_{k,h}) = a_k b_h$ has rank at most one. This is formalized in `MPOTensor.NonCartesianActiveSectorCandidate.full_lowLevel_counterexample`. Thus the local properties already inherited by an absorbed representative do not suffice to prove the factorization. Precisely, the formal theorem supplies injectivity, SAL (and hence MPDO generation), literal physical-trace idempotence, and an expression $\mathcal{K} = \mu A$ with A normal and $\mu \neq 0$. It does not package an ambient simple tensor in biCF whose canonical reconstruction contains $\mathcal{K}$. Moreover, for the singleton normalization produced here, $|\mu| = \sqrt{r}$, where r is the Perron value of the unnormalized transfer. The transfer endomorphism has trace 1353/5000. Its nonzero eigenvalues are roots of

$$\lambda^3 - \frac{1353}{5000}\lambda^2 - \frac{27}{5000}\lambda + \frac{9}{20000} = 0,$$

whose Perron root is $r \approx 0.284$. Hence $|\mu| \approx 0.533 < 1$. Since this normalization has only one coefficient, the line-246 convention requires that coefficient to have unit modulus and therefore fails. Rescaling to unit weight destroys literal ZCL. This is the scale tension recorded in `docs/paper-gaps/cpsv16_unit_weight_rfp_scale_tension.tex`. Consequently the construction does not resolve GitHub issue #6775 or refute the printed source-context assertion.

The non-Cartesian tensor is not itself the normal tensor assumed in Case I, and its normal representative loses literal ZCL. It therefore does not refute Lemma C.5 under its complete Case-I hypotheses. The completed normal Case-I corollary does not replace literal ZCL by TNLean's scale-invariant `MPOTensor.IsSourceZCL` predicate. The completed Case-II analytic inheritance likewise does not prove that coefficient absorption preserves normality.

By contrast, the earlier four-sector active-trace witness $\mathcal{K}$ has explicit trace-preserving completely positive renormalization maps, both before and after two-site blocking, but its sole BNT coefficient is strictly subunit. Its globally normalized representative $A = (4/\sqrt{5})\mathcal{K}$ retains SAL and TNLean's scale-invariant `MPOTensor.IsSourceZCL` relation, while $A^{[2]}$ fails the literal idempotence forced by conclusion (v). The normalized representative therefore refutes the broadened implication using that up-to-scalar predicate, but it is not an unqualified counterexample to Theorem 4.9 because it fails the paper's literal ZCL diagram. Independently, the repeated-copy theorem above refutes the exact printed implication (iv)$\Rightarrow$(v). The source-context (ii)$\Rightarrow$(iv) factorization remains GitHub issue #6775; the direct coarsening or sector-mixing construction of the blocked channel pair in GitHub issue #6793 is an alternative route to (ii)$\Rightarrow$(v). The conditional combination through the orthogonal outer physical sectors is formalized in GitHub issue #6632, and its trace-preserving completion outside the active support resolves GitHub issue #6807.

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.