

The Martingale Principal-Angle Estimate for MPS Parent Hamiltonians

2026-06-04

Verdict: scope-restriction (open).

1 The source assertion

The martingale-method paragraph in Cirac–Perez-Garcia–Schuch–Verstraete 2021 can be compared with the quantitative estimate isolated in the formal parent-Hamiltonian proof.¹ The relevant source passage is the discussion of gaps for MPS parent Hamiltonians in [CPGSV21, Section IV.C], where the local paragraph `sec:4:hams-gaps` develops the inequalities labelled `eq:4:martingale-1` and `eq:4:martingale-2`.

Let

$$H = \sum_i h_i$$

be a frustration-free Hamiltonian whose local terms are projections. The source recalls that a gap lower bound $\gamma > 0$ is equivalent, at the quadratic-form level, to

$$H^2 \geq \gamma H.$$

Since $h_i^2 = h_i$, the expansion of H^2 separates diagonal terms, overlapping products, and non-overlapping products. The non-overlapping products are treated as non-negative. Thus it is enough to control each overlapping pair by an anti-commutator estimate

$$h_i h_j + h_j h_i \geq -c_{ij}(1 - \gamma)(h_i + h_j), \quad (\text{eq:4:martingale-2})$$

where the coefficients c_{ij} are chosen so that the contributions in each row add up to one. The formal row-sum algebra uses the slightly more flexible condition that, for each fixed i , the sum of the coefficients attached to terms overlapping h_i is at most one.

The same paragraph interprets this estimate as a lower bound on the minimum non-zero angle between the ground spaces of overlapping regions. It then

¹For traceability, the finite-row reduction is recorded in issue #1044. The remaining blocked MPS ground-space estimate is recorded in issue #5455, and the broader parent-Hamiltonian theorem in issue #460.

cites the martingale method and the MPS parent-Hamiltonian results of Fannes–Nachtergaele–Werner and Nachtergaele, together with the later Kastoryano–Lucia formulation, for the conclusion that MPS parent Hamiltonians have a uniform positive spectral gap after a suitable blocking [FNW92, Nac96, KL18]. In the convention of this note, the Kastoryano–Lucia formulation cited there is the estimate for

$$\delta(\ell) := \|q_i q_j - P\|, \quad (1)$$

where q_i and q_j are the projectors onto the ground spaces $\ker h_i$ and $\ker h_j$, P is the projector onto $\ker h_i \cap \ker h_j = \ker(h_i + h_j)$, and ℓ is the size of the overlap region. The review writes the displayed product with the symbols h_i, h_j ; in the convention fixed above, those symbols must be read as the ground-space projectors q_i, q_j , not as the excitation projections h_i used in `eq:4:martingale-2`. If one inserted the excitation projections into $\|h_i h_j - P\|$, the norm would be at least one on the common ground space. With the ground-space projectors, the cited result states that $\delta(\ell)$ decays exponentially with universal constants whenever the frustration-free Hamiltonian is gapped, and that this implies the martingale condition because both quantities depend only on the principal angles between $\ker h_i$ and $\ker h_j$. The review text does not spell out the numerical constants or the precise finite cyclic blocking convention needed for the formal statement below. In particular, the intersection identity

$$\ker h_i \cap \ker h_j = \ker(h_i + h_j)$$

identifies the common local ground space, but it is not a substitute for the quantitative principal-angle estimate.

The source route that has to be matched is more precise than the review paragraph. Kastoryano–Lucia define, for finite regions A and B ,

$$\delta(A, B) = \|(P_A - P_{A \cup B})(P_B - P_{A \cup B})\| = \|P_A P_B - P_{A \cup B}\|$$

under frustration-freeness, and identify this with the cosine of the first principal angle between the two reduced ground-space projectors [KL18, Definition 2 and Remark 2]. Nachtergaele’s martingale hypotheses C3 and C3’ instead use the projection estimates

$$\|G_{\Lambda_{n+1} \setminus \Lambda_{n-l}} E_n\| \leq \varepsilon_l, \quad \|G_{\Lambda_{n+p} \setminus \Lambda_{n-l}} E_n^{(p)}\| \leq \eta_l,$$

with thresholds $\varepsilon_l < 1/\sqrt{l+1}$ and $\eta_l < 1/\sqrt{2}$, and Theorem 3 gives the corresponding gap lower bounds [Nac96, Equations (2.4)–(2.5) and Theorem 3]. In Section 6, Nachtergaele proves the GVBS/MPS gap by verifying C3’ using the projection estimates inherited from Fannes–Nachtergaele–Werner [Nac96, Section 6]. The formal development instead obtains a sufficient range-two estimate after physical blocking. The optimized FNW numerator and the identification of its source constants remain open.

1.1 Unit bound at the full-window volume

Fix an MPS tensor A and an interaction length $L > 0$. If $P_{\mathcal{G}_L^{\text{ES}}(A)}$ denotes the orthogonal projector onto the length- L MPS ground space, define the excitation projector by

$$P_L^{\text{ES}}(A) := \mathbb{1} - P_{\mathcal{G}_L^{\text{ES}}(A)}.$$

For a chain of length N and a start i with $i + L \leq N$, let $h_{i,\text{ES}}(A, L)$ be the copy of this excitation projector on the window $[i, i + L - 1]$, acting as the identity outside the window. Set

$$I_{L,N} = \{i \in \{0, \dots, N - 1\} : i + L \leq N\}$$

and define the compatible open Hamiltonian by

$$H_{N,L}^{\text{open}}(A) := \sum_{i \in I_{L,N}} h_{i,\text{ES}}(A, L).$$

At the full-window volume $N = L$, the compatible open Hamiltonian consists of one orthogonal projection.² Thus

$$H_{L,L}^{\text{open}}(A) = P_L^{\text{ES}}(A).$$

Consequently it preserves the norm on the orthogonal complement of its kernel. This identity verifies only the single-window case. It gives no gap estimate for $N > L$. The general-volume result instead requires the C1–C3 ingredients described below: the open-chain ground-space intersection, the quantitative projector-defect estimate, and the finite-row martingale reduction.

1.2 Open-chain projector geometry

The source-faithful open-chain geometry is now isolated before any cyclic specialization. In the common Hilbert space on Λ_{n+1} , the formal subspaces `openChainLeftGroundSpaceES` and `openChainTailGroundSpaceES` are the ranges of G_{Λ_n} and $G_{\Lambda_{n+1} \setminus \Lambda_{n-l}}$, expressed using the existing last-site and fixed-prefix restriction maps. Under L_0 -block injectivity, with $n = K + L_0$ and $l = L_0$, the theorem `openChainLeft_inf_tailGroundSpaceES` proves

$$U \cap V = \mathcal{G}_{K+L_0+1}^{\text{ES}}(A).$$

The generic theorem `isFriedrichsBound_of_norm_starProjection_com`
`p_sub_inf_starProjection_le` converts

$$\|P_U P_V - P_{U \cap V}\| \leq \eta$$

into a Friedrichs bound. Combining these two results with the complementary projection theorem gives the open-chain corollary `re_inner_openChain_antic`

²This is the Blueprint theorem “Full-window compatible interaction,” labelled `thm:open_parent_hamiltonian_full_window`; the formal declaration is `MPSTensor.openParentHamiltonianES_self_eq_parentInteractionES`.

`ommutator_ge_neg_of_groundProjection_defect`. Its hypothesis is precisely the C3 defect

$$\|G_{\Lambda_{n+1} \setminus \Lambda_{n-l}} G_{\Lambda_n} - G_{\Lambda_{n+1}}\| \leq \eta,$$

using the identity following condition C3 in [Nac96, Equation (2.4)]. Equation (2.5) is the separate batched condition C3' involving $E_n^{(p)}$.

The open-chain numerical step is also complete. The geometric estimate below is uniform in the left length, and `MPSTensor.IsPrimitiveMPS.exists_openChain_groundProjection_defect_lt_c3_threshold` chooses one block-injective $l > 1$, measured in sites of the input tensor, and one coefficient $\varepsilon_l \geq 0$ with

$$\varepsilon_l < \frac{1}{\sqrt{l+1}}.$$

The companion theorem `MPSTensor.IsPrimitiveMPS.exists_re_inner_openChain_anticommutator_ge_c3_threshold` combines Nachtergaele's projector-defect identity with the two-projection estimate proved here to give the corresponding anticommutator inequality for every left length. For the final range-two argument, exact blocked three-site identities now connect the three-block defect to the local terms of the blocked tensor. No excitation-compression estimate is used.

1.3 Three-site to cyclic transport

The independent locality step is now complete for range-two terms. If the two adjacent terms h_0, h_1 on the three-site Hilbert space satisfy

$$h_0 h_1 + h_1 h_0 \geq -c(h_0 + h_1)$$

in quadratic form, then `MPSTensor.re_inner_localTermES_adjacent_twoSite_forward_ge_of_threeSite` and `MPSTensor.re_inner_localTermES_adjacent_twoSite_backward_ge_of_threeSite` transport the same inequality, with the same coefficient c , to every forward and backward adjacent cyclic pair on every chain of length $N \geq 3$. The proof splits the configuration basis into a cyclic three-site window and its ordered complement, applies the three-site estimate fiberwise through cyclic restriction, and sums over the complement. Since the window may begin at any site, the result includes positions crossing the periodic cut; the backward statement begins the forward window at the cyclic predecessor.

The exact identification is `MPSTensor.blockTensor_threeSite_openChain_defect_norm_eq`. It equates the three-site open-chain projector defect of $B = \text{block}_p(A)$ with the whole-block defect on three consecutive p -site regions of A . Together with `MPSTensor.IsPrimitiveMPS.exists_threeBlock_wholeIncrement_defect_le_seven_sixteenths`, it supplies the coefficient $7/16$ for the adjacent three-site terms of B . The two transport theorems above then cover every overlapping off-diagonal cyclic pair, including the periodic cut. This argument does not claim the optimized FNW coefficient and does not use an excitation-compression surrogate.

2 Finite-row martingale reduction

The current development separates the general finite-row martingale algebra from the completed range-two blocked MPS specialization. Fix an interaction range $L > 1$ and an N -site periodic chain with $N \geq 2L$. Let p_i denote the Hilbert-space representative of the length- L local parent-Hamiltonian projection at the cyclic window starting at i . Write $i \sim_L j$ when the two length- L cyclic windows intersect and $i \neq j$. For each fixed i , the number of indices j with $i \sim_L j$ is at most

$$m = 2(L - 1).$$

The row-sum reduction specializes the source coefficients to

$$c_{ij} = \frac{1}{2(L - 1)}$$

for overlapping cyclic windows and to 0 otherwise. The row-cardinality bound then gives $\sum_j c_{ij} \leq 1$.

With this choice, the constant in the local statement is

$$\gamma_L = \frac{1}{4L}.$$

The ordered overlapping estimate sufficient for the gap bound is

$$\operatorname{Re}\langle p_i v, p_j v \rangle \geq - \left(1 - \frac{1}{4L}\right) \frac{1}{2(L - 1)} \operatorname{Re}\langle p_i v, v \rangle \quad (\text{F})$$

for every $N \geq 2L$, every off-diagonal overlapping pair $i \sim_L j$, and every vector v in the finite chain Hilbert space. Given this estimate, the already formalized row-sum argument proves $H_{N,L}^2 \geq \gamma_L H_{N,L}$, and hence a norm lower bound with constant γ_L on the orthogonal complement of the ground space.

The proof also isolates a sufficient norm-compression form of the same input:

$$\|p_i p_j v\| \leq \left(1 - \frac{1}{4L}\right) \frac{1}{2(L - 1)} \|p_i v\|. \quad (\text{N})$$

For a symmetric projection p_i one has

$$\operatorname{Re}\langle p_i v, p_j v \rangle = \operatorname{Re}\langle p_i v, p_i p_j v \rangle \geq -\|p_i v\| \|p_i p_j v\|.$$

Thus (N), together with $\operatorname{Re}\langle p_i v, v \rangle = \|p_i v\|^2$, implies (F). This is a sufficient projection-geometry comparison used in the proof: a norm-compression estimate for each overlapping pair is converted by Cauchy–Schwarz to the ordered cross-term estimate, and the finite row-sum argument converts those ordered estimates into the global gap bound. The Lean development also records the

source anticommutator row-sum reduction directly, without passing through this stronger one-sided estimate.³

The anticommutator row-sum formulation, including its finite-overlap cyclic-window specialization, is now the formal condition matching the source. The flexible constant formulation records a sufficient norm-compression variant: instead of requiring the special coefficient in (N), it assumes that for some $\eta \geq 0$,

$$\eta 2(L-1) < 1 \tag{E}$$

and that every overlapping cyclic-window pair satisfies

$$\|p_i p_j v\| \leq \eta \|p_i v\|. \tag{N_\eta}$$

The finite-row argument then gives a positive gap with lower bound $1 - \eta 2(L-1)$. The special displayed estimate (N) is one sufficient choice of η , but the source comparison need only supply a compression constant strictly smaller than $1/(2(L-1))$ for the blocking convention used here.

3 What remains to compare

The source discussion and the formal theorem are aligned at the qualitative level: both use the martingale expansion, non-negativity of non-overlapping terms, row-sum coefficients for overlapping terms, and principal-angle estimates for local ground spaces. The point still requiring a mathematical check is quantitative and conventional.

First, the source-matching formal theorem asks for the cyclic-window anticommutator estimate

$$p_i p_j + p_j p_i \geq - \left(1 - \frac{1}{4L}\right) \frac{1}{2(L-1)} (p_i + p_j)$$

³The current formal anchors include the anticommutator row-sum reductions `crossTerm_sum_bound_of_anticommutator_rowCol`, `quadraticForm_sum_projections_of_anticommutator_rowCol`, and `parentHamiltonianES_gap_bound_of_anticommutator_local_term_bounds`. The finite-overlap anticommutator reductions are `quadraticForm_sum_projections_of_finite_overlap_anticommutator`, `parentHamiltonianES_gap_bound_of_finite_overlap_anticommutator`, and `parentHamiltonianES_gap_bound_of_cyclic_window_overlap_anticommutator`. The corresponding final gap statements are `parentHamiltonianES_gap_bound_of_anticommutator` and `parentHamiltonian_gapped_of_anticommutator` in `TNLean/MPS/ParentHamiltonian/Martingale/Gap.lean`. The cyclic-window compression reductions are `parentHamiltonianES_gap_bound_of_cyclic_window_overlap_ordered_cross_term`, `parentHamiltonianES_gap_bound_of_cyclic_window_overlap_norm_bound`, and `parentHamiltonianES_gap_bound_of_cyclic_window_overlap_norm_bound_of_le` in `TNLean/MPS/ParentHamiltonian/Martingale/Reduction.lean`. The corresponding constant- η reductions are `parentHamiltonianES_gap_bound_of_cyclic_window_overlap_norm_bound_of_lt` in the same file and `parentHamiltonianES_gap_bound_of_overlap_norm_constant`, `parentHamiltonian_gapped_of_overlap_norm_constant` in `TNLean/MPS/ParentHamiltonian/Martingale/Gap.lean`. The projection-geometry reductions used by these statements are in `TNLean/Analysis/ProjectionGeometry.lean`. The declaration names identify the intermediate ordered cross-term estimate and the sufficient norm-compression estimates; this note uses “anticommutator estimate” for the source condition and “norm compression” for the local sufficient strengthening.

for overlapping off-diagonal pairs. The review text states the same martingale condition with general coefficients c_{ij} whose rows add to one, and it refers to blocking and principal angles to obtain a positive gap. It does not by itself identify the Fannes–Nachtergaele–Werner, Nachtergaele, or Kastoryano–Lucia estimates with this coefficient and this fixed cyclic-window convention. A proof may instead derive a row-sum anticommutator estimate with different positive constants and then specialize the final gap constant accordingly.

Second, the blocking parameter in the cited martingale argument must be matched with the formal convention. The formal statement is uniform for all $N \geq 2L$ and for the length- L cyclic windows used in the formal finite-chain Hilbert-space realization. A source proof that produces some positive gap after increasing or changing the blocking may still be mathematically sufficient for the existential MPS gap theorem, but it would not immediately be the displayed all- L estimate with constant $1/(4L)$ unless the constants are converted to the same blocking convention.

The general source comparison remains open at the level of the optimized FNW coefficient and its source constants. It is no longer needed for the proved range-two blocked theorem: the exact three-site identification gives coefficient $7/16$, and cyclic transport gives every overlapping pair. In particular, `re_inner_anticommutator_starProjection_orthogonal_ge_neg_of_friedrichs_bound` shows that a Friedrichs bound η for the two ground-space ranges, after removing their common intersection, gives

$$h_i h_j + h_j h_i \geq -\eta(h_i + h_j)$$

for the complementary excitation projections. What remains open is whether the cited FNW–Nachtergaele–Kastoryano estimates recover their optimized numerator and source constants in the present Gram and blocking conventions. The qualitative fact that two finite-dimensional subspaces have a well-defined principal-angle decomposition is not enough: the martingale reduction needs a uniform numerical bound after the same blocking convention has been fixed.

The all-vector norm-compression theorem remains available as a stronger conditional implication: any estimate (N_η) with $\eta 2(L-1) < 1$ gives a positive gap through the already formalized reduction. It is not the source principal-angle target. There is, however, a directional point. If (N_η) is imposed for all vectors v , then $p_i v = 0$ implies $p_i p_j v = 0$. Thus p_j maps $\ker p_i$ into $\ker p_i$. This is stronger than merely knowing that the smallest non-zero principal angle between the two local ground spaces is positive. A proof using compression must therefore identify a principal-angle estimate that gives the directed compressed-product bound (N_η) . A direct proof can instead target the anticommutator form

$$h_i h_j + h_j h_i \geq -c_{ij}(1 - \gamma)(h_i + h_j)$$

used in [CPGSV21, Section IV.C].

The development now also records the symmetric operator-product route. Instead of the directed bound (N_η) , suppose the operator-product bound

$$\|p_i(p_j v)\| \leq \eta \|p_j v\| \quad (O_\eta)$$

holds for all v , that is, a bound $\|p_i p_j\| \leq \eta$ on the norm of the product. The identity $\text{Re}\langle p_i v, p_j v \rangle = \text{Re}\langle p_i v, p_i p_j v \rangle$, Cauchy–Schwarz, and the arithmetic–geometric mean inequality give, for $0 \leq \eta$,

$$\text{Re}\langle p_i v, p_j v \rangle + \text{Re}\langle p_j v, p_i v \rangle \geq -\eta(\|p_i v\|^2 + \|p_j v\|^2),$$

which is exactly the symmetric anticommutator estimate with coefficient η . The difference from (N_η) is which kernel the hypothesis constrains: imposed for all v , (N_η) forces $p_j(\ker p_i) \subseteq \ker p_i$, whereas (O_η) imposes no such preservation of $\ker p_i$. Through this conversion any (O_η) with $\eta 2(L-1) < 1$ yields the same positive gap $1 - \eta 2(L-1)$.

This removes the kernel-preservation artefact of (N_η) , but it does not make the hypothesis achievable for the MPS application, and (O_η) must not be identified with the cited principal-angle estimate. The symbols $p_i = h_i$ here are the *excitation* projections (the local parent-Hamiltonian terms), not the ground-space projectors q_i of §1. The Kastoryano–Lucia quantity $\delta(\ell) = \|q_i q_j - P\|$ bounds the principal angle between the two local *ground* spaces, while $\|p_i p_j\|$ is the cosine of the smallest principal angle between the two *excitation* ranges; these are different quantities. If the two overlapping excitation ranges share any nonzero vector then $\|p_i p_j\| = 1$, so (O_η) with $\eta < 1$ is unsatisfiable, and the target $\eta < 1/(2(L-1))$ cannot follow from the ground-space intersection estimate without an additional quotient or disjointness argument. Like (N_η) , the operator-product bound is therefore a *conditional* reduction recorded in the formal development, not an achievable source-faithful estimate. The abstract conversion is `re_inner_anticommutator_ge_neg_of_norm_apply_le`; the cyclic-window entry points are `parentHamiltonianES_gap_bound_of_cyclic_window_overlap_operator_norm_of_le` and `parentHamiltonianES_gap_bound_of_cyclic_window_overlap_operator_norm_of_lt`, with public reformulations `parentHamiltonianES_gap_bound_of_overlap_operator_norm_constant` and `parentHamiltonian_gapped_of_overlap_operator_norm_constant`. For the blocked range-two theorem, the exact three-site projector identity and the forward and backward adjacent-pair transports replace this conditional route. The generic theorem `re_inner_anticommutator_starProjection_orthogonal_ge_neg_of_friedrichs_bound` then supplies the excitation-projection anticommutator estimate. This route uses the intersection identity $\ker h_i \cap \ker h_j = \ker(h_i + h_j)$ or its restatement `adjacent_localTermES_eq_zero_iff_mem_groundSpaceES_succ`; a uniform small bound on $\|p_i p_j\|$ itself is not available in the generic overlapping case, since that norm is 1 whenever the two excitation ranges have a common nonzero vector.

4 Relation to the blueprint

The blueprint item labelled `thm:martingale_mps` in `blueprint/src/chapter/ch13_parent_hamiltonian_spectral_gap_martingale.tex` records the completed blocked range-two theorem. For a primitive tensor with a positive-definite fixed point it chooses $p > 0$, sets $B = \text{block}_p(A)$, and proves a gap

lower bound $1/8$ for $H_{N,2}(B)$ for every $N \geq 4$. Its preceding item records the exact three-site projector identification and the transport of the $7/16$ adjacent anticommutator estimate to every cyclic overlap.

The source-matching theorem entry `thm:anticommutator_gap_bound` points to `parentHamiltonianES_gap_bound_of_anticommutator` in `TNLean/MPS/ParentHamiltonian/Martingale/Gap.lean`. Its hypothesis is the cyclic-window anticommutator lower bound with coefficient

$$\left(1 - \frac{1}{4L}\right) \frac{1}{2(L-1)}.$$

The theorem entries `thm:overlap_norm_gap_bound` and `thm:overlap_norm_gap_bound_constant` point respectively to the special coefficient theorem and to the constant- η theorem in `TNLean/MPS/ParentHamiltonian/Martingale/Gap.lean`. The special coefficient theorem is the specialization of the constant- η theorem with

$$\eta = \left(1 - \frac{1}{4L}\right) \frac{1}{2(L-1)}, \quad 1 - \eta 2(L-1) = \frac{1}{4L}.$$

The two norm-compression theorem entries are therefore conditional on the same stronger input, specialized either to this displayed coefficient or to an arbitrary coefficient satisfying (E). Thus all finite-row, cyclic-window, non-overlap, anticommutator, norm-compression, and spectral-theorem reductions are proved, and the blocked range-two MPS specialization is complete. The all-vector compression bound (N_η) , if separately assumed, remains only a stronger conditional estimate. The separate open problem is to compare the blocked Hamiltonian with a parent interaction for the original tensor in original-site units, as recorded under issue #5926.

5 Conclusion

The source states the martingale condition in terms of general row-sum coefficients and invokes principal-angle estimates to obtain a positive MPS parent-Hamiltonian gap. The formal proof contains the direct finite-overlap anticommutator reduction for cyclic windows and specializes it to the blocked range-two tensor with adjacent-pair coefficient $7/16$, giving gap $1/8$. It also records the stronger norm-compression route

$$\|p_i p_j v\| \leq \eta \|p_i v\|, \quad \eta 2(L-1) < 1.$$

The earlier coefficient

$$\left(1 - \frac{1}{4L}\right) \frac{1}{2(L-1)}$$

is a sufficient special case for this stronger route, not the only possible way to use the formal martingale reduction.

The Cirac–Perez-Garcia–Schuch–Verstraete 2021 statement is not being challenged. The formal result gives a sufficient explicit constant for the range-two parent Hamiltonian after blocking. Recovering the optimized FNW numerator and its source constants remains a separate comparison.

6 Inverse-Gram reduction and the remaining optimized FNW comparison

For $n \geq 1$, write $[n] := \{0, \dots, n-1\}$. Fix an MPS tensor $A = \{A^s\}_{s \in [d]}$ with $A^s \in \mathbb{C}^{D \times D}$. For $a, b \in [D]$, let E_{ab} be the matrix unit with its only nonzero entry, equal to one, in row a and column b . The transfer map and its L -fold composition are

$$\mathcal{E}_A(X) := \sum_{s \in [d]} A^s X (A^s)^\dagger, \quad \mathcal{E}_A^L := \underbrace{\mathcal{E}_A \circ \dots \circ \mathcal{E}_A}_{L \text{ factors}}. \quad (2)$$

For a linear map $T: \mathbb{C}^{D \times D} \rightarrow \mathbb{C}^{D \times D}$, this note uses the rectangular Choi convention

$$J(T)_{(i,a),(j,b)} := (T(E_{ij}))_{a,b}. \quad (3)$$

The open-chain boundary map is

$$\Gamma_L(X)(\sigma) := \text{tr}(A^\sigma X). \quad (4)$$

Let $U: \mathbb{C}^{D \times D} \rightarrow \ell^2([D] \times [D])$ be Frobenius vectorization, so that $U(E_{ab}) = \mathbf{e}_{(b,a)}$ for the standard unit vector $\mathbf{e}_{(b,a)}$. After the canonical physical ℓ^2 identification, write Γ_L^{ES} for the transported boundary map and

$$\mathcal{K}_L := (\Gamma_L^{\text{ES}})^\dagger \Gamma_L^{\text{ES}}. \quad (5)$$

Its finite-dimensional Hilbert realization, exact range, and Gram operator are recorded by `MPSTensor.groundSpaceMapES`, `MPSTensor.range_groundSpaceMapES`, and `MPSTensor.groundSpaceGram`. With the coordinate order just fixed, the exact identity is

$$\langle \mathbf{e}_{(b,a)}, \mathcal{K}_L(\mathbf{e}_{(e,c)}) \rangle_{\ell^2} = J(\mathcal{E}_A^L)_{(c,e),(a,b)} = (\mathcal{E}_A^L(E_{ca}))_{e,b}. \quad (6)$$

This is `MPSTensor.inner_single_groundSpaceGram_single_eq_rectangularChoi`. For any matrix endomorphism $\mathcal{L}$, let $\mathcal{R}(\mathcal{L})$ be the Euclidean operator with matrix entries

$$\mathcal{R}(\mathcal{L})_{(b,a),(e,c)} := J(\mathcal{L})_{(c,e),(a,b)}. \quad (7)$$

The algebraic matrix $R(\mathcal{L})$ with these entries is `Matrix.gramReshuffleMatrix`, and the operator $\mathcal{R}(\mathcal{L})$ is `Matrix.gramReshuffle`. The entrywise Frobenius squared sum is invariant under this reshuffling (`Matrix.gramReshuffleMatr`

`ix_norm_sq_eq_rectangularChoi_norm_sq`). The generic Euclidean estimate `Matrix.toEuclideanCLM_norm_sq_le_sum_norm_sq` proves, for every finite index set I and every complex matrix $B \in \mathbb{C}^{I \times I}$,

$$\|B\|_{\ell^2(I) \rightarrow \ell^2(I)}^2 \leq \sum_{i,j \in I} |B_{ij}|^2. \quad (8)$$

Combining these two facts gives the completed reshuffled-Choi corollary `Matrix.gramReshuffle_norm_sq_le_rectangularChoi_norm_sq`:

$$\|\mathcal{R}(\mathcal{L})\|_{\ell^2(I \times I) \rightarrow \ell^2(I \times I)}^2 \leq \sum_{u,v \in I \times I} |J(\mathcal{L})_{u,v}|^2. \quad (9)$$

The remaining generic conversion is also complete. Equip $\mathbb{C}^{I \times I}$ with its L^2 operator norm and write $\|\mathcal{L}\|_{\text{op}, L^2}$ for the resulting operator norm of $\mathcal{L}$. The column estimate `Matrix.sum_column_norm_sq_le_norm_sq`, its entrywise consequence `Matrix.sum_entry_norm_sq_le_card_mul_norm_sq`, and the matrix-unit normalization `Matrix.l2_opNorm_single_one` give `Matrix.gramReshuffle_norm_sq_le_card_cube_mul_opNorm_sq`:

$$\|\mathcal{R}(\mathcal{L})\|_{\ell^2(I \times I) \rightarrow \ell^2(I \times I)}^2 \leq |I|^3 \|\mathcal{L}\|_{\text{op}, L^2}^2. \quad (10)$$

Thus for $I = \{0, \dots, D-1\}$ the corresponding norm factor is $D^{3/2}$. This is a safe generic dimension conversion. It is not claimed sharp, and it is not identified with the coefficient in the FNW weighted-Gram estimate.

Reshuffling is still not operator-norm invariant and need not preserve singular values. The identity map gives the basic boundary: its operator norm on $\mathbb{C}^{D \times D}$ is one, whereas $\|\mathcal{R}(\text{id})\| = D$ for $D > 0$. Hence the preceding $D^{3/2}$ factor is an upper bound, not an invariance statement or a source constant.

`MPSTensor.groundSpaceGram_eq_gramReshuffle` upgrades the coordinate identity to the exact operator equality

$$\mathcal{K}_L = \mathcal{R}(\mathcal{E}_A^L). \quad (11)$$

The ordering is a Choi reshuffling. It must not be replaced by the generally false Hilbert–Schmidt expression $\langle E_{ab}, \mathcal{E}_A^L(E_{ce}) \rangle_{\text{HS}}$. Reshuffling commutes with subtraction by `Matrix.gramReshuffle_sub`. Physical Gram convergence in these coordinates is now complete under the exact complementary spectral hypothesis. Let $\rho \in \mathbb{C}^{D \times D}$ have nonzero trace, let $P_\rho(X) = \text{tr}(X)\rho / \text{tr}(\rho)$, and suppose that $\mathcal{E}_A$ is trace-preserving, $\mathcal{E}_A(\rho) = \rho$, and

$$\rho_{\text{spec}}(\mathcal{E}_A - P_\rho) < 1. \quad (12)$$

The generic prerequisite `geometric_bound_of_spectralRadius_lt_one` gives constants $C > 0$ and $0 < r < 1$ with $\|(\mathcal{E}_A - P_\rho)^n\| \leq Cr^n$ for every n . Combining the transfer-power decomposition, the exact Gram reshuffling identity, subtraction compatibility, and the D^3 squared norm conversion gives `MPSTensor.groundSpaceGram_sub_fixedPointProj_norm_sq_le_geometric`:

$$\|\mathcal{K}_n - \mathcal{R}(P_\rho)\|^2 \leq D^3 (Cr^n)^2, \quad n \geq 1. \quad (13)$$

Consequently, `MPSTensor.groundSpaceGram_tendsto_gramReshuffle_fixedPointProj` proves

$$\mathcal{K}_n \longrightarrow \mathcal{R}(P_\rho) \quad (14)$$

in operator norm. These theorems require exactly nonzero trace, trace preservation, the fixed-point equation, and complementary spectral radius strictly below one. They impose no positivity, injectivity, or uniqueness hypothesis.

The limiting Gram operator is $\mathcal{R}(P_\rho)$, generally not the identity in the current unweighted Frobenius coordinates. Its exact action is right multiplication in the Frobenius vectorization U defined above,

$$\mathcal{R}(P_\rho)U(X) = U\left(\frac{X\rho}{\text{tr } \rho}\right). \quad (15)$$

This identity is `Matrix.gramReshuffle_fixedPointProj_frobeniusEquivEuclidean_apply`. Hence a positive-definite fixed point makes the limiting Gram invertible. The correct perturbation is around $\mathcal{R}(P_\rho)$, not around the identity.

The generic identity

$$P_{\text{ran } T} = T(T^\dagger T)^{-1}T^\dagger \quad (16)$$

for injective finite-dimensional maps is `ContinuousLinearMap.injectiveRangeProjector_eq_starProjection`, and the denominator estimate $\|(1-t)^{-1}\| \leq (1-a)^{-1}$ under $\|t\| \leq a < 1$ is `NormedRing.norm_inverse_one_sub_le`. For an injective T , write $S_T := (T^\dagger T)^{-1}$ for its inverse Gram operator. The identities `ContinuousLinearMap.inverseGram_comp_adjoint_comp_self` and `ContinuousLinearMap.adjoint_comp_self_comp_inverseGram` show that S_T is both a left and right inverse of $T^\dagger T$. The theorem `ContinuousLinearMap.inverseGram_eq_ringInverse` identifies it with the canonical ring inverse. Consequently, `ContinuousLinearMap.norm_inverseGram_le_of_norm_one_sub_adjoint_comp_self_le` proves directly that

$$\|S_T\| \leq (1-a)^{-1} \quad (17)$$

whenever $\|1 - T^\dagger T\| \leq a < 1$.

Suppose that $L_0 > 0$ and that A is L_0 -block injective. Block injectivity then makes the Hilbert-space boundary map injective at every length $L \geq L_0$; this is recorded by `MPSTensor.groundSpaceMapES_injective_of_isNBlkInjective` and `MPSTensor.groundSpaceMapES_injective_of_isNBlkInjective_of_le`. Writing $\mathcal{G}_L^{\text{ES}}(A)$ for the physical local ground space, the specialization `MPSTensor.groundSpaceES_starProjection_eq_groundSpaceMapES_comp_inverseGram_comp_adjoint` therefore proves, for every $L \geq L_0$, the exact physical formula

$$P_{\mathcal{G}_L^{\text{ES}}(A)} = \Gamma_L^{\text{ES}} \mathcal{K}_L^{-1} (\Gamma_L^{\text{ES}})^\dagger. \quad (18)$$

This is the orthogonal projector onto $\text{ran } \Gamma_L^{\text{ES}}$. It is not the transfer fixed-point projector `fixedPointProj`; the latter acts on the virtual matrix space and cannot substitute for the inverse-Gram formula.

To state the source estimate, let $l \in \mathbb{N}$ be the overlap length and let $0 \leq \lambda < 1$ be an admissible transfer-mixing rate satisfying $|\mu| < \lambda$ for every non-unit eigenvalue μ of the relevant transfer operator. Let $c \geq 0$ be its prefactor; the source allows $c = k^2$ when k is the auxiliary, or bond, dimension. Choose l so that $c\lambda^l < 1$. These are parameters of the quoted FNW estimate, not bounds derived by the present formal development.

What remains open is the source-specific comparison of these exact inverse-Gram range projectors on the two overlapping physical windows. The completed exact and generic prerequisites now include:

- The Choi reshuffling is packaged as the exact operator identity `MPSTensor.gramSpaceGram_eq_gramReshuffle` (module `TNLean/MPS/ParentHamiltonian/GroundSpaceGram.lean`), and subtraction compatibility is `Matrix.gramReshuffle_sub`.
- The generic geometric power estimate is `geometric_bound_of_spectralRadius_lt_one`. Under nonzero trace, trace preservation, the fixed-point equation, and complementary spectral radius below one, the quantitative and limiting physical Gram theorems are `MPSTensor.gramSpaceGram_sub_fixedPointProj_norm_sq_le_geometric` and `MPSTensor.gramSpaceGram_tendsto_gramReshuffle_fixedPointProj`. Their limit is $\mathcal{R}(P_\rho)$, and the quantitative squared estimate has factor $D^3(Cr^n)^2$.
- The reconstructed quantitative inverse infrastructure is `MPSTensor.IsPrimitiveMPS.gramSpaceGram_geometric_inverse_bounds` and `MPSTensor.IsPrimitiveMPS.gramSpaceMapES_geometric_inverseGram_bounds`. Assume that the bond dimension is nonzero and that the primitive fixed-point witness ρ is explicitly positive definite. Set

$$\mathcal{K}_\infty = \mathcal{R}(P_\rho), \quad \mathcal{I}_\infty = \mathcal{K}_\infty^{-1}. \quad (19)$$

Then there are $g, c, r > 0$, with $r < 1$ and $c = \|\mathcal{I}_\infty\|g$, such that for every $n \geq 1$,

$$\|\mathcal{K}_n - \mathcal{K}_\infty\| \leq gr^n. \quad (20)$$

If $cr^n < 1$, then $\mathcal{K}_n$ is invertible and

$$\|\mathcal{K}_n^{-1}\| \leq \frac{\|\mathcal{I}_\infty\|}{1 - cr^n}, \quad (21)$$

$$\|\mathcal{K}_n^{-1} - \mathcal{I}_\infty\| \leq \frac{cr^n}{1 - cr^n} \|\mathcal{I}_\infty\|. \quad (22)$$

At the same lengths the physical boundary map is injective, its inverse Gram is $\mathcal{K}_n^{-1}$, and it obeys the same two bounds. The proof retains the available squared estimate $\|\mathcal{K}_n - \mathcal{K}_\infty\|^2 \leq D^3(Cr^n)^2$ and takes square roots directly, giving $g = \sqrt{D^3}C$ without any additional dimension loss. The base is the generally nonidentity operator $\mathcal{R}(P_\rho)$, not the identity.

- The elementary propagation theorem `MPSTensor.geometric_smallness_at_c3_lengths` shows that $cr^l < 1$ implies

$$cr^{l+1} < 1, \quad cr^{K+l} < 1, \quad cr^{K+l+1} < 1 \quad (23)$$

for every $K \in \mathbb{N}$. Thus one base-length smallness hypothesis supplies the three inverse conditions appearing in the C3 correction telescope. This is quantitative inverse infrastructure only: these results do not prove contraction of the full physical projector defect or the coefficient in Nachtergaele's equation (6.1).

- The generic operator-norm-to-entrywise-Frobenius estimate, its reshuffled-Choi corollary, and the safe dimension conversion are `Matrix.toEuclideanCLM_norm_sq_le_sum_norm_sq`, `Matrix.gramReshuffle_norm_sq_le_rectangularChoi_norm_sq`, and `Matrix.gramReshuffle_norm_sq_le_card_cube_mul_opNorm_sq`. The last theorem gives squared factor $|I|^3$ against the squared operator norm of the map when the matrix space carries its L^2 operator norm. For $I = \{0, \dots, D-1\}$ this is norm factor $D^{3/2}$; it is not claimed sharp and is not identified with the coefficient in the FNW weighted-Gram estimate.
- The spectator-indexed boundary maps for the tail and left windows are defined algebraically in `TNLean/MPS/ParentHamiltonian/SpectatorBoundary.lean`. Their finite-dimensional Hilbert realizations are `MPSTensor.tailBoundaryMapES` and `MPSTensor.leftBoundaryMapES` in `TNLean/MPS/ParentHamiltonian/SpectatorBoundaryGram.lean`. The latter module proves the exact physical ranges, the factorizations through Γ_{K+L}^{ES} , and injectivity from the corresponding local boundary map. It also computes the Grams as direct sums of $\mathcal{K}_L$ on tail fibers and $\mathcal{K}_K$ on left fibers: `MPSTensor.tailBoundaryMapES_adjoint_comp_self_eq_fiberwise_groundSpaceGram` and `MPSTensor.leftBoundaryMapES_adjoint_comp_self_eq_fiberwise_groundSpaceGram`. The inverse Grams are likewise fiberwise by `MPSTensor.inverseGram_tailBoundaryMapES_eq_fiberwise_inverseGram` and `MPSTensor.inverseGram_leftBoundaryMapES_eq_fiberwise_inverseGram`. These are exact unweighted identities, including empty spectator types; they assert no direct-sum norm estimate. The C3 specializations are stated next.
- The C3' whole-increment algebra and a quantitative prefix-uniform estimate are now complete. For a suffix increment Q , define the left virtual map by

$$J_Q^{\text{left}} X = (U(A^j X))_{j \in \Omega_Q}, \quad (24)$$

where $\Omega_Q = \{0, \dots, (-1\}d^Q$ is the set of length- Q physical words and U is Frobenius vectorization. Also set

$$\mathcal{K}_n = (\Gamma_n^{\text{ES}})^\dagger \Gamma_n^{\text{ES}}, \quad \mathcal{K}_\infty = \mathcal{R}(P_\rho), \quad (25)$$

where $P_\rho(X) = (\text{tr } X / \text{tr } \rho) \rho$. The declarations `MPSTensor.leftVirtualMapES_norm_map_of_leftCanonical` and `MPSTensor.leftVirtualMapES_norm_le_one_of_leftCanonical` prove

$$\|J_Q^{\text{left}} X\| = \|X\|, \quad \|J_Q^{\text{left}}\| \leq 1. \quad (26)$$

The reassociation and exact common-ambient projector identity are `MPSTensor.physicalReassocES`, `MPSTensor.physicalReassocES_comp_groundSpaceMapES`, `MPSTensor.reassocTailBoundaryMapES`, `MPSTensor.reassocTailBoundaryMapES_comp_tailVirtualMapES`, and `MPSTensor.whole_increment_injectiveRangeProjector_residual_eq`. They place the tail window of length $L + Q$, the left window of length $K + L$ with Q spectators, and the full window of length $K + L + Q$ in one Hilbert space.

The exact mixed-Gram and centered formulas are `MPSTensor.inner_tailBoundaryMapES_adjoint_leftBoundaryMapES_q`, whose named-reassociation bridge is `MPSTensor.inner_reassocTailBoundaryMapES_adjoint_leftBoundaryMapES`, and `MPSTensor.inner_reassocTailBoundaryMapES_adjoint_leftBoundaryMapES_centered`. The direct-sum factors `MPSTensor.wholeIncrementLeftOverlapFactorES` and `MPSTensor.wholeIncrementRightOverlapFactorES`, with their exact fiber formulas and norm bounds, factor `MPSTensor.wholeIncrementCenteredResidualES` through $\mathcal{K}_L - \mathcal{K}_\infty$. This is recorded by `MPSTensor.wholeIncrementCenteredResidualES_eq_overlapFactors`, `MPSTensor.wholeIncrementCenteredResidualES_norm_le`, `MPSTensor.wholeIncrementCenteredResidualES_norm_le_of_leftCanonical`, and `MPSTensor.IsPrimitiveMPS.wholeIncrementCenteredResidualES_norm_le`.

The whole-increment finite-Gram telescope and its three correction maps are now also explicit. The supporting Gram, inverse-Gram, cancellation, and limiting-product identities are `MPSTensor.reassocTailBoundaryMapES_adjoint_comp_self_eq_fiberwise_groundSpaceGram`, `MPSTensor.inverseGram_reassocTailBoundaryMapES_eq_fiberwise_inverseGram`, `MPSTensor.wholeIncrementLeftFiniteGramCancellation_eq_groundSpaceMapES_adjoint`, `MPSTensor.wholeIncrementTailFiniteGramCancellation_eq_groundSpaceMapES_adjoint`, and `MPSTensor.wholeIncrement_limitingOverlapProduct_eq`. The declarations `MPSTensor.wholeIncrementTailFiniteGramCorrectionES`, `MPSTensor.wholeIncrementFullInverseGramCorrectionES`, and `MPSTensor.wholeIncrementLeftFiniteGramCorrectionES` package the corrections. The exact virtual and physical telescopes are `MPSTensor.wholeIncrement_virtualResidual_eq_centered_sub_gramErrors` and `MPSTensor.wholeIncrement_injectiveRangeProjector_residual_eq_centered_sub_corrections`; the latter uses `MPSTensor.wholeIncrementCenteredProjectorResidualES`. The reassociated tail pseudoinverse bound is `MPSTensor.norm_reassocTailBoundaryMapES_comp_inv`

erseGram_le_sqrt. The centered physical sandwich is controlled by `MPSTensor.wholeIncrementCenteredProjectorResidualES_norm_le`. The three correction bounds are `MPSTensor.wholeIncrementTailFiniteGramCorrectionES_norm_le`, `MPSTensor.wholeIncrementFullInverseGramCorrectionES_norm_le`, and `MPSTensor.wholeIncrementLeftFiniteGramCorrectionES_norm_le`.

The assembled quantitative estimate is `MPSTensor.IsPrimitiveMPS.wholeIncrement_groundProjection_defect_le_geometric`. For a primitive tensor with positive-definite fixed point and nonzero bond dimension, it gives constants $C, c, r > 0$, with $r < 1$, such that for every prefix length K , overlap length $L \geq 1$, and suffix length Q satisfying $cr^L < 1$. Write

$$P_{K;L,Q}^{\text{tail}} := P_{\text{ran}(\tilde{\Gamma}_{K;L,Q}^{\text{tail}})}, \quad P_{K+L,Q}^{\text{left}} := P_{\text{ran}(\Gamma_{K+L,Q}^{\text{left}})}. \quad (27)$$

These are the orthogonal projectors onto the ranges of the reassocated tail and left boundary maps. Then

$$\left\| P_{K;L,Q}^{\text{tail}} P_{K+L,Q}^{\text{left}} - P_{\mathcal{G}_{K+L+Q}^{\text{ES}}(A)} \right\| \leq \frac{Cr^L}{1 - cr^L}. \quad (28)$$

The constants are independent of both K and Q . This estimate is reconstructed from the exact inverse-Gram telescope and is not attributed to FNW. It supplies the quantitative C3'-type estimate needed here, but it does not claim the optimized numerator in Nachtergaele's equation (6.1).

The theorem `MPSTensor.IsPrimitiveMPS.exists_uniform_wholeIncrement_defect_le_seven_sixteenths` then chooses one integer $p > 0$, measured in sites of the input tensor, such that A is p -block injective and the preceding defect with $L = Q = p$ is at most $7/16$ for every prefix length K . Its specialization `MPSTensor.IsPrimitiveMPS.exists_threeBlock_wholeIncrement_defect_le_seven_sixteenths` takes $K = p$ and gives the same bound for three consecutive blocks of p original sites. Thus prefix-uniform quantitative smallness and the three-block coefficient $7/16$ are resolved in original-site units.

- The mixed Gram of the C3 tail map at $(K, l+1)$ and left map at $(K+l, 1)$ is computed exactly by `MPSTensor.inner_tailBoundaryMapES_adjoin_t_leftBoundaryMapES`. Write $\Gamma_{K,l+1}^{\text{tail}}$ and $\Gamma_{K+l,1}^{\text{left}}$ for these two Hilbert-space maps. For boundary-family vectors x and y , write X_u and Z_j for their matrix fibers. Then

$$\langle x, (\Gamma_{K,l+1}^{\text{tail}})^\dagger \Gamma_{K+l,1}^{\text{left}} y \rangle = \sum_{u \in \{0, \dots, d-1\}^K} \sum_{j \in \{0, \dots, d-1\}^1} \langle U(A^j X_u), \mathcal{K}_l U(Z_j A^u) \rangle. \quad (29)$$

This formula is an exact algebraic identity reconstructed from the boundary-map definitions, word factorization, and trace cyclicity. It is

not stated in this form by FNW or Nachtergaele and is not attributed to them. The sourced target remains the projector defect in Nachtergaele's equation (2.4), while the rational estimate is quoted after that equation and in equation (6.1). No mixed-Gram norm estimate follows from this identity alone.

- For injective maps with common factorizations $TJ_1 = C = LJ_2$, `ContinuousLinearMap.injectiveRangeProjector_comp_sub_of_factorizations` factors the exact projector residual through

$$T^\dagger L - (T^\dagger T)J_1 S_C J_2^\dagger (L^\dagger L). \quad (30)$$

The concrete specialization `MPSTensor.c3_injectiveRangeProjector_residual_eq` uses the tail map at $(K, L_0 + 1)$, the left map at $(K + L_0, 1)$, the common map $\Gamma_{K+L_0+1}^{\text{ES}}$, and the corresponding virtual maps. Its three inverse-Gram range projectors have the open-chain tail, open-chain left, and full ground-space ranges by the exact range theorems above. All three injectivity assumptions are explicit. This is an algebraic identity only; a source-specific estimate of the mixed-Gram residual is still required.

- For an injective finite-dimensional map T , with $S_T = (T^\dagger T)^{-1}$, the two pseudoinverse factors in its exact range projector satisfy

$$\|TS_T\| = \sqrt{\|S_T\|}, \quad \|S_T T^\dagger\| = \sqrt{\|S_T\|}. \quad (31)$$

These identities are `ContinuousLinearMap.norm_T_comp_inverseGram_eq_sqrt` and `ContinuousLinearMap.norm_inverseGram_comp_adjoint_eq_sqrt`.

- If A is a primitive MPS tensor with nonzero virtual dimension, write J_K^{tail} for its length- K tail virtual map. Then `MPSTensor.IsPrimitiveMPS.tailVirtualMapES_norm_uniform` gives a constant $C > 0$, independent of K , such that

$$\|J_K^{\text{tail}}\| \leq C \quad (32)$$

for every $K \in \mathbb{N}$. No positive-definiteness assumption on ρ is used. The proof preserves the fourth-power intermediate normalization, first establishing a uniform bound for $\|J_K^{\text{tail}}\|^4$, and then obtains the norm bound. It reconstructs the required norm infrastructure from the exact tail adjoint-self estimate, the fixed-point/complementary-power decomposition, the geometric power bound, and the dimension conversion from the L^2 operator norm to entrywise mass. This is not an FNW or Nachtergaele contraction estimate.

- Let A be an MPS tensor, let Ω_K be the set of length- K physical words, and let $\mathcal{B}_K = \bigoplus_{u \in \Omega_K} M_D(\mathbb{C})$. Write $J_1^{\text{left}}(X) = (A^j X)_{j \in \Omega_1}$ for the one-site left virtual map and $J_K^{\text{tail}}(X) = (XA^u)_{u \in \Omega_K}$ for the length- K tail

virtual map. The two common direct-sum overlap factors `MPSTensor.c3LeftOverlapFactorES` and `MPSTensor.c3RightOverlapFactorES` have the exact fiber formulas

$$(F_K X)_{u,j} = A^j X_u, \quad (G_K Z)_{u,j} = Z_j A^u, \quad (33)$$

recorded by `MPSTensor.boundaryFamilyEquiv_c3LeftOverlapFactorES_apply` and `MPSTensor.boundaryFamilyEquiv_c3RightOverlapFactorES_apply`. Their definitions and fiber formulas are valid for arbitrary finite, possibly empty, configuration sets and require no nonzero-cardinality assumption. Moreover, `MPSTensor.c3LeftOverlapFactorES_norm_le` and `MPSTensor.c3RightOverlapFactorES_norm_le` give

$$\|F_K\| \leq \|J_1^{\text{left}}\|, \quad \|G_K\| \leq \|J_K^{\text{tail}}\|, \quad (34)$$

with no loss depending on the number of prefix configurations.

When $D \geq 1$ and ρ is explicitly positive definite, set $\mathcal{K}_\infty = \mathcal{R}(P_\rho)$. The centered residual `MPSTensor.c3CenteredVirtualResidualES` factors exactly as

$$R_{K,l} = F_K^\dagger (\mathcal{K}_l - \mathcal{K}_\infty)^{\oplus(\Omega_K \times \Omega_1)} G_K \quad (35)$$

by `MPSTensor.c3_centeredVirtualResidualES_eq_overlapFactors`, and `MPSTensor.c3_centeredVirtualResidualES_norm_le` gives

$$\|R_{K,l}\| \leq \|J_1^{\text{left}}\| \|\mathcal{K}_l - \mathcal{K}_\infty\| \|J_K^{\text{tail}}\|. \quad (36)$$

Thus, if A is primitive with fixed-point witness ρ , if ρ is explicitly positive definite, and if $l \geq 1$, then `MPSTensor.IsPrimitiveMPS.c3_centeredVirtualResidualES_norm_sq_le_geometric` supplies constants $C > 0$ and $0 < r < 1$, independent of K , for which

$$\|R_{K,l}\|^2 \leq D^3 (Cr^l)^2 \quad (37)$$

for every $K \in \mathbb{N}$. These formulas are reconstructed algebra, not formulas quoted from FNW or Nachtergaele. They control only the centered term and prove neither the full projector defect nor the rational coefficient in equation (6.1).

- For each length n , let $\mathcal{K}_n$ be the ground-space Gram, and whenever it is invertible set $S_n = \mathcal{K}_n^{-1}$. Let $\mathcal{K}_\infty = \mathcal{R}(P_\rho)$, with ρ positive definite, and set $S_\infty = \mathcal{K}_\infty^{-1}$. The exact C3 virtual residual is telescoped by `MPSTensor.c3_virtualResidual_eq_centered_sub_gramErrors` into its limiting-center residual and three corrections containing, in order,

$$\mathcal{K}_{l+1} - \mathcal{K}_\infty, \quad (38)$$

$$S_{K+l+1} - S_\infty, \quad (39)$$

$$\mathcal{K}_{K+l} - \mathcal{K}_\infty. \quad (40)$$

The theorem `MPSTensor.c3_injectiveRangeProjector_residual_eq_centered_sub_gramErrors` places this identity between the two local pseudoinverse factors in the exact tail-after-left physical projector residual.

Under the explicit hypotheses that the virtual dimension is nonzero and ρ is positive definite, `MPSTensor.inner_limitingMixedGramIntertwining` proves the limiting intertwining identity. Under the same hypotheses, `MPSTensor.inner_c3_centeredVirtualResidual_eq_gramError` identifies the centered residual pairing with the length- l Gram error. The cancellation uses the nonidentity limiting metric, its inverse, and the virtual-map adjoints; no transfer fixed-point equation is needed.

- The three physical correction sandwiches are `MPSTensor.c3TailFiniteGramCorrectionES`, `MPSTensor.c3FullInverseGramCorrectionES`, and `MPSTensor.c3LeftFiniteGramCorrectionES`. Write

$$T = \Gamma_{K,l+1}^{\text{tail}}, \quad (41)$$

$$L = \Gamma_{K+l,1}^{\text{left}}, \quad (42)$$

$$C = \Gamma_{K+l+1}^{\text{ES}}. \quad (43)$$

The exact factorizations `MPSTensor.c3TailFiniteGramCorrectionES_eq_factored`, `MPSTensor.c3FullInverseGramCorrectionES_eq_factored`, and `MPSTensor.c3LeftFiniteGramCorrectionES_eq_factored` give

$$Q_1 = TS_{l+1}^{\oplus \Omega_K} (\mathcal{K}_{l+1} - \mathcal{K}_\infty)^{\oplus \Omega_K} J_K^{\text{tail}} S_{K+l+1} C^\dagger, \quad (44)$$

$$Q_2 = TS_{l+1}^{\oplus \Omega_K} \mathcal{K}_\infty^{\oplus \Omega_K} J_K^{\text{tail}} S_\infty (\mathcal{K}_\infty - \mathcal{K}_{K+l+1}) S_{K+l+1} C^\dagger, \quad (45)$$

$$Q_3 = TS_{l+1}^{\oplus \Omega_K} \mathcal{K}_\infty^{\oplus \Omega_K} J_K^{\text{tail}} S_\infty (J_1^{\text{left}})^\dagger (\mathcal{K}_{K+l} - \mathcal{K}_\infty)^{\oplus \Omega_1} S_{K+l}^{\oplus \Omega_1} L^\dagger. \quad (46)$$

The middle equality uses the right-oriented identity `MPSTensor.inverseGram_sub_limitingInverse_eq_resolvent`,

$$S_{K+l+1} - S_\infty = S_\infty (\mathcal{K}_\infty - \mathcal{K}_{K+l+1}) S_{K+l+1}. \quad (47)$$

The fiberwise subtraction identity is `MPSTensor.boundaryFiberwiseMap_sub`, and invertibility of a finite Gram follows from injectivity by `MPSTensor.groundSpaceGram_isUnit_of_groundSpaceMapES_injective`.

The spectator pseudoinverse estimates `MPSTensor.norm_tailBoundaryMapES_comp_inverseGram_le_sqrt` and `MPSTensor.norm_inverseGram_comp_leftBoundaryMapES_adjoint_le_sqrt`, together with the uniform fiberwise norm bound, yield `MPSTensor.c3TailFiniteGramCorrectionES_norm_le`, `MPSTensor.c3FullInverseGramCorrectionES_norm_le`, and `MPSTensor.c3LeftFiniteGramCorrectionES_norm_le`:

$$\|Q_1\| \leq \sqrt{\|S_{l+1}\|} \|\mathcal{K}_{l+1} - \mathcal{K}_\infty\| \|J_K^{\text{tail}}\| \sqrt{\|S_{K+l+1}\|}, \quad (48)$$

$$\|Q_2\| \leq \sqrt{\|S_{l+1}\|} \|\mathcal{K}_\infty\| \|J_K^{\text{tail}}\| \|S_\infty\| \|\mathcal{K}_{K+l+1} - \mathcal{K}_\infty\| \sqrt{\|S_{K+l+1}\|}, \quad (49)$$

$$\|Q_3\| \leq \sqrt{\|S_{l+1}\|} \|\mathcal{K}_\infty\| \|J_K^{\text{tail}}\| \|S_\infty\| \|J_1^{\text{left}}\| \|\mathcal{K}_{K+l} - \mathcal{K}_\infty\| \sqrt{\|S_{K+l}\|}. \quad (50)$$

No spectator-cardinality factor appears. These factorizations and bounds are exact consequences of the inverse-Gram formulas for the physical range projectors. FNW and Nachtergaele state overlapping-projector estimates, while Cirac–Perez-Garcia–Schuch–Verstraete state the martingale anticommutator criterion and its principal-angle interpretation. None of those sources gives the three displayed correction formulas. Their uniform combination is given below; comparison with the optimized coefficient in Nachtergaele’s equation (6.1) remains open.

Fix $K, L_0 \in \mathbb{N}$ and let $\mathcal{H}_{K+L_0+1}$ be the $(K + L_0 + 1)$ -site Euclidean space $\ell^2(\{0, \dots, d-1\}^{K+L_0+1})$. Specializing the tail transport theorem to $(K, L_0 + 1)$ identifies its range with the open-chain tail subspace $\mathcal{V}_{K, L_0+1}(A)$, while specializing the left transport theorem to $(K + L_0, 1)$ (the only length at which the left boundary map matches the Nachtergaele C3 left-window convention) identifies its range with the open-chain left subspace $\mathcal{U}_{K+L_0}(A)$, both in the common ambient $\mathcal{H}_{K+L_0+1}$. The relevant projector defect is

$$\|P_{K, L_0+1}^{\text{tail}}(A) P_{K+L_0}^{\text{left}}(A) - P_{\mathcal{G}_{K+L_0+1}^{\text{ES}}(A)}\| \quad (51)$$

in the common ambient $\mathcal{H}_{K+L_0+1}$ (Nachtergaele C3, arXiv:cond-mat/9410110, eq. (2.4)). Here $\mathcal{U}_{K+L_0}(A)$ and $\mathcal{V}_{K, L_0+1}(A)$ are the left and tail ground subspaces of Definition 4.2 (`MPSTensor.openChainLeftGroundSpaceES` and `MPSTensor.openChainTailGroundSpaceES`), their intersection equals $\mathcal{G}_{K+L_0+1}^{\text{ES}}(A)$ under block injectivity with $D \geq 1$ and $L_0 > 0$ (`MPSTensor.openChainLeft_inf_tailGroundSpaceES`). The symbols $P_{K+L_0}^{\text{left}}(A)$ and $P_{K, L_0+1}^{\text{tail}}(A)$ denote the orthogonal projectors onto $\mathcal{U}_{K+L_0}(A)$ and $\mathcal{V}_{K, L_0+1}(A)$, respectively, as in Definition 4.2, while $P_{\mathcal{G}_{K+L_0+1}^{\text{ES}}(A)}$ denotes the orthogonal projector onto the full ground subspace.

The two ranges are transported into this common ambient by `MPSTensor.tailBoundaryMap_range_map_symm_eq_openChainTailGroundSpaceES` (specialized to $(K, L_0 + 1)$) and `MPSTensor.leftBoundaryMap_range_map_symm_eq_openChainLeftGroundSpaceES` (specialized to $(K + L_0, 1)$).

The present formal development now controls the deviation from the correct limiting Gram operator:

$$\|\mathcal{K}_n - \mathcal{R}(P_\rho)\|^2 \leq D^3 (Cr^n)^2. \quad (52)$$

It does not control $\|1 - \mathcal{K}_n\|$ in these unweighted coordinates, because $\mathcal{R}(P_\rho)$ is generally not the identity. Under the explicit additional hypothesis that ρ is positive definite, the limiting Gram is invertible. The reconstructed pointwise estimates `MPSTensor.IsPrimitiveMPS.groundSpaceGram_geometric_inverse_bounds` and `MPSTensor.IsPrimitiveMPS.groundSpaceMapES_geometric_inverse_bounds` perturb around the nonidentity operator $\mathcal{R}(P_\rho)$ and give the inverse and inverse-Gram bounds with denominator $1 - cr^n$ whenever $cr^n < 1$. The theorem `MPSTensor.geometric_smallness_at_c3_lengths` propagates the base condition $cr^l < 1$ to the three correction lengths $l+1$, $K+l$, and $K+l+1$. The normalization still descends from the squared estimate with

factor D^3 , yielding the norm constant $\sqrt{D^3}C$ without additional dimension loss; it is not replaced by an identity-centered source constant. Furthermore, `MPSTensor.IsPrimitiveMPS.groundSpaceGram_ringInverse_tendsto` proves

$$\mathcal{K}_n^{-1} \longrightarrow \mathcal{R}(P_\rho)^{-1}. \quad (53)$$

The uniform combination is now complete. Put

$$\mathcal{K}_\infty = \mathcal{R}(P_\rho), \quad \mathcal{I}_\infty = \mathcal{K}_\infty^{-1}. \quad (54)$$

Choose $g, c, r > 0$ from `MPSTensor.IsPrimitiveMPS.groundSpaceMapES_geometric_inverseGram_bounds`, so that $c = \|\mathcal{I}_\infty\|g$ and $0 < r < 1$, and choose $T > 0$ from `MPSTensor.IsPrimitiveMPS.tailVirtualMapES_norm_uniform`. Set

$$\begin{aligned} J &= \|J_1^{\text{left}}\|, & M &= \|\mathcal{K}_\infty\|\|\mathcal{I}_\infty\|, \\ C_* &= 4\|\mathcal{I}_\infty\|Tg(J+1)(M+1). \end{aligned} \quad (55)$$

Then `MPSTensor.IsPrimitiveMPS.openChain_groundProjection_defect_leave_geometric` proves, for every K and every $l \geq 1$ with $cr^l < 1$,

$$\left\| P_{K,l+1}^{\text{tail}} P_{K+l}^{\text{left}} - P_{\mathcal{G}_{K+l+1}^{\text{ES}}(A)} \right\| \leq \frac{C_* r^l}{1 - cr^l}. \quad (57)$$

The proof uses the exact physical four-term assembly `MPSTensor.c3_injectiveRangeProjector_residual_eq_centered_sub_corrections` and the centered-sandwich estimate `MPSTensor.c3CenteredProjectorResidualES_norm_le`. Its virtual input is bounded by `MPSTensor.c3_centeredVirtualResidualES_norm_le`; the other inputs are the uniform tail bound, the pointwise inverse-Gram bounds and their three-length smallness theorem, and the three correction estimates `MPSTensor.c3TailFiniteGramCorrectionES_norm_le`, `MPSTensor.c3FullInverseGramCorrectionES_norm_le`, and `MPSTensor.c3LeftFiniteGramCorrectionES_norm_le`. The exact range identities for the tail, left, and full boundary maps give the displayed physical projectors. No physical projector is identified with a transfer fixed-point projector.

This coefficient is a reconstructed TNLean substitute sufficient for Nachtergaele's condition C3. Indeed, `MPSTensor.IsPrimitiveMPS.exists_openChain_groundProjection_defect_lt_c3_threshold` uses

$$\sqrt{n+1}r^n \longrightarrow 0 \quad (58)$$

to choose $l > 1$ above a positive block-injectivity length such that $cr^l < 1$ and

$$\frac{C_* r^l}{1 - cr^l} < \frac{1}{\sqrt{l+1}}. \quad (59)$$

The same l works for every spectator length K , and `MPSTensor.IsPrimitiveMPS.exists_re_inner_openChain_anticommutator_ge_c3_threshold` then

gives the open-chain anticommutator inequality by the two-projection argument. The length l is measured in sites of the input tensor. If that tensor is already a blocked representative, one such site is one blocked-chain filtration step.

The coefficient is not the optimized FNW expression

$$c\lambda^l \frac{1 + c\lambda^l}{1 - c\lambda^l}, \quad c\lambda^l < 1, \quad (60)$$

quoted as equation `boundAm` in [Nac96, Section 6]. Recovering that numerator and its source constants still requires the weighted Gram estimates of FNW, Lemmas 5.2, 5.3, and 6.2. No identification of the reconstructed constant with a source dimension-dependent value is made; the available unweighted Gram conversion retains its proved $D^{3/2}$ norm factor inside g .

7 Verdict

Verdict: scope restriction. The exact whole-increment algebra, the reconstructed prefix-uniform geometric C3'-type estimate, the three-block coefficient $7/16$, and the resulting cyclic range-two gap $1/8$ for a suitable blocked tensor are proved. All lengths before blocking are measured in sites of the input tensor, while the final chain length counts blocked sites. These bounds come from the formal inverse-Gram reconstruction and are not attributed to FNW. The optimized FNW numerator and its source constants remain open.

The blocked-to-original comparison is now proved on periodic lengths divisible by the chosen block length. Under the canonical configuration isometry, every blocked range-two term is exactly the original range- $2p$ term beginning at a block boundary. The conjugated blocked Hamiltonian is therefore the sparse sum over aligned starts, and the omitted original local projections give a positive quadratic-form remainder. Block injectivity identifies the blocked and original kernels with the same periodic MPS line under this isometry. Hence `MPSTensor.IsPrimitiveMPS.exists_parentHamiltonianES_gap_eighth_mul` retains the gap constant $1/8$ for $H_{Np,2p}(A)$, $N \geq 4$.

This finite cyclic sparse-sum argument is a TNLean reconstruction. Nachtergaele, [Nac96, equations (3.12)–(3.16)], formulates compatible open intervals and the proof of Theorem 2.1(ii) uses the rescaled sequence $\tilde{\Lambda}_n = \Lambda_{ln}$; the source does not state the finite periodic operator comparison above. The precise remaining leaf for the unrestricted normal-MPS theorem is the remainder-length case: chains $Np + r$ with $0 < r < p$. It requires a boundary/remainder comparison preserving the common ground-space complement and a uniform positive constant. No all-length claim is made here.

References

- [CPGSV21] J. Ignacio Cirac, David Pérez-García, Norbert Schuch, and Frank Verstraete. Matrix product states and projected entangled pair

- states: Concepts, symmetries, theorems. *Reviews of Modern Physics*, 93(4):045003, 2021.
- [FNW92] Mark Fannes, Bruno Nachtergaele, and Reinhard F. Werner. Finitely correlated states on quantum spin chains. *Communications in Mathematical Physics*, 144(3):443–490, 1992.
- [KL18] Michael J. Kastoryano and Angelo Lucia. Divide and conquer method for proving gaps of frustration free spin systems. *Journal of Statistical Mechanics: Theory and Experiment*, 2018(3):033105, 2018.
- [Nac96] Bruno Nachtergaele. The spectral gap for some spin chains with discrete symmetry breaking. *Communications in Mathematical Physics*, 175(3):565–606, 1996.