

Inhomogeneous Matrix Product States: Block Length

2026-10-02

Verdict: scope-restriction (open).

1 What the source states

The paragraph “Inhomogeneous short-range correlated MPS” of [MSWC24] considers a sequence of matrix product states $|\phi_N\rangle$ “with bond dimension at most D ” that are not translation invariant. The sequence has finite correlation length if, after blocking $q = O(\log N)$ sites, the state $|\phi_{\text{pos}}\rangle$ of the positive parts P_i of the polar decompositions of the blocked tensors (eq. (18)) is approximated by a product of nearest-neighbour pairs $|\Omega\rangle = \bigotimes_{i=1}^{N/q} |\omega^i\rangle_{R_i L_{i+1}}$ with $\varepsilon(\Omega, \phi_{\text{pos}}) \rightarrow 0$. The paragraph then asserts that preparing $|\Omega\rangle$ and applying the isometries of the polar decompositions gives depth $O(\log(N/\varepsilon))$ with error $\varepsilon(\Omega, \phi_{\text{pos}})$.

2 The restriction adopted

The formal statements¹ work on a ring of N sites cut into blocks of lengths ℓ_k , with bond dimensions $D_0, \dots, D_{N-1} \leq D$ varying along the ring as in the source, and blocked tensors that need not be injective. They assume that every block has at least $3D$ sites, which the circuit needs to hold the two bond indices of a block in registers of D sites each and to route them together. The first statement, with error exactly $\varepsilon(\Omega, \phi_{\text{pos}})$, also assumes $D^2 \leq d^{3D}$, which holds for every D when $d \geq 2$. For blocks of $q \geq 3D \geq 1$ sites it follows from the source’s requirement $d^q \geq D^2$ recalled below, since $d^q \geq 1$ forces $d \geq 1$ and for $d = 1$ both conditions read $D \leq 1$; it is therefore no further restriction. The second statement, with error at most δ under the approximation hypothesis, holds for every d .

The source has no bound $3D$, but it is not silent on the block length. Its polar decompositions need “ q large enough that $d^q \geq D^2$ ” (footnote to the paragraph “Approximation through the fixed-point state”), and its circuit holds the legs

¹Formal declarations: `MPSPreparation.exists_isPreparedInDepth_inhomogeneous_isometricExtension` and `MPSPreparation.exists_isPreparedInDepth_of_isPairApproximable` in `TNLean/MPS/Preparation/PartialIsometryPreparation`.

R_i and L_{i+1} , of dimension D each, on the sites of the blocks, with the central sites in $|0\rangle$ (eq. (10) and the paragraph “Preparing the approximate state”). A block must therefore hold two registers of dimension at least D , that is, at least $2\lceil\log_d D\rceil$ sites. The bound $3D$ is a coarser sufficient condition: each register uses D sites, as $D \leq d^D$, and a third stretch of D sites leaves room for routing the left register next to the right one. For $q = O(\log N)$ growing with N both bounds hold for large N .

The finite-correlation assumption is stated for one ring: some unit pairs ω^k on $\mathbb{C}^{D_j} \otimes \mathbb{C}^{D_j}$, for the bond j joining block k to block $k+1$, have $\varepsilon(\Omega, \phi_{\text{pos}}) \leq \delta$. The formal statements prove depth at most CL for blocks of lengths at most L , with C depending only on d and D . The source’s total depth $O(\log(N/\varepsilon))$ is the choice $L = O(\log(N/\varepsilon))$, which is not made here; the source’s definition fixes blocks of length $q = O(\log N)$.

3 Partial isometries

Without injectivity the isometric factor V_k of the polar decomposition $B_k = V_k P_k$ is a partial isometry, $V_k^\dagger V_k = \Pi_k$ with Π_k the projector onto the range of P_k , as in the Supplemental Material, “Proof of Lemma 1 and extension to non-normal tensors” of [MSWC24]; the footnote to the paragraph “The sequential-RG circuit” states that the derivation of the circuit “remains valid also for non-injective tensors B ”, with P^{-1} read as the pseudo-inverse. A unitary circuit cannot prepare $(\bigotimes_k V_k)|\Omega\rangle$ itself, whose norm is that of $(\bigotimes_k \Pi_k)|\Omega\rangle$ and may be smaller than 1. The formal statements prepare instead $(\bigotimes_k W_k)|\Omega\rangle$ for isometries W_k with $W_k \Pi_k = V_k$, isometric extensions of V_k . They agree with V_k on the range of Π_k , and $W_k P_k = V_k P_k = B_k$, so $|\phi_N\rangle = (\bigotimes_k W_k)|\phi_{\text{pos}}\rangle$ and the overlap of $(\bigotimes_k W_k)|\Omega\rangle$ with the normalized $|\phi_N\rangle$ is that of $|\Omega\rangle$ with the normalized $|\phi_{\text{pos}}\rangle$. The error is therefore exactly $\varepsilon(\Omega, \phi_{\text{pos}})$, the error the source states. The block unitary implementing W_k applies, after routing the two input registers together, a unitary $T_k^\dagger$ on them whose conjugate T_k has its first columns spanning the range of Π_k , and then the staircase of the sequential factorization of $V_k T_k$ on those columns. This is a choice of the prepared state, not a restriction: the source names the isometry V_k and its circuit, and W_k is the isometry that this circuit implements.

When every blocked tensor is injective in the sense of the next section, Π_k is the coordinate projector onto the rectangle of the bonds of block k , $|\Omega\rangle$ lies in the range of $\bigotimes_k \Pi_k$, and $(\bigotimes_k W_k)|\Omega\rangle = (\bigotimes_k V_k)|\Omega\rangle$; the corresponding statements² prepare $(\bigotimes_k V_k)|\Omega\rangle$ itself and keep the injectivity hypothesis, together with the block length $3D$, for that purpose.

²Formal declarations: `MPSPreparation.exists_isPreparedInDepth_inhomogeneous` in `TNLean/MPS/Preparation/InhomogeneousPreparation`, and `MPSPreparation.exists_isPreparedInDepth_chainBlockIsometryState_of_isInjectiveOn` and `MPSPreparation.exists_isPreparedInDepth_chainBlockIsometryState` in `TNLean/MPS/Preparation/DepthUpperBound`.

4 Varying bond dimensions by zero padding

The rectangular matrices A_k^i of size $D_k \times D_{k+1}$ are padded with zeros to $D \times D$ matrices³. The padded chain has the same state, which is proved. For a block of at least one site, its blocked tensor B_k vanishes outside the upper-left corner $[0, D_{o_k}) \times [0, D_{o_{k+1}})$, which is also proved, and its entries on that corner are those of the rectangular blocked tensor. This identification is proved by `VaryingBondChain.zeroPad_rectangularBlockTensor` in `TNLean/MPS/Preparation/RectangularBlocks`, together with the ordered composition formula for the actual rectangular site transfer maps. The injectivity hypothesis of the statement for injective blocked tensors is stated directly for B_k , as spanning every matrix unit in the corner. For a block of length zero the two readings differ, since B_k is then the identity matrix of size D ; the statement using the hypothesis assumes blocks of at least one site. Under the hypothesis, the support projector of the polar decomposition $B_k = V_k P_k$ is the coordinate projector Π_k onto the corner, so V_k is an isometry on it and vanishes on the other inputs. The pairs ω^k are padded in the same way, and $|\Omega\rangle$ is supported on the rectangles. The objects $|\phi_{\text{pos}}\rangle$, $(\bigotimes_k V_k)|\Omega\rangle$ and $\varepsilon(\Omega, \phi_{\text{pos}})$ of the formal statements are those of the padded chain; they agree with the source's after the zero coordinates are dropped. This is the reading adopted, not a restriction: the theorems are stated for every chain with bond dimensions at most D .

5 Nonzero state of the positive parts

The error $\varepsilon(\Omega, \phi_{\text{pos}})$ compares $|\Omega\rangle$ with the normalized state $|\phi_{\text{pos}}\rangle/\|\phi_{\text{pos}}\|$, which presupposes $|\phi_{\text{pos}}\rangle \neq 0$. The formal approximation condition⁴ states $|\phi_{\text{pos}}\rangle \neq 0$ explicitly; without it the normalized vector would be read as 0 and the condition would hold for every $\delta \geq 1$. This is a local fix of the intended reading, not a restriction.

6 Elimination

1. *Varying bond dimensions.* Done by zero padding, as described above; the earlier restriction to a common bond dimension D is removed.
2. *Injectivity.* Done, as described in the section on partial isometries: the block unitary implements an isometric extension W_k of the partial isometry V_k , and the error is exactly $\varepsilon(\Omega, \phi_{\text{pos}})$ without injectivity.
3. *Block length.* The bound $3D \leq \ell_k$ is needed only by the circuit, not by the state. Registers of $\lceil \log_d D \rceil$ sites in place of D sites would lower it to a multiple of $\log_d D$, the order of the source's own requirement $d^q \geq D^2$;

³Formal declarations: `VaryingBondChain.zeroPad` and `VaryingBondChain.coeff_zeroPad` in `TNLean/MPS/Chain/VaryingBondChain`.

⁴Formal declaration: `VaryingBondChain.IsPairApproximable`.

this is not done. The upper bound $q = O(\log N)$ alone allows bounded block lengths, so the sequence-level statement reads the source's blocking as $q_N \geq c \log N$ for some $c > 0$, a block length growing with N as in the translation-invariant case. Then $q_N \geq 3D$ once $N \geq N_0$, with N_0 depending only on D and c ; each of the finitely many rings with $N < N_0$ is prepared exactly by a circuit of depth bounded in terms of N_0 and d , and the largest of these depths is absorbed into the constant of the depth bound. This argument is formalized in `MPSPreparation.exists_isPreparedInDepth_inhomogeneous_of_log_lower` in `TNLean/MPS/Preparation/InhomogeneousSequence`, including the one-site ring with depth zero. The positive lower-growth hypothesis remains explicit; it is not inferred from the upper bound.

4. *Sequence-level statement.* With unit pairs satisfying $\varepsilon(\Omega, \phi_{\text{pos}}) \leq \delta_N$ and $\delta_N \rightarrow 0$, together with the preceding lower-growth hypothesis, `MPSPreparation.exists_isPreparedInDepth_inhomogeneous_sequence` gives prepared unit states on every positive ring length with depth at most CL_N and error tending to zero. Thus $L_N = O(\log N)$ gives logarithmic depth. This removes the restriction to one ring under that hypothesis, but it does not give the depth $O(\log(N/\varepsilon))$: the source supplies no rate for δ_N , and the choice of a block length L with error at most ε and $L = O(\log(N/\varepsilon))$ needs one. In the translation-invariant case the rate is proved: the error after blocking q sites of a ring of N sites is $O((N/q)e^{-\gamma q/\xi})$ for every $\gamma < 1/2$, for one fixed normal tensor, with constants depending on that tensor ([MSWC24], Lemma 1 and eq. (17))⁵. For the inhomogeneous sequence the tensors change with N , and the same choice of L requires a bound $\delta(N, q) \leq CN^k e^{-cq}$ with $C, c > 0$ and k independent of N . The source states no such uniform bound and does not derive one from $\delta_N \rightarrow 0$. A formal statement with this uniform rate as a hypothesis would be a further scope-restricted theorem, not the source's; with $\delta_N \rightarrow 0$ alone the sequence-level depth $O(\log(N/\varepsilon))$ remains open.

A conditional variant is now formalized in `MPSPreparation.exists_isPreparedInDepth_inhomogeneous_le_log_of_uniform_rate` in `TNLean/MPS/Preparation/InhomogeneousUniformRate`. It assumes a rate $KN^k e^{-rq}$, with $K, r > 0$ and $k \geq 0$ independent of N , at every scale $1 \leq q \leq N$ for some partition with $q \leq \ell_j \leq 2q$. Choosing $q = \lceil (\max(k, 1)/r) \log(N/\varepsilon) + b \rceil$ gives the accuracy-dependent depth for every positive ring length. If $q > N$, the ring is prepared exactly in depth $O(N)$ by `MPSPreparation.exists_isPreparedInDepth_normalizedChainState` in `TNLean/MPS/Preparation/InhomogeneousExactPreparation`: the whole ring is one block, and its normalized positive-part state is the input pair. Generic finite-dimensional synthesis is used only for the fixed finite

⁵Formal declaration: `MPSTensor.exists_approximationError_le_mul`, which proves the faster rate $e^{-2\gamma q/\xi}$ for $\gamma < 1$; it gives the homogeneous depth in `docs/paper-gaps/mswc24_depth_upper_bound_divisible_length.tex`.

set $N < 3D$. This proves the stated conditional consequence; it does not derive the uniform rate from the source's definition.

7 A quantitative sufficient condition on actual blocks

A separate structural result assumes a common faithful fixed state σ of trace one and normalized Choi domination $J(F_j) \geq (\eta/D)(\sigma \otimes I)$, $0 \leq \eta \leq 1$, for the trace-preserving transfer maps F_j of actual contiguous blocks of at most s sites. It proves

$$\begin{aligned} \|B^\dagger B - \sigma^\top \otimes I\| &\leq 2D(1 - \eta)^{\lfloor N/s \rfloor}, \\ \|P_B - (\sqrt{\sigma})^\top \otimes I\| &\leq K_\sigma(1 - \eta)^{\lfloor N/s \rfloor}, \end{aligned} \quad (1)$$

for the whole blocked chain, with K_σ independent of its tensors and length. The declarations are `MPSChainTensor.transferMap_blockTensor`, `MPSChainTensor.traceNorm_transferMap_blockTensor_sub_le_of_choi_domination`, and `MPSPreparation.exists_norm_polarPos_chainBlockTensor_sub_le_of_choi_domination` in `TNLean/MPS/Chain/Transfer`, `TNLean/MPS/Preparation/InhomogeneousDoeblin`, and `TNLean/MPS/Preparation/InhomogeneousPositivePartRate`.

This condition is imposed on ordered products over actual intervals, not on powers of the individual site maps. Separate sitewise spectral gaps are insufficient to justify that replacement. The result is restricted to a common bond space and a common faithful fixed state. It gives a uniform Gram and polar estimate under those additional structural hypotheses, not the full source preparation theorem: varying rectangular bond spaces, varying fixed states, and passage from qualitative convergence to a uniform quantitative hypothesis remain separate requirements.

A second sufficient-condition result now starts with actual ordered transfer products on each nonwrapping contiguous interval in the chosen ring cut:

$$\|T_{E_{A_i} \circ \dots \circ E_{A_{i+\ell-1}}} - T_\sigma\| \leq K e^{-r\ell},$$

with $K, r > 0$ independent of the ring length and interval, in the L^2 operator norm on the transfer matrices. The common limiting state is faithful and has trace one. The declarations `MPSPreparation.exists_norm_trace_prod_polarPos_sub_one_le` and `MPSPreparation.exists_one_sub_norm_inner_chainPosState_le` prove the ordered mixed-overlap and normalized pair estimates: transfer error at most δ on every selected block and on the whole ring gives normalized pair error at most $C_\sigma M \delta$. The constants are chosen before the tensors, physical dimension, ring, and partition. The whole-ring transfer controls the actual normalization $c_N^2 = \text{Tr}(T_{B_{\text{whole}}})$; the local Gram estimates are not used as a substitute for that normalization.

Consequently `MPSPreparation.exists_isPreparedInDepth_inhomogeneous_le_log_of_ordered_mixing` in `TNLean/MPS/Preparation/Inhomogeneous`

`eousMixingPreparation` derives pair error at most $C_\sigma K N e^{-r^q}$ and physical depth $O(\log(N/\varepsilon))$. It assumes nonzero periodic targets at every positive ring length, as required to interpret the source's normalized quantum states. Mixing alone guarantees nonvanishing only eventually; it does not silently remove the finite short-ring condition. Exact preparation covers scales longer than the ring. This closes the uniform overlap and normalization step under the stated ordered mixing hypothesis. It does not derive that hypothesis from qualitative finite correlation, and it does not cover varying reference states or rectangular bond spaces.

The direct consumer `MPSPreparation.exists_isPreparedInDepth_in_homogeneous_le_log_of_choi_domination` in `TNLean/MPS/Preparation/InhomogeneousDoeblinPreparation` now connects the actual-site sufficient condition above to physical preparation. With trace-preserving site maps fixing the same faithful σ , and uniform Choi domination $0 < \eta \leq 1$, inverse Gram reshuffling yields ordered-transfer error $K(1-\eta)^\ell$ on each actual interval. Weakening to $\eta/2$ retains the endpoint $\eta = 1$ while giving the finite positive exponent $-\log(1 - \eta/2)$. Strict domination itself proves nonvanishing on every positive ring: `MPSPreparation.le_norm_chainState_sq_of_choi_domination` gives $\|\phi_N\|^2 \geq \eta > 0$ by splitting off the first site and using nonnegativity of the supertrace of the completely positive residual. Thus this concrete corollary needs no additional nonzero-target hypothesis. The general ordered-mixing theorem retains that physical normalization premise. This is a quantitative sufficient condition, not a replacement of ordered mixing by separate sitewise spectral gaps.

8 Short rings and undersized physical blocks

The algebraic pair estimate at a small blocking scale does not assert an isometry from $\mathbb{C}^{D^2}$ into an undersized physical space. The physical compiler chooses $q \geq 3D$ first. For $d \geq 2$, this implies $D^2 \leq d^q$, and it constructs support-preserving isometric extensions $W_j \Pi_j = V_j$ even for rank-deficient Gram matrices. For $d \leq 1$ it prepares the normalized target directly in depth zero. When the selected $q > N$, exact preparation handles the whole ring; generic synthesis is needed only for the fixed finite range $N < 3D$.

For example, AKLT dimensions $d = 3, D = 2$ give $d < D^2$, so no full-input one-site block isometry exists. The general ordered-mixing estimate permits this; circuit synthesis uses longer blocks or the exact short-ring branch. By contrast, strict faithful Choi domination at a microscopic site forces Choi rank D^2 , whereas its Kraus rank is at most d . The direct actual-site Choi corollary therefore does not apply to unblocked AKLT. Such domination may hold after a fixed blocking; it is not supplied by separate sitewise eigenvalue gaps.

The usual AKLT transfer spectrum is $1, -1/3, -1/3, -1/3$, giving periodic squared norm $1 + 3(-1/3)^N$. In particular the one-site periodic vector vanishes. The all-positive-length normalized theorem correctly keeps its explicit nonzero premise. The separate result `MPSPreparation.exists_isPreparedInDepth_`

`inhomogeneous_le_log_eventually_of_ordered_mixing` removes it from an eventual-length conclusion. Whole-ring convergence chooses a uniform N_0 with $\|\phi_N\|^2 \geq 1/2$ for $N \geq N_0$, then proves physical preparation for the original family on that range. It uses `MPSPreparation.exists_isPreparedInDepth_inhomogeneous_le_log_of_uniform_rate_from`; the old all-positive-length rate interface is its $N_0 = 0$ specialization. No finite short-ring replacement is made in this proof. This is an application under uniform quantitative ordered mixing, not a derivation of that assumption from qualitative finite correlation.

9 Varying positive semidefinite references on a square bond space

A separate sufficient-condition result permits density matrices $\sigma_{N,i}$ that vary with the ring and site and may be singular, while keeping the common square bond dimension D . It assumes the actual ordered transfer matrix on every positive nonwrapping interval starting at a is within $Ke^{-r\ell}$ of $X \mapsto \text{Tr}(X)\sigma_{N,a}$, with $K, r > 0$ independent of all intervals and rings. There is no separate transport or stationarity hypothesis: the global estimate is proved directly from these quantitative tensor-product bounds.

The tensor-level comparison `MPSTensor.exists_norm_gram_transferMatrix_sub_le` in `TNLean/MPS/Preparation/PositivePartRate` uses one constant $K_D > 0$ chosen before the physical dimension, tensor, and positive semidefinite reference. It bounds the Gram error by the transfer-matrix error and conversely, both in the L^2 operator norm. No trace-one, stationarity, or faithfulness assumption is needed for this comparison. Positivity is needed only to identify the fixed-point tensor's transfer map with the algebraic trace-reset map. This finite-dimensional norm comparison supplies no mixing rate by itself.

The global square-root Hölder bound replaces a reference-dependent Lipschitz constant by a dimension-dependent factor and halves the exponential rate. Compactness of the set of density matrices makes the mixed-transfer constants uniform, including at singular references. Since $R_\sigma R_\tau = R_\sigma$ for trace-one references, ordered reference products stay bounded. The positive reference tensor of block j contributes the pair with coefficient $(\sqrt{\sigma_j})_{r_{j-1}, l_j}$, so the outgoing pair from block j uses σ_{j+1} . A two-block regression with distinct rank-one references detects this shift.

The resulting normalized pair error is at most $C_D M \sqrt{\delta}$. For exponential transfer error this gives $C_D \sqrt{K} N e^{-rq/2}$ and logarithmic physical preparation of the original family beyond a uniform ring cutoff. Whole-ring transfer convergence derives eventual nonvanishing; no assumption is made on short zero rings. The common faithful-reference result retains its sharper exponent r .

These declarations⁶ remove the common-reference and faithfulness restric-

⁶Formal declarations: `MPSPreparation.exists_norm_trace_prod_polarPos_sub_one_le_varying_reference`, `MPSPreparation.exists_one_sub_norm_inner_chainPosState_le_varying_reference`, and `MPSPreparation.exists_isPreparedInDepth_inhomogeneous_le_log_eventually_of_varying_reference`.

tions from this fixed-square-bond sufficient condition. They still do not derive a uniform rate from the paper’s qualitative finite-correlation definition. Rectangular bonds require the corresponding corner-supported reset maps; zero-padded rectangular site maps are not generally trace preserving on the full square space, so the square-space hypotheses are not silently applied to them.

10 A local window criterion with derived references

A quantitative sufficient condition is now proved for actual square-bond chains whose site transfer maps are trace preserving. Fix an integer $s \geq 1$ and a strength $\eta > 0$ uniformly across the family. Every actual nonwrapping s -window must satisfy the normalized output-first Choi bound $J(F) \geq (\eta/D)(\tau \otimes I)$ for some density matrix τ . The minorizer can depend on the ring and window and can be singular. The actual window at $N = s$ implies $\eta \leq 1$ internally.

The proof constructs a fixed density of the whole-ring channel and propagates it through ordered suffixes, deriving compatible references $\sigma_N = \sigma_0$ and $E_{[a,a+n)}(\sigma_{a+n}) = \sigma_a$. These references need not equal the local Choi minorizers. Subtracting each reset contribution leaves completely positive residual maps with trace scales $1 - \eta$. Their ordered product yields Gram error $2D(1 - \eta)^{\lfloor n/s \rfloor}$, and a fixed-dimensional reshuffling converts it to the required transfer operator norm. The shorter final window uses only complete positivity and trace preservation.

Weakening to $\eta/2$ before the logarithm covers the original $\eta = 1$ endpoint with a finite positive rate. The resulting theorem ⁷ gives physical logarithmic-depth preparation of the original family above a uniform cutoff $N_0 \geq s$. Global mixing, a compatible reference family, faithfulness, a common stationary state, and short-ring nonvanishing are derived or unnecessary, rather than additional input hypotheses.

A concrete two-dimensional alternating unitary/reset family satisfies the criterion with $s = 2$ and $\eta = 1$, with distinct singular minorizers. The regression includes a three-site remainder and a zero one-site periodic vector with no available two-site window. This tests the necessity of the eventual conclusion without patching the family.

This establishes a useful local sufficient criterion. It does not prove that the source’s qualitative finite-correlation definition implies uniform window minorization. Rectangular varying bonds remain outside this theorem: zero-padding does not preserve full-space trace preservation.

⁷Formal declaration: `MPSPreparation.exists_isPreparedInDepth_inhomogeneous_le_1og_eventually_of_window_domination`.

11 Actual rectangular transfer rates

The rectangular dimensions need not be replaced by a fixed square-bond hypothesis. For $F_{a:b} : M_{D_b} \rightarrow M_{D_a}$, use the Hilbert–Schmidt induced operator norm in the matrix-unit bases and assume a uniform rate $\|F_{a:b} - \text{Tr}(\cdot)\sigma_a\| \leq K e^{-r(b-a)}$ for every nonempty block of every partition, with $K, r > 0$ chosen before the ring length. The references σ_a are positive semidefinite of trace one on the actual bonds and may be singular.

The padded reference is $X \mapsto \text{Tr}(P_b X) J_a \sigma_a J_a^\dagger$, where J_a is the coordinate isometry and $P_b = J_b J_b^\dagger$. The full-space trace reset would be incorrect on inputs outside the outgoing corner. Its Gram reference is $(J_a \sigma_a J_a^\dagger)^\top \otimes P_b$. The shared square-root and ordered-product estimates give pair error $C_D \sqrt{K} N e^{-rq/2}$ for blocks of lengths between q and $2q$. Whole-ring mixing supplies nonvanishing beyond a fixed cutoff. The theorem `MPSPreparation.exists_isPreparedInDepth_le_log_eventually_of_rectangular_mixing` in `TNLean/MPS/Preparation/RectangularPreparation` therefore gives physical depth $C \log(N/\varepsilon)$ for all $N \geq N_0$ and $0 < \varepsilon \leq 1$, with N_0, C selected first. This is a quantitative sufficient condition, and does not derive the required rate from qualitative convergence alone.

Dimension-one bonds are included. Equation (S39) concerns the paper’s numerical open-boundary study; the formal physical circuit conclusion uses ring adjacency. It does not assert open-line locality without a proof that the wrap-around edge is unused.

12 Rectangular local-window criterion

The preceding actual rectangular rate can be derived from a local condition. Fix $s \geq 1$ and $\eta > 0$, uniformly across a family with $1 \leq D_i \leq D$. Require every available window $F_{a:a+s} : M_{D_{a+s}} \rightarrow M_{D_a}$ to satisfy

$$J(F_{a:a+s}) \geq \frac{\eta}{D_{a+s}} (\tau_a \otimes I_{D_{a+s}}),$$

where the normalized output-first Choi convention includes the actual input dimension, and τ_a is a density that may be singular and vary with the window. The original rectangular site maps must be completely positive and trace preserving. No full-space trace-preservation claim is made about zero-padding.

A fixed density of the full-cycle channel and suffix transport derive compatible cyclic references. Subtracting the local trace-prepare contributions gives completely positive residual maps. Peeling complete windows from the right leaves a shorter left remainder. The residual product has trace scale $t = (1 - \eta)^{\lfloor n/s \rfloor}$ and unnormalized Choi trace D_{a+nt} . This bounds the actual Hilbert–Schmidt induced interval error by $2Dt$ directly, without introducing a common enlarged channel or changing the physical tensors.

The family window at $N = s$ gives $\eta \leq 1$. Weakening to $\eta/2$ yields $K = 4D$ and $r = -\log(1 - \eta/2)/s > 0$. The theorem `MPSPreparation.exists_isPr`

paredInDepth_le_log_eventually_of_rectangular_window_domination in `TNLean/MPS/Preparation/RectangularWindowPreparation` then gives the original physical ring preparation conclusion with $N_0 \geq s$ and C chosen before N, ε . Dimension-one bonds and singular references are included. This local condition remains additional quantitative input; the paper’s qualitative finite-correlation discussion does not provide its uniform window length and strength.

References

- [MSWC24] Daniel Malz, Georgios Styliaris, Zhi-Yuan Wei, and J. Ignacio Cirac. Preparation of matrix product states with log-depth quantum circuits. *Phys. Rev. Lett.*, 132:040404, 2024.