

The First MPU Source Cut and the Orientation of the Source Gates

2026-09-01

Verdict: local-correction (resolved).

1 Source-cut orientation

Let $\mathcal{U}_{j,\alpha,\beta}^i$ be a matrix product operator tensor, with upper and lower physical indices i, j and left and right auxiliary indices α, β . Section III of Ref. [CPGSV17] introduces two matrix flattenings. The diagonal cut in the first diagram separates the upper and right legs from the left and lower legs. The second cut separates the left and upper legs from the lower and right legs. Hence

$$(M_1)_{(i,\beta),(\alpha,j)} = \mathcal{U}_{j,\alpha,\beta}^i, \quad (1)$$

$$(M_2)_{(\alpha,i),(j,\beta)} = \mathcal{U}_{j,\alpha,\beta}^i. \quad (2)$$

The same leg assignment is explicit in equation `eq:3-leg` of Ref. [FRBC25]. In particular, X_1 carries (i, β) and Y_1 carries (α, j) in the factorization $M_1 = X_1 Y_1$.

An earlier formal reading used the transpose of Eq. 1, with row (α, j) and column (i, β) . That reading preserved the rank of M_1 , but it changed the types of X_1 and Y_1 and therefore changed every source contraction in which their open legs mattered. The corrected source cuts are `MPOTensor.sourceCutM` and `MPOTensor.sourceCutM`.

2 Weight and adjunction

The row space of M_1 is ordered as the upper physical factor followed by the right auxiliary factor. The weighted normalization in Ref. [CPGSV17] is therefore represented without an additional swap:

$$W_\rho = \text{Id}_d \otimes \rho, \quad (3)$$

$$X_1^\dagger W_\rho X_1 = \text{Id}_r. \quad (4)$$

These are `MPOTensor.sourceWeight` and `MPOTensor.sourceX_weighted_isometry`. Physical adjunction exchanges the two cuts by conjugate transpose,

$$M_1(\mathcal{U}^\sharp) = M_2(\mathcal{U})^\dagger, \quad (5)$$

$$M_2(\mathcal{U}^\sharp) = M_1(\mathcal{U})^\dagger. \quad (6)$$

The corresponding formal declarations are `MPOTensor.sourceCutM_physicalAdjointTensor` and `MPOTensor.sourceCutM_physicalAdjointTensor`.

3 Source gates

With $M_1 = X_1 Y_1$ and $M_2 = X_2 Y_2$, the gates in equations `uuvv` of Ref. [CPGSV17] and `eq:uv` of Ref. [FRBC25] have entries

$$u_{(l,r),(i_1,i_2)} = \sum_{\beta} (Y_2)_{l,(i_1,\beta)} (Y_1)_{r,(\beta,i_2)}, \quad (7)$$

$$v_{(j_1,j_2),(r,l)} = \sum_{\alpha} (X_1)_{(j_1,\alpha),r} (X_2)_{(\alpha,j_2),l}. \quad (8)$$

They are formalized by `MPOTensor.SourceFactors.sourceU` and `MPOTensor.SourceFactors.sourceV`. The periodic two-site operator then factors as

$$U_{(i_1,i_2),(j_1,j_2)}^{(2)} = \sum_{l,r} v_{(i_1,i_2),(r,l)} u_{(l,r),(j_2,j_1)}. \quad (9)$$

This is `MPOTensor.mpo_two_pair_entry_eq_sourceV_mul_sourceU_swap`. The open two-site forms use X_2, u, X_1 or v, Y_1, Y_2 , as recorded by `MPOTensor.SourceFactors.mul_apply_eq_sum_X_mul_sourceU_mul_X` and `MPOTensor.SourceFactors.mul_apply_eq_sum_sourceV_mul_Y_mul_Y`.

The contractions Y_1-X_2 and X_1-Y_2 remain useful auxiliary kernels, but they are not the gates u and v . Treating either mixed kernel as a paper gate produces the wrong physical and auxiliary leg assignment.

4 Weight orientation and the transfer fixed pair

Let ρ be the fixed point of the normalized transfer map, $\frac{1}{d} \sum_{i,j} \mathcal{U}^{ij} \rho (\mathcal{U}^{ij})^\dagger = \rho$, and let $|\rho\rangle$ be its column-stacking vector, so that the doubled-bond coordinate (α, α') carries the entry $\rho_{\alpha'\alpha}$, with α the index of the conjugated layer. Inserting $|\rho\rangle\langle\Phi|$ between two double-layer letters therefore pairs the conjugated leg with the column index of ρ :

$$(W^{p_1 q_1} |\rho\rangle) (\Phi | W^{p_2 q_2})_{(a_1, b_1), (a_2, b_2)} = \sum_{\alpha, \alpha', \beta, j_1, j_2} \overline{\mathcal{U}_{p_1, a_1, \alpha}^{j_1}} \mathcal{U}_{q_1, b_1, \alpha'}^{j_1} \rho_{\alpha' \alpha} \overline{\mathcal{U}_{p_2, \beta, a_2}^{j_2}} \mathcal{U}_{q_2, \beta, b_2}^{j_2}. \quad (10)$$

The printed weight $W_\rho = \text{Id}_d \otimes \rho$ of Eq. 3 instead pairs the conjugated leg with the row index of ρ :

$$(Y_1^\dagger Y_1)_{(a,p),(b,q)} = \sum_{i,\beta,\beta'} \overline{\mathcal{U}_{p,a,\beta}^i} \rho_{\beta\beta'} \mathcal{U}_{q,b,\beta'}^i. \quad (11)$$

The two readings agree exactly when $\rho = \rho^\top$. In Ref. [CPGSV17] the fixed point is diagonal, so the printed identity $X_1^\dagger(\text{Id} \otimes \rho)X_1 = \text{Id}$ and the diagram **X1X2b** coincide there. For a general positive definite ρ the two coordinate expressions differ by a transpose; the source does not make that generalization.

The formalization keeps the printed weight `MPOTensor.sourceWeight` and records the transpose on the inserted vector. The source- u Gram closes to a trace with $|\rho^\top\rangle(\Phi|$ in `MPOTensor.SourceFactors.sourceU_gram_eq_transpose_fixed_pair_trace`, and the source- v dressed Gram equation is equivalent to the two-letter identity with the insertion $|\rho^\top\rangle(\Phi|$ in `MPOTensor.sourceVDressedGram_iff_simple2_transpose_fixed_pair`. The source gates remain the gates built from the printed weight ρ ; replacing that weight by $\rho^\top$ recomputes X_1, Y_1 and gives a different gate when ρ is not symmetric. The source's diagonal convention makes the two inserted vectors equal, and `MPOTensor.IsMPUCanonicalFormII.sourceV_isIsometry` proves the forward direction of Theorem III.8 on precisely that convention. The converse direction in `MPOTensor.IsMPU.isMPUSimple_of_sourceV_isIsometry` produces the witnesses $(\Phi|$ and $|\rho^\top\rangle$ for the same weight ρ . Together with the diagonal source- u isometry, `MPOTensor.IsMPUCanonicalFormII.isMPUSimple_tfae` gives Theorem III.8 under the source's standing convention; the transport of source cuts to a reduced representative of an arbitrary ambient MPU remains recorded in `mpu_canonical_form_full_support.tex`. No source-factor definition is changed.

For the source- u gate, the transpose closure and the source's diagonal coordinates give a direct retained-window contraction. If the normalized transfer power supplies

$$E^J = |\rho\rangle(\Phi|, \quad (12)$$

then the reflected transfer identity at $K = JD^2$ and cyclicity of the trace give

$$\sum_{l,r} u_{(l,r),q} \overline{u_{(l,r),p}} = \text{tr}(W^{p_1 q_1} W^{p_2 q_2} E^K). \quad (13)$$

Expanding this trace produces one K -site input tail between two unblocked one-site endpoints. The p entries are conjugated and the q entries are not, so MPU tail isometry gives $u^\dagger u = \text{Id}$ and $d^2 \leq r\ell$. The arbitrary- K declarations are `MPOTensor.SourceFactors.sourceU_gram_eq_closed_doubleLayer_trace_of_isSymm`, `MPOTensor.SourceFactors.sourceU_gram_eq_normalized_input_tail`, `MPOTensor.SourceFactors.sourceU_isIsometry_of_isMPU`, and `MPOTensor.SourceFactors.mul_self_le_rightRank_mul_leftRank_of_isMPU`. The compact-SVD wrapper is `MPOTensor.IsMPUCanonicalFo`

`rmII.sourceU_isIsometry`, while `MPOTensor.SourceFactors.sourceU_gram_eq_normalized_mpo_input_tail_of_isDiag` converts the raw trace-one transfer power at J into the direct fixed pair at $K = JD^2$. The independent right-shift checks `MPOTensor.sourceU_rightShiftPaperSourceFactors_gram` and `MPOTensor.sourceU_rightShiftPaperSourceFactors_isIsometry` fix the same normalization and orientation explicitly. These conclusions assume the source’s diagonal fixed point; they do not assert a source- u theorem for arbitrary nonsymmetric ρ .

5 Dressings acting on the corrected cut

Any operator written against a row or a column space of M_1 inherits the leg order of Eq. 1. The $\mathbb{Z}_2$ appendix of Ref. [FRBC25] considers a tensor whose adjoint is a unitary virtual gauge conjugate of itself, $\mathcal{U}_{ij}^\dagger = T^\dagger \mathcal{U}_{ij} T$, and dresses the second cut into the first,

$$M_1(\mathcal{U}) = A M_2(\mathcal{U})^\dagger B, \quad (14)$$

with A acting on the row index (i, β) and B on the column index (α, j) . The printed candidate factors place the identity factor first in both dressings.¹ The row index is physical followed by auxiliary, so the left dressing keeps that order, whereas the column index is auxiliary followed by physical and reverses it:

$$A = \text{Id}_d \otimes \overline{T}, \quad (15)$$

$$B = T^\top \otimes \text{Id}_d. \quad (16)$$

These are the same two operators as the printed ones; only the order in which their tensor factors are written follows Eq. 1. No further correction is involved, and the corresponding declarations are `MPOTensor.inverseCompatibleLeftGauge` and `MPOTensor.inverseCompatibleRightGauge`.

6 Marker boundary for the complete source- u network

The transposed first cut was a local correction, and it is resolved. The separate supplied-fixed-pair restriction of the complete source- u module is resolved as well. Its generic theorems still take a stabilized fixed pair as an explicit hypothesis, but they are stated as reusable steps for an arbitrary positive diagonal weight and an arbitrary stabilization exponent, and Lemma III.7 itself is now stated for a tensor in canonical form II, where the weight and the stabilized power are both recorded by the convention. The source itself fixes ρ to be diagonal with positive weights $\rho_n > 0$, so diagonality is not a restriction relative to

¹Source traceability: `Papers/2502.20257/main.tex`, line 5432.

the cited lemma. Accordingly, the module carries no scope marker for the supplied fixed pair. The remaining source-data restriction, namely the transport of source cuts from an arbitrary ambient MPU to its reduced representative, is recorded in `mpu_canonical_form_full_support.tex`.

The elimination step for the supplied-data restriction was the bundling of the standing canonical-form-II convention and the restatement of the MPU-level results through it, and it is complete. Issue GitHub issue #7653 instead asks for a nonsymmetric ambient-weight strengthening that CPSV17 does not claim. It is not the elimination step for the source-faithful theorem.

7 Affected claims and audit boundary

The transposed first cut affected source-factor recovery, physical adjunction, physical reindexing, tensor products of supplied source factors, open and closed source networks, source- v dressing, and the explicit shift examples. Rank-only statements could remain numerically true because transposition preserves rank, while entrywise and open-leg statements required new proofs. The corrected declarations now use the leg assignments in Eqs. 1– 8.

Any source- u or source- v result written before this correction must be revalidated from its current declaration rather than accepted from an older compiled artifact. In particular, a proof must expose the Y_2 – Y_1 contraction for u or the X_1 – X_2 contraction for v , with the product orders shown above. The supplied-fixed-pair source- u contraction has now been revalidated in the source’s diagonal canonical-form-II coordinates. Results outside this diagonal source- u boundary still require the same statement-by-statement check.

8 Verdict

The former definition of M_1 was the transpose of the source diagram. The correction changes no source rank, but it changes the typed source factors and all order-sensitive contractions. The first cut, weight, adjunction identities, paper gates, two-site factorizations, tensor products, and shift examples now use one consistent leg assignment. The correction is resolved, while the revalidation rule remains applicable to any future recovery of source-gate work predating this audit. The printed weight and the column-stacking transfer fixed pair differ by a transpose for non-diagonal ρ ; the source- u Gram and source- v network theorems carry this transpose on the inserted vector without changing the source gate’s weight. The source’s diagonal fixed point removes the mismatch in Theorem III.8 and proves the supplied-fixed-pair source- u isometry.

References

- [CPGSV17] J. Ignacio Cirac, David Pérez-García, Norbert Schuch, and Frank Verstraete. Matrix product unitaries: Structure, symmetries, and

topological invariants. *Journal of Statistical Mechanics: Theory and Experiment*, 2017(8):083105, 2017.

- [FRBC25] Adrián Franco Rubio, Arkadiusz Bochniak, and J. Ignacio Cirac. Symmetry defects and gauging for quantum states with matrix product unitary symmetries, 2025.