1. Isometric PEPS: Motivation, Moses move

Canonical forms for 1D MPS are very useful: since constituents are isometries, contractions are trivial:

Standard tool for obtaining canonical form: SVD 

Can we similarly express PEPS state purely through isometries to obtain an 'isometric PEPS' ? YES!

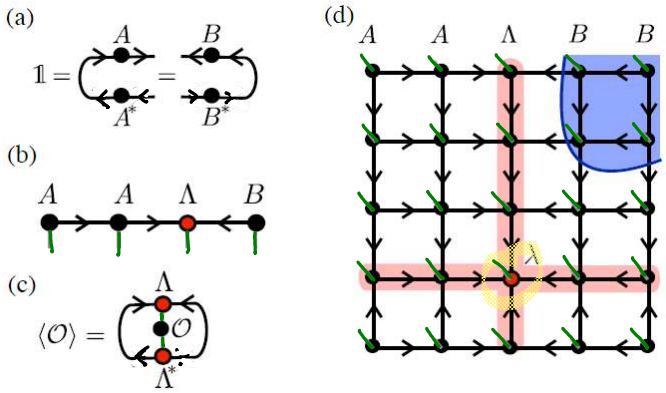

FIG. 1. Schematic representation of the canonical form in 1D and 2D. (a) Left and right isometries are represented by arrows whose orientation indicates whether $A^\dagger A = BB^\dagger = \mathbb{1}$. We view the isometry as an RG-like procedure from the large Hilbert space (incoming arrows) to the smaller one (outgoing arrows). In the case of higher-rank tensors, the contraction $A^\dagger A = \mathbb{1}$ is always over all the incoming arrows. (b) A 1D MPS can be brought into a mixed canonical form with orthogonality center Λ. Note that each dangling physical index implicitly has an incoming arrow. (c) Expectation values of local operators can be directly obtained from Λ. (d) 2D canonical form with "orthogonality hypersurfaces" Λ (column and row highlighted in red). The orthogonality center λ is marked by a red dot. In blue we indicate an example of a subregion with only outgoing arrows, whose boundary map is consequently an isometry.

MPS^2: all tensors except those in 'orthogonality hypersurfacer' Λ (red row, column) are isometries:

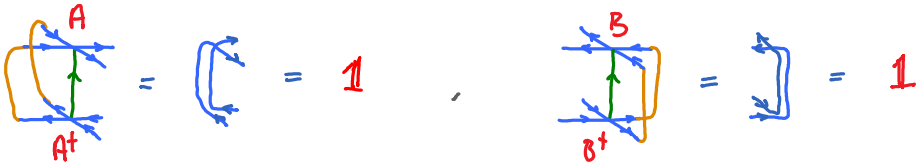

(1)

Λ is wave function of the system in an orthonormal basis, because its exterior is an isometry from the physical bonds to the incoming virtual bonds of Λ.

This allows local expectation values to be computed as $\langle\psi|\hat{O}|\psi\rangle = \langle\Lambda|\hat{O}|\Lambda\rangle$ (2) (compare (c))

since tensors A, B outside of orthogonality center Λ contract to $\mathbb{1}$ by isometry condition (1).

Interesting open question: how does variational power of isometric PEPS differ from general PEPS? One restriction is: many of its correlations decay exponentially, because any two-point function along orthogonality center can be reduced to that of the MPS Λ, whose correlations decay exponentially.

Shifting orthogonality center: Moses move (MM)

$\Lambda^\ell B^{\ell+1} \overset{(a)}{=} A^\ell \Lambda B^{\ell+1} = A^\ell \Lambda^{\ell+1}$ (3)

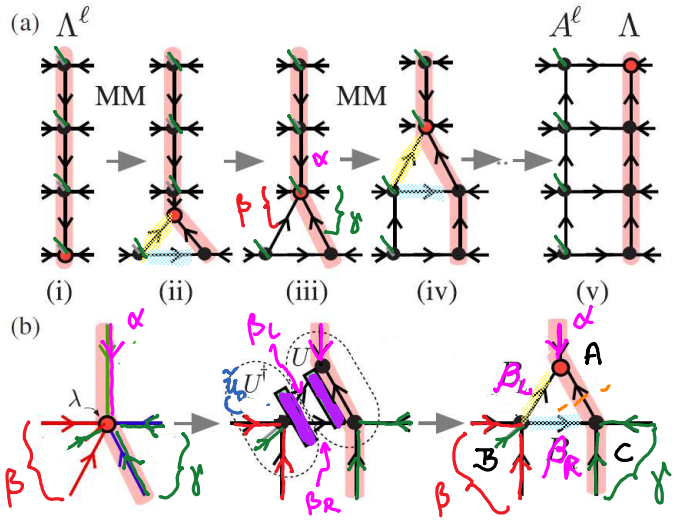

FIG. 2. The Moses move. (a) The orthogonality hypersurface Λ^ℓ is split into the product of a left isometry A^ℓ and a zero-column state Λ with no physical indices. The unzipping is performed by successively applying the splitting procedure shown in panel (b). The legs of the center site λ are grouped into a tripartite state $|\alpha\beta\gamma\rangle$ which is "split" into three tensors in two steps: first find $|\alpha\beta\gamma\rangle \approx \tilde{u}_0 U^\dagger |\alpha\beta_L, \beta_R\gamma\rangle$ for an initial guess of the isometry $\tilde{u}_0$ and unitary U which minimizes the entanglement across the vertical bond highlighted in red; second set $a = \tilde{u}_0 U^\dagger$ and split $|\alpha\beta_L, \beta_R\gamma\rangle$ in two via SVD. The resulting a comprise the tensors in A^ℓ, and the choice of U will produce a Λ with minimal vertical entanglement.

At central cite (red dot), three 'regions' meet: top (α), left/bottom (β), bottom/right (γ). We need to split three-partite entanglement into two-partite entanglement, top/left (α, β_L) vs. bottom-right (β_R, γ).

Needed: the optimal isometry B which 'splits' the composite index β into two sub-indices β_L, β_R

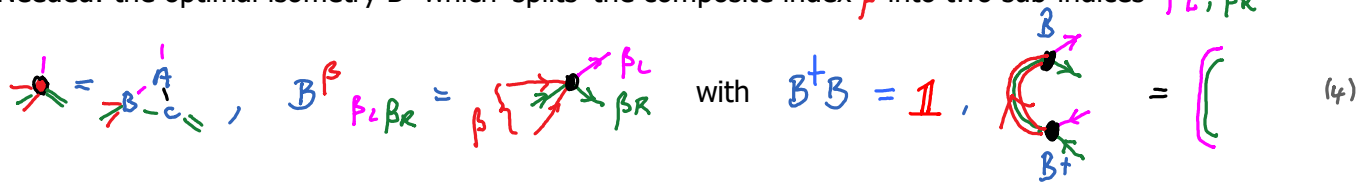

with $B^\dagger B = \mathbb{1}$, (4)

while minimizing $S_{\alpha\beta_L,\beta_R\gamma}$, the entanglement entropy between top/left and bottom/right.

Choice of B: (5)

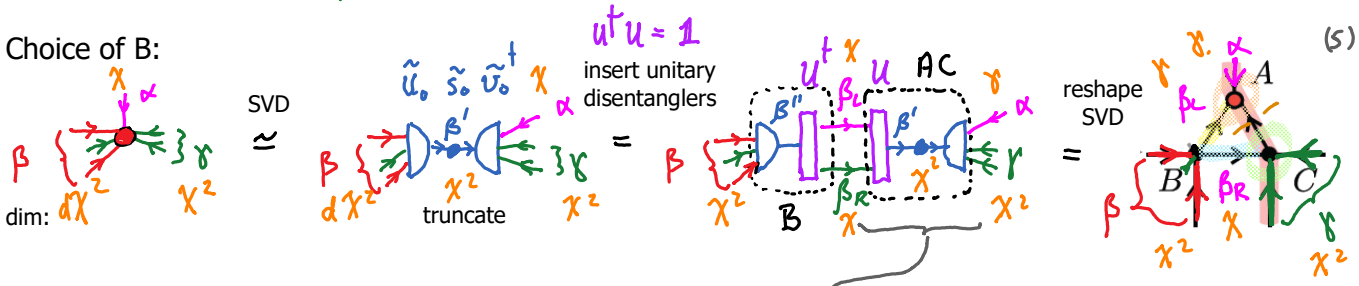

with disentangler U chosen to minimize $S^{(2)}_{\alpha\beta_L,\beta_R\gamma}$ (using projective truncation methods of Evenbly & Vidal).

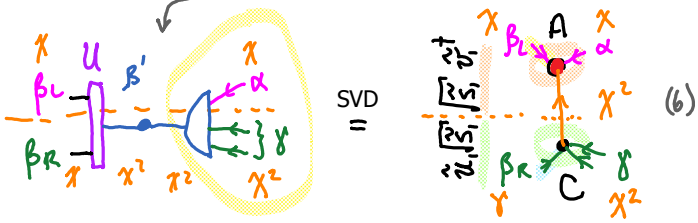

(6)

General question: find U which splits given system into two minimally entangled subsystems, via $\beta' = (\beta_L, \beta_R)$

$$|\psi\rangle = |\gamma\rangle_R |\beta'\rangle |\alpha\rangle_L \psi^{\alpha\beta'\gamma} \xrightarrow{?} |\tilde{\psi}\rangle = U|\psi\rangle = |\gamma\rangle_R |\beta_R\rangle |\beta_L\rangle |\alpha\rangle_L \psi^{\alpha\beta_L,\beta_R\gamma}$$

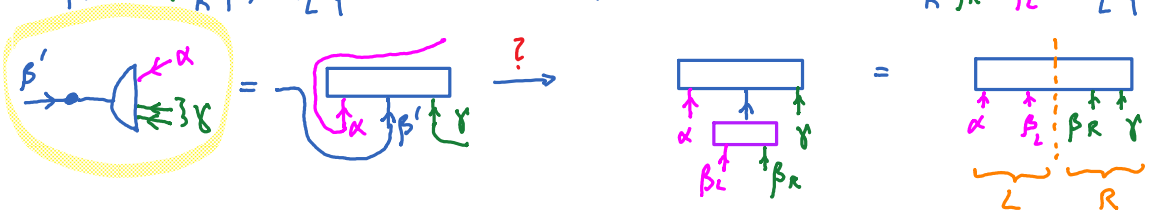

(7)

Tool: construct reduced density matrix, $\tilde{\rho}_L = \mathrm{Tr}_R |\tilde{\psi}\rangle\langle\tilde{\psi}|$ (8)

and choose U such that it minimizes von Neumann (or Shannon) entropy:

$$S^{(1)}_{LR}(\tilde{\rho}_L) = -\mathrm{Tr}\, \tilde{\rho}_L \log_2 \tilde{\rho}_L \quad (9)$$

$S^{(1)}_{LR}$ is always ≥ 0

is maximal if all weights are equal, $\Rightarrow |\tilde{\psi}\rangle$ is 'maximally entangled'

is minimal if one weight = 1, all others zero, $\Rightarrow |\tilde{\psi}\rangle$ is 'product state' (10)

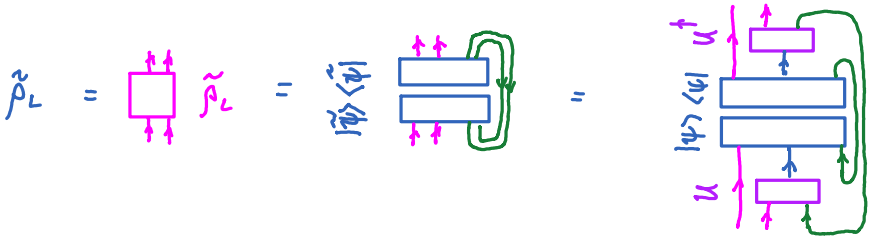

(11)

Practical problem: von Neumann entropy is a highly non-linear function of reduced density matrix, there is no straightforward way to minimize it. Alternative: consider Rényi entropies!

Rényi entropy [Rényi1961] : $S^{(\alpha)}(\rho) = \frac{1}{1-\alpha} \log_2 [\mathrm{Tr}\, \rho^\alpha], \quad 0 \le \alpha \le \infty$ (12)

Rényi entropy has properties similar to the von Neumann entropy. It reduces to the latter for $\alpha \to 1$:

Proof: $S^{(\alpha)}_{L/R}(\rho) = \frac{1}{1-\alpha} \log_2 [\mathrm{Tr}(\rho \cdot \rho^{\alpha-1})] = \frac{1}{1-\alpha} \log_2 [\mathrm{Tr}(\rho\, e^{(\alpha-1)\ln \rho})]$

$\xrightarrow{\alpha \to 1} \frac{1}{1-\alpha} \log_2 [\mathrm{Tr}(\rho \cdot [\underbrace{1}_{=1} + (\alpha-1)\ln\rho + \mathcal{O}((\alpha-1)^2)])]$

[use $\mathrm{Tr}\,\rho = 1$] $= \frac{1}{1-\alpha} \log_2 [1 + (\alpha-1)\,\mathrm{Tr}\,\rho \cdot \ln\rho + \mathcal{O}((\alpha-1)^2)]$

$S^{(1)}_{L/R}(\rho) = -\mathrm{Tr}[\rho \log_2 \rho] \Rightarrow$ [we used: $\log_2 x = \frac{\ln x}{\ln 2}$, $\ln(1+y) \xrightarrow{y \to 0} y$]

Properties (10) hold also for Rényi entropies, hence they also serve as useful entanglement measures.

To minimize L/R entanglement, we may hence minimize $S^{(\alpha)}_{L/R}(\tilde\rho_L)$ for some α , e.g. $\alpha = 2$

Minimization of 'second Rényi entropy', [Hauschild2018, Sec.III.A]

$S^{(2)}_{L/R}(\tilde\rho_L) = -\ln \underbrace{\mathrm{Tr}(\rho_L)^2}_{Z_2}$ (13)

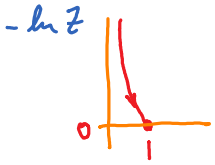

We seek a unitary 'disentangler' U, (14)

$|\tilde\psi\rangle = U|\psi\rangle$

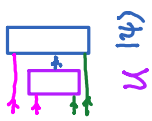

with constraint $U^\dagger U = U U^\dagger = 1$

which minimizes $S^{(2)}_{L/R}(\tilde\rho_L)$, i.e. which maximizes $Z_2(\tilde\rho_L) = \mathrm{Tr}\,\tilde\rho_L^2$ (15)

with $\tilde\rho_L =$

$=$

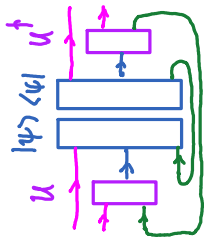

(16)

This constrained nonlinear optimization problem can not be solved directly.

But it can be solved iteratively via linearization: [Evenbly2009, Sec. IV], [Evenbly2017, Sec. III], TNR.3

Define environment of U via $Z_2 = t\,\mathrm{Tr}[\Gamma_U \otimes U]$ (17)

Do SVD on environment of U, $\Gamma_U = \tilde U \tilde S \tilde V^\dagger$ (18)

and update U by the prescription $U' = \tilde V \tilde U^\dagger$: (19)

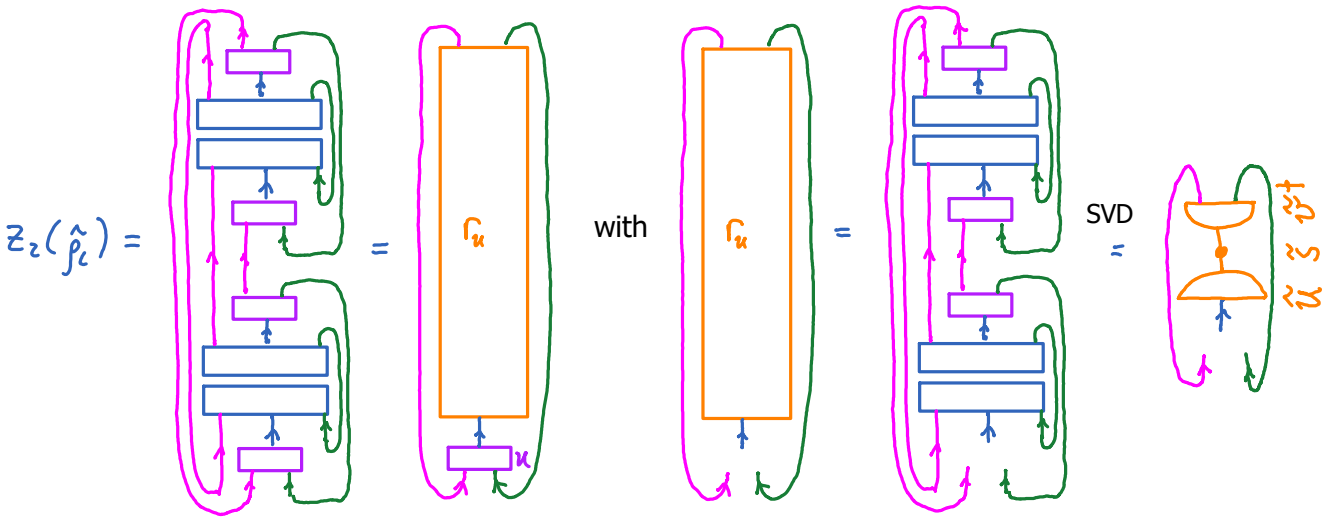

$$Z_2 = tTr\left[\Gamma_u \otimes u\right] = tTr\left[\tilde{u}\,\hat{s}\,\tilde{v}^\dagger\, u\right] \xrightarrow{\text{update}} tTr\left[\tilde{u}\,\hat{s}\,\tilde{v}^\dagger\, \underbrace{\tilde{v}\,\tilde{u}^\dagger}_{u'}\right] = tTr\left[\hat{s}\right] \tag{20}$$

Iterate until convergence of singular value spectrum of Γ_u .

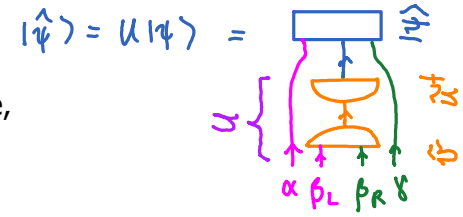

(21)

Final U is desired disentangler. Use it in (6) to achieve Moses move,

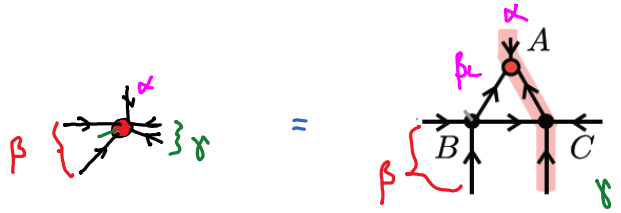

(23)

with A, B, C given by:

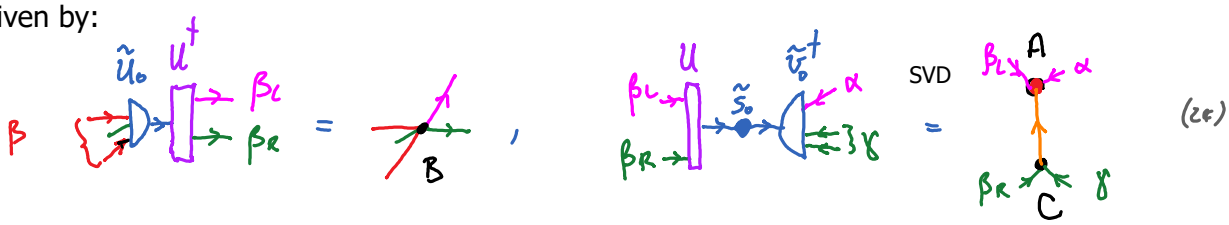

(24)

TEBD^2 algorithm (square lattice, nearest neighbors)

$$H = \sum_{c=1}^{L_x} H_c + \sum_{r=1}^{L_y} H_r \quad , \quad e^{-\beta H} = \prod_r e^{-\tau H_r} \prod_c e^{-\tau H_c} \qquad (1)$$

acts on column c acts on row r

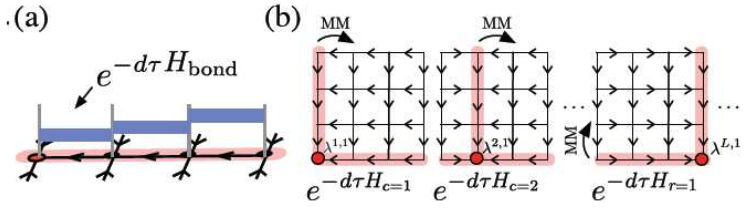

Transverse field Ising model.
FIG. 4. The TEBD^2 algorithm: (a) Trotterization of $e^{-\tau H_{r/c}}$ into a product of two-site terms acting on a single row/column of an orthogonality center. The gates are applied on the up sweep, and the MM on the down sweep. (b) To complete one time step, the 1D update is applied to all row/columns by sequentially shifting the orthogonality center $\lambda^{c,r}$ using the MM.

Begin with orthogonality center (OC) a lower left corner, $\lambda^{1,1}$.

Up-sweep along first row: apply $e^{-\tau H_{c=1}}$ to column Λ^1 by calling standard 1D TEBD.

Down-sweep along first row: Moses move to shift OC to the right by one column, to $\lambda^{2,1}$

Repeat up-down sweeps until reaching column Λ^{L_x} on the right, with OC at $\lambda^{L_x,1}$

Then successively apply $e^{-\tau H_r}$ to rows, and MM to shift OC upwards, until it reaches λ^{L_x,L_y}.

Repeat to bring OC λ counterclockwise around the four corners back to $\lambda^{1,1}$ to complete time step.

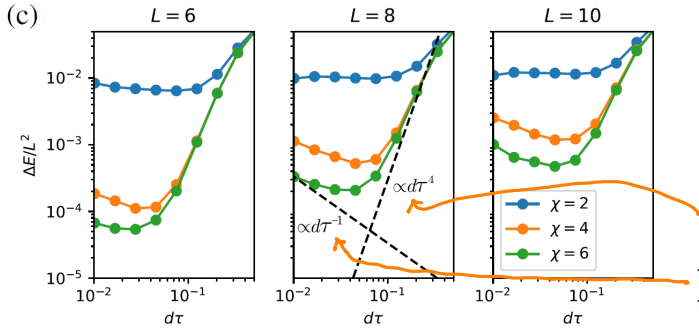

(c) Error densities of the energy of the transverse field Ising model with $g = 3.5$ for different system sizes and maximal bond dimensions χ as function of the Trotter step size $d\tau$.

Trotter error increases as $(d\tau)^4$

truncation error due to MM decreases as $\frac{1}{(d\tau)}$, since larger time steps mean less MM moves.

MPS to MPS^2 conversion: entanglement reshuffling

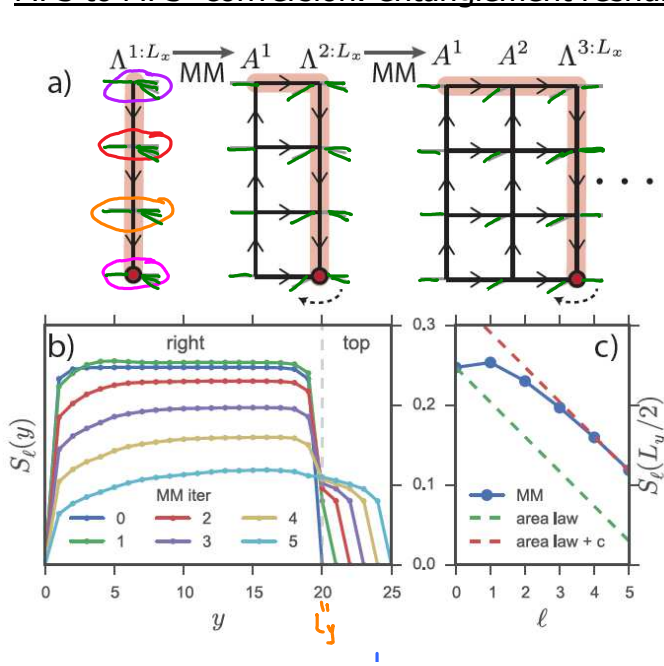

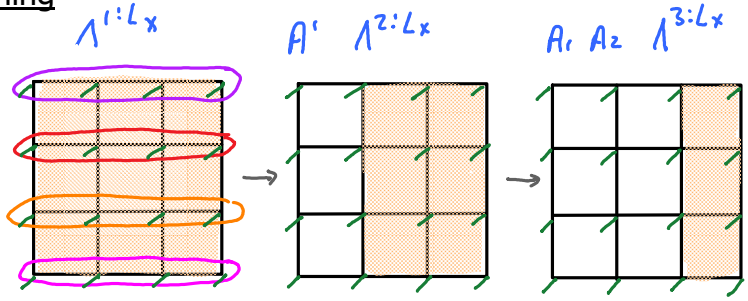

FIG. 3. The MPS to MPS^2-algorithm: (a) The MPS $\Lambda^{1:L_x}$ for an $L_x \times L_y$ strip is fed into the MM by treating the legs of the first column as the left ancilla and the remaining columns as the right ancilla to obtain $\Lambda^{1:L_x} = A^1\Lambda^{2:L_x}$. The renormalized wavefunction $\Lambda^{2:L_x}$ is then reshaped by viewing the legs of the second column as physical, and the remaining as right-ancilla. Applying the MM again, we can repeat to obtain a canonical TNS. (b) Entanglement entropy S_ℓ for the sequence of orthogonality centers (highlighted in pink) after ℓ iterations. y runs from bottom right, to top right, to top left. (c) S_ℓ for a cut at $y \sim L_y/2$, compared against the bulk area law determined from DMRG. $L_x = 6$, $L_y = 20$

y L_y ℓ

left. (c) S_ℓ for a cut at $y \sim L_y/2$, compared against the bulk area law determined from DMRG. $L_x = 6$, $L_y = 20$

(a) Start with 'fat' MPS $\Lambda^{1:L_x}$, where each 'site' contains the L_x spins of the corresponding row.

Use MM to iteratively 'peal off' columns of the wavefunction, $\Lambda^{\ell:L_x} = A^\ell \Lambda^{\ell+1:L_x}$, (2)

producing an MPS^2.

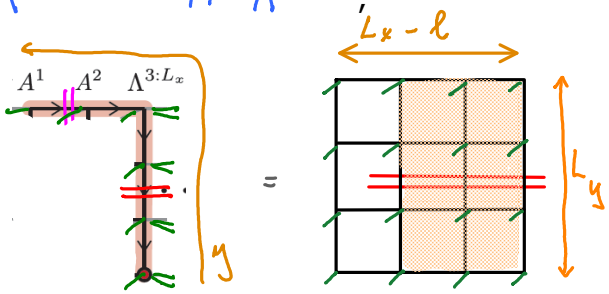

(b) Entanglement entropy $S_\ell(y)$ for cuts across Λ

(cut runs at fixed y), after iteration step ℓ.

In 'right-region' ($y < L_y$), S_ℓ decreases with ℓ, since $\Lambda^{\ell:L_x}$ becomes 'thinner'.

(c) After initial delay, S_ℓ reflects area law: $S_\ell(y = L_y/2) \sim s|\partial A| \sim s(L_x - \ell)$ (3)

Thus pealing off an isometric column from Λ removes from it a 'unit' s of horizontal entanglement.

Initial delay is expected, because any two vertically-entangled degrees of freedom will individually have some entanglement with their horizontal neighbors. The isometries A can fully remove them from right-region of Λ only after their entire support is to the left of Λ.

The entanglement removed from right-region of Λ by the isometries A is redistributed to the top-region of Λ, where it is encoded as horizontal entanglement. Its magnitude, given by $S_\ell(y)$ for a vertical cut through the top-region of Λ, is of order of the entanglement 'unit' s.

For $\ell = L_x$ ($= 6$), smoothly matches up between right/top regions (despite anisotropic algorithm!)

3. Canonical form for bond in 2D tensor network

CanF.3

[Evenbly2018]

For 1D MPS, there is a canonical form for a virtual bond:
Schmidt decomposition into left and right parts,
yielding diagonal bond matrix:

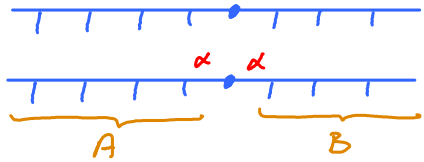

Method to obtain Schmidt decomposition: SVD

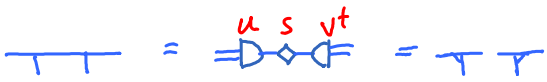

(1)

Schmidt decomposition also serves to 'fix gauge' on bond:
bond matrix should be diagonal, with only positive diagonal entries, arranged from large to small.

Generic question for 2D tensor network: can a canonical form for bonds be defined?
The question is nontrivial due to presence of loops, so Schmidt decomposition does not apply.

Evenbly proposes a generic answer in terms of 'bond environment'. [Evenbly2018, Sec. II]
(a) Consider bond between tensors A and B, described by bond matrix σ_{AB}
(b,c) Its 'bond environment', Y_{AB} , is 4-leg tensor obtained by deleting this bond (twice) from $\langle\psi|\psi\rangle$

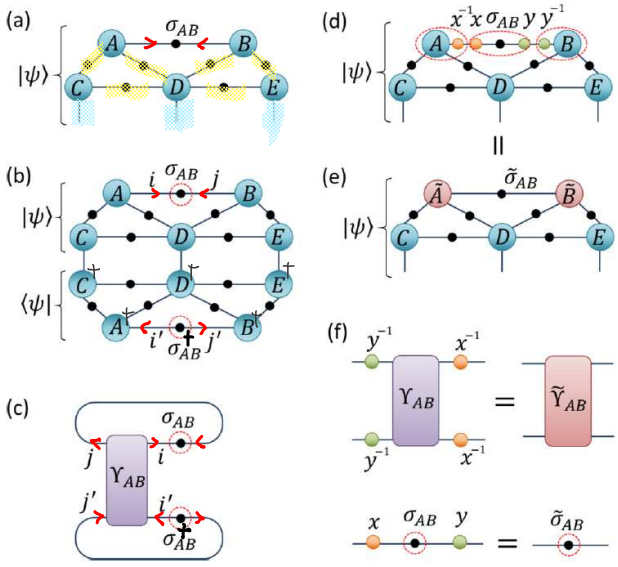

$\langle\psi|\psi\rangle = \sigma^{\dagger}_{j'i'}\, {}^{j'}_{j}\Upsilon^{i'}_{i}\, \sigma^{ij}$ (2)

gauge change: $= \tilde{\sigma}^{\dagger}_{j'i'}\, {}^{j'}_{j}\tilde{\Upsilon}^{i'}_{i}\, \tilde{\sigma}^{ij}$ (3)

FIG. 1. (a) Quantum state $|\psi\rangle$ defined from a network of tensors $\{A, B, C, D, E\}$ with bond matrices σ sitting between pairs of tensors. (b) Tensor network for $\langle\psi|\psi\rangle$. (c) The bond environment Υ_{AB} is defined by contracting $\langle\psi|\psi\rangle$ while leaving the index between A and B (and their conjugates) open. (d)–(e) A change of gauge, which leaves the state $|\psi\rangle$ invariant, is enacted on the index between A and B via matrices x and y together with their inverses. (f) Depiction the new bond environment $\tilde{\Upsilon}_{AB}$ and associated bond matrix $\tilde{\sigma}_{AB}$ from the gauge change in (e).

Useful properties of bond environment:

E.1 Bond environment of an index is invariant under gauge changes on all other internal indices of network.

E.2 All bond environments of all indices are invariant under unitary transformations on external indices of network (because external indices are contracted out when constructing Y).

E.3 If bond is a 'bridge',  , its bond environment factorizes:

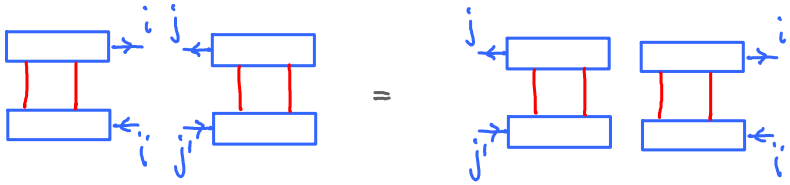

$${}^{j'}_{j}Y^{i'}_{i} = Y_R{}^{i'}_{i}\, Y_L{}^{j'}_{j}$$ (4)

Gauge fixing [Evenbly2018, Sec. III]

Define left and right 'boundary matrices' by contracting $\sigma^\dagger\sigma$ or $\sigma\sigma^\dagger$ onto Υ from left or right:

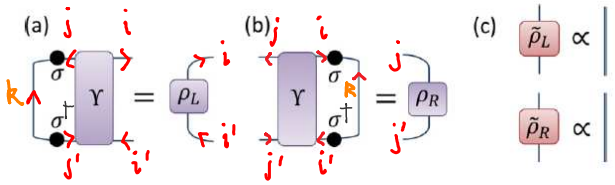

FIG. 2. (a)–(b) The left boundary matrix ρ_L is formed from contracting a bond environment Υ with (two copies of) the associated bond matrix σ. (b) The right boundary matrix ρ_R. (c) The weighted trace gauge (WTG) is the choice of gauge that yields trivial environment matrices, $\tilde{\rho}_L \propto \mathbb{I}$ and $\tilde{\rho}_R \propto \mathbb{I}$.

$$\sigma^\dagger_{j'k}\,\sigma^{ki}\,\Upsilon^{j'i'}_{ji}\,\sigma^\dagger_{i'i} = \rho_L{}^{j'}_{j} \qquad (5)$$

contract left index of σ with matching index of $\sigma^\dagger$

$$\Upsilon^{j'i'}_{ji}\,\sigma^{ik}\,\sigma^\dagger_{ki'} = \rho_R{}^{j'}_{j} \qquad (6)$$

contract right index of σ with matching index of $\sigma^\dagger$

Index is in 'weighted trace gauge' (WTG) if both left and right boundary matrices are proportional to $\mathbb{I}$ and bond matrix σ itself is diagonal, with positive diagonal elements, arranged from large to small.

If bond is bridge index, so that bond environment factorizes, WTG implies

(7)

which is reminiscent of MPS canonical form:

(8)

How to transform into WTG: right/left-multiply by inverse of dominant left/right transfer operators

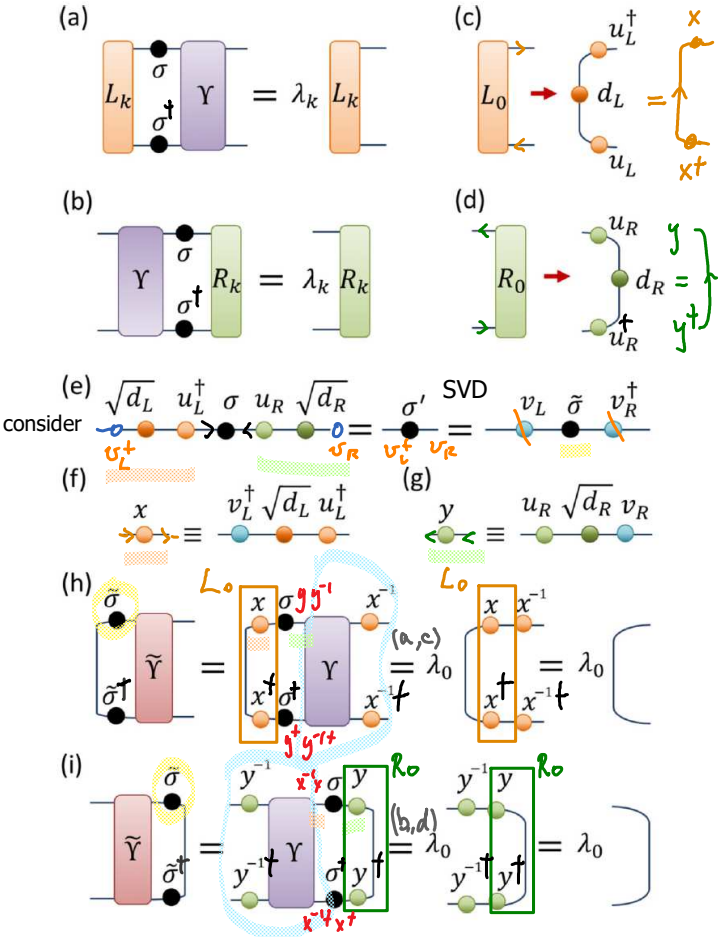

(a) Left transfer operator: $L_k(\sigma^\dagger\sigma\Upsilon) = \lambda_k L_k$

(c) SVD of dominant one: $L_0 = u_L d_L u_L^\dagger$

(b) Right transfer operator: $(\Upsilon\sigma\sigma^\dagger)R_k = \lambda_k R_k$

(d) SVD of dominant one: $R_0 = u_R d_R u_R^\dagger$

(e) definition of modified bond matrix, and its SVD.

Writing (e) as 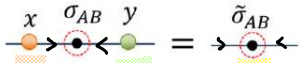 , we identify x and y as given by (f). By construction, they satisfy

$$x^\dagger x = (u_L\sqrt{d_L}\,v_L)(v_L^\dagger\sqrt{d_L}\,u_L^\dagger) = u_L d_L u_L^\dagger = L_0$$

$$y y^\dagger = (u_R\sqrt{d_R}\,v_R)(v_R^\dagger\sqrt{d_R}\,u_R^\dagger) = u_R d_R u_R^\dagger = R_0$$

These properties ensure that $\tilde{\Upsilon}$ and $\tilde{\sigma}$ satisfy WTG, as demonstrated in (h) = (a)$^{x^{-1}}_{x^{-1\dagger}}$ and (i) = $^{y^{-1}}_{y^{-1\dagger}}$(b).

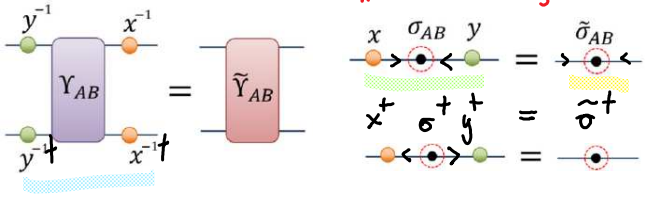

4. Full environment truncation (FET)

[Evenbly2018, Sec. V]

CanF.4

Goal: replace $|\psi\rangle$, with bond dimension χ on bond between A and B,
by $|\tilde{\psi}\rangle$, with bond dimension $\tilde{\chi} < \chi$ on that bond, while keeping their overlap maximal.

Strategy: replace bond matrix  by 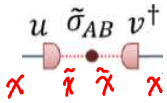 (1)

and chose isometries, u, v, with $u^\dagger u = \mathbb{1}$, $v^\dagger v = \mathbb{1}$, (2)

such that they maximize the cost function $F(\psi, \tilde{\psi}) = \dfrac{\langle\tilde{\psi}|\psi\rangle\langle\psi|\tilde{\psi}\rangle}{\langle\tilde{\psi}|\tilde{\psi}\rangle\langle\psi|\psi\rangle}$ (3)

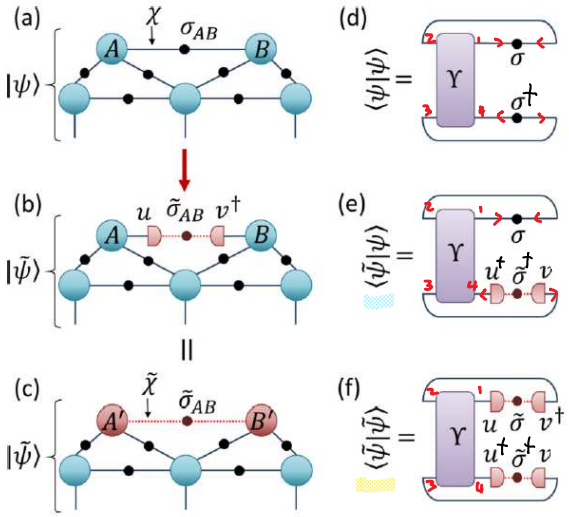

(d,e,f): this cost function can be expressed through the bond enviroment, Υ. Hence the name: 'full environment trunctation' (FET).

FIG. 5. (a)–(c) The index connecting tensors A and B is truncated to smaller dimension, $\tilde{\chi} < \chi$, by replacing σ_{AB} with the product $u(\tilde{\sigma}_{AB})v^\dagger$, where u, v are isometries and $\tilde{\sigma}_{AB}$ is a $\tilde{\chi} \times \tilde{\chi}$ matrix. (d)–(f) The overlaps $\langle\psi|\psi\rangle$, $\langle\tilde{\psi}|\psi\rangle$, $\langle\tilde{\psi}|\tilde{\psi}\rangle$ of the states from (a)–(b), which have been expressed in terms of the bond environment Υ.

Initialize optimization protocol by doing SVD on bond matrix: 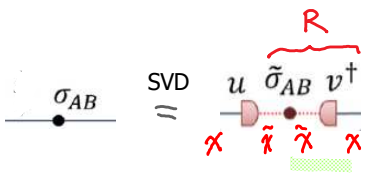 (4)

Then iteratively compute (b)-(g).

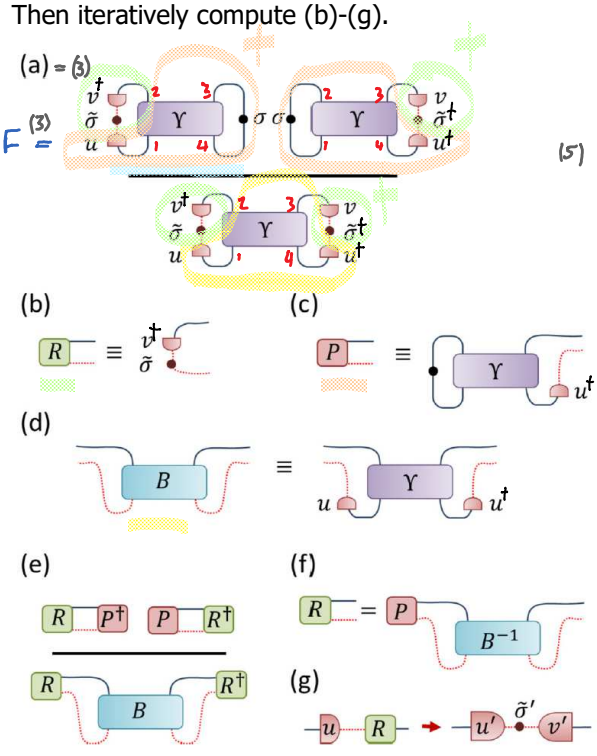

FIG. 11. Diagrams relating to the full environment truncation (FET) algorithm. (a) The fidelity F between an initial and a truncated state expressed in terms of a bond environment Υ, see Fig. 5. (b)–(d) Definitions of tensors R, P, and B. (e) The fidelity F can be expressed as a generalized eigenvalue problem in R as $F = (RP^\dagger PR^\dagger)/(RBR^\dagger)$. (f) The fidelity is maximized with the choice $R = PB^{-1}$. (g) Updated tensor u', $\tilde{\sigma}'$, and v' are obtained from the SVD of the product uR.

(b) Initialize $R = \tilde{\sigma}v^\dagger$ using SVD from (4).

Then seek to solve for optimal R while u is held fixed.

(c-d) Define tensors P, B, such that cost function takes the form

(e) $F = \dfrac{RAR^\dagger}{RBR^\dagger}$, with (6) $A = P^\dagger P$ (7)

(f) Since A is simply outer product of two vectors, optimal choice for R is known analytically: $R = PB^{-1}$ (8)

(g) Do SVD on uR (= (4)) to obtain updated $u', \tilde{\sigma}', v'$

Next, update the product $L = (u\tilde{\sigma})$ in similar way, with v held fixed.

Iterate these two steps until convergence is reached (typically 'less than 20 iterations are needed).

Benchmark results for FET

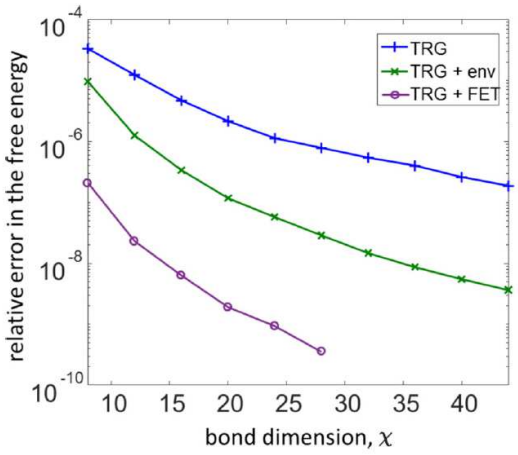

FIG. 13. Relative error in the free energy per site of the classical Ising model on a $2^{16}\times 2^{16}$ lattice of spins at critical temperature, comparing (i) tensor renormalization group [32] (TRG), (ii) tensor renormalization group with enlarged environment [36] (TRG + env) and (iii) tensor renormalization group that includes full environment truncations (TRG + FET).

SRG of [Zhao2010]

FET does better than TRG or SRG.

Evenbly claims (without showing data) that performance of FET is comparable to Gilt.

Gilt is 'smarter' than FET, because Gilt does more than optimize the overlap $\langle\tilde{\psi}|\psi\rangle$

Instead, Gilt eliminates all information flowing out from bond that cannot traverse environment and reach outside world. Gilt and FET actually use same bond environment, but treat its indices differently:

Gilt:

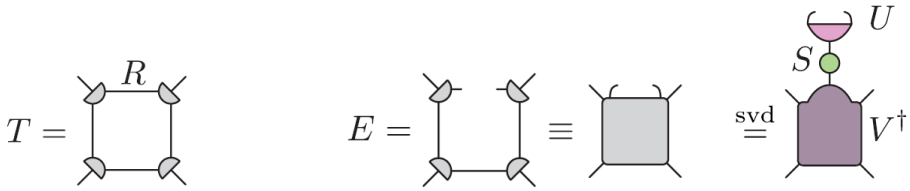

Gilt uses SVD between indices defining upper and lower gap:

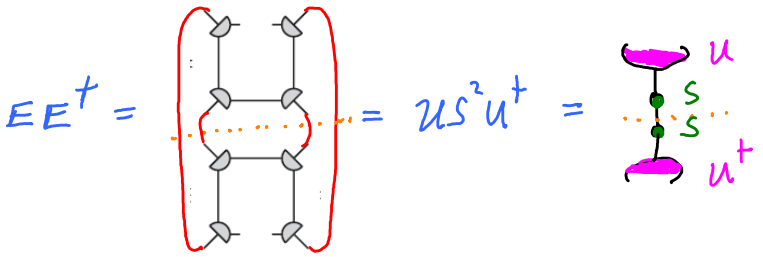

FET:

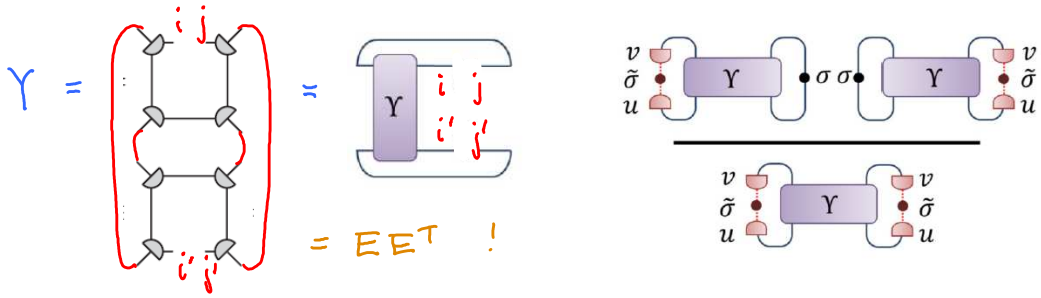

FET successively optimizes σ bond, then $\tilde{\sigma}$ bond, iteratively until convergence.
But it never makes a 'horizontal SVD' through the middle of Υ, which governs how information flows from inside to outside. For this reason, FET is less efficient than Gilt.