1. Original formulation of DMRG

[White1992], [White1993], [Schollwöck2011, Sec 2.2]

Goal: finding ground state of infinite chain

Infinite-size DMRG (iDMRG)

Diagonalize small system (e.g. 2 sites), write ground state in the form

$$|\psi\rangle = \sum_{\alpha\beta} |\beta\rangle_R |\alpha\rangle_L \psi^{\alpha\beta} \qquad (1)$$

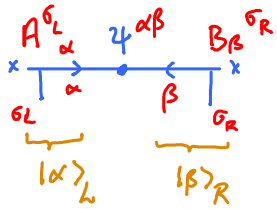

'Block L' describes left part of system, with basis $\{|\alpha\rangle_L\}$

'Block R' describes right part of system, with basis $\{|\beta\rangle_R\}$

Now add two sites between blocks L and R, and seek new ground state of $H_{L\bullet\bullet R}$ of the form

$$|\psi\rangle = |\beta\rangle_R |\sigma_R\rangle |\sigma_L\rangle |\alpha\rangle_L \psi^{\alpha\sigma_L,\sigma_R\beta} \qquad (2)$$

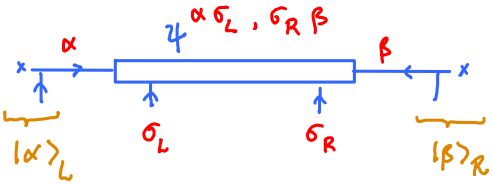

by minimizing (Lanzcos) $\dfrac{\langle\psi|\hat{H}_{L\bullet\bullet R}|\psi\rangle}{\langle\psi|\psi\rangle}$

Bond dimension has grown from $D\times D$ for $\psi^{\alpha\beta}$ to $Dd\times Dd$ for $\psi^{\alpha\sigma_L,\sigma_R\beta}$, so truncation is needed.
Split enlarged system in the middle, and call left side (new) block L, right side (new) block R.

Write ground state in the form

$$|\psi\rangle = |b\rangle_R |a\rangle_L \psi^{a,b} \qquad (3)$$

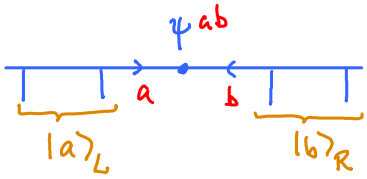

with composite indices $a = (\alpha, \sigma_L), b = (\beta, \sigma_R)$

of dimension $D_a = D_\alpha d$. White's truncation prescription: compute reduced DM of $L\bullet$,

$$\rho_{L\bullet} = \mathrm{Tr}_{\bullet R}|\psi\rangle\langle\psi| = |a'\rangle_L \underbrace{\psi^{a'b'}\overline{\psi}_{b'a}}_{(\rho_{L\bullet})^{a'}{}_a} {}_L\langle a| = \sum_c |\tilde{c}\rangle_L \rho_c {}_L\langle\tilde{c}| \qquad (4)$$

diagonalize

Construct truncated basis for block L , using the D eigenvectors $|\tilde{c}\rangle_L$ with the largest eigenvalues ρ_c . Rename: $|\alpha\rangle_L^{new} := |\tilde{c}\rangle_L$.

here truncation happens

Ditto for block R.

Then iterate,: add two more sites, etc.

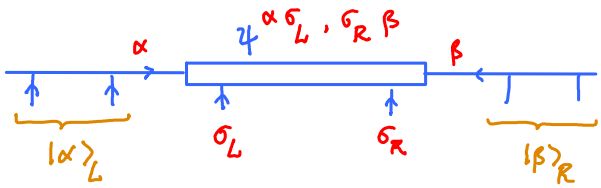

(5)

Then iterate,: add two more sites, etc.

Remark: we established early on (see MPS-II.2) that the eigenvalues ρ_c of reduced density matrix of L, are obtained by SVD of ψ^{ab}

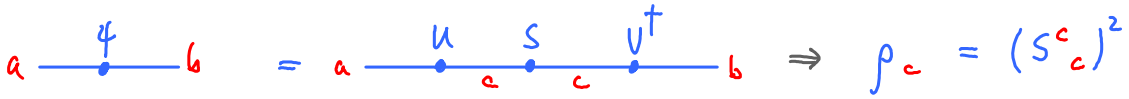

So retaining the ρ_L-eigenstates with largest ρ_c is equivalent to just doing SVD-truncation on ψ^{ab}.

Modern formulation

Start with MPS in bond-canonical form:

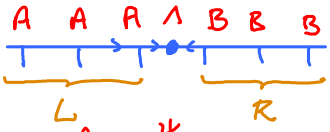

Add two central sites and find ground state (Lanczos):

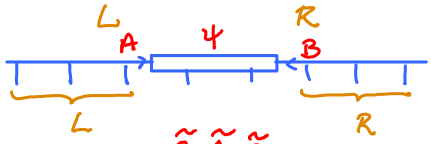

Do SVD to split chain into two larger blocks, and truncate:

$\psi = U S V^\dagger = \tilde{A}\,\tilde{\Lambda}\,\tilde{B}$

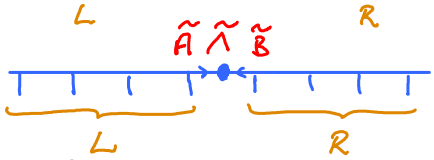

Iterate: make chain longer and longer, until ground state energy per site converges.

'iDMRG state prediction' [McCulloch2008], [Schollwöck2011, Sec. 10.1]

To speed up Lanczos search for ground state, construct initial guess for ψ from previous data:

initial Ansatz

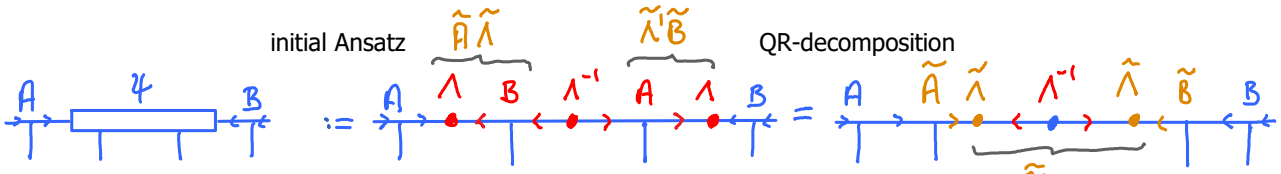

Logic: let A be followed by ΛB, and B preceded by $A\Lambda$

to reverse arrows between B and A, use Λ^{-1}

use this to initialize Lanczos

This leads to 'dramatic speedup' of iDMRG.

Finite-size DMRG

Grow chain to some length N using infinite-size DMRG algorithm.

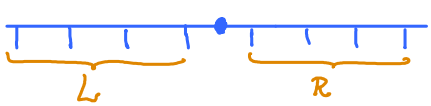

Then reduce L, enlarge R, optimize ψ:

Diagonalize ρ_L, truncate.

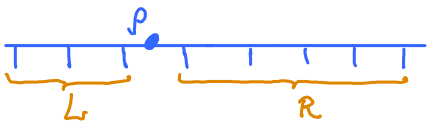

Iterate: sweep back and forth until convergence.

This is conceptually identical to variational optimization with two-site update.

Single-site DMRG is also possible ⟷ variational single-site update.

2. Subspace expansion for 1-site DMRG

'strictly single-site DMRG (DMRG3S)'

DMRG-II.2

Problem: when exploiting quantum numbers, 1-site DMRG performs poorly, because it does not explore subspaces with different quantum numbers. An early remedy for this is 2-site DMRG, but that is computationally much more expensive than 1-site DMRG. Subsequent suggestions for 1-site DMRG with symmetries are 'density matrix perturbation' [White2005], the 'center matrix wave function formalism [McCulloch2007], and, most recently, 'subspace expansion' [Hubig2015], which performs best.

Reminder of 1-site DMRG,
in site-canonical representation:
Local basis: $|\alpha,\sigma,\beta\rangle := |\beta\rangle_{\ell+1}|\sigma\rangle|\alpha\rangle_{\ell-1}$

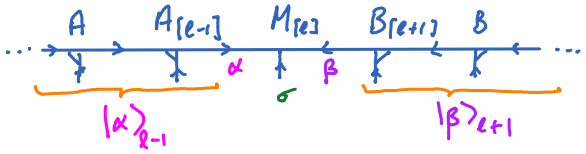

(1)

Minimize energy with constraint of fixed normalization, 1 site at a time:

$$\frac{\partial}{\partial B^{\dagger}_{[\ell]}}\left[\langle\psi|\hat{H}|\psi\rangle - \lambda\langle\psi|\psi\rangle\right] = 0 \qquad (2)$$

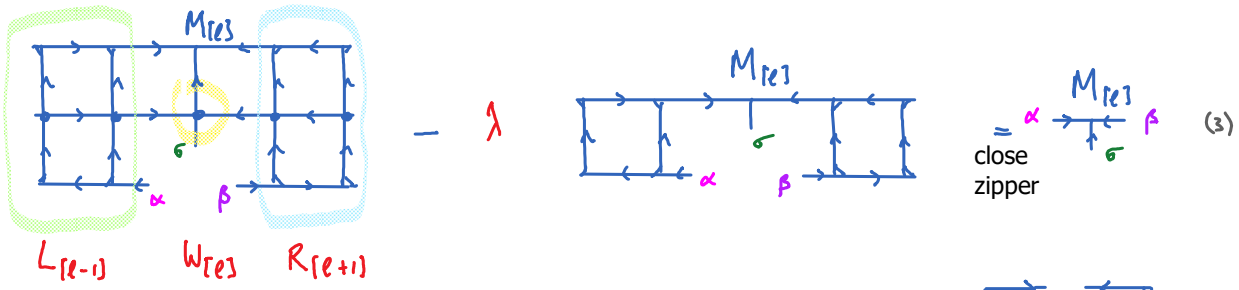

(3)

$$H^{a'}_{[\ell]\,a} M^{a}_{[\ell]} = \lambda M^{a'}_{[\ell]} \qquad a = (\alpha,\sigma,\beta)$$

with $H_{[\ell]} = L_{[\ell-1]} W_{[\ell]} R_{[\ell+1]} =$

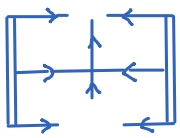
(4)

Solve for 'eigenvector' with lowest eigenvalue, say $\tilde{M}_{[\ell]}$, then do SVD on it to move to next site:

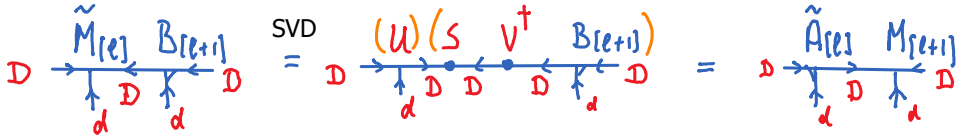

(5)

Important: dimensions of $\tilde{M}_{[\ell]}$ are fixed, hence truncation is neither needed nor possible!

Reminder of 2-site DMRG,
in site-canonical representation:
Local basis: $|\alpha,\sigma,\bar{\sigma},\gamma\rangle := |\gamma\rangle_{\ell+2}|\bar{\sigma}\rangle|\sigma\rangle|\alpha\rangle_{\ell-1}$

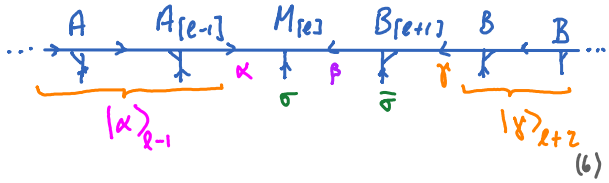

(6)

Minimize energy two sites at a time:

$$H^{a'}_{[\ell,\ell+1]\,a}\left(M_{[\ell]}B_{[\ell+1]}\right)^{a} = \lambda\left(M_{[\ell]}B_{[\ell+1]}\right)^{a'} \qquad a = (\alpha,\sigma,\bar{\sigma},\gamma)$$

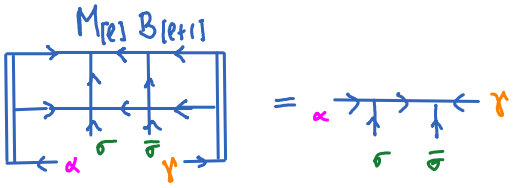

(7)

Solve for 'eigenvector' with lowest eigenvalue, $(\tilde{MB})$, then do SVD and truncate (!) to move to next site:

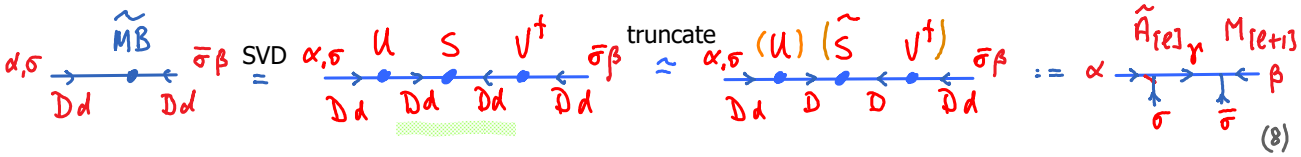

(8)

Problem of single-site optimization: it is constrained to a variational space defined by outgoing state spaces $|\beta\rangle_{\ell+1}\,|\alpha\rangle_{\ell-1}$. If the ranges of quantum numbers Q_α and Q_β for these spaces are too small to accurately represent the ground state, single-site DMRG has no way to enlarge them.

Two-site optimization does not have this problem: the action of H on two sites enlarges bond dimension in between, adding the full range of quantum numbers needed on that bond. If a certain quantum number was missing on that bond before the action of H, but appears afterwards with non-negligible weight, it will survive after SVD and truncation. Hence: two-site optimization can add missing quantum numbers, if needed.

But this comes at a cost: effective two-site Hamiltonian has dimension $D^2d^2 \times D^2d^2$.
By contrast, effective one-site Hamiltonian has dimension $D^2d \times D^2d$. (9)

Example: 6-site chain of S=1/2 spins, exploiting Sz quantum numbers: $\sigma_\ell \in \{+1, -1\}$

(10)

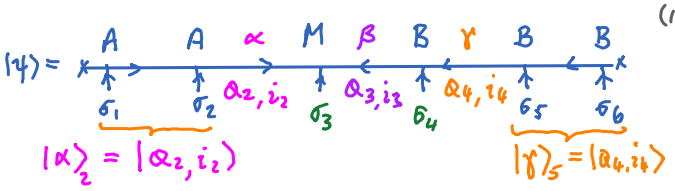

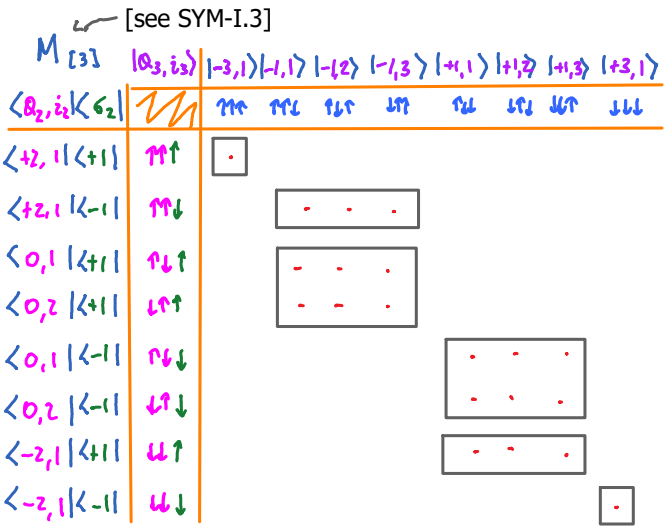

spin conservation: $-(Q_2 + \sigma_3) = Q_3 = \sigma_4 + Q_4$ (11)

$$|\psi\rangle = \sum_{\{\vec\sigma\}} |\vec\sigma\rangle (A^{\sigma_1} A^{\sigma_2})^{\prime} \propto M^{\alpha\sigma_3\beta}_{[3]} B_{[4]\beta}{}^{\sigma_4\gamma} (B^{\sigma_5} B^{\sigma_6})_{\gamma}{}^{\prime}$$

$$= |\gamma\rangle_5 |\sigma_4\rangle |\sigma_3\rangle |\alpha\rangle_2 M^{\alpha\sigma_3\beta}_{[3]} B_{[4]\beta}{}^{\sigma_4\gamma}$$

$$= |Q_4, i_4\rangle |\sigma_3\rangle |Q_2, i_2\rangle M^{Q_2,i_2,\sigma_3,Q_3,i_3}_{[3]} B_{[4]\,Q_3,i_3}{}^{\sigma_4,Q_4,i_4}$$

(12)

In general, the product $M_{[3]}\,B_{[4]}$ has block matrix product structure, where sum over bond index $\beta = (Q_3, i_3)$ matches blocks with same Q_3 :

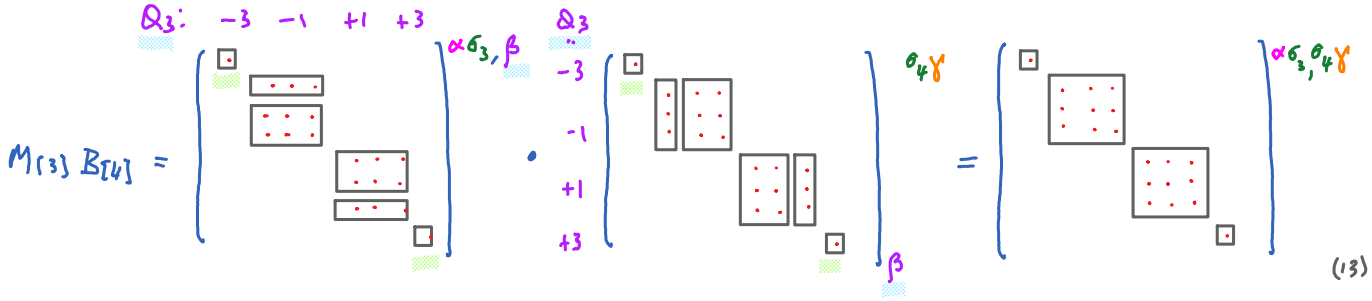

(13)

Bond dimensions: $D_2 \cdot d_3 \times D_3 = 4\cdot 2 \times 8$ $\quad D_3 \times d_4 \cdot D_4 = 8 \times 2\cdot 4$ $\quad 4\cdot 2 \times 4\cdot 2$

But in practice, we use truncated state spaces. Suppose $D_{max} = 4$,

with $D_2 = 4$, $d_3 = 2$, $D_3 = 4$ and that random initial state does not contain $Q_3 = \pm 3$

Then the corresponding two blocks are missing in both $M_{[3]}$ and $B_{[4]}$:

Truncated:

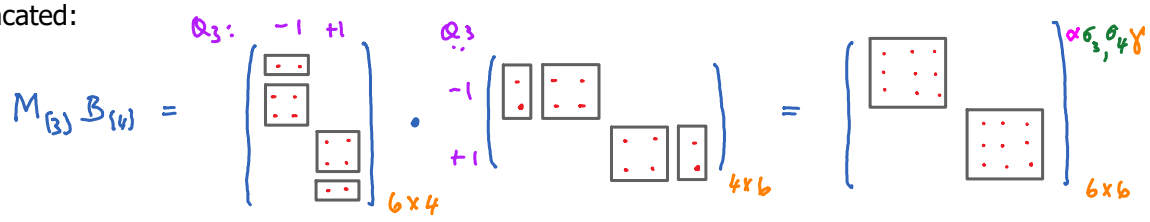

(14)

1-site optimization of truncated $M_{[3]}$ will never find a good ground state if the latter has non-negligible contributions from missing blocks.

2-site optimization reinstates missing block: E.g.:

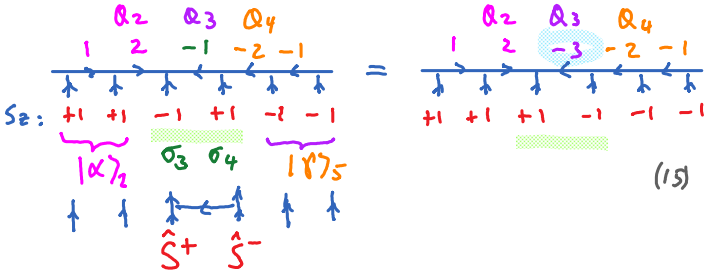

(15)

action of spin-flip term on sites 3 and 4 ⇝ can increases range of Q_3

Subspace expansion

1-site scheme that can 'expand' the variational space to recover missing blocks!

Suppose site ℓ has been optimized by 1-site DMRG, yielding $M_{[\ell]}$, defining 'current ground state' $|\psi_0\rangle$

Conceptual idea: if the variational space $\mathrm{span}\{|\psi\rangle\}$ is too small, expand it to $\mathrm{span}\{|\psi\rangle, \hat{H}|\psi_0\rangle\}$

current approx. for ground state

all MPS with specified bond dimensions and quantum numbers

explores more quantum numbers (except if $|\psi_0\rangle$ is an exact eigenstate)

Concretely: define $|\tilde{\psi}\rangle = |\psi_0\rangle + c\,|\psi_\alpha\rangle$ ($\sim \hat{H}|\psi_0\rangle$) and adjust c to further minimize the energy. (16)

'mixing coefficient'

$|\psi_0\rangle = |\gamma\rangle_{\ell+2}\,|\bar{\sigma}\rangle|\sigma\rangle\,|\alpha\rangle_{\ell-1}\,(M_{[\ell]})^{\alpha\sigma\beta}\,B_{[\ell+1]\beta}{}^{\bar{\sigma}\gamma} =$

shorthand ⇝ $|\alpha,\sigma,\bar{\sigma},\gamma\rangle$

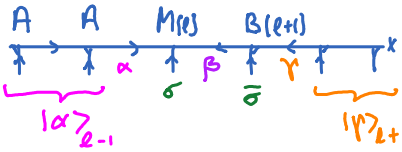

(17)

$|\psi_\alpha\rangle = \hat{P}\hat{H}|\psi_0\rangle =$

projector into local basis $\sum |\alpha,\sigma,\bar{\sigma},\gamma\rangle\langle\alpha,\sigma,\bar{\sigma},\gamma|$

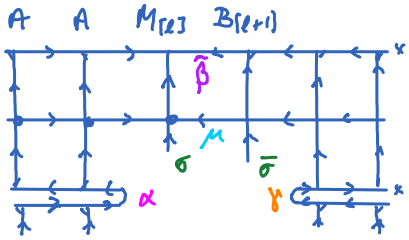

=:

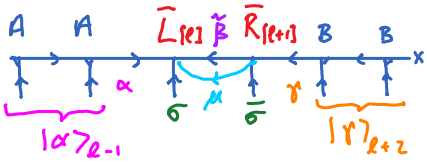

(18)

Crucial point: $\tilde{\beta}$ enumerates all states possibly generated by action of H on sites $\ell, \ell+1$, hence its range can be larger than that of β in (17) for $|\psi_0\rangle$, possibly including new quantum numbers!

Thus, virtual state space on this bond has been 'expanded': $\dim(\tilde{\beta}) = \dim(\alpha)\cdot\dim(\sigma)$

[for an example, see (15)]

$\tilde{D} = Dd$ (19)

For 6-site chain discussed earlier:

for each fixed μ :

For 6-site chain discussed earlier:

$(M_{[\ell]})^{\alpha\sigma\beta} \cdot (B_{[\ell+1]})_{\beta}^{\bar\sigma\gamma}$ and

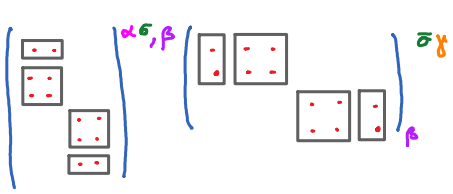

for each fixed μ :

$(\bar L_{[\ell]}^{\mu})^{\alpha\sigma\tilde\beta} \cdot (\bar R_{[\ell+1]\,\mu})_{\tilde\beta}^{\bar\sigma\gamma}$

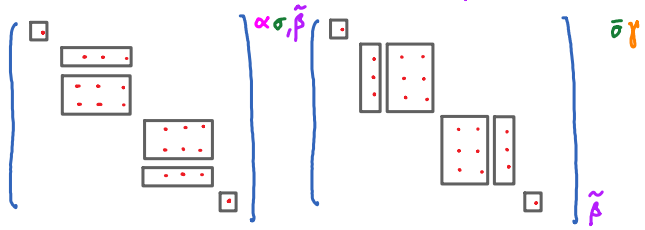

(20)

In (18), we identified coefficients tied to $|\sigma\rangle|\kappa\rangle_{\ell-1}$ and to $|\gamma\rangle|\bar\sigma\rangle$:

$\bar L_{[\ell]}^{\alpha\sigma\mu\tilde\beta} :=$

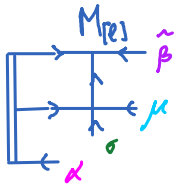

(21a),

$\bar R_{[\ell+1]\,\tilde\beta\mu}^{\bar\sigma\gamma} :=$

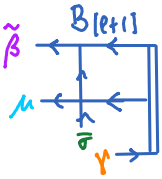

(21b)

$= (L_{[\ell-1]} W_{[\ell]} M_{[\ell]})^{\alpha\sigma,\mu\tilde\beta}$ (22a)

$= (B_{[\ell+1]} W_{[\ell+1]} R_{[\ell+2]})_{\tilde\beta\mu}^{\bar\sigma\gamma}$ (22b)

dimensions: $(D\cdot d)\times(d_W\cdot\tilde D)$ (23)

dimensions: $(D\cdot d)\times(d_W\cdot\tilde D)$

Now we add the two MPS in (16), by yet again expanding virtual state space and block-multiplying their tensors:

$$|\tilde\psi\rangle = |\psi_0\rangle + c|\psi_a\rangle = |\alpha,\sigma,\bar\sigma,\gamma\rangle \underbrace{\left[(M_{[\ell]})^{\alpha\sigma\beta}(B_{[\ell+1]})_{\beta}^{\bar\sigma\gamma} + c\left(\tilde L_{[\ell]}\right)^{\alpha\sigma\mu\tilde\beta}(\tilde R_{[\ell+1]})_{\tilde\beta\mu}^{\bar\sigma\gamma}\right]}_{\left[M_{[\ell]}\;\; c\bar L_{[\ell]}\right]^{\alpha\sigma}_{\tilde{\tilde\beta}}\cdot\begin{bmatrix}B_{[\ell+1]}\\ \bar R_{[\ell+1]}\end{bmatrix}^{\bar\sigma\gamma}_{\tilde{\tilde\beta}} =: (\mathcal{M}\cdot\mathcal{B})^{\alpha\sigma\bar\sigma\gamma}} \quad (24)$$

(25)

$(Dd)\times(Dd)$

$\dim(\tilde{\tilde\beta}) = \dim(\beta) + \dim(\tilde\beta)\cdot\dim(\mu)$ (26)

6-site example:

$\mathcal{M}\cdot\mathcal{B} =$

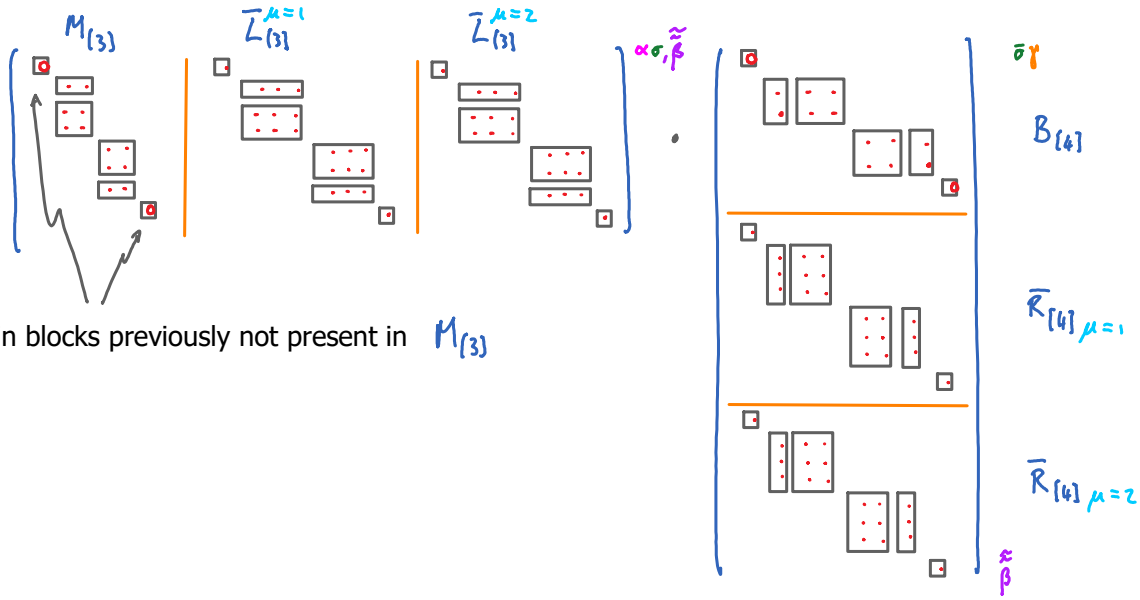

Insert zeros in blocks previously not present in $M_{[3]}$

To deal with large-dimension bond $\tilde{\beta}$, use SVD on $M = U S V^\dagger$ and truncate back to (!!)

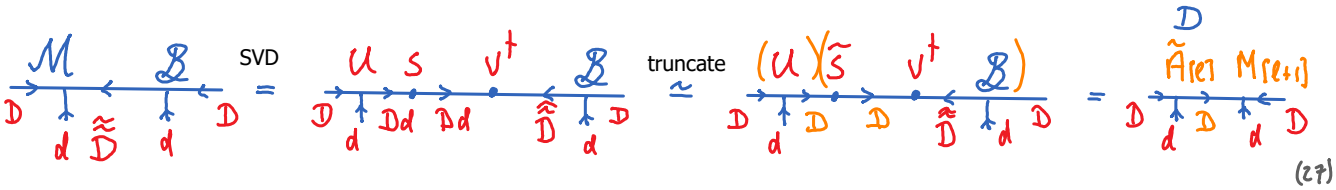

Left-normalized $\tilde{A}_{[\ell]}$ is final result of optimization of site ℓ, and $M_{[\ell+1]}$ initializes optimization of site $\ell+1$.

Crucial point: as mentioned earlier, $\tilde{\beta}$ may include larger range of quantum numbers than β.
If these contribute significant weight to $|\psi_a\rangle$, then they will have significant weight in SVD of (27), hence will survive the truncation in (27), and thus be present in $\tilde{A}_{[\ell]}$.
In this way, subspace expansion succeeds in 'reinstating missing quantum numbers' (without ever needing to do Lanczos with huge effective 2-site Hamiltonian of 2-site DMRG!).

6-site example: if $Q_3 = \pm 3$ carries significant weight, then SVD, followed by trunction with dim(β) = 4, will reinstate those blocks:

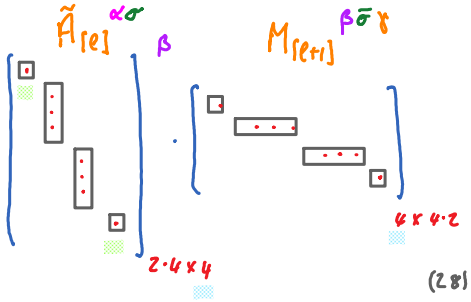

Remark: [Hubig2015] choose $\bar{R}_{[\ell+1]} = 0$ [instead of (22b)], arguing that it only affects $M_{[\ell+1]}$, which will be updated anyway when optimizing site $\ell+1$ via Lanczos ground state solver in next iteration. But using (21b) speeds up convergence for that Lanczos solver [Weichselbaum (private communication)].

Summary of subspace expansion algorithm 'strictly single-site DMRG (DMRG3S)'

Left-to-right sweep:

1. In site-canonical basis for site ℓ, use 1-site DMRG to find optimized $M_{[\ell]}$

   with $H_{[\ell]} M_{[\ell]} = E M_{[\ell]}$ (29)

   ↖ new current energy estimate

2. Expand subspace, $M_{[\ell]} B_{[\ell+1]} \to \mathcal{M}_{[\ell]} \cdot \mathcal{B}_{[\ell+1]} = \begin{pmatrix} M_{[\ell]} & c\bar{L}_{[\ell]} \end{pmatrix} \begin{pmatrix} B_{[\ell+1]} \\ \bar{R}_{[\ell+1]} \end{pmatrix} \approx \tilde{A}_{[\ell]} M_{[\ell+1]}$ (30)

   do SVD + truncation:

3. Obtain 'post-expansion' energy estimate $\tilde{E} =$ (31)

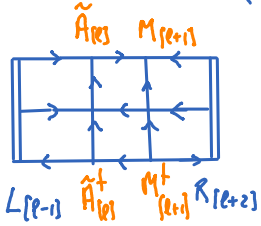

Back to 1. and repeat.

Right-to-left sweep is analogous, with

2. Expand subspace, $A_{[\ell-1]} M_{[\ell]} \to \mathcal{A}_{[\ell-1]} \mathcal{M}_{[\ell]} = \begin{pmatrix} A_{[\ell-1]} & c\bar{L}_{[\ell-1]} \end{pmatrix} \begin{pmatrix} M_{[\ell]} \\ \bar{R}_{[\ell]} \end{pmatrix} \approx M_{[\ell-1]} \tilde{B}_{[\ell]}$ (32)

   do SVD + truncation:

3. Obtain 'post-expansion' energy estimate $\tilde{E}$ =

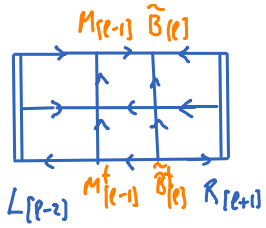

(33)

Choice of mixing coefficient c :

ΔE_0 energy change from 1-site optimization, usually < 0

ΔE_T energy change from truncation, can be either < 0 or > 0.

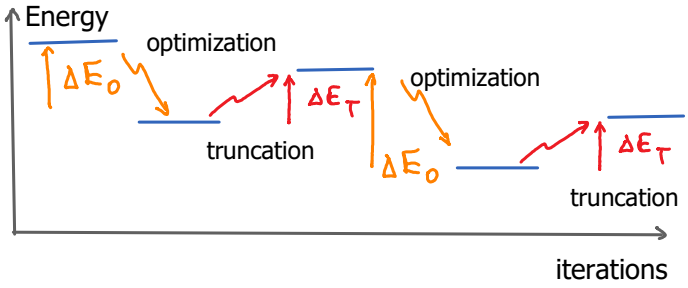

Behavior of energy changes depends on c :

If c is too small, one can get stuck in local minimum, i.e. $\Delta E_0 \simeq 0$

If c is too large, truncation destroys changes made by previous local optimization, and $0 < \Delta E_T \simeq |\Delta E_0|$ (34)

Then reduce c for next iteration on next site.

If ΔE_T is much smaller than $|\Delta E_0|$ or even negative, then one can increase c for next iteration.

In practice, start with $c = 1$ and aim for $\Delta E_T \simeq 0.3 |\Delta E_0|$ (for scheme of [Hubig2015], with $\bar{R} = 0$) and change c by multiplying with some factor greater/smaller than 1 (rather than subtracting/adding). After several DMRG sweeps back and forth through the chain, the 'correct' quantum numbers will typically have been found. So, one can eventually take $c \to 0$, returning to pure 1-site DMRG.

Remark: there may be better ways to choose c : build the 2-dimensional Krylov space $K = \{|\psi_0\rangle, |\psi_a\rangle\}$ and choose c such that $|\tilde{\psi}\rangle = |\psi_0\rangle + c|\psi_a\rangle$ is the ground state of in this space. (35)

This corresponds to the first step of a Lanczos iteration, cf. DMRG-I.2, Eqs. (6-11).

Orthogonalize: $b_1 |\psi_1\rangle := |\tilde{\psi}_1\rangle := \hat{P}\hat{H}|\psi_0\rangle - |\psi_0\rangle \underbrace{\langle\psi_0|\hat{P}\hat{H}|\psi_0\rangle}_{a_0}$ (36)

$= |\psi_a\rangle - a_0 |\psi_0\rangle$ (37)

Computing $|\tilde{\psi}_1\rangle$ involves addition of two MPS, can be done as in (24,25) [no need for SVD + truncation here], to arrive at an MPS with same bond dimensions as $|\psi_0\rangle$.

$b_1 = \sqrt{\langle\tilde{\psi}_1|\tilde{\psi}_1\rangle} = \langle\psi_1|H|\psi_0\rangle$, $\quad a_1 = \langle\psi_1|H|\psi_1\rangle$ (38)

Then diagonalize $H_K = \begin{pmatrix} a_0 & b_1 \\ b_1 & a_1 \end{pmatrix}$, its lower-energy eigenvector yields optimal c. (39)

Then use that to compute a new $|\tilde{\psi}\rangle$, and from that compute $M \cdot B = \tilde{A}_{[l]} M_{[l+1]}$ as described in (24,25,27). Explore whether this yields better convergence behavior!

How can one estimate / control error during extrapolations with increasing bond dimension? See [Hubig2018].