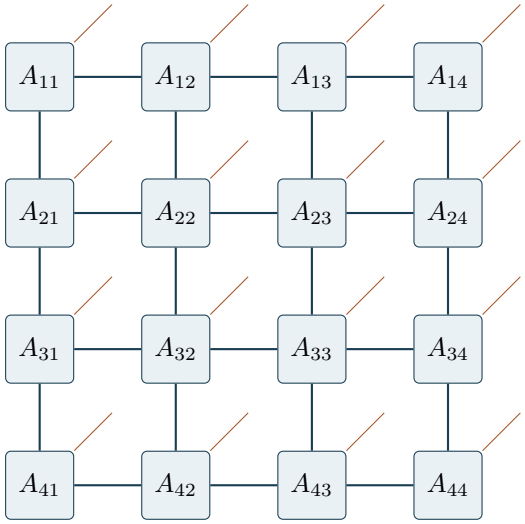
Tensor Network Methods

Andrew Hallam



Contents

1 Why should anyone care about tensor networks? 1

2 Tensor contraction and tensor-networks 4

3 AKLT and the matrix product representation 7

4 Properties of MPS 11

5 Contractions and transfer operators 19

References 26

1 Why should anyone care about tensor networks?

Tensor-network methods are currently one of the most powerful approaches for numerically simulating and analytically understanding complex quantum systems. The central idea behind tensor networks is simple: break the large vectors that represent quantum wavefunctions into smaller objects called *tensors*, and then repeatedly *locally* optimise these tensors to achieve the particular task you have in mind. This simple idea has proven to be incredibly flexible, leading to many tensor networks of different shapes and sizes and efficient numerical algorithms for ground-state optimisation, non-equilibrium quantum dynamics, simulations of quantum circuits, and even applications in areas such as machine learning and time evolution of PDEs (wow!).

Before jumping into tensor networks themselves, I'd like to give a few pieces of motivation for where tensor networks came from and why people in physics and beyond care about them.

1.1 Getting beyond Landau–Ginzburg theory

The traditional approach to studying equilibrium condensed matter systems was built upon the work of Landau and Ginzburg, where phases are understood through their underlying symmetries and how those symmetries are spontaneously broken. A common starting point is to expand an order parameter around its mean-field value and then study fluctuations perturbatively:

$$\widehat{O} = \langle \widehat{O} \rangle + \delta \widehat{O}. \tag{1.1}$$

When fluctuations are sufficiently small, this expansion gives an accurate description of the phase. Mean-field approaches, together with effective quasiparticle descriptions, were highly successful from the 1950s to the 1970s, with BCS theory of superconductivity and Fermi-liquid theory of metals as notable examples.

It later became clear that many interesting condensed matter phases could not be described by symmetry breaking alone. A striking example is topological order, which cannot be characterised by a local symmetry-breaking order parameter and therefore lies beyond the traditional Landau–Ginzburg classification. The AKLT state was introduced in 1987 by Affleck, Kennedy, Lieb and Tasaki as a concrete realisation of a particular type of symmetry-protected topological order. The use of tensor-networks as concrete realisations of exotic quantum phases remains one of the key motivations for studying them: they provide explicit, analytically tractable wavefunctions that capture physics that semiclassical methods based on local order parameters can miss.

1.2 Entanglement area laws

A second motivation for tensor-network states comes from applying ideas from quantum-information theory to large quantum systems. A generic state of N spin-1/2 particles requires 2^N coefficients to be specified. It therefore seems like it should be completely hopeless to even begin to analyse such a system. However, *physically interesting* states are *not* generic: they are highly unusual! We can see this by asking how entangled a region A is with the rest of the system. For a pure state $|\psi\rangle$, we obtain the reduced density matrix of A by tracing out its complement $\bar{A}$, and define the entanglement entropy:

$$\rho_A = \mathrm{Tr}_{\bar{A}} |\psi\rangle\langle\psi|, \qquad S(A) = -\mathrm{Tr}(\rho_A \log \rho_A). \tag{1.2}$$

It turns out that generic states are highly entangled: you would expect every site of region A to be highly entangled with sites outside of region A. In these states $S(A)$ grows with the number of sites in A and we say that it satisfies an entanglement *volume law*. If only sites close to the

boundary between A and its surroundings are entangled, then the entanglement will instead grow like the size of the boundary of A, and we say that this state satisfies an entanglement *area law*. In Figure 1 an illustrative example of the two scalings for a $2D$ system is shown. Let's briefly dwell on the case of one-dimensional systems. In this instance, when A is a contiguous l-site region, an entanglement volume law means $S(A) \propto l$ and an entanglement area law means $S(A)$ is constant, regardless of the size of A!

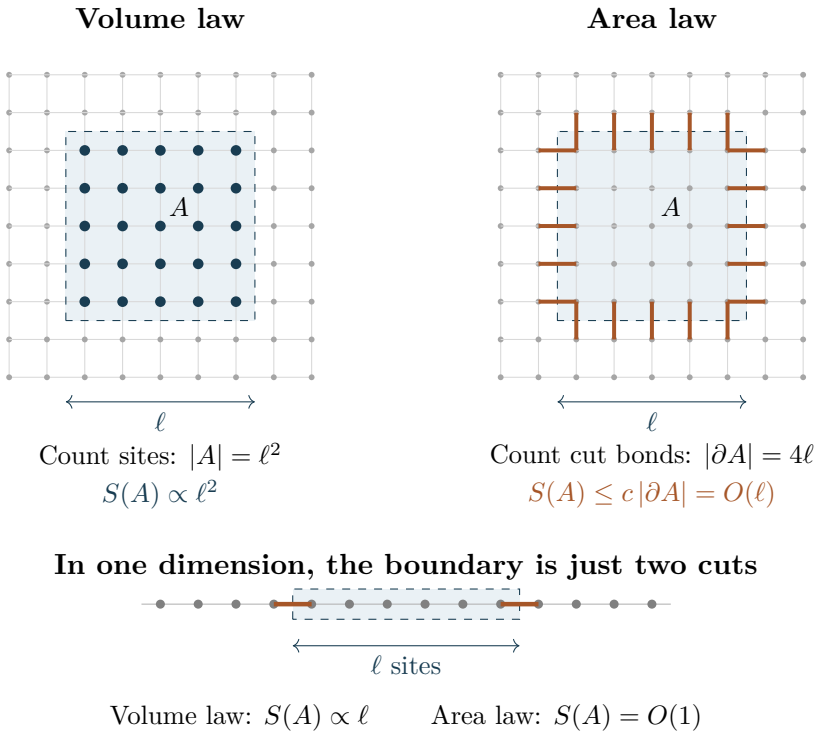


Figure 1: Volume versus boundary scaling of entanglement. The shaded region A is the same in both square lattices: a volume law scales with all its sites (left), whereas an area law is bounded by the number of bonds crossing its boundary (right), with a size-independent constant c. In two dimensions, "area" means the size of the perimeter. For a one-dimensional system, an area-law is size independent.

As I said above, we know that many physically interesting states satisfy entanglement area laws. In particular, it has been shown that the ground states of gapped, one-dimensional systems with finite-range interactions have to satisfy an entanglement area law [9]. In general, similar statements are also believed to be true in higher dimensions.

This insight is so important because tensor-networks naturally represent states with a limited amount of entanglement. Matrix product states (MPS) with bounded bond dimension obey a one-dimensional entanglement area law and provide accurate approximations to ground states of gapped local chains [9, 10]. They are therefore an incredibly powerful tool for describing the ground states of complex quantum systems.

1.3 What is DMRG anyway?

We have seen that tensor networks can represent important classes of states, but this is only half of the motivation for their use. Just as significantly, there are powerful algorithms for manipulating and optimising these states. The simplest numerical method for finding the ground state of a many-body quantum system, exact diagonalisation (ED), simply involves representing

the particular Hamiltonian as a matrix which is then, in the most naive case, diagonalised in order to obtain all of the energy eigenstates. However, since the Hamiltonian is exponentially large in the number of sites N, ED methods for studying quantum systems rapidly become completely intractable.

For this reason, there has been a long-standing interest in more sophisticated methods for numerically studying many-body quantum systems that sidestep this exponential scaling. One such method introduced by Steven White is the density matrix renormalisation group (DMRG) algorithm for finding ground states of one-dimensional spin chains. Remarkably, it transformed the exponential scaling of ED to *linear* scaling in the system size N! It was soon realised that DMRG can be thought of as a variational method in the space of matrix-product states, which was the basis for its huge success. The ideas used to develop the DMRG algorithm were subsequently generalised to a wide range of new tensor-network algorithms—all with equally unintuitive acronyms—which are probably the main reason for the interest in tensor-network algorithms today.

2 Tensor contraction and tensor-networks

For a system of N sites, a general pure state can be written as

$$|\psi\rangle = \sum_{s_1,\ldots,s_N} C_{s_1\cdots s_N} |s_1\rangle \otimes \cdots \otimes |s_N\rangle, \tag{2.1}$$

where s_i labels a basis state at site i. These site labels do not assume any particular spatial geometry.

We usually think of the coefficients $C_{s_1\cdots s_N}$ as the entries of a single vector in the many-body Hilbert space. Alternatively, we can view them as an N-dimensional array, with d entries along each axis, which we call a *tensor*. It's important to remember that this is just a different way of organising the same coefficients.

A *tensor network* is simply a decomposition of C into smaller tensors. Just as matrices can be multiplied to produce new matrices, tensors can be *contracted* to produce new tensors. In a tensor-network representation of a wavefunction, various tensors are contracted to produce C. Writing these tensor contractions explicitly can quickly become cumbersome, so a helpful diagrammatic notation that makes the tensors and their manipulations easier to see has been introduced.

2.1 Tensor-networks and tensor-network diagrams

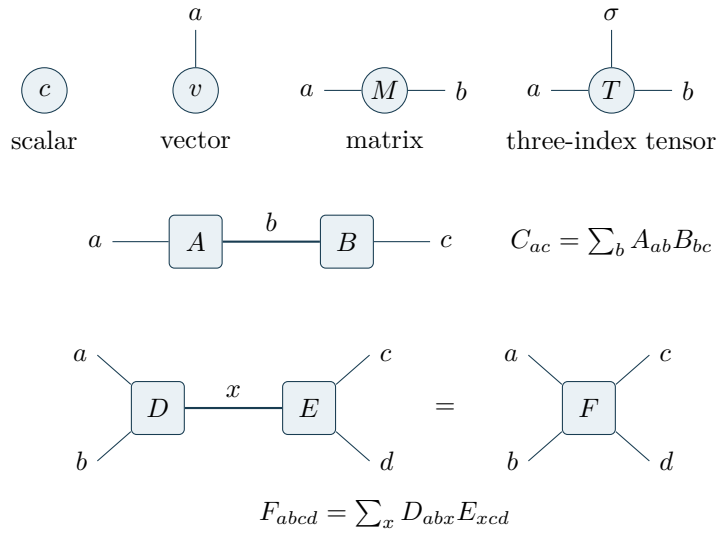


Figure 2: Tensor notation and contraction. The middle row shows matrix multiplication. In the bottom row, contracting the shared index x of two three-leg tensors produces a four-leg tensor F; the open indices a, b, c, d remain.

In a tensor-network diagram, tensors are represented as nodes with a number of legs sticking out, each of the legs representing an individual index of the tensor (see Fig. 2). A scalar has no indices, a vector has one, and a matrix has two. The number of legs counts the indices; the dimension of each leg counts the possible values of its index. Tensor contraction is summing over legs of particular tensors in order to produce new tensors, for example multiplying two matrices A and B to produce a new matrix C can be written as

$$C_{ac} = \sum_b A_{ab}B_{bc}. \tag{2.2}$$

In a tensor-network diagram this is depicted as a two-leg A tensor with one of its legs connected to a two-leg B tensor, producing an overall two-leg object, C. Of course, it would be total overkill to introduce a whole new graphical notation to just do matrix-matrix and matrix-vector multiplication, so let's think about contracting one leg of two three-leg tensors D and E,

$$F_{abcd} = \sum_x D_{abx} E_{xcd}. \tag{2.3}$$

Since four legs of D and E remain uncontracted, this produces F, a four-leg tensor.

In practice, how do we actually do tensor contractions? We are always free to reshape tensors by grouping (fusing) legs to create a larger leg whose dimension is the product of all the fused legs. Similarly, we can always separate (split) one leg into several smaller legs, provided the product of the dimensions of the split legs matches the dimension of the original leg. Therefore, we can always reshape tensors with a high number of legs into matrices. With this tool at our disposal it's clear how to perform tensor contractions on a computer. For each tensor, group together the legs to be contracted and, separately, the uncontracted legs. The tensor contraction is then simply matrix multiplication. More generally, this means we can perform *any* matrix operation on a tensor, such as singular value decomposition (SVD), or eigenvalue decomposition when the reshaped matrix is square. We will use SVD to construct MPS and explain their entanglement in Section 4.2.

2.2 Contraction order matters

An important detail to remember when performing tensor contraction is that the cost can depend upon the order in which we contract. Consider the three-tensor network (see Fig. 3)

$$T_{pq} = \sum_{a,b,c} X_{ab} Y_{apc} Z_{bcq}, \tag{2.4}$$

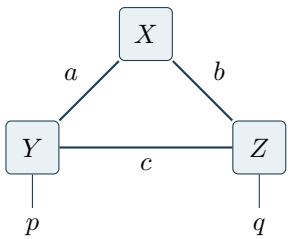
with every index of dimension D. Contracting X and Z first gives

$$K_{acq} = \sum_b X_{ab} Z_{bcq}, \qquad T_{pq} = \sum_{a,c} Y_{apc} K_{acq}. \tag{2.5}$$

Both steps naively cost $O(D^4)$ with ordinary dense contractions, and the largest intermediate has D^3 entries. Contracting Y and Z first instead produces

$$Q_{apbq} = \sum_c Y_{apc} Z_{bcq}, \qquad T_{pq} = \sum_{a,b} X_{ab} Q_{apbq}. \tag{2.6}$$

The first step costs $O(D^5)$ and stores D^4 entries; the final step costs $O(D^4)$. Despite producing the same result, this tensor contraction is significantly less efficient.



First contract X with Z:
$O(D^4) + O(D^4)$.

First contract Y with Z:
$O(D^5) + O(D^4)$.

Figure 3: Two different ways of contracting produce the same output but can have different scaling.

2.3 Choosing a network geometry

We can now return to representing quantum many-body wavefunctions in terms of tensor-networks, by choosing a representation of the N-leg tensor C in terms of smaller tensors. While we're free to choose any graph we like, typically we want to choose a regular structure for theoretical and numerical reasons. The simplest choice is a matrix-product state (MPS) which is a chain of 3-leg tensors, with two legs being contracted to the left and the right, and a free leg being one of the free legs of C. For open boundary conditions, this tensor network is terminated by a two-leg tensor at each end of the chain; for periodic boundary conditions, the first and last three-leg tensors are simply contracted.

An MPS is the most natural tensor network for 1D systems, and the one we will consider for the rest of the notes, but other choices of tensor-network exist. A tree tensor network has a hierarchical structure of tensors instead, and is often useful for representing states that are more entangled than an area law. A projected entangled-pair state (PEPS) is the generalisation of MPS to higher-dimensional lattices. For example, on a square lattice each bulk PEPS tensor has five legs: four connect to its neighbours, and one is a physical leg (see Fig. 4).

Each internal index has a dimension, called its *bond dimension*. These dimensions determine the number of coefficients stored in the tensors and constrain the entanglement the network can represent. Choosing a network geometry is therefore only part of specifying the ansatz: we must also choose its bond dimensions.

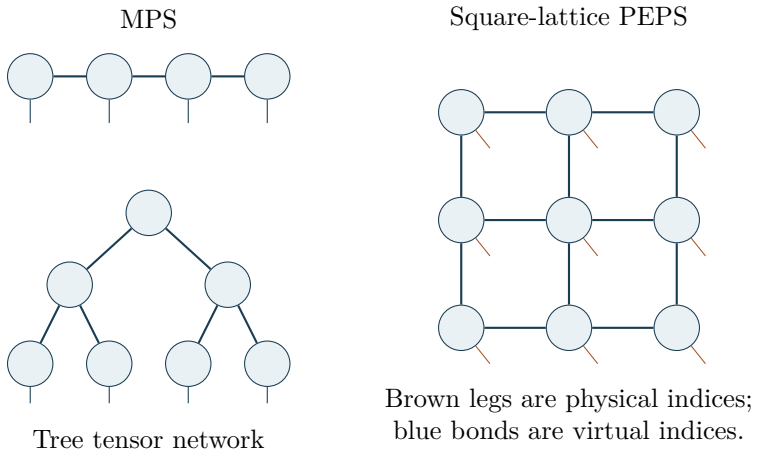


Figure 4: There are many different types of tensor-networks, including matrix-product states (MPS), projected entangled-pair states (PEPS), and tree tensor networks (TTNs).

3 AKLT and the matrix product representation

3.1 The Haldane phase and AKLT model

As discussed above, tensor networks were originally introduced to study exotic quantum phases. We will now consider the case of the *Haldane phase*, one of the first examples of symmetry-protected topological (SPT) order. This story begins with a debate about the ground state of the uniform antiferromagnetic Heisenberg chain,

$$H = J\sum_i \mathbf{S}_i \cdot \mathbf{S}_{i+1}, \qquad J > 0. \tag{3.1}$$

It was well known that the ground state of the model was gapless for spin-1/2, due to the Bethe ansatz, and higher half-integer spins, due to the Lieb-Schultz-Mattis theorem. For a long time it was unknown if the integer-spin Heisenberg model was also gapless. This was resolved in 1983 when Haldane predicted that integer-spin systems would instead be gapped but would *not* display conventional antiferromagnetic order [2]. The spin-1 Haldane phase has no conventional magnetic long-range order, but has nontrivial symmetry-protected topological structure, leading to exotic properties such as ground-state degeneracy that depends upon the topology of the system due to fractionalised gapless edge modes. Haldane's calculation was based on a field theory calculation, so it remained unclear how to understand the Haldane phase on the lattice. For this reason the AKLT model was introduced as an exactly solvable spin-chain representation of the Haldane phase [3].

To find the AKLT Hamiltonian, we begin by noticing that two neighbouring spin-1 sites can combine to give total spin $J_{\text{tot}} = 0, 1, 2$. Let $P^{(2)}_{i,i+1}$ be the orthogonal projector onto their total spin-2 subspace. We can then define the AKLT Hamiltonian in terms of this projector:

$$H_{\text{AKLT}} = \sum_i P^{(2)}_{i,i+1}. \tag{3.2}$$

Each term is positive and penalises the total spin-2 component on every neighbouring pair. Because the Hamiltonian is a sum of these projector terms, if we can construct a state with zero total spin-2 on each neighbouring pair of sites, it has zero energy and must be a ground state.

How does this Hamiltonian relate to the Heisenberg model? To relate these models, write $x_i = \mathbf{S}_i \cdot \mathbf{S}_{i+1}$. Since each site has $\mathbf{S}_i^2 = 2$,

$$x_i = \frac{1}{2}\left[\mathbf{J}^2 - \mathbf{S}_i^2 - \mathbf{S}_{i+1}^2\right] = \tfrac{1}{2}\left[(\mathbf{S}_i + \mathbf{S}_{i+1})^2 - 4\right]. \tag{3.3}$$

Its eigenvalues in the $J_{\text{tot}} = 0, 1, 2$ sectors are therefore $-2, -1, 1$, respectively. To relate this to the total spin-2 projector, we need a polynomial that vanishes at -2 and -1 and equals one at 1:

$$P^{(2)}_{i,i+1} = \frac{(x_i + 2)(x_i + 1)}{6} = \frac{1}{2}\left(x_i + \frac{1}{3}x_i^2 + \frac{2}{3}\right). \tag{3.4}$$

Thus, up to a positive overall scale and an additive constant, the AKLT Hamiltonian is the Heisenberg Hamiltonian with an additional biquadratic interaction:

$$H_{\text{bilinear-biquadratic}} = J\sum_i \left[\mathbf{S}_i \cdot \mathbf{S}_{i+1} + \frac{1}{3}(\mathbf{S}_i \cdot \mathbf{S}_{i+1})^2\right] = 2JH_{\text{AKLT}} - \frac{2J}{3}N_b, \qquad J > 0, \tag{3.5}$$

where N_b is the number of bonds. In particular, the two conventions have the same ground states. We use the projector convention so that the ground state has zero energy.

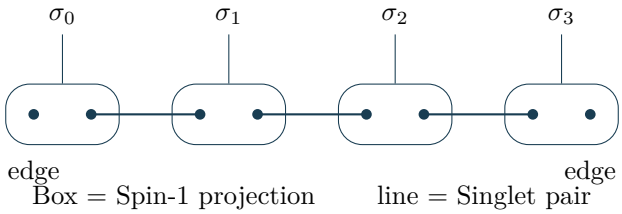


Figure 5: AKLT valence-bond construction. Each box combines two virtual spins into a physical spin 1. Open boundaries leave two effective spin-1/2 edge degrees of freedom.

3.2 Constructing the AKLT ground state

We will now try to construct the ground state of the AKLT model by finding a state that is annihilated by every local projector term. To do this, we imagine that each spin-1 is secretly two spin-1/2 degrees of freedom that are then fused together to create the physical spin-1 states. Now let's place a singlet between one virtual spin-1/2 at site i and one at site $i+1$:

$$|\omega\rangle = \frac{|\uparrow\downarrow\rangle - |\downarrow\uparrow\rangle}{\sqrt{2}}. \tag{3.6}$$

Then we want to fuse together these singlets so that at each physical site we get a spin-1. To do this, we use the following map:

$$|+1\rangle = |\uparrow\uparrow\rangle, \quad |0\rangle = \frac{|\uparrow\downarrow\rangle + |\downarrow\uparrow\rangle}{\sqrt{2}}, \quad |-1\rangle = |\downarrow\downarrow\rangle. \tag{3.7}$$

The important thing to remember is that singlets are introduced *between* sites, and the projectors are introduced *on* sites, as shown in Fig. 5. Remarkably, it turns out this state is annihilated by all of the total spin-2 projectors. To see this, remember that this big spin is constructed of four of the underlying spin-1/2 degrees of freedom we introduced. But we've put the central two spin-1/2 degrees of freedom into a singlet, which has total spin zero. The remaining two virtual spins can carry at most total spin 1. Also, projecting each site onto its symmetric spin-1 subspace commutes with simultaneous spin rotations, so it cannot create a total-spin-2 component. Consequently,

$$P^{(2)}_{i,i+1}|\psi_{\mathrm{AKLT}}\rangle = 0 \quad \text{for every bond,} \qquad H_{\mathrm{AKLT}}|\psi_{\mathrm{AKLT}}\rangle = 0. \tag{3.8}$$

The state is therefore the exact ground state of the AKLT Hamiltonian. This Hamiltonian is *frustration free* in the sense that this state minimises every local term separately.

3.3 MPS representation of the AKLT ground state

We will now turn the singlets and on-site projections into tensors, then contract them to obtain the coefficients of the AKLT ground state. This will be our first explicit example of an MPS.

Start with a single singlet. Its coefficients can be arranged in a matrix:

$$|\omega\rangle = \sum_{\alpha,\beta} Q_{\alpha\beta}|\alpha\rangle|\beta\rangle, \qquad Q = \frac{1}{\sqrt{2}}\begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}. \tag{3.9}$$

Rows and columns are ordered as $(\uparrow,\downarrow)$. The row index labels the virtual spin on the left end of the bond, and the column index labels the one on the right.

At each site, the map from the two virtual spin-1/2 degrees of freedom to the physical spin-1 is

$$\mathcal{P} = |+1\rangle\langle\uparrow\uparrow| + \frac{|0\rangle}{\sqrt{2}}(\langle\uparrow\downarrow| + \langle\downarrow\uparrow|) + |-1\rangle\langle\downarrow\downarrow|. \tag{3.10}$$

With rows ordered as $(+1, 0, -1)$ and columns as $(\uparrow\uparrow, \uparrow\downarrow, \downarrow\uparrow, \downarrow\downarrow)$, it is the matrix

$$\mathcal{P} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1/\sqrt{2} & 1/\sqrt{2} & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}. \tag{3.11}$$

Alternatively, keep the two virtual indices separate and define $P^\sigma_{\alpha\beta} = \langle\sigma|\mathcal{P}|\alpha\beta\rangle$. For each physical value σ, reshaping the corresponding row into a 2×2 matrix gives

$$P^+ = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}, \qquad P^0 = \frac{1}{\sqrt{2}}\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \qquad P^- = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}. \tag{3.12}$$

Together, these three matrices are one three-leg tensor: two virtual legs and one physical leg.

We can now construct the overall wavefunction through the tensor contraction of P and Q everywhere. The AKLT wavefunction is found by contracting the virtual α, β legs of $P^\sigma_{\alpha\beta}$ with two different Q matrices. The uncontracted legs of the Q matrices are then contracted with different P tensors to build up a large physical state, shown in Fig. 6. We can write the coefficient matrix $C_{\sigma_1 \cdots \sigma_N}$ through repeated contractions of the Q and P tensors. For any fixed choice of the physical indices $\sigma_1, \ldots, \sigma_N$, which selects a particular coefficient of C, we are simply taking the *product* of P and Q matrices. By working with a periodic system, we can close the final virtual bond by summing the first and last indices together, giving the following trace:

$$C_{\sigma_1 \cdots \sigma_N} = \mathrm{Tr}(P^{\sigma_1} Q \cdots P^{\sigma_N} Q). \tag{3.13}$$

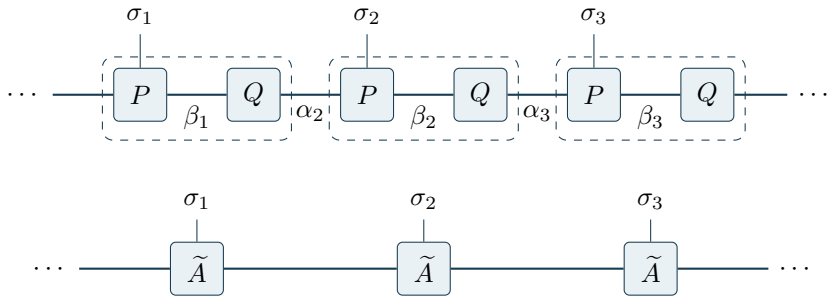


Figure 6: From valence bonds to an MPS. Each dashed box groups an on-site projection tensor P with a singlet tensor Q, contracting the index β_i. The resulting tensor $\tilde{A}$ has one physical leg and two virtual legs of dimension two.

We can absorb the singlet to the right of each site into its projection tensor:

$$\tilde{A}^\sigma = P^\sigma Q, \qquad \tilde{A}^+ = \frac{1}{\sqrt{2}}\begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}, \quad \tilde{A}^0 = \frac{1}{2}\begin{pmatrix} -1 & 0 \\ 0 & 1 \end{pmatrix}, \quad \tilde{A}^- = -\frac{1}{\sqrt{2}}\begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix}. \tag{3.14}$$

Finally, for reasons we'll explain below in Section 5.4, we use the conventionally normalised tensors $A^\sigma = (2/\sqrt{3})\tilde{A}^\sigma$:

$$A^+ = \sqrt{\frac{2}{3}}\begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}, \quad A^0 = \frac{1}{\sqrt{3}}\begin{pmatrix} -1 & 0 \\ 0 & 1 \end{pmatrix}, \quad A^- = -\sqrt{\frac{2}{3}}\begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix}. \tag{3.15}$$

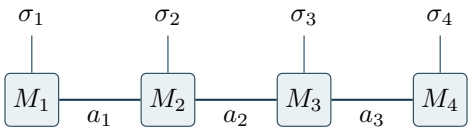
Therefore, we have written the AKLT ground state for a periodic system in terms of these locally defined three-leg A tensors:

$$|\psi_{\mathrm{AKLT}}\rangle = \sum_{\sigma_1,\ldots,\sigma_N} \mathrm{Tr}(A^{\sigma_1}\cdots A^{\sigma_N})|\sigma_1\cdots\sigma_N\rangle, \tag{3.16}$$

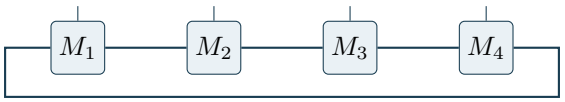
This is a *matrix product state*: each coefficient is the trace of a product of matrices, with one matrix selected by the physical spin at each site. The size of the matrices is called the *bond dimension*, which for the AKLT ground state is two because each virtual index takes the values $\uparrow$ and $\downarrow$.

4 Properties of MPS

4.1 From AKLT to a general MPS



Open chain: boundary bonds have dimension one.



Periodic chain: the virtual indices are traced.

Figure 7: MPS for open and periodic boundary conditions. Physical legs are uncontracted while all virtual bonds are contracted.

In the previous section, we saw that the AKLT ground state could be written as a tensor contraction of a collection of rank-3 tensors $A^{\sigma_i}_{\alpha\beta}$, each tensor representing one physical site, i, of the system. A particular coefficient of the wavefunction was found by fixing the physical legs of the rank-3 tensors and then taking the product of N matrices. To generalise the AKLT construction, we allow the rank-3 tensors $A^{\sigma_i}_{\alpha\beta}$ to vary from site to site, both in their size and in the value of each coefficient. An open-boundary MPS can be written as

$$|\psi\rangle = \sum_{\boldsymbol{\sigma}} A_1^{\sigma_1} A_2^{\sigma_2} \cdots A_N^{\sigma_N} |\boldsymbol{\sigma}\rangle, \tag{4.1}$$

where $A_i^{\sigma_i}$ is a $\chi_{i-1} \times \chi_i$ matrix and $\chi_0 = \chi_N = 1$. The product is therefore a scalar. Boundary vectors can equivalently be written explicitly as $v_L^\dagger A_1^{\sigma_1} \cdots A_N^{\sigma_N} v_R$. For periodic boundaries a common form is

$$|\psi\rangle = \sum_{\boldsymbol{\sigma}} \mathrm{Tr}(A_1^{\sigma_1} A_2^{\sigma_2} \cdots A_N^{\sigma_N}) |\boldsymbol{\sigma}\rangle. \tag{4.2}$$

The trace closes the virtual chain into a loop (see Fig. 7). OBC MPS are much more convenient numerically, so we will focus on open chains in the rest of the notes. The cost of storing an MPS is $\sum_i d\chi_{i-1}\chi_i = O(Nd\chi^2)$ for bounded χ. This is linear in N provided that we can represent the state we want at a fixed bond dimension. As we will see below, whether this is possible depends upon how entangled the particular state is.

4.2 Bond dimension, entanglement, and the Schmidt decomposition

What is the physical meaning of the bond dimension of a matrix-product state? It's clear that as you let the bond dimension become larger and larger you are representing a larger class of states, but we can be more quantitative by calculating the entanglement entropy of MPS of a particular bond dimension. We begin by cutting an open-chain MPS between sites n and $n+1$. Contracting all tensors to the left gives a family of states $|L_a\rangle$, one for each value of the exposed virtual index. The right half similarly gives $|R_a\rangle$. Reconnecting the bond sums over that index:

$$|\psi\rangle = \sum_{a=1}^{\chi_n} |L_a\rangle |R_a\rangle. \tag{4.3}$$

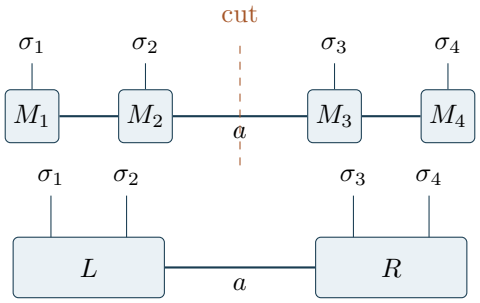


Figure 8: Cutting an MPS bond, illustrated for four sites with $n = 2$. Grouping the tensors on each side produces two block tensors.

For example, the coefficients of $|L_a\rangle$ are $(A_1^{\sigma_1} \cdots A_n^{\sigma_n})_{1a}$. The right block is constructed analogously from sites $n+1, \ldots, N$. This can be shown graphically in Fig. 8. These block states need not be normalised, orthogonal, or even linearly independent.

Can we instead choose orthonormal block states and assign a weight to each pair? To do this, combine the physical indices on the left into $\ell = (\sigma_1, \ldots, \sigma_n)$ and those on the right into $r = (\sigma_{n+1}, \ldots, \sigma_N)$. The coefficient tensor becomes a $d^n \times d^{N-n}$ matrix $C_{\ell r}$. Its compact singular value decomposition (SVD) is

$$C = USV^\dagger, \qquad S = \operatorname{diag}(s_1, \ldots, s_{r_n}), \qquad s_1 \geq \cdots \geq s_{r_n} > 0, \tag{4.4}$$

where $r_n = \operatorname{rank} C$ and $U^\dagger U = V^\dagger V = \mathbb{I}_{r_n}$. Define

$$|a_L\rangle = \sum_\ell U_{\ell a} |\ell\rangle, \qquad |a_R\rangle = \sum_r (V^\dagger)_{ar} |r\rangle. \tag{4.5}$$

The orthonormal columns of U and rows of $V^\dagger$ give orthonormal block states, so

$$|\psi\rangle = \sum_{a=1}^{r_n} s_a |a_L\rangle |a_R\rangle, \qquad \langle a_L | b_L\rangle = \langle a_R | b_R\rangle = \delta_{ab}. \tag{4.6}$$

This is the *Schmidt decomposition*. The Schmidt states generally differ from the original $|L_a\rangle, |R_a\rangle$ in Eq. (4.3). Since that expression already writes the state as a sum of χ_n product states,

$$r_n \leq \min(\chi_n, d^n, d^{N-n}). \tag{4.7}$$

For a normalised state, $\sum_a s_a^2 = 1$. Tracing out the right block gives the reduced density matrix for the left $\rho_L = \sum_a s_a^2 |a_L\rangle\langle a_L|$, from which we can calculate the von Neumann entanglement entropy

$$S_E(L) = -\sum_a s_a^2 \log s_a^2 \leq \log r_n \leq \log \chi_n, \tag{4.8}$$

where we saturate this bound on the entropy when all nonzero Schmidt weights are equal. We have now arrived at a physical intuition as to the meaning of the bond dimension of an MPS: a bounded bond dimension bounds the entanglement across a single cut logarithmically. In particular, the bond dimension determines the number of non-zero Schmidt coefficients across any particular cut of the system. A product state has $\chi = 1$ because it has just one Schmidt coefficient; a Bell pair has two Schmidt coefficients with value $a = 1/\sqrt{2}$, hence the AKLT state has bond dimension $\chi = 2$. An interior interval cuts two virtual bonds: its Schmidt rank is at most the product of their dimensions, giving $S_E \leq 2 \log \chi$ when both dimensions are at most χ. The same argument applies to an interval on a periodic chain and establishes that MPS do indeed represent states which satisfy the one-dimensional entanglement area law.

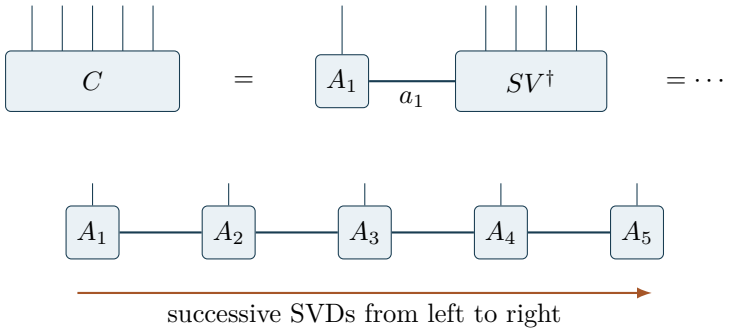


Figure 9: An MPS can be constructed through repeated SVDs. Each SVD separates one physical index and introduces a virtual bond.

4.3 Turning an arbitrary state into an MPS

We have seen how an MPS of fixed bond dimension has a limited amount of entanglement. We will now show that any state can be represented as an MPS, provided the bond dimension is big enough. We will do this by applying the SVD repeatedly across the system, following the construction shown in Fig. 9. First reshape the coefficient tensor as $C_{\sigma_1,[\sigma_2\cdots\sigma_N]}$, separating the first leg from all the others which are grouped together, and take its SVD. Set the first tensor equal to the U obtained from the SVD, $U = A_1^{\sigma_1}$, and leave the remaining parts of the SVD $SV^\dagger = \widetilde{C}_{a_1\sigma_2\cdots\sigma_N}$ to the right:

$$C_{\sigma_1\cdots\sigma_N} = \sum_{a_1} (A_1^{\sigma_1})_{1a_1} \widetilde{C}_{a_1\sigma_2\cdots\sigma_N}. \tag{4.9}$$

For the second site, reshape the remaining tensor $\widetilde{C}$ into a matrix with (a_1, σ_2) as its row index and $(\sigma_3, \ldots, \sigma_N)$ as its column index. Taking its SVD gives

$$\widetilde{C}_{[a_1\sigma_2],[\sigma_3\cdots\sigma_N]} = \sum_{a_2} (U_2)_{[a_1\sigma_2],a_2} (S_2 V_2^\dagger)_{a_2,[\sigma_3\cdots\sigma_N]}. \tag{4.10}$$

We reshape U_2 into the second site tensor and carry $S_2V_2^\dagger$ to the right:

$$(A_2^{\sigma_2})_{a_1a_2} = (U_2)_{[a_1\sigma_2],a_2}, \qquad \widetilde{C}^{(2)}_{a_2\sigma_3\cdots\sigma_N} = (S_2V_2^\dagger)_{a_2,[\sigma_3\cdots\sigma_N]}. \tag{4.11}$$

Substituting this into the first decomposition makes the first two MPS tensors explicit:

$$C_{\sigma_1\cdots\sigma_N} = \sum_{a_1,a_2} (A_1^{\sigma_1})_{1a_1} (A_2^{\sigma_2})_{a_1a_2} \widetilde{C}^{(2)}_{a_2\sigma_3\cdots\sigma_N}. \tag{4.12}$$

By repeating this process until only the final site remains, we will have transformed the state into an MPS. At each step, reshape U into a site tensor with one physical and two virtual indices, and pass the remaining factor to the next site. It is important to note that sweeping from left to right and keeping the U matrices as the MPS tensors was an arbitrary choice. We could equally well sweep from right to left, keeping the V matrices as MPS tensors.

Since we're keeping every nonzero singular value this construction is exact, with the minimal bond dimension at cut n equal to its Schmidt rank r_n. For N sites of local dimension d, the first SVD has d rows and d^{N-1} columns, so $\chi_1 \leq d$. The second has $d\chi_1$ rows and d^{N-2} columns, giving $\chi_2 \leq \min(d\chi_1, d^{N-2}) \leq \min(d^2, d^{N-2})$. More generally, at the cut after site n, the left and right Hilbert spaces have dimensions d^n and d^{N-n}. The Schmidt rank cannot exceed either dimension, so the exact construction gives

$$\chi_n = r_n \leq \min(d^n, d^{N-n}) = d^{\min(n,N-n)}. \tag{4.13}$$

Thus $d^{\min(n,N-n)}$ is a sufficient bond dimension at each cut to represent any state. A generic state saturates these bounds, with the largest bond near the middle of the chain:

$$\chi_{\max} = d^{\lfloor N/2 \rfloor}. \tag{4.14}$$

This is consistent with the maximum entanglement entropy of a d-dimensional spin chain being $S_E \leq N/2 \log d$. Every finite pure state therefore has an exact MPS representation, but often this will not be an efficient description because a generic state requires exponentially large bond dimension. Fortunately, as we discussed earlier, many physically interesting states can be captured efficiently by MPS.

4.4 Gauge freedom and canonical form

The MPS tensors are not uniquely determined by the physical state. We can always insert $XX^{-1} = \mathbb{I}$ on any internal bond, as long as X is invertible, and absorb its factors into the adjacent MPS tensors:

$$A_i^\sigma \longmapsto A_i^\sigma X, \qquad A_{i+1}^\sigma \longmapsto X^{-1} A_{i+1}^\sigma \tag{4.15}$$

This leaves the wavefunction unchanged while transforming the underlying tensors. The gauge freedom of MPS is very important, because we can use this to choose the most convenient form of the tensors in which to do analytic and numerical calculations. In particular, we want to choose a representation in which the block states we've seen above are orthonormal. This is called a *canonical form*.

The successive SVDs used to construct an MPS already provide the essential approach. Because each U has orthonormal columns, the tensors obtained from a left-to-right SVD obey

$$\sum_\sigma A_L^{\sigma\dagger} A_L^\sigma = \mathbb{I}_{\chi_i}. \tag{4.16}$$

They are *left-canonical*. The identity follows from $U^\dagger U = \mathbb{I}$, with (a_{i-1}, σ) treated as a row index. Starting at the left boundary, this identity ensures that the corresponding left-block states are orthonormal at every cut. A right-to-left decomposition produces *right-canonical* tensors,

$$\sum_\sigma A_R^\sigma A_R^{\sigma\dagger} = \mathbb{I}_{\chi_{i-1}}, \tag{4.17}$$

and orthonormal right-block states. Instead of making the whole system left canonical or right canonical, we can partially sweep from left to right and right to left and meet somewhere in the middle to create a *mixed canonical form*. In a mixed canonical form centred on a bond,

$$C_{\boldsymbol{\sigma}} = A_L^{\sigma_1} \cdots A_L^{\sigma_n} \Lambda_n A_R^{\sigma_{n+1}} \cdots A_R^{\sigma_N}. \tag{4.18}$$

The A_L tensors represent the orthonormal left Schmidt states and the A_R tensors represent the orthogonal right Schmidt vectors. Therefore this is a Schmidt decomposition of the system and $\Lambda_n = \mathrm{diag}(s_1, \ldots, s_{r_n})$ is a matrix with the Schmidt coefficients of the system along its diagonal (see Fig. 10).

Alternatively, absorb Λ into a neighbouring tensor to obtain a site centre A_C at site n. Then

$$|\psi\rangle = \sum_{a,\sigma,b} (A_C^\sigma)_{ab} |a_L\rangle |\sigma\rangle |b_R\rangle. \tag{4.19}$$

Because both block bases are orthonormal, $\langle\psi|\psi\rangle = \sum_{a,\sigma,b} |(A_C^\sigma)_{ab}|^2$. To move the centre one site to the right, reshape A_C as $(A_C)_{(a,\sigma),b} = (USV^\dagger)_{(a,\sigma),b}$, retain U as the new A_L tensor at site n, and multiply $SV^\dagger$ into site $n+1$. Moving left produces an A_R tensor using the analogous decomposition with the physical index grouped on the right.

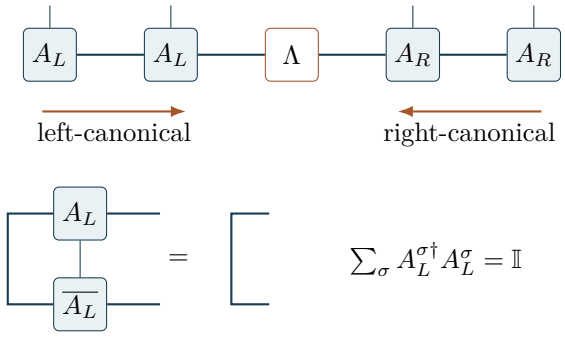


Figure 10: Mixed canonical form and cancellation of a left-canonical tensor with its conjugate. The upper diagram places the Schmidt coefficients on the central bond.

4.5 Truncation and its meaning

As we've seen, the bond dimension can in principle grow exponentially with the system. While we've seen special cases where the Schmidt coefficients across any cut of the system are strictly bounded, this is clearly the result of fine-tuning. In practice, even most states that satisfy an entanglement area law aren't like this. Instead, they have many non-zero Schmidt coefficients, but they drop exponentially and are therefore essentially irrelevant to properties of the quantum wavefunction. From this, it is clear that these small Schmidt coefficients can easily be *truncated* with little cost, but how do we do that? As we've seen, the Schmidt decomposition of an MPS is essentially an SVD of the wavefunction. From an SVD we obtain the singular values of a matrix, which are positive numbers arranged in descending order $s_i \geq s_{i+1}$; the larger the singular value, the more it contributes to the matrix. In computer science, one way to compress information is through a low-rank approximation, where only the first r singular values are preserved, and the rest are set to zero.

A greyscale image provides a visual example: its pixel intensities form a matrix I. Writing $I = USV^\dagger$, we retain only the largest r singular values to obtain a low-rank approximation. Figure 11 shows how increasing r recovers progressively finer detail. Alternatively, we may keep singular values up to some error $\epsilon = \sum_{i>r} s_i$.

The same truncation of singular values is the basis for MPS compression. In mixed canonical form, the central bond at site n is simply the Schmidt coefficients, so we can reduce its dimension by keeping only the largest χ coefficients.

$$|\widetilde{\psi}_\chi\rangle = \sum_{a=1}^{\chi} s_a |a_L\rangle |a_R\rangle. \tag{4.20}$$

This is optimal in the sense that it minimises the Hilbert-space distance to $|\psi\rangle$ among states with Schmidt rank at most χ at this cut. Its squared error is the discarded weight,

$$\epsilon_{\text{trunc}} = \sum_{a>\chi} s_a^2 = \| |\psi\rangle - |\widetilde{\psi}_\chi\rangle \|^2. \tag{4.21}$$

After renormalisation, the squared overlap with the original state is $1 - \epsilon_{\text{trunc}}$. If a state is area-law entangled, and a sufficiently large bond dimension is chosen, ϵ_{trunc} should be exponentially small. Turning a state into an MPS of fixed bond dimension χ everywhere requires sweeping through the system and performing this truncation at every bond. In this sense an MPS is simply a generalised low-rank approximation where we have used the underlying tensor-product structure to further simplify our representation of the state.

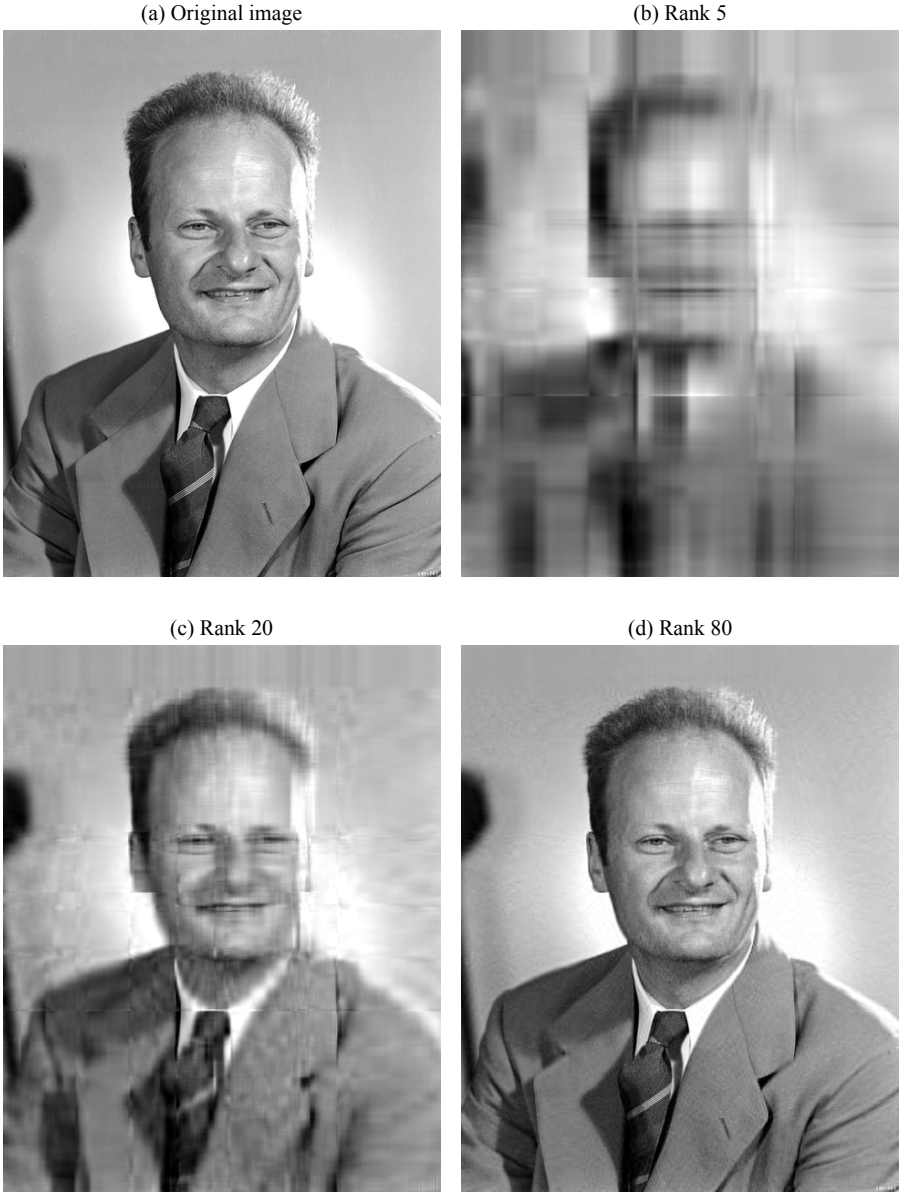


Figure 11: Low-rank approximation of an image. The original photograph is a 600×480 greyscale matrix; the other panels retain 5, 20 and 80 singular values.

4.6 Symmetry on the virtual indices

Given that MPS decompose a large wavefunction into local objects which have a complex virtual description, it is not clear how the symmetries of the wavefunction manifest at the level of the tensors. Let's suppose that a wavefunction is invariant under a particular representation $u(g)$ of a group with elements g. Then we know that $\otimes_n u(g)|\psi\rangle = e^{i\Theta(g)}|\psi\rangle$. For an injective uniform MPS in canonical form, this manifests at the level of individual tensors as

$$\sum_{\sigma'} u(g)_{\sigma\sigma'} A^{\sigma'} = e^{\mathrm{i}\theta(g)} V(g)^\dagger A^\sigma V(g) \tag{4.22}$$

where $V(g)$ is a unitary matrix acting on the virtual indices and may form a projective representation. As with the gauge transformations discussed above, neighbouring virtual transformations cancel (see Fig. 12). On a periodic chain, the trace also cancels the transformations at the ends; on an open chain, they act on the boundary indices. For the AKLT matrices in Eq. (3.15) and

$u(\theta) = e^{\mathrm{i}\theta S^z}$,

$$V(\theta) = e^{-\mathrm{i}\theta\sigma^z/2}, \qquad A^{\pm} \mapsto e^{\pm\mathrm{i}\theta} A^{\pm}, \quad A^0 \mapsto A^0. \tag{4.23}$$

The symmetry on the physical level has spin-1 character, while the virtual action has spin-1/2 character. More generally, the virtual matrices may compose *projectively*, $V(g)V(h) = \omega(g,h)V(gh)$. For spin rotations, the physical symmetry group is $SO(3)$: a 2π rotation is the identity on a spin-1, but gives a minus sign on a virtual spin-1/2. This is the fractionalisation of the symmetry. A nontrivial projective class on the virtual indices distinguishes an SPT phase even when the physical state preserves the symmetry. In the infinite AKLT chain, it also gives two equal Schmidt coefficients across a single cut [13].

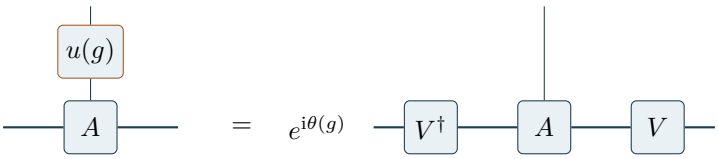


Figure 12: An on-site physical symmetry becomes a change of basis on the virtual bonds. Neighbouring virtual transformations cancel inside a chain. The orientation shown matches Eq. (4.22).

Another consequence appears at the boundaries. In the bulk of the AKLT chain, each virtual spin-1/2 is paired into a singlet with a neighbouring site. An open chain leaves one unpaired virtual spin at each end, as in Fig. 5. The cancellation of the V matrices stops at these two exposed indices, so a physical spin rotation acts on the edges as spin-1/2 rotations.

We can see the resulting ground states directly in the MPS. Leave the two boundary indices open and define the unnormalised states

$$|\psi_{ab}\rangle = \sum_{\boldsymbol{\sigma}} (A^{\sigma_1} \cdots A^{\sigma_N})_{ab} |\boldsymbol{\sigma}\rangle, \qquad a, b = 1, 2. \tag{4.24}$$

Every choice of a, b has the same singlet structure on the internal bonds. Therefore

$$H^{\text{open}}_{\text{AKLT}} |\psi_{ab}\rangle = 0, \qquad H^{\text{open}}_{\text{AKLT}} = \sum_{i=1}^{N-1} P^{(2)}_{i,i+1}. \tag{4.25}$$

For $N \geq 2$, these four states are linearly independent and define the entire AKLT ground-state space. Therefore the open AKLT chain has an *exactly fourfold-degenerate* ground state [3]. The two edge spins combine as

$$\tfrac{1}{2} \otimes \tfrac{1}{2} = 0 \oplus 1, \tag{4.26}$$

so this space consists of one total-spin singlet and one triplet. The bulk remains gapped; the degeneracy comes from the edges.

Closing the chain adds the projector on the bond between sites N and 1. The two exposed virtual spins are then paired into a singlet, giving the trace MPS constructed earlier. For a periodic chain with $N \geq 3$, this is the unique ground state. We have therefore found a system whose ground-state properties in the thermodynamic limit depend crucially on the *topology* of the spin chain, which in one dimension simply means its boundary conditions.

4.7 Parent Hamiltonians

The AKLT construction suggests a second question: given an MPS, can we find a local Hamiltonian for which it is an exact ground state? In fact, the answer is yes. This can be shown through the finitely correlated states construction of Fannes, Nachtergaele, and Werner, which generalises the finite-dimensional virtual construction behind AKLT [4].

This *parent Hamiltonian* for an MPS can be found by returning to the Schmidt decomposition of an MPS. The reduced density matrix of a finite l-site block $\rho_{[i,i+\ell-1]}$ has at most χ^2 non-zero eigenvalues of its d^l total eigenvalues. If we make l big enough that $d^\ell > \chi^2$ we can define h_i as the orthogonal projector onto the space of zero eigenvalues of $\rho_{[i,i+\ell-1]}$. Then

$$h_i^2 = h_i \geq 0, \qquad h_i|\psi\rangle = 0, \qquad H_{\text{parent}} = \sum_i h_i. \tag{4.27}$$

The second equation follows from $\langle\psi|h_i|\psi\rangle = \text{Tr}(h_i\rho_{[i,i+\ell-1]}) = 0$ and positivity. Thus $|\psi\rangle$ is a zero-energy, frustration-free ground state.

This construction alone does not guarantee a unique ground state or a gap, but such properties can be established in many cases. These results, combined with the discussions of symmetries above, allow MPS to be used to characterise a wide variety of gapped quantum phases in one dimension.

5 Contractions and transfer operators

We have seen how an MPS represents a wavefunction compactly. We now ask how to use that representation to calculate overlaps and expectation values without storing the full wavefunction. The key is to choose the contraction order carefully. For a four-site example, see Fig. 13.

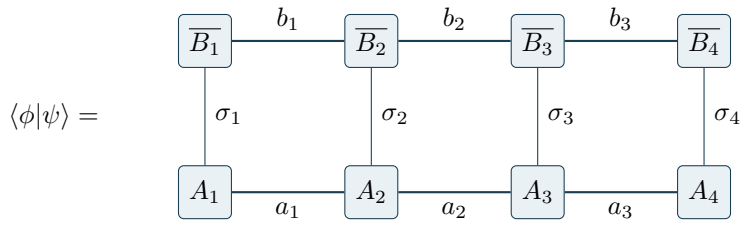


Figure 13: The full overlap of two four-site MPS with open boundaries. The upper row contains the conjugate tensors of $|\phi\rangle$ and the lower row those of $|\psi\rangle$. All physical and virtual indices are contracted, so the result is a scalar.

5.1 Overlaps and contraction order

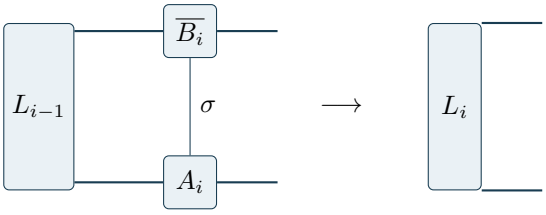
Consider two finite open-chain MPS, $|\psi\rangle$ with tensors A_i^σ and $|\phi\rangle$ with tensors B_i^σ. Their overlap contracts each physical index between the ket and bra:

$$\langle\phi|\psi\rangle = \sum_{\sigma_1,\ldots,\sigma_N} (\overline{B_1^{\sigma_1}\cdots B_N^{\sigma_N}})(A_1^{\sigma_1}\cdots A_N^{\sigma_N}). \tag{5.1}$$

One possible contraction order is to contract all virtual bonds within the ket and bra first, but this simply reproduces the two full coefficient tensors, each containing d^N entries. This loses all of the underlying MPS structure, so manipulating these tensors is exponentially expensive, even when the original MPS has small bonds.

How should we do this instead? We will take inspiration from the sweeping algorithms we used to construct MPS by contracting from one end of the double-layer network, summing each physical index as soon as its site is reached. Let's imagine starting from the leftmost site of the system $i = 1$. We can imagine that there is an auxiliary 1×1 matrix $L_0 = 1$ to the left of the first A and B tensors. We can then contract the left virtual legs of $A_1^{\sigma_1}$ and $B_1^{\sigma_1}$ with L_0 before contracting their physical legs. This creates an output matrix L_1 whose legs are the uncontracted virtual legs of $A_1^{\sigma_1}$ and $B_1^{\sigma_1}$. We can repeat this iteratively, as is shown in Fig. 14, sweeping through the system like so:

$$L_i = \sum_\sigma B_i^{\sigma\dagger} L_{i-1} A_i^\sigma. \tag{5.2}$$



Absorb one site; only two virtual indices remain open.

Figure 14: One step of the overlap sweep. Contracting physical and virtual indices together keeps the intermediate object small, whereas contracting each full ket or bra first does not.

At the final boundary L_N is a scalar and $L_N = \langle\phi|\psi\rangle$. For comparable bond dimensions χ, first form $L_{i-1}A_i^\sigma$, then multiply by $B_i^{\sigma\dagger}$ and sum over σ. Each step of this process scales like $O(d\chi^3)$, giving $O(Nd\chi^3)$ scaling, which is linear in N. This is the standard efficient contraction order for a dense open-chain MPS [12]. Contracting from right to left is equally efficient:

$$R_{i-1} = \sum_\sigma A_i^\sigma R_i B_i^{\sigma\dagger}, \qquad R_N = 1. \tag{5.3}$$

Sweeps from opposite ends can meet at any bond, where $\langle\phi|\psi\rangle = \mathrm{Tr}(L_i R_i)$. One simple application of this idea is calculating the norm of a wavefunction, which can then be used to appropriately normalise the state.

5.2 One-site expectation values and canonical form

We will now determine how to calculate single-site expectation values for MPS. Let's consider an operator O at site n, with matrix elements $O_{\sigma\tau} = \langle\sigma|O|\tau\rangle$, where σ and τ label the physical legs of the bra and ket. The essential idea is that when we calculate this expectation value, $B_i^\sigma = A_i^\sigma$ everywhere except at site n, where $B_n^\sigma = \sum_\tau O_{\sigma\tau}A_n^\tau$. We can therefore utilise the same approach introduced for calculating overlaps of MPS. In particular, we can contract all other sites into L_{n-1} and R_n. Writing $\mathcal{N} = \langle\psi|\psi\rangle$, the exact finite-chain expectation value is

$$\langle O_n\rangle = \frac{1}{\mathcal{N}} \sum_{\sigma,\tau} O_{\sigma\tau}\,\mathrm{Tr}[L_{n-1}A_n^\tau R_n A_n^{\sigma\dagger}]. \tag{5.4}$$

Canonical form simplifies the L and R tensors, which define the *environment* for this expectation value. If the tensors to the left are left-canonical, we simply find $L_{n-1} = \mathbb{I}$. For a left-canonical tensor at site n,

$$\langle O_n\rangle = \sum_{\sigma,\tau} O_{\sigma\tau}\,\mathrm{Tr}[A_L^\tau R_n A_L^{\sigma\dagger}]. \tag{5.5}$$

In left-canonical form the right environment will still be a non-trivial matrix. Similarly, in right-canonical form $R_n = \mathbb{I}$, giving

$$\langle O_n\rangle = \sum_{\sigma,\tau} O_{\sigma\tau}\,\mathrm{Tr}[L_{n-1}A_R^\tau A_R^{\sigma\dagger}]. \tag{5.6}$$

In mixed canonical form centred on site n, both environments are identities, so expectation values just involve A_C:

$$\langle O_n\rangle = \sum_{\sigma,\tau} O_{\sigma\tau}\,\mathrm{Tr}[A_C^\tau A_C^{\sigma\dagger}], \qquad \sum_\sigma \mathrm{Tr}[A_C^\sigma A_C^{\sigma\dagger}] = 1. \tag{5.7}$$

This reduces the calculation to the centre tensor and the local operator (see Fig. 15).

5.3 The MPS transfer matrix

The norm and expectation-value calculations above repeatedly update an environment matrix. For the norm, the bra and ket tensors are the same, and each update is a linear map. Define the right-to-left map and its left-to-right adjoint by

$$\mathcal{E}_i(X) = \sum_\sigma A_i^\sigma X A_i^{\sigma\dagger}, \qquad \mathcal{E}_i^*(Y) = \sum_\sigma A_i^{\sigma\dagger} Y A_i^\sigma. \tag{5.8}$$

An operator insertion is also a linear map:

$$\mathcal{E}_{n,O}(X) = \sum_{\sigma,\tau} O_{\sigma\tau} A_n^\tau X A_n^{\sigma\dagger}. \tag{5.9}$$

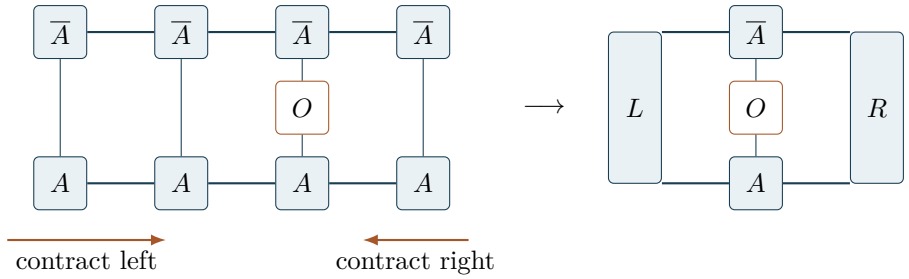


Figure 15: A local expectation value. All sites except the operator site are absorbed into left and right environments. In mixed canonical form, norm environments become identities.

In this notation, the one-site expectation value in Eq. (5.4) is $\mathrm{Tr}[L_{n-1}\mathcal{E}_{n,O}(R_n)]/\mathrm{Tr}[L_n R_n]$.

Thus $R_{i-1} = \mathcal{E}_i(R_i)$ and $L_i = \mathcal{E}_i^*(L_{i-1})$. If we combine the two open virtual indices into one, we can think of this map as an ordinary matrix:

$$(E_i)_{(a,a'),(b,b')} = \sum_\sigma (A_i^\sigma)_{ab}\overline{(A_i^\sigma)_{a'b'}}. \tag{5.10}$$

This is the *transfer matrix*. Its action on the flattened matrix $|X)$ is $E_i|X) = |\mathcal{E}_i(X))$. For bond dimension χ it has $\chi^2 \times \chi^2$ entries. It is useful conceptually, but we should normally use the tensor contraction we saw above because forming the full transfer matrix costs $O(d\chi^4)$ per site, and a dense matrix-vector multiplication costs $O(\chi^4)$, compared with $O(d\chi^3)$ that we saw above. We can also reshape $\mathcal{E}_{n,O}$ into a different $\chi^2 \times \chi^2$ matrix, $E_{n,O}$. In this notation, an expectation value is simply $(L_{n-1}|E_{n,O}|R_n)/(L_n|R_n)$.

We've seen that MPS allow properties of the system to be stored locally in particular MPS tensors. The transfer matrix is important because it allows this local information to be translated along the system to the left or right. If the tensors are the same along a stretch of a finite chain, the sweep repeatedly applies the same map $\mathcal{E}$. This is a power iteration method, and we can extract useful information about the sweep through the spectrum of the transfer matrix. For this reason, we can introduce the eigenvalue decomposition of the transfer matrix,

$$E = \sum_n \lambda_n |r_n)(l_n| \tag{5.11}$$

We can assume without loss of generality that $|\lambda_n| \geq |\lambda_{n+1}|$.

5.4 Expectation values in the AKLT ground state

Let's consider some concrete calculations using the AKLT ground state. First we can explain the normalisation introduced in Section 3.3. The unnormalised tensors $\widetilde{A}^\sigma = P^\sigma Q$ of Eq. (3.14) give

$$\widetilde{E} = \sum_\sigma \widetilde{A}^\sigma \otimes \overline{\widetilde{A}^\sigma} = \frac{1}{4}\begin{pmatrix} 1 & 0 & 0 & 2 \\ 0 & -1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 2 & 0 & 0 & 1 \end{pmatrix}, \tag{5.12}$$

using the basis $(\uparrow\uparrow, \uparrow\downarrow, \downarrow\uparrow, \downarrow\downarrow)$. Its eigenvalues are $3/4, -1/4, -1/4, -1/4$. Rescaling each tensor by $2/\sqrt{3}$ multiplies E by $4/3$, so the tensors in Eq. (3.15) have dominant eigenvalue one, with spectrum $1, -1/3, -1/3, -1/3$. This explains why this normalisation was chosen. They satisfy both canonical identities,

$$\sum_\sigma A^{\sigma\dagger}A^\sigma = \sum_\sigma A^\sigma A^{\sigma\dagger} = \mathbb{I}_2. \tag{5.13}$$

The transfer matrix is used to translate information across the system. To see this, we note that the normalised transfer map acts particularly simply on the identity and Pauli matrices:

$$\mathcal{E}(\mathbb{I}) = \mathbb{I}, \qquad \mathcal{E}(\sigma^\alpha) = -\tfrac{1}{3}\sigma^\alpha, \qquad \alpha = x, y, z. \tag{5.14}$$

By this, I mean the transfer matrix E acting on a vectorised version of these operators, *not* physical Pauli matrices since the AKLT model is a spin-1 model. The factors of 1 and $-1/3$ here arise due to the spectrum of the transfer matrix. The identity $\mathbb{I}$ corresponds to the eigenvalue 1 space and the Pauli matrices span the eigenvalue $-1/3$ space. Every 2×2 environment X is a linear combination of these four matrices. Let's suppose I perform an MPS contraction like the calculation of an expectation value from left to right and get a $\chi \times \chi$ matrix X. If I now want to translate this matrix m steps, I find:

$$\mathcal{E}^m(X) = \frac{\operatorname{Tr} X}{2}\mathbb{I} + \left(-\frac{1}{3}\right)^m \left(X - \frac{\operatorname{Tr} X}{2}\mathbb{I}\right). \tag{5.15}$$

Here $\mathcal{E}^* = \mathcal{E}$, so sweeping right to left obeys the same formula. We see that the traceless part of X decays exponentially with m due to the factor $(-1/3)^m$. Let's consider the effect this has on gapless edge modes we discussed above for the AKLT model. We can begin and end an OBC AKLT ground state with arbitrary two-dimensional vectors, $v_L^\dagger A^{\sigma_1} \cdots A^{\sigma_N} v_R$ so $L_0 = v_L v_L^\dagger$ and $R_N = v_R v_R^\dagger$. But what the calculation above shows is that, up to an overall normalisation factor, L_m and R_{N-m} will approach $\mathbb{I}$ exponentially. Any prefactors cancel in Eq. (5.4), and we may use

$$L = \mathbb{I}, \qquad R = \mathbb{I}/2, \qquad \operatorname{Tr}(LR) = 1. \tag{5.16}$$

What we've seen is that changes from the bulk properties of the AKLT ground state disappear exponentially as you move away from the boundary. We've also seen that this exponential decay is controlled by the subleading eigenvalues of the MPS transfer matrix. We can now calculate one-site spin expectation values explicitly using matrices. In the physical basis $(+1, 0, -1)$, the spin-1 operators are

$$S^x = \frac{1}{\sqrt{2}}\begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix}, \quad S^y = \frac{1}{\sqrt{2}}\begin{pmatrix} 0 & -\mathrm{i} & 0 \\ \mathrm{i} & 0 & -\mathrm{i} \\ 0 & \mathrm{i} & 0 \end{pmatrix}, \quad S^z = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & -1 \end{pmatrix}. \tag{5.17}$$

For a physical operator O, inserting it between the bra and ket tensors gives

$$E_O = \sum_{\sigma,\tau} O_{\sigma\tau} A^\tau \otimes \overline{A^\sigma}. \tag{5.18}$$

Using the same paired virtual basis as above, $(\uparrow\uparrow, \uparrow\downarrow, \downarrow\uparrow, \downarrow\downarrow)$, the three spin-1 operator transfer matrices are

$$E_{S^x} = \frac{1}{3}\begin{pmatrix} 0 & -1 & -1 & 0 \\ 1 & 0 & 0 & 1 \\ 1 & 0 & 0 & 1 \\ 0 & -1 & -1 & 0 \end{pmatrix}, \tag{5.19}$$

$$E_{S^y} = \frac{\mathrm{i}}{3}\begin{pmatrix} 0 & 1 & -1 & 0 \\ 1 & 0 & 0 & 1 \\ -1 & 0 & 0 & -1 \\ 0 & 1 & -1 & 0 \end{pmatrix}, \tag{5.20}$$

$$E_{S^z} = \frac{2}{3}\begin{pmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ -1 & 0 & 0 & 0 \end{pmatrix}. \tag{5.21}$$

In this convention a virtual matrix X is vectorised as $|X) = (X_{\uparrow\uparrow}, X_{\uparrow\downarrow}, X_{\downarrow\uparrow}, X_{\downarrow\downarrow})^{\mathsf{T}}$. The converged environment matrices $L = \mathbb{I}$ and $R = \mathbb{I}/2$ therefore become

$$(L| = \begin{pmatrix} 1 & 0 & 0 & 1 \end{pmatrix}, \qquad |R) = \frac{1}{2}\begin{pmatrix} 1 \\ 0 \\ 0 \\ 1 \end{pmatrix}, \qquad (L|R) = 1. \tag{5.22}$$

Multiplying the matrices by the right environment gives

$$E_{S^x}|R) = \frac{1}{3}\begin{pmatrix} 0 \\ 1 \\ 1 \\ 0 \end{pmatrix}, \quad E_{S^y}|R) = \frac{\mathrm{i}}{3}\begin{pmatrix} 0 \\ 1 \\ -1 \\ 0 \end{pmatrix}, \quad E_{S^z}|R) = \frac{1}{3}\begin{pmatrix} 1 \\ 0 \\ 0 \\ -1 \end{pmatrix}. \tag{5.23}$$

Each has zero contraction with $(L|$. Consequently, for a site in the bulk of an open chain,

$$\langle S_n^\alpha \rangle \simeq (L|E_{S^\alpha}|R) = 0, \qquad \alpha = x, y, z. \tag{5.24}$$

This result depends on the left/right environment tensors, so particular choices of boundary edge spins can produce nonzero magnetisation near the ends of the chain.

5.5 Two-point functions and exponential decay

Let's now think about correlation functions between two local operators. For operators P_i and Q_j with $i < j$, we can find their expectation value by finding $E_{i,P}$ and $E_{j,Q}$ and placing ordinary transfer matrices on the sites between them. For a normalised state, we find that the expectation value is

$$\langle P_i Q_j \rangle = (L_{i-1}|E_{i,P}E_{i+1}\cdots E_{j-1}E_{j,Q}|R_j). \tag{5.25}$$

For a uniform MPS with converged left/right environments and separation $r = j - i$ (see Fig. 16),

$$\langle P_i Q_{i+r} \rangle \simeq (L|E_P E^{r-1} E_Q|R), \qquad r \geq 1. \tag{5.26}$$

The only dependence on the separation is the power E^{r-1}. To interpret this result, we assume that E is diagonalisable and has a unique leading eigenvalue $\lambda_1 = 1$, with all others strictly inside the unit circle $|\lambda_a| < 1$. We then use the eigenvalue decomposition of E in which left and right eigenvectors satisfy $(l_a|r_b) = \delta_{ab}$:

$$E^m = |R)(L| + \sum_{a \geq 2} \lambda_a^m |r_a)(l_a|. \tag{5.27}$$

It's important to remember that the transfer matrix need not be Hermitian and that its left and right eigenvectors can differ. The leading term depends only upon $|R)$ and $(L|$, giving the product of the one-site expectation values. Subtracting it leaves the connected correlation function

$$\langle P_i Q_{i+r} \rangle_{\mathrm{c}} \equiv \langle P_i Q_{i+r} \rangle - \langle P_i \rangle \langle Q_{i+r} \rangle = \sum_{a \geq 2} \lambda_a^{r-1} (L|E_P|r_a)(l_a|E_Q|R). \tag{5.28}$$

Since all of the subleading eigenvalues are less than 1, this means that the correlations decay exponentially. The longest correlation length, given by the largest subleading eigenvalue λ_2, is

$$\xi = -\frac{1}{\log|\lambda_2|}, \qquad 0 < |\lambda_2| < 1. \tag{5.29}$$

This shows that for MPS where the transfer matrix has a unique largest eigenvalue, correlations decay exponentially. Hastings famously proved that the ground states of locally interacting

gapped Hamiltonians exhibit exponentially decaying correlations [8], showing once again the close connection between MPS and this class of states. MPS transfer matrices do not need to have a unique largest eigenvalue. A famous example is the GHZ state

$$|\mathrm{GHZ}\rangle = \frac{|00\ldots\rangle + |11\ldots\rangle}{\sqrt{2}}. \tag{5.30}$$

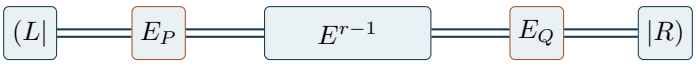
This can be defined in terms of MPS tensors

$$A^0 = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}, \qquad A^1 = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}, \tag{5.31}$$

with boundary vectors

$$\ell = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \end{pmatrix}, \qquad r = \begin{pmatrix} 1 \\ 1 \end{pmatrix}. \tag{5.32}$$

A simple calculation will show you that the transfer matrix of this state has a twofold-degenerate subspace with eigenvalue 1.



Each double line is a paired ket–bra virtual index.

Figure 16: Two-point expectation value with converged outer environments. Replacing the middle transfer by $|R)(L|$ gives the disconnected product; subleading modes give the connected correlation.

5.6 Correlation functions in the AKLT ground state

Let's look at correlation functions for the AKLT ground state. We already know most of what we need from the previous calculations we did for AKLT. The (vectorised) identity $|\mathbb{I})$ has eigenvalue one, while the Pauli matrices $|\sigma^\alpha)$ have eigenvalue $-1/3$. This is all we need to calculate the S_z correlation function since

$$E_{S^z}|R) = \tfrac{1}{3}|\sigma^z), \tag{5.33}$$

$$E^{r-1}|\sigma^z) = \left(-\tfrac{1}{3}\right)^{r-1}|\sigma^z), \tag{5.34}$$

$$E_{S^z}|\sigma^z) = -\tfrac{2}{3}|\mathbb{I}). \tag{5.35}$$

Thus the correlation function between two separated S_z operators is

$$\langle S^z_i S^z_{i+r}\rangle = (L|E_{S^z}E^{r-1}E_{S^z}|R) = -\frac{4}{9}\left(-\frac{1}{3}\right)^{r-1} \tag{5.36}$$

Since the one-site spin averages vanish in this bulk approximation, this is also the connected correlation. This correlation function has correlation length $\xi = 1/\log 3$, which is consistent with the correlation length we'd extract directly from the transfer matrix eigenvalues. Rotational symmetry means that each spin operator yields the same result, giving $\langle \mathbf{S}_i \cdot \mathbf{S}_{i+r}\rangle \simeq 4(-1/3)^r$, but this can also be calculated directly from the transfer matrix using a similar method.

5.7 Hidden string order in the AKLT model

We saw above that ordinary AKLT correlations decay exponentially, but SPT phases such as the Haldane phase also exhibit a type of long-range order called *string order* that reveals their

topological nature. To reveal this string order we insert an operator on *every* site between the endpoints of a correlation function. This defines the string correlator

$$C^z_{\text{string}}(i,j) = \left\langle S^z_i \left(\prod_{k=i+1}^{j-1} e^{\mathrm{i}\pi S^z_k} \right) S^z_j \right\rangle. \tag{5.37}$$

With $G = e^{\mathrm{i}\pi S^z} = \mathrm{diag}(-1,1,-1)$, the expectation value of the string correlator for the AKLT ground state is

$$C^z_{\text{string}}(i,j) = (L_{i-1}|E_{i,S^z}E_{i+1,G}\cdots E_{j-1,G}E_{j,S^z}|R_j) \tag{5.38}$$

Using the AKLT MPS tensors, we find the transfer matrix with this string operator inserted to be

$$E_G = \sum_\sigma G_{\sigma\sigma} A^\sigma \otimes \overline{A^\sigma} = \frac{1}{3}\begin{pmatrix} 1 & 0 & 0 & -2 \\ 0 & -1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ -2 & 0 & 0 & 1 \end{pmatrix}. \tag{5.39}$$

Crucially, we find that $E_G|\sigma^z) = |\sigma^z)$. Thus $|\sigma^z)$ is an eigenvector of E_G just as it was an eigenvector of E but it now has eigenvalue 1. We can now evaluate the string correlator in the bulk of the AKLT ground state to be

$$C^z_{\text{string}}(i,i+r) \simeq (L|E_{S^z}E_G^{r-1}E_{S^z}|R) = -\frac{4}{9} \tag{5.40}$$

This gives a finite value even as $r \to \infty$! All of the local correlation functions we calculated were zero for the AKLT ground state, suggesting it lacked any kind of local magnetic order. We now see that the model possesses a hidden kind of *topological* order which can be characterised by this string operator. Once again, S_z was chosen for simplicity, and any choice of spin operators will yield an identical result due to the rotational symmetry of the model. The conventional string order parameter includes a minus sign:

$$\mathcal{O}^\alpha_{\text{string}} = -\lim_{j-i\to\infty} \left\langle S^\alpha_i \exp\left(\mathrm{i}\pi \sum_{k=i+1}^{j-1} S^\alpha_k \right) S^\alpha_j \right\rangle = \frac{4}{9} \quad \text{for AKLT.} \tag{5.41}$$

This string order can act as an order parameter for the Haldane phase, just as local operators served as order parameters in the Landau paradigm of traditional condensed matter physics [15].

References

[1] E. Lieb, T. Schultz, and D. Mattis, *Two soluble models of an antiferromagnetic chain*, Ann. Phys. **16**, 407–466 (1961). doi:10.1016/0003-4916(61)90115-4.

[2] F. D. M. Haldane, *Nonlinear field theory of large-spin Heisenberg antiferromagnets: semiclassically quantized solitons of the one-dimensional easy-axis Néel state*, Phys. Rev. Lett. **50**, 1153 (1983). doi:10.1103/PhysRevLett.50.1153.

[3] I. Affleck, T. Kennedy, E. H. Lieb, and H. Tasaki, *Rigorous results on valence-bond ground states in antiferromagnets*, Phys. Rev. Lett. **59**, 799 (1987). doi:10.1103/PhysRevLett.59.799.

[4] M. Fannes, B. Nachtergaele, and R. F. Werner, *Finitely correlated states on quantum spin chains*, Commun. Math. Phys. **144**, 443–490 (1992). doi:10.1007/BF02099178.

[5] K. G. Wilson, *The renormalization group: critical phenomena and the Kondo problem*, Rev. Mod. Phys. **47**, 773 (1975). doi:10.1103/RevModPhys.47.773.

[6] S. R. White, *Density matrix formulation for quantum renormalization groups*, Phys. Rev. Lett. **69**, 2863 (1992). doi:10.1103/PhysRevLett.69.2863.

[7] S. Óstlund and S. Rommer, *Thermodynamic limit of density matrix renormalization*, Phys. Rev. Lett. **75**, 3537 (1995). arXiv:cond-mat/9503107.

[8] M. B. Hastings, *Locality in quantum and Markov dynamics on lattices and networks*, Phys. Rev. Lett. **93**, 140402 (2004). arXiv:cond-mat/0405587.

[9] M. B. Hastings, *An area law for one-dimensional quantum systems*, J. Stat. Mech. (2007) P08024. arXiv:0705.2024.

[10] F. Verstraete and J. I. Cirac, *Matrix product states represent ground states faithfully*, Phys. Rev. B **73**, 094423 (2006); preprint 2005. arXiv:cond-mat/0505140.

[11] T. Kuwahara and K. Saito, *Area law of noncritical ground states in 1D long-range interacting systems*, Nature Communications **11**, 4478 (2020). doi:10.1038/s41467-020-18055-x.

[12] U. Schollwöck, *The density-matrix renormalization group in the age of matrix product states*, Ann. Phys. **326**, 96–192 (2011). arXiv:1008.3477.

[13] F. Pollmann, E. Berg, A. M. Turner, and M. Oshikawa, *Entanglement spectrum of a topological phase in one dimension*, Phys. Rev. B **81**, 064439 (2010). arXiv:0910.1811.

[14] TeNPy, *Uniform matrix product states in the thermodynamic limit: Exercises*, documentation version 1.0.7. Official exercise notes. Fixed points and AKLT checks.

[15] F. Pollmann, E. Berg, A. M. Turner, and M. Oshikawa, *Symmetry protection of topological phases in one-dimensional quantum spin systems*, Phys. Rev. B **85**, 075125 (2012). arXiv:0909.4059.