Freek Witteveen
Department of Mathematical Sciences
University of Copenhagen
fw@math.ku.dk

Research

Quantum information theory and many-body physics

Freek Witteveen, currently a postdoctoral researcher at the University of Copenhagen, won the Stieltjes Prize for the academic year 2021–2022 with his dissertation *Quantum Information Theory and Many-Body Physics*. In this work, which he conducted at the CWI under the supervision of Michael Walter and Eric Opdam, he explored how representation theory of tensor networks can be used to shed light on the newly emerging field of quantum information theory. In this article, he explains what this field is about and how his research has contributed to it.

Many-body physics broadly concerns physical systems which are made up of a large number of subsystems. While the fundamental laws and principles of quantum mechanics are well known and can be formulated in compact equations, understanding quantum many-body physics and the emergent phenomena related to it gives rise to a whole new set of challenges. One approach to understanding large and complex physical system is by simulating them with the help of a computer. As scientists have been advancing the usage of computer simulations, they realized that the *theory* of computation as well as the closely related field of *information theory* are very useful to obtain a qualitative understanding of these complex systems. I will give an overview of the usage quantum information science in many-body physics, as developed over the last two decades. We will also see the notion of a *tensor network*, which is a convenient mathematical language for describing many-body quantum states. This article is partially based on the introduction of my thesis [11].

Quantum many-body physics

At small scales, our best fundamental theories of reality are quantum mechanical. What is quantum mechanics? Let us have a look at the simplest nontrivial quantum mechanical system: the qubit. A quantum mechanical system consists of a complex Hilbert space. In the case of the qubit this Hilbert space is $\mathbb{C}^2$ with the standard inner product. The *state* of the system is a normalized vector in the Hilbert space $\mathcal{H}$. It is common to use so-called bra-ket notation for vectors. A 'ket' $|\psi\rangle$ is simply a vector in the Hilbert space, and the 'bra' $\langle\psi|$ is the dual vector, which is such that $\langle\phi|\psi\rangle$ is precisely the inner product between the vectors $|\phi\rangle$ and $|\psi\rangle$. For the standard basis we write

$$|0\rangle = \begin{pmatrix}1\\0\end{pmatrix}, \qquad |1\rangle = \begin{pmatrix}0\\1\end{pmatrix},$$

so an arbitrary state can be written as $|\psi\rangle = \alpha\,|0\rangle + \beta\,|1\rangle$ where the normalization condition means that $|\alpha|^2 + |\beta|^2 = 1$. To get a sense of what this means physically you can think of a tiny magnet, which can either point upwards (state 0) or downwards (state 1). In this interpretation, the qubit is often called a *spin* particle which can be in state up $|\uparrow\rangle = |0\rangle$ or down $|\downarrow\rangle = |1\rangle$. Another possibility is that there is a particle (say an electron) which is either present (state 1) or absent (state 0); or in general any two-level physical system. The quantum mechanical state $|\psi\rangle$ allows 'superpositions' of these two states. This is similar (but different in a subtle way)

Illustration: Ryu Tajiri

to a system which is in state 0 with probability $p_0 = |\alpha|^2$ and in state 1 with probability $p_1 = |\beta|^2$, which would be a classical bit.

If one *measures* the system to see whether it is in state 0 or 1, one finds outcomes 0 and 1 precisely with probabilities p_0 and p_1. What makes the quantum state $|\psi\rangle$ different from a probability distribution over outcomes 0 and 1 is that the measurement depends on a choice of basis. This basis does not need to be the standard basis, but one could also express $|\psi\rangle$ in a different basis. This would give a different measurement procedure, with different outcomes. For this reason, a collection of qubits exhibits different behavior than a collection of classical bits.

Now that we have introduced the basic structure of quantum mechanics, we turn to more interesting systems than a single qubit. Most real quantum mechanical systems do not consist of a single particle with two states, but rather of many particles. In chemistry, a molecule has multiple electrons, each of which can be in different states. Similarly, in a solid we have a crystal structure with many electrons. In general, multiple particles (or systems) are modeled by the *tensor product*. To make this concrete, let us assume that we have N qubits. Then the Hilbert space of these N qubits is given by $(\mathbb{C}^2)^{\otimes N}$. The standard basis consists of tensor products of a choice of $|0\rangle$ or $|1\rangle$ for each tensor factor. In bra-ket notation we omit tensor product symbols and write $|\phi\rangle \otimes |\psi\rangle = |\phi\rangle|\psi\rangle$ or even $|\phi\psi\rangle$. This gives basis vectors

$$|i_1\rangle \cdots |i_N\rangle = |i_1 \cdots i_N\rangle \quad \text{for } (i_1,\dots,i_N) \in \{0,1\}^N.$$

In other words, a basis for N qubits is given by bitstrings of length N, which means that the Hilbert space has dimension 2^N. An arbitrary state can be expanded in this basis as

$$|\psi\rangle = \sum_{i_1,\dots,i_N} \psi_{i_1\cdots i_N} |i_1 \cdots i_N\rangle.$$

The physics of a quantum mechanical system is described by a *Hamiltonian*. The Hamiltonian describes the interactions between different particles (or with external fields). In general, it is a self-adjoint linear operator H acting on the Hilbert space. Since H is self-adjoint, it has real eigenvalues with orthogonal eigenvectors. The eigenvalues are the energies $E_0 \leq E_1 \leq E_2 \leq \cdots$ of the system, and the corresponding eigenvectors $|\psi_k\rangle$ are states with energy E_k. At zero temperature, the system will be in the *ground state* $|\psi_0\rangle$ with ground state energy E_0. The ground state energy could be degenerate (but in the following we assume for convenience that it is unique). The difference $\Delta = E_1 - E_0$ is the *ground state energy gap*. We say that a family of Hamiltonians of increasing system sizes is *gapped* if Δ is lower bounded by a constant. In many situations for electronic structure in chemistry and solids, the energy gap Δ to the first excited state is relatively large, compared to the temperature. This means that the behavior even at for instance room temperature is captured by ground state physics. For this reason, it makes sense to restrict our attention to ground state physics, even if we care about processes at finite temperature.

The precise form of the Hamiltonian is determined by the details of the physical system, but a common characteristic is that it arises from *local* interactions. This means that the Hamiltonian can be written as

$$H = \sum_X H_X$$

where the sum runs over some subsets $X \subset \{1,\dots,N\}$ of the qubits and H_X acts as the identity on all qubits except the ones in X. An even stronger constraint is *spatial* locality. Here we assume that the qubits are ordered in a lattice structure, and that interactions are only between nearby particles. For example, the sum only runs over pairs of nearest neighbors in the lattice.

To make this concrete, we can have a look at a simple model for magnetic interactions for a lattice of qubits (or spins). A basic example is the one-dimensional Ising model, which acts on a chain of qubits $n = 1,\dots,N$. We let

$$X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \qquad Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix},$$

be the Pauli matrices, and write X_n and Z_n for the operator which acts as X or Z on qubit n and as the identity on all other qubits. A term like $-Z_n Z_{n+1}$ is such that it has minimal energy if qubits n and $n+1$ are aligned (so they are in state $|00\rangle$ or $|11\rangle$), modelling a magnetic interaction. We may also add an external field in a transverse direction to get the Ising Hamiltonian

$$H = -J \sum_{n=1}^{N} Z_n Z_{n+1} + hX_n, \tag{1}$$

where h and J are real parameters and where we can take periodic boundaries. Since the operators X and Z do not commute, it is not obvious what the eigenvectors and thereby energies of this Hamiltonian are. The details of this model are not very important for the remainder of this article; I only introduce it here as a concrete example of what a many-body Hamiltonian may look like.

Ground state physics and computation

If a Hamiltonian is local, this means that it has an 'efficient' description: instead of an arbitrary $2^N \times 2^N$ self-adjoint matrix (with an exponential number of parameters) one only has to prescribe the (polynomially many) local terms. This leads to the formulation of one of the main computational problems in quantum many-body physics: given a local Hamiltonian H, what is the ground state energy of H? Certain Hamiltonians have 'easy' solutions that can be calculated exactly (this is, for example, the case for the one-dimensional Ising model). These are valuable for analytical understanding of the relevant physics and are good approximations in many instances. However, in other instances this approach does not suffice. Because the different local terms H_X in the Hamiltonian do not need to commute, it is not clear how to diagonalize the full Hamiltonian H. How would one go about (numerically) computing ground state energies? An immediate obstruction is that if the system consists of N qubits, the Hilbert space has dimension 2^N. While the Hamiltonian is local and has an efficient description, the ground state needs an exponential number of parameters to describe it! Even for modest N this means that one runs out of memory to store the full state on a computer. This barrier is one of the original motivations [4] for the development of *quantum computers* which perform computations on registers of qubits rather than bitstrings. A quantum computer with N qubits can, by definition, store the state of an N qubit quantum system, so quantum computers can at least *store* ground states of (local) Hamiltonians.

This is a good first step, but one still needs to *find* the ground state. It turns out this is a hard problem! It is known that determining the ground state energy of an arbitrary local Hamiltonian is a QMA-complete problem. QMA is a complexity class which is the quantum equivalent of the class NP, and it is strongly believed that there are no efficient (polynomial time) quantum algorithms for QMA-complete problems. While this may seem like bad news for the use of quantum computers for ground state physics, it does not mean that the problems one encounters in practice are necessarily intractable: perhaps certain subclasses of instances of the local Hamiltonian that occur in electronic structure computations in existing molecules are not computationally hard? There is in fact a powerful class of quantum algorithms for ground state problems, based on *quantum phase estimation*. These use the fact that a quantum computer is able to efficiently transform the state of the quantum computer by (an approximation of) the unitary operator $U=\exp(iH)$. Quantum phase estimation is an algorithm which computes the eigenvalues of the unitary U, in this case given by $\exp(iE_k)$. Of course, there is an exponential number of such eigenvalues. Which ones you find, will depend on the choice of initial state. To obtain E_0, one is required to prepare an initial state $|\phi\rangle$ on the quantum computer which is such that the overlap with the ground state $|\langle\phi|\psi_0\rangle|$ is not too small (the runtime of the algorithm scales with the inverse of this number). Since the full Hilbert space has exponential size, a random guess $|\phi\rangle$ will have exponentially small overlap, leading to an inefficient algorithm. Nevertheless, in many cases one can use classical methods to obtain reasonable initial states. Given such a good initial state, the resulting algorithm is polynomial in the system size N and the required precision of the approximation of E_0. The precise practical usefulness of such methods, compared to the best classical approaches, is still subject of a lively debate [8].

Quantum information theory

The mathematical theory of information was developed by Shannon in his landmark 1948 paper [10]. Shannon proposed that information is quantified by the *entropy* of a source. A source is modeled to be a probability distribution $p(x)$ with outcomes in an alphabet $x\in X$. The Shannon entropy of the probability distribution is given by

$$S(p)=-\sum_x p(x)\log p(x).$$

Since we want to measure in terms of bits, the logarithm is to base 2. The Shannon entropy $S(p)$ takes its maximum value $\log|X|$ when the distribution is uniform, and equals zero for a deterministic source. Shannon showed that when we are given independent samples from a source with distribution p, the outcomes can be compressed at a rate of $S(p)$ bits per outcome. This gives an operational meaning to $S(p)$ as the amount of information (measured in bits) of the source.

It was only much later that a fully quantum mechanical theory of information was pursued. Quantum mechanics allows a type of correlations which are fundamentally different from classical probability theory. Roughly speaking, going from probability theory to quantum theory one replaces probability distributions by normalized complex vectors

$$p_i\geq 0,\ \sum_i p_i=1 \Rightarrow \psi_i\in\mathbb{C},\ \sum_i|\psi_i|^2=1.$$

The difference becomes relevant when we study correlations between two parties. We consider a state $|\psi_{AB}\rangle$ shared by Alice and Bob, who hold quantum systems A and B. This means that $|\psi_{AB}\rangle$ is a vector in the tensor product Hilbert space $\mathcal{H}_A\otimes\mathcal{H}_B$. By a singular value decomposition (known as the Schmidt decomposition in this context) we may write

$$|\psi_{AB}\rangle=\sum_{i=1}^{r}s_i\,|a_i\rangle|b_i\rangle \tag{2}$$

where $|a_i\rangle$ and $|b_i\rangle$ form collections of orthonormal vectors in $\mathcal{H}_A$ and $\mathcal{H}_B$ respectively, and the s_i are positive numbers such that $s_1^2+\cdots+s_r^2=1$. The numbers s_i capture all the correlation between Alice and Bob. Such quantum correlations are known as *entanglement*. If $r=1$ and $s_1=1$ the state is a product state $|\psi_{AB}\rangle=|a_1\rangle|b_1\rangle$ and there are no correlations. On the other hand, if r is the minimum of the dimensions of the two Hilbert spaces and the s_i are uniformly equal to $s_i=r^{-1/2}$ the state is called *maximally entangled*. A special case is a pair of maximally entangled qubits

$$|\psi_{AB}\rangle=\frac{1}{\sqrt{2}}(|00\rangle+|11\rangle).$$

The maximally entangled pair of qubits is one of the fundamental units in quantum information theory.

As an analog to Shannon's noiseless coding theorem we would like to measure the entanglement in an arbitrary state $|\psi_{AB}\rangle$ in terms of maximally entangled qubits. The answer is given by the entropy of the probability distribution with probabilities $p_i=s_i^2$,

$$H(A)_\psi=-\sum_{i=1}^{r}s_i^2\log s_i^2.$$

This is known as the *entanglement entropy*. It equals zero if and only if the state is a product state, and equals the maximal value $\log(\dim(\mathcal{H}_A))$ is the state is maximally entangled. As an analog to the noiseless coding theorem the entanglement entropy is the optimal rate at which we can convert maximally entangled qubits to and from copies of a bipartite pure state. This corresponds to a scenario where we are given a large number of copies N of $|\psi_{AB}\rangle$. How many maximally entangled qubits can we distill using only local operations and classical communication? Conversely, suppose that we would like to create N copies of $|\psi_{AB}\rangle$ using only local operations and classical communication with some initial maximally entangled qubits. How many maximally entangled qubits do we need? For both questions we would like to know the optimal rate: as N goes to infinity, what is the optimal number of maximally entangled qubits per copy of $|\psi_{AB}\rangle$? The answer to both of these questions is that the optimal rate is the entanglement entropy $H(A)_\psi$. This sets the entanglement entropy as the operationally correct measure for entanglement.

This is only the beginning of the story: we restricted the discussion to 'pure' quantum states, but one can also model 'mixed' quantum states (which are a probabilistic mixture of pure quantum states). In this case the theory of entanglement becomes much more subtle. Quantum information theory studies many scenarios beyond entanglement. Important further topics are determining the possible rates of communication over noisy quantum channels and measuring the cryptographic properties of quantum states [9].

Ground state physics and entanglement

The entanglement entropy is a measure of entanglement in quantum states. What does the entanglement of ground states in many-body physics look like? We take a lattice of qubits, and we let A be a subregion of the lattice containing at most half of the sites of the full lattice

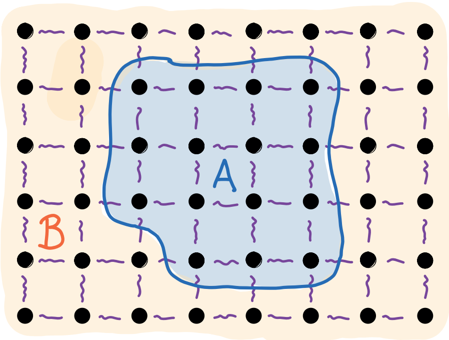

If we pick a random state (by taking a uniformly random point in the unit sphere of the Hilbert space), the state will be close to maximally entangled with high probability. That is, we have entanglement entropy $H(A) \approx |A|$, which equals $\log(\dim(\mathcal{H}_A))$. More generally, if the entanglement entropy scales with the size of the region $|A|$, this is known as *volume-law* entanglement.

Ground states of spatially local Hamiltonians behave in a fundamentally different way! An important feature of the Hamiltonian is its locality: the Hamiltonian is a sum of terms, each of which acts only on a few particles. One might expect that this carries over the ground state, and the ground state also has a 'local' nature. To some extent this is true, and the correlations in ground states are of a local nature. This is captured by the so-called *area law*. A quantum state on a lattice in D spatial dimensions satisfies an area law if the scaling of the entanglement entropy of any subregion is with the size of the boundary of the subregion, rather than the volume of the subregion. If we take a (large) region A in the lattice then this would mean that the entropy $H(A)$ is proportional to $|\partial A|$, the size of the boundary of A.

A closely related fact is that we expect exponentially decaying correlations between different sites. It is known that all local lattice Hamiltonians which are gapped have exponential decay of correlations [6]. Exponential decay of correlations is (at least on an intuitive level) closely related to area laws: if we have exponential decay of correlations one may expect that sites far away from the boundary of a region A do not contribute to the entanglement. However, exponential decay of correlations does not directly imply an area law, and the area law is only rigorously known to hold in special cases (for example, in one spatial dimension [5]), and it can be violated under certain circumstances.

Tensor network states

The area law captures the fact that in ground states of local Hamiltonians correlations are of a local nature. A next question is whether it is possible to represent the ground state itself in a local way. Suppose we have a quantum state $|\psi\rangle$ on a tensor product Hilbert space $\mathcal{H} = \mathcal{H}_1 \otimes \mathcal{H}_2 \otimes \cdots \otimes \mathcal{H}_n$, where $\mathcal{H}_k = \mathbb{C}^{d_k}$. We may expand $|\psi\rangle$ in a product basis:

$$|\psi\rangle = \sum_{\{i_k\}} \psi^{i_1 i_2 \ldots i_n} |i_1\rangle \otimes |i_2\rangle \otimes \cdots \otimes |i_n\rangle .$$

The collection of numbers $\psi_{i_1 i_2 \ldots i_n}$ defines a *tensor*. Of course, in general, the size of this collection is exponential in n. Tensor networks provide a method to parametrize a relevant subset of tensors efficiently, by 'breaking up ψ into smaller tensors.' We may represent a tensor graphically as

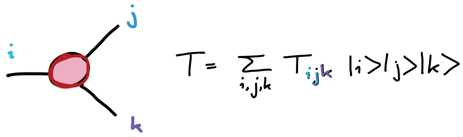

Given two tensors S and T with coefficients $S_{i_1,i_2,\ldots,i_n}$ and $T_{j_1,j_2,\ldots,j_m}$ we may for instance contract S and T along the first indices (provided the corresponding dimensions are equal) to get a tensor with coefficients

$$\sum_{i_1} S_{i_1,i_2,\ldots,i_n} T_{i_1,j_2,\ldots,j_m}.$$

Graphically this simply corresponds to connecting the two tensors

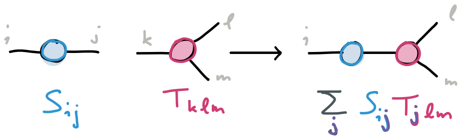

If the two tensors are 2-tensors, this simply corresponds to matrix multiplication.

Now, if we are given a collection of tensors, then we may contract indices along a graph which is defined by letting the tensors correspond to vertices with a number of dangling half-edges corresponding to the number of indices the tensor has; we then indicate which indices are contracted by connecting half-edges to form an edge in the graph. The resulting tensor has uncontracted indices on all unconnected half-edges

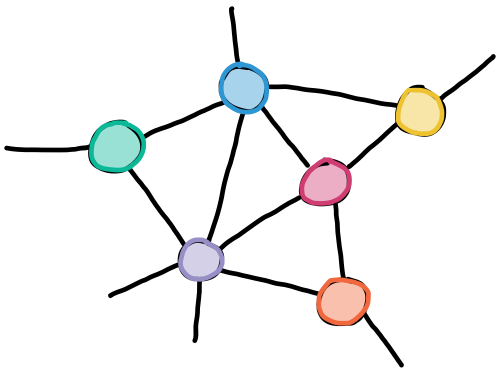

The dimensions along the contracted edges are called *bond dimensions*, while the dimensions along the uncontracted edges are the *physical dimensions*.

The decomposition in equation (2) is a special case of a tensor network decomposition:

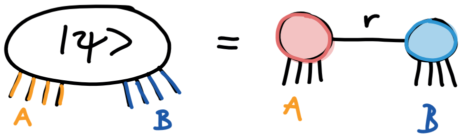

where the two tensors are connected by a an edge with bond dimension r. One way to think of tensor network states is that they are tensor decompositions which are a generalization of the singular value decomposition, and where the bond dimensions generalize the rank of the decomposition.

A different perspective on tensor network states is that one starts with a collection of maximally entangled states along the edges of the graph, and then applies linear maps at each of the vertices. For this reason, tensor network states are also known as Projected Entangled Pair States (PEPS). This means that tensor network states are quantum states where the entanglement is of a local nature, with respect to some given graph. For a state on a lattice, it is natural to take a tensor network that corresponds to the lattice geometry

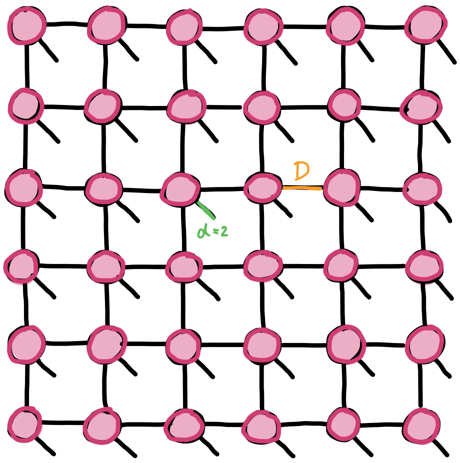

The dangling edges correspond to the physical degrees of freedom, which for example are two-dimensional (qubits), whereas the contracted vertical and horizontal edges have a bond dimension D. Typically, one takes this bond dimension D to be system size independent. This gives a huge reduction in the number of parameters: each tensor has one edge of dimension 2, and four (or fewer on the boundary of the lattice) of dimension D, giving $2D^4$ complex parameters. This leads to a state with at most $2ND^4$ parameters instead of the 2^N parameters required for arbitrary states of N qubits. Tensor network states by construction satisfy constraints on their entanglement entropy: if A is a subsystem and B is its complement, then the rank r of the state $|\psi_{AB}\rangle$ as in equation (2) is bounded by the product of the bond dimensions of any set of edges one has to cut to separate A from its complement. If the graph one starts with is based on a lattice, this directly imposes an area law, since in that case we cut a number of bonds proportional to the boundary ∂A, each with bond dimension D. This gives $H(A) \leq \log(r) \leq |\partial A| \log(D)$.

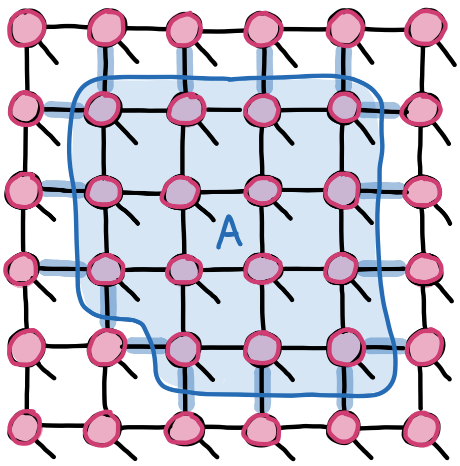

To make the tensor network approach more concrete, we start with *matrix product states* (MPS), which are tensor networks in one spatial dimension. In this case we consider a one-dimensional chain of N qubits and for convenience we take periodic boundary conditions and make the state translation invariant. We fix a bond dimension D. Then an MPS state is defined by a tensor of size $2 \times D \times D$, which we write as $M^{(i)}_{jk}$ and $j,k = 0,\dots,D-1$. Instead of considering a $2 \times D \times D$ tensor, we may also think of M as a pair of $D \times D$ matrices, and we denote by $M^{(i)}$ the $D \times D$ matrix with entries $\{M^{(i)}_{jk}\}_{j,k}$.

Then the associated MPS state is a state of N qubits, defined by the following product of matrices

$$|\psi\rangle = \sum_{i \in \{0,1\}^N} \mathrm{tr}\,[M^{(i_1)} M^{(i_2)} \dots M^{(i_N)}]\,|i_1 i_2 \dots i_N\rangle. \tag{3}$$

Graphically, this corresponds to

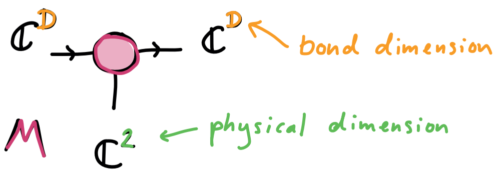

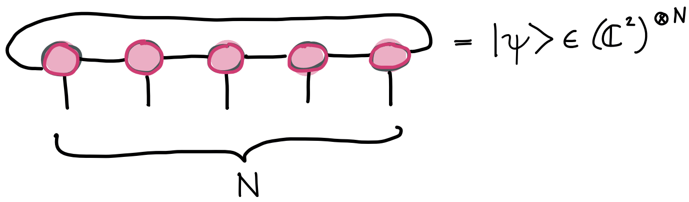

For MPS, if we let A be an interval, then the entanglement entropy $H(A)_\psi$ is upper bounded by $2\log(D)$ (since we only need to cut two edges to separate A from its complement):

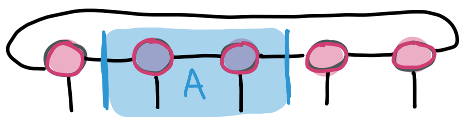

This again is consistent with an area law, since A has a constant size boundary, no matter the length of the interval! The total number of parameters in the description of the MPS state $|\psi\rangle$ is only $2ND^2$. For fixed D this number of parameters is *linear* in the system size rather than exponential, so this is a huge reduction in parameter space. Additionally, from equation (3) we can efficiently compute the coefficients of the state. If one increases the bond dimension D to a size exponential in N one can write any quantum state as an MPS (but one is usually interested in the regime where D is constant or polynomial in N).

For one-dimensional spin systems, the space of ground states of gapped Hamiltonians is well understood. If H is gapped, it is known that the ground state satisfies both an area law and exponential decay of correlations. One can approximate the ground state to precision ε by an MPS state with bond dimension $D = O(\mathrm{poly}(N, \frac{1}{\varepsilon}))$ [5]. Moreover, such an approximation can be found in $O(\mathrm{poly}(N, \frac{1}{\varepsilon}))$ by a classical computer [7]. Historically, the development of tensor networks started with the *density matrix renormalization group* (DMRG) algorithm which, in hindsight, is an algorithm for finding MPS approximations to ground states and gives excellent results in practice. From this we conclude that, while in general the problem of finding ground states of local Hamiltonians is QMA-hard, when we restrict to one-dimensional gapped Hamiltonians the problem can be solved in polynomial time on a classical computer.

Moreover, the fact that ground states of local gapped Hamiltonians can be approximated by MPS gives us a wealth of information about the structure of such ground states. An important tool to leverage this structure is a fact about the algebraic structure of MPS. Any MPS has a *gauge symmetry*, where one can change the individual tensors without changing the resulting MPS state. It is easy to see that if we replace $M^{(i)} \mapsto gM^{(i)}g^{-1}$ for $i=0,1$ and for a $D\times D$ invertible matrix g this does not change the resulting state in equation (3) since each g cancels against its inverse:

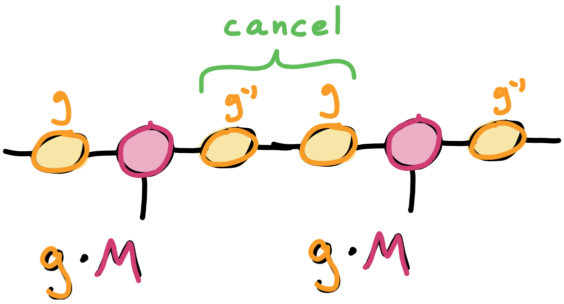

In other words, there is an action by the group $\mathrm{GL}(D)$ on the tensor which keeps the contracted state $|\psi\rangle$ invariant. The *Fundamental Theorem of MPS* states that this is (essentially) the only redundancy in the description. One way to formulate this is that one can always bring the tensor into a *canonical form* by (limits of) invertible transformations g, such that the only remaining symmetries are unitary $D\times D$ matrices. If we have two different tensors giving rise to the same MPS state for all system sizes, they can be brought in the same canonical form. There are various canonical forms, one may for example impose the condition

$$(M^{(0)})^{\dagger}M^{(0)}+(M^{(1)})^{\dagger}M^{(1)}=\mathrm{Id}$$

which (if we think of a 3-tensor of size $2\times D\times D$ again, instead of a pair of matrices) we can write graphically as

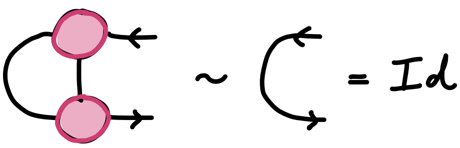

where we use the convention that rotating the tensor corresponds to complex conjugation. This is important for numerics: ignoring the gauge symmetry can lead to matrix computations with large condition numbers. It is also important for variational optimization algorithms (for approximating ground states) to fix redundancies in the description.

As a theoretical application of the Fundamental Theorem, one can study states which have a symmetry. Consider an MPS state $|\psi\rangle$ on N sites with single site tensor M which is invariant under unitaries $U_h^{\otimes N}$, where $h\mapsto U_h$ is a unitary representation of a group G. For example, on a spin chain we could consider states which are invariant under exchanging $|0\rangle$ and $|1\rangle$. This is a representation of $\mathbb{Z}/2\mathbb{Z}$ where the generator is mapped to the Pauli operator X. The Ising model is an example of a system which has this symmetry. Any such symmetry also gives rise to an action of G on the tensor M, mapping the tensor to a different tensor $U_h\cdot M$. By the Fundamental Theorem and the invariance of the state, we can deduce that M and $U_h\cdot M$ must be related by the action of some $D\times D$ unitary g_h:

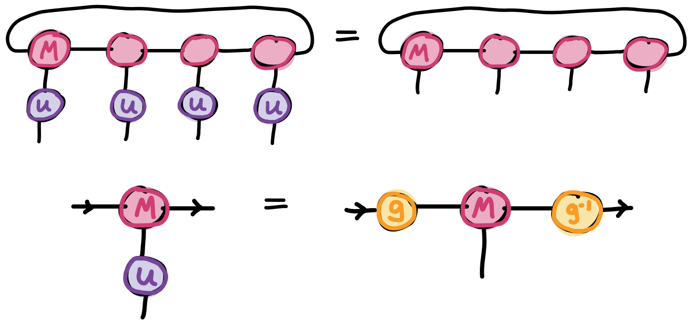

It turns out that the g_h form a projective representation of G, and one can classify equivalence classes of states with symmetry group G under continuous deformations ('symmetry protected topological phases') by the second cohomology group $H^2(G,U(1))$. This example illustrates the fact that global properties of the quantum many-body state are encoded in the local tensors.

In higher spatial dimensions tensor networks (which here often go under the name PEPS) can also be used to approximate ground states of lattice Hamiltonians. While the theory of MPS is relatively well understood, the general theory of PEPS is much more complicated. Similarly, numerical methods using PEPS are much more challenging.

The fact that numerical methods are more challenging is corroborated by results in complexity theory, which state that tensor network contraction (a necessary ingredient in any tensor network algorithms) is $\#P$-hard on two-dimensional lattices. This is closely related to the fact that the algebraic structure of PEPS is more intricate, and the involved multilinear algebra is more complicated. In relation to many-body physics, there are no rigorous approximation results available as in the case of MPS: there are no guarantees for approximation of ground states of local Hamiltonians by PEPS (even though one would expect this to be the case in most situations).

Given all these hurdles, why would one care about PEPS? A first motivation is given by numerical evidence: even though simulations are computationally challenging, they give good results for strongly interacting and complicated systems. Secondly, while we do not have rigorous guarantees, it is widely believed that in practice ground states of gapped Hamiltonians should have PEPS representations. This means that PEPS is a useful abstraction to reason about general ground states. Another good reason is that for many paradigmatic models there are *exact* PEPS ground states with a known description. These models provide important intuition for many-body physics. Amongst others, one can realize states with so-called topological order as PEPS. Finally, since PEPS states can be used as a proxy for ground states, they are useful in designing and understanding quantum algorithms for preparing ground states.

As an example of the application of techniques from multilinear algebra, one can define a canonical form for PEPS in higher spatial dimensions, extending the canonical form for MPS. PEPS tensors have a similar symmetry as MPS tensors. One can show [1] that one can bring a PEPS tensor into a canonical form according to the conditions

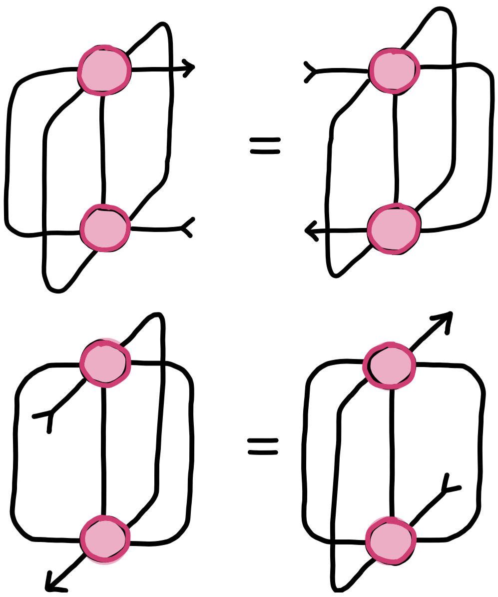

such that the only remaining freedom is given by the action of unitary matrices. This can be done using ideas from (geometric) invariant theory, which studies group actions (in this case a group action of $\mathrm{GL}(D) \times \mathrm{GL}(D)$). It follows that two tensors have the same canonical form if and only if they give rise to the same state, not only on every two-dimensional lattice but on any possible contraction graph:

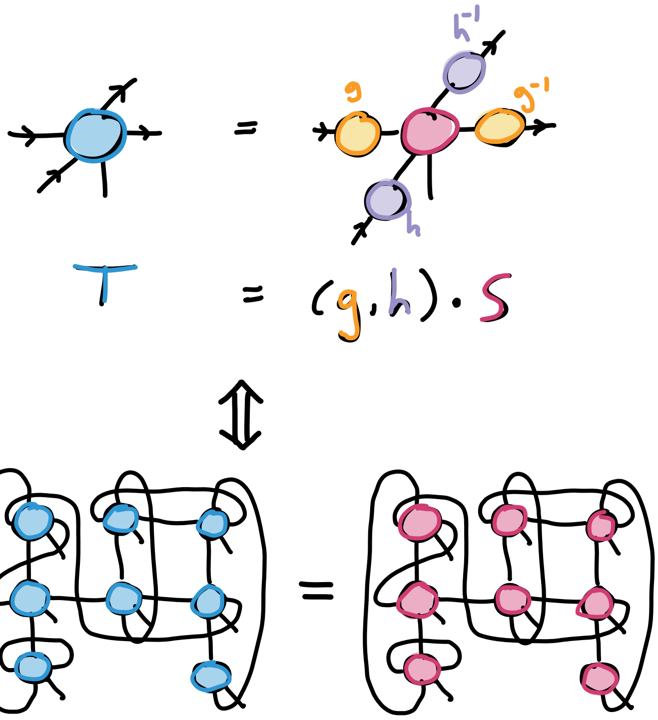

Note that in the special case of one spatial dimension there is only a single connected graph of size N according which one can contract the tensor.

Conclusions

I have tried to give an impression of interactions between quantum information science and quantum many-body physics. We have seen two perspectives:

- *A computational perspective:* one can try to classify the computational hardness of finding ground states of local Hamiltonians. The problem is computationally hard in general.
- *An information theoretic perspective:* ground states of spatially local Hamiltonians have very specific entanglement behavior. One can use this to decompose a global state into a tensor network, where all information is encoded in local tensors.

Both of these perspectives have been very fruitful over the last two decades, and many interesting (mathematical) questions remain, see for example [3]. Amongst these are: rigorous understanding of ground state approximation properties of PEPS, proving area laws for higher dimensional systems and obtaining a more general understanding of the (symmetry protected) topological order. Besides this, there are other intriguing interactions between quantum information theory and many-body physics, such as the intimate relation between quantum error correction and topological order, and the role quantum information theory plays in holographic theories of quantum gravity. If you are interested in learning more about these topics, I recommend the modern classic [9] for quantum information theory, the book [12] for an overview of quantum information theory in quantum many-body physics and [2] for an up-to-date review of the mathematics of tensor networks and many-body physics.

References

1 Arturo Acuaviva, et al., The minimal canonical form of a tensor network, in *2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS)*, IEEE, 2023, pp. 328–362.

2 J. Ignacio Cirac, David Pérez-García, Norbert Schuch and Frank Verstraete, Matrix product states and projected entangled pair states: Concepts, symmetries, theorems, *Reviews of Modern Physics* 93(4) (2021), 045003.

3 Juan Ignacio Cirac, José Garre-Rubio and David Pérez-García, Mathematical open problems in projected entangled pair states, *Revista Matemática Complutense* 32 (2019), 579–599.

4 Richard P. Feynman, Simulating physics with computers, *International Journal of Theoretical* Physics 21(6/7) (1982).

5 Matthew B. Hastings, An area law for one-dimensional quantum systems, *Journal of Statistical Mechanics: Theory and Experiment* 2007(08) (2007), P08024.

6 Matthew B. Hastings and Tohru Koma, Spectral gap and exponential decay of correlations, *Communications in Mathematical Physics* 265(3) (2006), 781–804.

7 Zeph Landau, Umesh Vazirani and Thomas Vidick, A polynomial time algorithm for the ground state of one-dimensional gapped local Hamiltonians, *Nature Physics* 11(7) (2015), 566–569.

8 Seunghoon Lee, et al. Evaluating the evidence for exponential quantum advantage in ground-state quantum chemistry, *Nature Communications* 14(1) (2023), 1952.

9 Michael A. Nielsen and Isaac Chuang, *Quantum Computation and Quantum Information,* American Association of Physics Teachers, 2002.

10 Claude Elwood Shannon, A mathematical theory of communication, *The Bell System Technical Journal* 27(3) (1948), 379–423.

11 Freek Witteveen, *Quantum Information Theory and Many-body Physics*, PhD thesis, University of Amsterdam, 2022.

12 Bei Zeng, Xie Chen, Duan-Lu Zhou and Xiao-Gang Wen, *Quantum Information Meets Quantum Matter*, Springer, 2019.