Basics of Quantum Phases and Entanglement
2020 UQM school at Perimeter Institute

Lecturer: McGreevy
Please email corrections and comments to mcgreevy at physics dot ucsd dot edu.

Last updated: 2020/12/28, 13:24:19

Contents

1 States of matter, classified by level of desperation **3**

2 Adiabatic continuation and local unitary circuits **10**

3 Entanglement, short and long **12**

4 Entanglement in groundstates of local Hamiltonians **14**

5 Renormalization group monotones from entanglement measures **19**

6 Omissions **21**

This is an introductory lecture on basic notions of quantum phases of matter, and how to think about them in terms of entanglement. For further reading, I suggest:

- Zeng, Chen, Zhou, Wen, *Quantum Information Meets Quantum Matter: From Quantum Entanglement to Topological Phase in Many-Body Systems*
- These TASI lectures.
- For more on applications of quantum information theory to many body physics from a similar viewpoint, see the lecture notes and problems posted here.

1 States of matter, classified by level of desperation

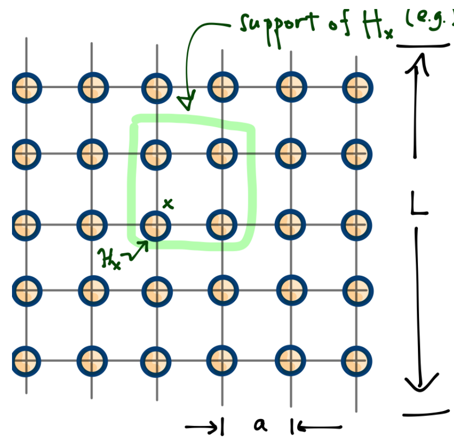

In this lecture we are going to talk about *extensive quantum systems*. A quantum system can be specified by its Hilbert space and its Hamiltonian. By the adjective *extensive* I mean that the Hilbert space is defined by associating finite-dimensional Hilbert spaces $\mathcal{H}_x$ to chunks of *space*, labelled by[1] some coordinates x. Then couple them by a local Hamiltonian, $H = \sum_x H_x$, where H_x acts only on the patch at x and not-too-distant patches (and as the identity operator on the other tensor factors in $\mathcal{H}$).

The phenomena whose study we will find most fulfilling only happen in the *thermodynamic limit*, where the number of patches grows without bound. I will use L to denote the linear size of the system. For a cubic chunk of d-dimensional hypercubic lattice, there are $\left(\frac{L}{a}\right)^d$ patches, where a is the size of the patches. So the thermodynamic limit is $L \to \infty$, or more precisely $L \gg a$. In the mysterious first sentence of this paragraph, I am referring to *emergent* phenomena: qualitatively new effects which can never be accomplished by small systems, such as spontaneous symmetry breaking (magnetism, superconductivity, the rigidity of solids), phase transitions, topological order, and all the other things we have not thought of yet because we are not very

[1] We can allow the local Hilbert space to be infinite-dimensional (as for rotors or bosons) if we add terms to the Hamiltonian that leave a finite-dimensional set of low-energy states. In practice, when simulating such systems for example, we can always truncate the Hilbert space to some finite value.

smart.[2] [3]

Perhaps the most basic question we can ask about such a system is: how many degrees of freedom are there at the lowest energies (lower than any interesting scale in the problem, in particular in the Hamiltonian)? By degrees of freedom (dofs) I mean excitations that can be created by a local operator, as in an experiment where we scatter particles (neutrons, photons...) off the material. There are essentially three possibilities:

1. **None.**
2. **Some.**
3. **A lot.**

As we proceed down this classification, our level of understanding rapidly decreases.

A more informative tour through that list goes like this. To get started let me make the assumption that the system has (at least discrete) translation invariance, so we can label the excitations by momentum.

1. **None:** Such a system has an energy gap ('is gapped'): the energy difference $\Delta = E_1 - E_0$ between the first excited state and the groundstate is nonzero, even in the thermodynamic limit. Note that Δ is almost always nonzero in finite volume. (Recall, for example, the spectrum of the electromagnetic field in a box of linear size L: $E_n \sim \frac{n}{L}$.) The crucial thing here (in contrast to the case of photons) is that this energy stays finite even as $L \to \infty$.

   The excitations of such a system are generally massive particles[4].

   Actually, it is useful to allow a finite number of states below the gap (which become degenerate in the thermodynamic limit) in our definition of a gapped system.

[2] In case you doubt that characterization, ask yourself this: How many of the items on this list were discovered theoretically before they were found to occur in Earth rocks by our friends who engage in experiments? The answer is **none**. Not one of them! Let us be humble. On the other hand: this is a source of hope for more interesting physics, in that the set of Earth rocks which have been studied carefully so far is likely to represent a very small sample of the possible emergent quantum systems.

[3] Can you think of other elements I should add to this list? One possibility (thanks to Ibou Bah for reminding me) can be called *gravitational order* – the emergence of dynamical space (or spacetime) (and hence gravity) from such emergent quantum systems. The best-understood example of this is AdS/CFT, and was discovered using string theory. I was tempted to claim this as a victory for theorists, but then I remembered that we discovered gravity experimentally quite a while ago.

[4] Ref. [1] proves a version of this statement. I think it is worth looking for loopholes here.

2. **Some:** An example of what I mean by 'some' is that the system can have excitations which are massless particles, like the photon.

   The lowest energy degrees of freedom occur at isolated points in momentum space: The dispersion relation of the photon $\omega(k) = c\sqrt{\vec{k} \cdot \vec{k}}$ vanishes at $\vec{k} = 0$.

   In this category I also put the gapless fluctuations at a critical point. It's not necessarily true that $\omega \sim k^{\text{integer}}$ and those excitations are not necessarily *particles*. But they are still at $k = 0$[5].

3. **A lot:** What I mean by this is Fermi surfaces, but importantly, not just free fermions or adiabatic continuations of free fermions (Landau Fermi liquid theory). Such systems exist, for example in the half-filled Landau level and in the strange metal regime of cuprate superconductors.

Let's reconsider the case of gapped systems and define the notion of a quantum phase. Different gapped states are in different phases if we can't deform the Hamiltonian to get from one to the other without closing the gap. So a gapped phase is an equivalence class of Hamiltonians.

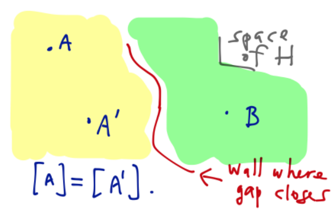

[6] [7]You might be bothered by the following: it is hard to imagine checking that there is no way around the wall of gaplessness. It is therefore important to find sharp characterizations of such states, like integer labels, which cannot change smoothly. This is the very definition of topology. An important goal in condensed matter physics is to figure out labels that can be put on states which can distinguish them in this way as distinct phases of matter.

Here are two classes of examples (even in the absence of symmetry) of topological labels.

Topological Order. Even the lowest-energy (even below the gap) physics of gapped systems can be deeply fascinating. For example, it may be that the *number* of groundstates depends on the topology of the space on which we put the system. Since this is an integer, it cannot vary continuously and can only jump when the gap closes.

[5]or some other isolated points in momentum space.

[6]Actually, there is an important extra equivalence relation that we must include: We don't care if on top of some nontrivial phase of matter someone sprinkles a dust of decoupled qubits which are totally inert and do nothing at all. This represents the same phase of matter. Then, further, we are allowed to adiabatically deform the hamiltonian, including these decoupled bits, so that they can interact with the original degrees of freedom. So: in addition to allowing adiabatic variation of couplings, we also allow the addition of decoupled qubits.

[7]Note that the closing of the gap does not by itself mean a quantum critical point: at a first order transition, just the lowest two levels cross each other.

This is a symptom of the phenomenon is called *topological order*[8]. Other necessary symptoms are excitations which cannot be created by any local operator – when these are particles, they are called *anyons*.

As an example of a state with topological order, consider the toric code, aka $\mathbb{Z}_2$ gauge theory. A representative groundstate wavefunction (the fixed point one) locally has the form

$$|\text{gs}, 0\rangle = \sum_{\text{closed loops},C} |C\rangle =$$

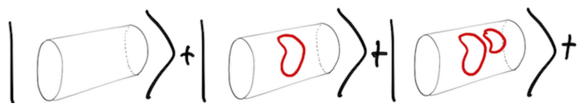

$\cdots$.

I say 'locally' because on a space with non-contractable loops, we get orthogonal groundstates by including (or not) the loops which wind around the non-contractible cycles. In the picture above, I tried to indicate that space is a cylinder, and I included only contractable loops in the sum. An orthogonal groundstate is

$$|\text{gs}, 1\rangle =$$

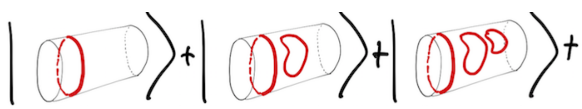

$\cdots$.

So the number of groundstates depends on the topology of space. States in these different sectors are only related by an operator that creates a whole large loop winding around the nontrivial cycle, hence not any local operator:

$$|\text{gs}, 1\rangle = \mathcal{O}\Big($$ 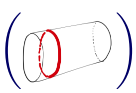 $$\Big) |\text{gs}, 0\rangle$$

The anyonic excitations arise by allowing the loops to end,

$$|\text{anyons at } x \text{ and } y\rangle = \mathcal{O}\Big($$ 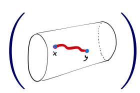 $$\Big) |\text{gs}\rangle$$

(where we can act on any of the groundstates). Such a state can arise as the groundstate of a spin system – it is a (gapped) *spin liquid* state; a Hamiltonian whose gapped groundstate this is will surely be written many times this week.[9]

---

[8]In order for this degeneracy to be stable, it must be that no local operator maps one of these groundstates to another. Suppose our hamiltonian H_0 happens to annihilate two states $|\psi_{1,2}\rangle$ (set the groundstate energy to zero). If we perturb H with any ΔH such that $\langle\psi_1| \Delta H |\psi_2\rangle$, the degeneracy will be split – the levels repel each other. More precisely, the matrix elements must vanish in the thermodynamic limit.

This property makes the groundstate subspace of a system with topological order into a quantum error-correcting code, with a *code distance* (the number of errors that can be corrected) of order the system size L.

[9]To see that topological order is a subtle thing that one might miss if asking the wrong questions,

The low-energy physics of gapped phases is often described by a (unitary) topological field theory, this is a theory of groundstates, and it can provide a way to distinguish states of matter. When this is case, the phase is called a *liquid*. Exceptions include fracton topological phases, where the lattice is not forgotten by the topological groundstates.

Edge modes. Another (distinct!) possibility is that even if the system in infinite space has an energy gap, if we cut the space open, new stuff can happen; for example there may be gapless edge modes. There is a class of topological states which are characterized by their edge modes, generally called *invertible states* (special cases include Symmetry-Protected Topological states (SPTs) and Topological Insulators). The edge modes of such a state must carry some property which cannot be recreated locally on the surface; such a property is called an *anomaly*[10]

Both of these phenomena (topological order and edge states) happen in (fractional) quantum Hall systems. A reason to think that an interface between the vacuum and a gapped state of matter which is distinct from the trivial one might carry gapless modes is that the couplings in the hamiltonian are forced to pass through the wall where the gap closes. In fact there are important exceptions to this conclusion, which you can learn about *e.g.* here [2, 3]. For example, the toric code admits gapped boundaries (interface with vacuum).

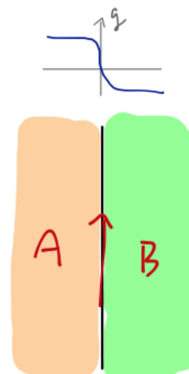

Gaplessness is something special that needs to be explained. An energy gap (and no topological order or special edge modes) should probably be the generic expectation for what happens if you pile together a bunch of degrees of freedom and couple them in some haphazard (even translation invariant) way. At the very least this follows on general grounds of pessimism: if you generically got something interesting by

consider the following model of spins at the sites of the square lattice: $\mathbf{H} = \sum_i X_i$, where X_i is the Pauli σ^x operator.

Its groundstate, written in the σ^z eigenbasis, is (the product state) $\otimes_i \left|\rightarrow\right\rangle_i \propto \prod_i \sum_{s_i=\uparrow,\downarrow} \left|\{s_i\}\right\rangle$. We can visualize this as a sum over closed loops: draw a loop around each region of $s_i = \downarrow$ (the red dots in the figure at right). They are closed loops because they are the boundaries of a region. Isn't this a toric code groundstate? No! One way to see the difference is that since these loops are defined as the boundaries of regions, they are always contractable.

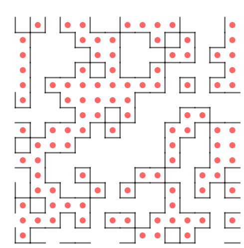

[10]This definition of SPT as a state characterized by its anomalous edge modes may be imperfect. There are examples of distinct states protected by lattice symmetries which seem not to have interesting edge modes. See appendix A of this paper or this paper. Thanks to Mike Hermele for bringing this phenomenon to my attention. A related exception would seem to be "higher-order topological insulators," which have no edge states, but have excitations localized to sharp corners of the boundary.

doing this, physics would be a lot easier (or more likely: we wouldn't find it interesting anymore). More seriously, gaplessness is an extreme case of a finite degeneracy: if allowed local operators could mix the the low-lying states, the levels would repel and create a gap under generic perturbations of the Hamiltonian.

Here is a list of some possible reasons for gaplessness (if you find another, you should write it down):

1. tuning to a critical point – notice that this requires some agent to do the tuning, and will only occur on some subspace of the space of couplings of nonzero codimension.

2. broken continuous symmetry (Goldstone bosons).

3. continuous unbroken gauge invariance (*e.g.* photons). Actually, [4] this is a special case of item 2 for continuous one-form symmetries.

4. Fermi surface (basically only in this case do we get gapless degrees of freedom at some locus of dimension greater than one in momentum space)

5. edge of an invertible topological phase. Here the gaplessness or degeneracy is protected by an anomaly.

6. a symmetry which forbids mass terms in some other way. This is called 'technical naturalness'. An example is unbroken chiral symmetry, which forbids fermion masses. Supersymmetry (where chiral symmetry prevents fermion masses, and supersymmetry relates boson masses to fermion masses) goes in this category.

7. CFT with no relevant operators. I am not sure if there are examples of this which are not examples of item 3. Sometimes this is called self-organized criticality. See here for a recent search.

Each entry in this list is something to be understood[11]. If you encounter a gapless model and it does not fit into this list then I will bet you $5 that it is fine tuned, meaning that its creator simply didn't add enough terms to the Hamiltonian.

We can no longer define the boundary of a gapless phase by a wall of gap-closing. Instead, a useful definition is that perturbation theory (in the difference of Hamiltonians) works within the phase. A phase is an equivalence class of hamiltonians such

[11] Note that the masslessness of the graviton is a mystery not obviously solved by an element of this list.

that within the phase, physics (thermodynamics, local operator expectations) varies smoothly[12].

Refinement by symmetry. Another important axis along which we may organize states of matter is by symmetry. Specifically, we can label states according to the symmetry group G that acts on their Hilbert space and restrict the space of Hamiltonians to those which are preserved by G. Here I am speaking about what are called *global symmetries*, that is, symmetries (not redundancies of our labelling, like gauge transformations).

There are many refinements of this organization. We can ask how the symmetry G is *realized*, in at least three senses:

1. most simply, what representations of the group appear in the system?
2. is the symmetry preserved by the groundstate? If not, this is called 'spontaneous symmetry breaking'.
3. is it 'on-site'? Alternatively, is it 'anomalous'? What are its anomaly coefficients? I'll postpone the explanation of these terms. The keywords associated with them are SPT (symmetry-protected topological) and SET (symmetry-enriched topological) phases. The latter is what happens when there is both topological order and some anomalous realization of symmetry.

Entanglement. An important perspective for organizing our understanding of quantum phases of matter – on which we focus for the rest of this lecture – is the amount and structure of entanglement in the groundstate.

A pure state is completely unentangled if it is a product state, $\otimes_x |s_x\rangle$. In a phase with a product-state representative, mean field theory applies. (And when it applies it is very useful: it predicts phase transitions and the associated critical theory of the order parameter, and even the excitation spectrum (*e.g.* spin waves).) Such states can be distinguished only by symmetries acting independently on each site. This problem is approximately solved (from the point of view of the experiment-free discussion of condensed matter physics we are having) by the representation theory of G. Elsewhere in such a phase, the effects of entanglement are a perturbative correction to non-universal quantities.

[12] An annoying fact is that sometimes within a phase there are observables which vary non-analytically across a point where the thermodynamics and all local observables are perfectly smooth. A classic example is the roughening transition of Wilson loops in lattice gauge theory. So when I say 'physics varies smoothly' I really mean local observables. For non-local Hamiltonians, we do not even know how to define a notion of phase.

We've just defined what is 'unentangled'. We will want to be more quantitative about entanglement between A and its complement $\bar{A}$; here $\mathcal{H} = A \otimes \bar{A}$ is any bipartition of the Hilbert space, such as a region of space. Since we are talking about pure states (as opposed to *e.g.* thermal density matrices), we can do this using the entanglement entropy:

$$\boldsymbol{\rho}_A \equiv \mathrm{tr}_{\bar{A}} \left|\psi\right\rangle \left\langle\psi\right| \qquad S_A \equiv -\mathrm{tr}\boldsymbol{\rho}_A \log \boldsymbol{\rho}_A \quad .$$

So: highly-entangled and mean-field are antonyms. The description in terms of weakly-interacting waves above an ordered groundstate breaks down when the entanglement matters. The frontier of our understanding is states of matter where quantum mechanics is essential, not just a correction that can be included perturbatively. This is now a big industry (*e.g.* [5, 6]) and I will try to give some flavor of it. The states of interest here are distinguished instead by their patterns of quantum entanglement. Furthermore, since such new states of matter are distinguished by different new kinds of orders, the phase transitions which separate them go beyond those described by fluctuations of local symmetry-breaking order parameters. This leads to new renormalization-group fixed points and new conformal field theories (CFTs).

2 Adiabatic continuation and local unitary circuits

[Zeng, Chen, Zhou, Wen, chapter 7] A useful alternative characterization of a gapped phase motivates this entanglement-based point of view.

First, a quantum phase is actually a property of the groundstate. For example, in the case of topological order, all of the data about the characteristic anyon excitations are encoded in the groundstate wavefunctions on a torus (see here) or indeed even a single wavefunction (see here or, most elegantly, here). We'll see some evidence below.

Claim: Two groundstates are representatives of the same phase[13] iff there is a quasi-local unitary circuit U of finite depth (the depth of a circuit is the (maximum) number of elementary gates acting on each site, 'finite' means independent of L, and I will explain quasi-local below) which maps one state to the other. In symbols, $[\mathbf{H}_0] = [\mathbf{H}_1] \Leftrightarrow \left|\psi(1)\right\rangle = U \left|\psi(0)\right\rangle$.

$\boxed{\Rightarrow}$ Suppose there is a path $\mathbf{H}(s)$ in the space of Hamiltonians starting at $\mathbf{H}_0$ (whose groundstate is $\left|\psi(0)\right\rangle$) and ending at H_1 (whose groundstate is $\left|\psi(1)\right\rangle$), with a gap for every s in between. In finite volume, the adiabatic theorem says we can construct

[13] In this discussion, we assume that the hamiltonians have a unique groundstate. So if we are talking about a phase with TO, we study it on a simply-connected space. The notion of phase is a local property.

a unitary which *probably* maps $|\psi(0)\rangle$ to $|\psi(1)\rangle$, namely slow-enough time-evolution along the path $\mathbf{H}(s)$,

$$\mathcal{T}e^{\mathbf{i}\int_0^1 dt\tilde{\mathbf{H}}(t)} |\psi(0)\rangle \propto |\psi(0)\rangle + \cdots . \tag{2.1}$$

Since the gap is independent of L, the required duration is too. The failure rate (the amplitude for the $\cdots$ in (2.1)), however, is extensive. This problem can be fixed by a procedure called *quasi-adiabatic filtering* introduced by Hastings (a review is here) – one can construct a modified family of Hamiltonians $\tilde{\mathbf{H}}(s)$ which are almost as local[14] but precisely map groundstates to groundstates (the idea is to filter out the contributions from the excited states to which non-adiabatic transitions can happen)

$$|\psi(1)\rangle = \mathcal{T}e^{\mathbf{i}\int_0^1 dt\tilde{\mathbf{H}}(t)} |\psi(0)\rangle .$$

(So really the title of this section should have been 'quasi-adiabatic continuation...'.)

Now this continuous time evolution can be *Trotterized*. That is, we can approximate it by a circuit, by breaking the time evolution into tiny steps; the range of the terms in the Hamiltonian then determines the range of the individual unitary gates. The crucial point is that finite time evolution (independent of L) means a finite number of layers of elementary gates – this is a finite-depth circuit.

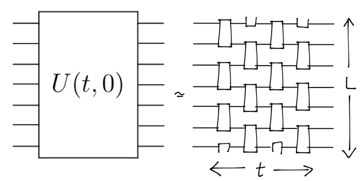

I will regard circuits and continuous time evolution unitaries as equivalent.

$\boxed{\Leftarrow}$ Given a circuit $U = \mathcal{T}e^{\mathbf{i}\int_0^1 dt\tilde{\mathbf{H}}(t)}$ that accomplishes $|\psi(1)\rangle = U|\psi(0)\rangle$, we can define $U(s) \equiv \mathcal{T}e^{\mathbf{i}\int_0^s dt\tilde{\mathbf{H}}(t)}$ and a family of states $|\psi(s)\rangle = U(s)|\psi(0)\rangle$. These states are the gapped groundstates of

$$\tilde{\mathbf{H}}(s) = \sum_x U(s)\tilde{\mathbf{H}}_x U(s)^\dagger$$

(gapped because the spectrum is independent of s) where $\mathbf{H}(0) = \sum_x \mathbf{H}_x$ is local, meaning that each $\mathbf{H}_x$ has finite range ξ (independent of L); the range $\tilde{\xi}$ of the terms in the filtered Hamiltonian $\tilde{\mathbf{H}} = \sum_x \tilde{\mathbf{H}}_x$ is still finite. But then the range of $U(s)\tilde{\mathbf{H}}_x U(s)^\dagger$ is bounded by $\tilde{\xi}+sv_{\max}$, where $v_{\max}$ is the maximum speed of propagation of correlations

[14] I've oversimplified the discussion here. Actually, there is a trade-off between locality of the filtered $\tilde{\mathbf{H}}$ and the precision with which the groundstates are mapped to each other. In fact, in order to precisely map the groundstates to each other, the operators $\tilde{\mathbf{H}}_x$ must have some tails, that is they have a profile which behaves like $e^{-r^{1-\delta}}$ where r is the distance from the point x – not quite exponential decay. This is the meaning of the modifier 'quasi-local'. Approximations to the exact map which are just as good for practical purposes can be made with $\tilde{H}_x$ which are strictly local. Thanks to Subir Sachdev for noticing this.

via $\tilde{\mathbf{H}}(t \leq s)$, which is again (according to the Lieb-Robinson bound) independent of L. ■

Notice that there is a lot of freedom in defining the unitary U which relates the two groundstates – we're actually only specifying its action on a single vector. What it does to the excited states (for example, the fact that it preserves the spectrum) is largely meaningless.

By a trivial phase we'll mean one with a representative groundstate which is a product state. This result implies that any groundstate in a nontrivial phase *cannot* be made from a product state by a finite-depth circuit. An example is a toric code ground state $\sum_{\text{loops},C} |C\rangle$.

3 Entanglement, short and long

Mean field theory is product states, which means there is no entanglement between regions of space at all. The next level of complication and interest to consider for possible groundstates of quantum many body systems is the case of states obtained by acting with a short-ranged quantum circuit of small depth on a product state. Let us consider such states, which are called short-range-entangled. What does their entanglement entropy of subregions look like and how do we distinguish which bits might be properties of a phase?

Let us focus on $d = 2$ space dimensions for definiteness. If the entanglement is short-ranged, we can construct a local 'entanglement entropy density' which is supported along the boundary of the region A [Grover-Turner-Vishwanath]:

$$S_A = \oint_{\partial A} s d\ell = \oint \left(\Lambda + bK + cK^2 + ...\right) d\ell = \Lambda \ell(\partial A) + \tilde{b} + \frac{\tilde{c}}{\ell(\partial A)} + ...$$

In the first step, we use the fact that the entanglement is localized at the boundary between the region and its complement. In the second step we parametrize the local entropy density functional in a derivative expansion; K is the extrinsic curvature of the boundary. Since the total system is in a pure state, $S(A) = S(\bar{A})$, which implies $b = 0$: since interchanging A and $\bar{A}$ reverses the orientation of the boundary, the extrinsic curvature cannot contribute. This means that the subsystem-size-independent term cannot come from terms local on the boundary; it is universal in the sense that it cannot be changed by changing the UV regulator (*e.g.* by rearranging lattice details).

Where can such a term come from? For the example of the groundstate of $\mathbb{Z}_2$ gauge theory (the toric code), a closed string that enters the region A must leave again. This is one missing bit of freedom for the reduced density matrix of A, which means a contribution to the EE that is independent of the size of A:

$$S_A = |\partial A|\Lambda - \log 2 \equiv |\partial A|\Lambda - \gamma \tag{3.1}$$

where the area-law coefficient Λ is some short-distance-dependent junk and γ is a universal characterization of the nature of the topological order.

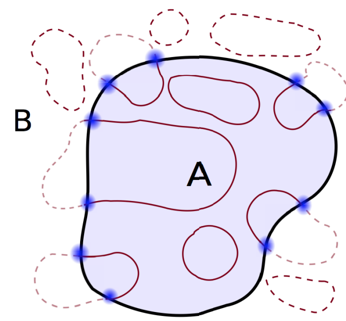

[fig: Tarun Grover]

This is true for each component of the boundary of A individually, so the generalization of (3.1) to regions with $b_0(\partial A)$ boundary components is $S(A) = |\partial A|\Lambda - \gamma b_0(\partial A)$.

The universal constant term γ is called the topological entanglement entropy (TEE)[15]. For more general topological orders, γ can related to the spectrum of anyons; for Abelian states γ is $\frac{1}{2}\log$ (#torus groundstates). A beautiful argument for this is the Kitaev-Preskill wormhole construction (see their Fig. 2).

It is instructive to try to combine entropies of different regions to isolate the TEE from the area-law junk.

If the entanglement is indeed all short-ranged, then for collections of regions where the boundaries cancel out, $\partial(AB) + \partial(BC) = \partial(B) + \partial(ABC)$, (such as in the figure at right) nothing will be left. Let $S(x)$ be the EE of the subregion x in the state in question.

$$I(A : C|B) := S(AB) + S(BC) - S(B) - S(ABC)$$

is the conditional mutual information – correlations between variables A and C if we knew B. In general this combination of entropies satisfies a deep inequality called Strong Subadditivity (SSA), $I(A : C|B) \geq 0$. In general gapped phases in 2d, for the arrangement of regions at right, $I(A : C|B) = 2\gamma$, where γ is the subleading term to the area law defined in (3.1). The area-law contributions cancel out pairwise (notice that the corners cancel too).

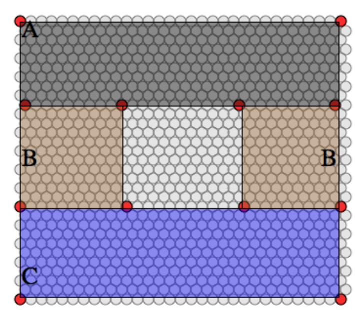

The regions should be large compared to the lattice spacing and the correlation length.

When $\gamma = 0$, SSA is saturated. $I(A : C|B) = 0$ means $\boldsymbol{\rho}_{ABC}$ is a 'quantum Markov

[15] It was introduced for $d = 2$ by Hamma-Ionicioiu-Zanardi, Kitaev-Preskill, Levin-Wen; the higher-dimensional generalizations are explained in the Grover et al paper linked above.

chain,' a state which can be reconstructed from its marginals $\boldsymbol{\rho}_A, \boldsymbol{\rho}_B, \boldsymbol{\rho}_C$ (by a formula due to Petz). So the quantity γ is an obstruction to this automatic reconstruction of the global state from local data.

The above argument shows that the TEE is not a short-distance artifact, but is it a property of a phase for any choice of A, B, C? And is it only nonzero for states with topological order? Almost. The papers linked above argue – assuming that the system is a liquid – that the TEE is independent of small changes in the regions (using $S_A = S_{\bar{A}}$ for pure states) and therefore insensitive to changes in the Hamiltonian that keep the correlation length short. There is, however, an important exception if the phase is not a liquid, whereby small changes of the regions lead the TEE to jump, and to give nonzero answers in states without TO.

In $d = 3$, ∂A is characterized by its number of components b_0 and its number of noncontractable loops b_1; these are related by $\chi = 2b_0 - b_1 = V - E + F = \frac{1}{2\pi}\int_{\partial A} R$ (the Gauss-Bonnet theorem) to the integral of a local density. The EE of A is linear in b_0 and b_1 (see Appendix E of the Grover-Turner-Vishwanath paper) but only one combination of them is a signature of long-range entanglement. Again this 3d TEE can be extracted by combining regions whose boundaries and corners cancel.

The TEE is only one number characterizing the nature of the topological order, and by no means uniquely characterizes it. For example, the double semion state is a distinct topological order from the toric code in $d = 2$, whose representative wavefunction is $\sum_{\text{closed loops},C}(-1)^{b_0(C)} |C\rangle$ (where $b_0(C)$ is the number of components of the loops). As you can see from the form of the wavefunction it also has four groundstates on the torus and hence the same TEE. However, by now humans have learned to extract a great deal of the data specifying a given topological order from the entanglement properties of a single wavefunction, the most advanced incarnation of which is the *entanglement bootstrap*.

4 Entanglement in groundstates of local Hamiltonians

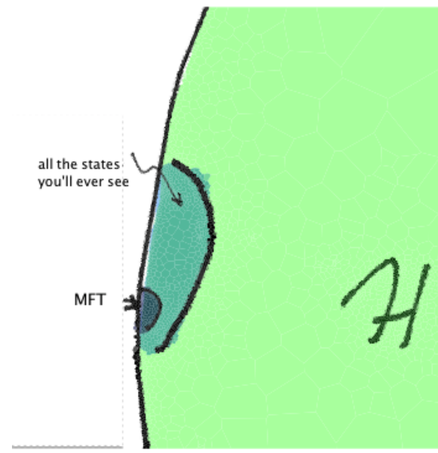

We began our discussion by saying that $\mathcal{H} = \otimes_x \mathcal{H}_x$ is the Hilbert space. However, most of many-body Hilbert space is fictional, in the sense that it cannot be reached from a product state by time evolution with local Hamiltonians in a time polynomial in the system size. The representation of the wavefunction

$$|\psi\rangle = \prod_i \sum_{s_i=\pm} c_{s_1 \ldots s_{2L^d}} \left| s_1 \ldots s_{2L^d} \right\rangle$$

as a vector of $e^{L^d \log 2}$ complex numbers $c_{s_1 \ldots s_{2L^d}}$ is not useful. It's too many numbers, and you can't get there from here. (For more rhetoric along these lines, I recommend *e.g.* [7].)[16]

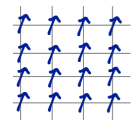

Mean field theory means unentangled states of the sites, of the form $|\psi_{\rm MF}\rangle = \otimes_i \left(\sum_{s_i=\pm} c_{s_i} |s_i\rangle\right)$. Such a state depends on only L^d numbers c_{s_i}; this is too far in the other direction.

There is more in the world: entangled groundstates (even if only short-range-entangled) mean new phenomena, *e.g.* SPT states (whose entanglement looks some-

[16]The result of [7] is that the overwhelming majority of states cannot be reached from a product state by time evolution with a k-local Hamiltonian in a time polynomial in L. Here k-local is a generalization of local which merely demands that each term in the Hamiltonian only act on $\leq k$ sites at a time, but does not require them to be nearby in any metric. The argument is a simple counting argument, similar to the proof that most boolean functions on n variables are not efficiently computable. For a summary, see here, §1.3.

The property that the inaccessible states have is not so much that they have volume-law entanglement, but rather that they have high *complexity*. You may be tempted to argue with this definition of 'accessible' in terms of time-evolution from a product state. However, I claim that all experiments actually proceed in this way: the experimenter decides (to some extent) what is the Hamiltonian, and then waits while the system (defined to include all degrees of freedom entangled with the system) evolves itself unitarily. If we could implement other operations, such as euclidean time evolution, we could do more powerful things, in fact so powerful as to bump up against complexity-theoretic lore suggesting that it is impossible.

Notice also that I am not saying that all volume law states are inaccessible! We expect that time evolution by a chaotic Hamiltonian for a time polynomial in L will produce a state which locally looks thermal. The statement is merely that these accessible volume-law states are special, low-complexity states.

Thanks to Tarun Grover and Dam Son for discussions of this subject.

thing like this: 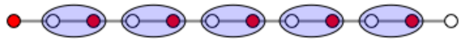) topological order, quantum critical points, ... An energy gap implies correlations of local operators are short-ranged, $\langle \mathcal{O}(x)\mathcal{O}(0)\rangle \sim e^{-|x|/\xi}$. But the entanglement needn't be, as we've seen in the toric code example.

> Problem: characterize the physical corner of $\mathcal{H}$ by its entanglement properties and parameterize it efficiently, in particular in a way which allows observables to be efficiently calculated given the wavefunction.

Groundstates of local Hamiltonians are special: generically (with few, well-understood exceptions) the entanglement entropy of large-enough subregions satisfies an area law. (Large-enough means large compared to the correlation length, so that the massive particles can't get involved.) This means S_A scales like the area of the boundary of the region in question. The idea is just that minimizing a local Hamiltonian demands that the strongest entanglement is between nearest neighbors. I will state it as a
Basic expectation (**area law**): In groundstates of local, motivic[17] Hamiltonians ,

$$S_A = aR^{d-1} + \text{smaller}$$

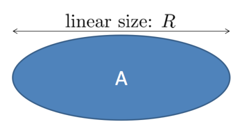

where R is the linear size of the subregion A, say its diameter. This statement is supported by a great deal of evidence and has been rigorously proved for gapped systems in 1d [8, 9]. The proof uses the Lieb-Robinson bound on the spread of correlations. We will give a more general explanation below, following [10]. The area law has been essential in identifying efficient numerical representations of groundstates in terms of tensor networks, and in the development algorithms for finding them (a useful introduction is [11, 12] or [13]).

In $d = 1$, the area law implies a MPS (matrix product state) representation of the groundstate (for a review of this subject, I recommend Ref. [11]). This is the form of the groundstate that is output by DMRG algorithms (for more explanation, see *e.g.* [14]).

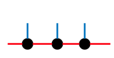 $= \sum_{a_{1,2\ldots}=1}^{\chi} M^{\sigma_1}_{a_1 a_2} M^{\sigma_2}_{a_2 a_3} \cdots |\sigma_1, \sigma_2 \cdots\rangle$ 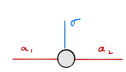 $\equiv M^{\sigma}_{a_1 a_2}$

χ, the range of the auxiliary index, is called the *bond dimension*. This encodes the groundstate in $L^{d=1}\chi^2$ numbers. In such a state, each site is manifestly entangled with the rest of the system only through its neighbors.

[17] By this term I mean that even if there is no translation invariance, the form of the Hamiltonian is the same in any region of space.

In $d > 1$ more can happen. A direct generalization of an MPS to $d > 1$ is called a 'PEPS' and looks like this: 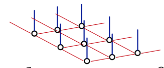. Such a state manifestly has an area law: the entanglement entropy of a subregion of such a state is bounded above by the number of (red) bonds cut, times $\log \chi$.

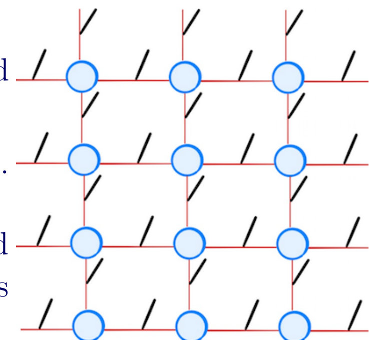

Such a state is not necessarily short-range entangled, however. For example, at right is a PEPS for the toric code groundstate. Here the physical degrees of freedom live on the links; each link is a two state system, where $i, j = 0, 1$ indicate the absence or presence of a string, respectively. The tensors involved are just 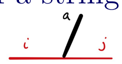 $= \delta_{ij}\delta_{ia}$ and 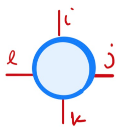 $= 1$ if $i + j + k + l$ is even, and 0 otherwise. So the tensor at the vertices guarantees that only closed strings appear. This tensor has a simple representation as 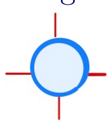 $=$ 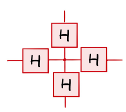 where $\mathsf{H}_{ij} = (-1)^{ij}$ is proportional to the Hadamard gate, and the tensor at the center again sets all the incoming legs equal in the given basis.

The entangling power of PEPS can be thought of as arising from *post-selection* – we prepare a state of an auxiliary, larger system, and then project it into the physical Hilbert space ('PEPS' stands for 'projected entangled pair states'). Realizing a state by this method in an experiment requires doing it over and over until the desired (unlikely) outcome obtains. The computational power of such methods is quite strong[18].

Even if the bond dimension χ is small, such a network is slow to contract, which one needs to do to compute matrix elements, and to determine the values of the tensors, for example in a variational calculation. Worse, even with a gap, rigorous results only show that there exists a PEPS with $\chi \sim e^{\log^d(L)}$, growing with linear system size L.

So numerical methods which incorporate only this datum (the area law) struggle with gapless states in $d = 1$ and even with gapped states in $d > 1$. A more refined statement takes into account how much entanglement there is at each length scale. Incorporating this extra data allows one to make efficiently-contractible networks. The process of organizing our understanding of the entanglement in the state scale-by-scale is called *entanglement renormalization*. The best-developed implementation of this

[18]This fact has an interesting geometric interpretation via holographic duality.

idea is MERA (the multicale entanglement renormalization ansatz) [15], which is the state-of-the-art method for the study of 1d quantum critical points [16, 17].

Exceptions to the area law. The 'expectations' above are often correct (even beyond the examples where they are precisely true), but there are some real exceptions to the area law expectation, even for groundstates of local Hamiltonians: groundstates at quantum critical points in $d = 1$ have

$$S(A) = \frac{c}{3} \log R/\epsilon \tag{4.1}$$

(here A is a single interval of length R) whereas the $d = 1$ area law would be independent of R. ϵ is a short-distance cutoff, and c, the 'central charge', is a measure of the number of critical degrees of freedom. This includes the well-studied case of $1+1$-dimensional conformal field theory, where much can be said (if you are impatient, look here). A positive way to look at this is that the entanglement entropy of subregions can diagnose a continuous phase transition. Another class of examples of area law violation in $d = 1$ arises from highly disordered systems, namely random singlet states.

In $d > 1$, even critical points (are expected to) satisfy the area law[19]. An important class of exceptions to the area law in any dimension is metallic groundstates of fermions, such as free fermions in partially-filled bands. This also leads to super-area-law scaling:

$$S_{\text{metal}}(R) \sim (k_F R)^{d-1} \log(k_F R) \tag{4.2}$$

– a logarithmic violation, where the Fermi momentum k_F makes up the dimensions. This result can be understood from the $1 + 1$-d CFT case, as explained here. The idea is that each point on the Fermi surface behaves like a $1 + 1$d CFT; (4.2) results from adding up these contributions.[20]

[19] There is nice way to visualize the difference between $d = 1$ and $d \geq 1$ here. (I learned this argument from here.) Critical points are scale-invariant, so let's demand a scale-invariant representation of the groundstate. Such a thing is a MERA network, which is like a discretization of hyperbolic space. The entanglement entropy of a subregion is bounded above by the number of bonds in the network that we have to cut to separate the subregion from the rest of the network. It is a purely geometric property of hyperbolic space that for $d = 1$, this number grows like $\log R$, while in $d > 1$ it is just R^{d-1}, the area law term.

For circular regions in Lorentz-invariant theories (not necessarily critical), there is a better proof of the area law, which I'll mention at the end of 5.

[20] More precisely, (4.2) was shown for free fermions in [18, 19], which conjectured a nice expression, called the Widom formula, for general shape of Fermi surface and of the region. The appealing picture of the violation in terms of 1d systems at each point on the Fermi surface was developed in [20, 21]. A strong numerical test of this picture was made in [22]. This allows for an extension to non-Fermi liquids (where the CFT at each $\vec{k}_F$ is not just free fermions), a result which was confirmed numerically in [22].

Non-groundstates. And of course there is more in the world than groundstates. The first excited state of a many body system has zero energy density in the thermodynamic limit. (Often it is a single particle.) Such states will still have an area law if the groundstate did. But in general, states with finite energy density and finite temperature states will have volume-law behavior of the entanglement entropy. The coefficient of the volume-law is the thermal entropy density, which grows with temperature.

Is this because thermal states are highly entangled states, becoming more quantum at high temperature? No. Look at a purification of thermal state:

$$|\sqrt{\boldsymbol{\rho}}\rangle = Z^{-1/2} e^{-\frac{1}{2}\beta \mathbf{H} \otimes \mathbb{1}} \sum_i \frac{1}{\sqrt{|\mathcal{H}|}} |i\rangle_1 |i\rangle_2 = Z^{-1/2} \sum_E e^{-\frac{1}{2}\beta E} |E\rangle_1 |E\rangle_2 . \tag{4.3}$$

By 'purification' I mean that if we trace out one of the copies we get back $\text{tr}_2 |\sqrt{\boldsymbol{\rho}}\rangle\langle\sqrt{\boldsymbol{\rho}}| = \boldsymbol{\rho}$. The maximally entangled state here can be in any basis; we can choose it to be a local basis: each site is strongly entangled with its purifying partner. The ancilllary Hilbert space doing the purifying is really just a proxy for a thermal bath and we are using our freedom to mess with the purification to make it look nice (like a copy of our system). (Beware that the individual object in the intermediate step I've written in (4.3) are maybe not so well-defined in the thermodynamic limit.) The lesson I am trying to convey is: Volume law entanglement entropy of thermal states can be regarded as (short-ranged) entanglement with the thermal bath, rather than entanglement between parts of the system itself. For such mixed states, it is necessary to use more sophisticated measures to isolate the quantum entanglement from this kind of entropy of mixture, such as the entanglement negativity (for a recent application, see here).

s-sourcery. Suppose you had a finite-depth unitary circuit $\mathbf{U}$ which doubles the system size. Notice that being unitary, it will have to act on not just the system itself, but also some initially-unentangled ancilla bits. For systems without much entanglement, we can find such a $\mathbf{U}$ which just starts with the unentangled bits and produces the groundstate at linear size $2L$, but for nontrivial (liquid) phases we need a copy of the original system at size L. Rather amazingly, there are some systems for which we need $s = 2$ copies at size L to produce a single copy at size L. This is the case for fracton topological phases, as shown here and here.

Such a circuit produces a counting of groundstates as a function of system size. The groundstate degeneracy satisfies: $G(2L) = G(L)^s$. This means that states with $s > 1$ must necessarily have strange growth of G with system size, as do fracton phases.

The existence of such a circuit $\mathbf{U}$ controls the growth of entanglement with system size. With sufficient locality properties, it implies recursive bounds on the entropy of subregions:

$$S(2R) \leq sS(R) + kR^{d-1}$$

$$S(2R) \geq sS(R) - k'R^{d-1}$$

for some constants k, k'. (The proof uses the Small Incremental Entangling result of [23].) For $s \leq 2^{d-1}$ and $d > 1$, (4.4) implies the area law. In fact, the existence of such a $\mathbf{U}$ with $s = 1$ can be shown using quasiadiabatic continuation for all known gapped liquid states, so this is a proof of the area law for all those states.

Finally, the s-sourcery circuit $\mathbf{U}$ can be used to construct a MERA, by Trotterization.

5 Renormalization group monotones from entanglement measures

The renormalization group (RG) is a procedure by which degrees of freedom are thinned; starting from a microscopic theory of all degrees of freedom, it is possible to coarse-grain our description in order to obtain a (different) theory of just the long-wavelength degrees of freedom. This procedure is hard to do in practice, and it is useful to know about quantities which behave monotonically under this process; such a quantity is then naturally regarded as a measure of the number of degrees of freedom.

The quantity c appearing in (4.1), the entanglement entropy of a single interval in the groundstate of a 1+1d CFT is such a quantity. It was proved to be an RG monotone by Zamolodchikov long ago. A proof of its monotonicity using Lorentz invariance and SSA was found by Casini and Huerta.

Proof of entropic c-theorem in $d = 1 + 1$. We will show that the entanglement entropy of an interval of length r satisfies

$$0 = \left(\frac{d}{d \log r} \right)^2 S(r). \tag{5.1}$$

This says that $c(r) \equiv 3r\partial_r S(r)$ satisfies $c'(r)/3 = rS'' - S' \leq 0$. In particular, if $S(r) = \frac{c}{3} \log r/\epsilon$, then $c(r) = c$ is the central charge.

First a Minkowski space geometry exercise. Consider an interval d, and the inward light rays from its endpoints b_1, c_1. Extend these for a time t. The interval connecting the endpoints of the lightrays is a. Let b and c be the spacelike surfaces in the figure. Lorentzian geometry then implies $bc = ad$, where the letters indicate the proper lengths of the associated segments. This says $c = \lambda a, d = \lambda b, \lambda = \frac{c}{a} = \frac{d}{b} \geq 1$.

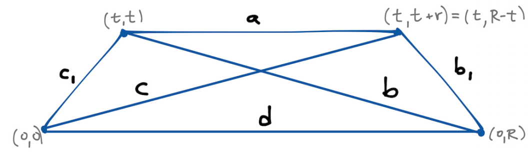

More explicitly, if $d \equiv |d^\mu d_\mu| = R, a = r$, then $c^\mu = (t, t+r)^\mu = (t, R-t)^\mu$ has $c^2 = |c^\mu c_\mu| = r(r-2t) = rR = |b^2| = bc$.

SSA says

$$\underbrace{S(c_1 a)}_{=S(c)} + \underbrace{S(ab_1)}_{=S(b)} \geq S(a) + S(d).$$

The underbraced equations are consequences of Lorentz invariance. Then

$$S(b) - S(a) \geq S(\lambda b) - S(\lambda a).$$

Notice that the area law term cancels in the differences. If we set $\lambda = 1 + \epsilon$, (5.1) follows.

This paper shows that the assumption of Lorentz invariance is necessary.

In $d = 2, 3$, universal terms in the entanglement entropy of subregions (in particular the quantity γ studied above in $d = 2$) have also been shown to be RG monotones, by a more complicated argument along the same lines. This paper is a good place to look.

What's the big deal about RG monotones? Often we are interested in some particular microscopic model (*e.g.* a lattice model or a field theory) and would like to know what is a useful description of its long-wavelength physics. Knowing an RG monotone can put useful constraints on what physics can emerge. A nice example is this paper which uses these results to constrain the emergence of algebraic spin liquids. In high-energy physics terms, the question is: how many flavors of charged Dirac fermion are required to prevent $\mathsf{U}(1)$ gauge theory in $2+1$d from confining.

By the way, the proof of monotonicity of γ in $d = 2$ for Lorentz-invariant theories implies that they satisfy the area law[21]. It shows, for $A =$ a round disk, that $\frac{\partial^2 S}{\partial R^2} \leq 0$, which means that $S(R)$ cannot grow faster than R, the area law behavior in $d = 2$.

[21] Thanks to Tarun Grover for reminding me.

6 Omissions

- $T \neq 0$ and entanglement of mixed states. Nearly all of the long-range entangled phases we know have a phase transition at $T = 0$ as the temperature is raised, to a phase which is adiabiatically connected to $T = \infty$ – an unentangled state. It is sad.

- Non-equilibrium matter can also be ultra-quantum (see Matthew Fisher's lecture).

- Entanglement and higher form symmetries.

- Entanglement is a useful tool for labelling gapless topological phases (for example, here).

Acknowledgements. Thanks to Tarun Grover, Mike Hermele, Ethan Lake, Subir Sachdev, Dam Son, Liujun Zou for helpful questions and comments, and to Nat Tantivasakadarn for fielding many questions during the lecture. This work was supported in part by funds provided by the U.S. Department of Energy (D.O.E.) under cooperative research agreement DE-SC0009919, and by the Simons Collaboration on Ultra-Quantum Matter, which is a grant from the Simons Foundation (651440, JM).

References

[1] J. Haegeman, S. Michalakis, B. Nachtergaele, T. J. Osborne, N. Schuch, and F. Verstraete, "Elementary Excitations in Gapped Quantum Spin Systems," *Physical Review Letters* **111** (Aug., 2013) 080401, 1305.2176. 4

[2] E. Dennis, A. Kitaev, A. Landahl, and J. Preskill, "Topological quantum memory," *J. Math. Phys.* **43** (2002) 4452–4505, quant-ph/0110143. 7

[3] A. Kitaev, "Toward Topological Classification of Phases with Short-range Entanglement," *unpublished* (2011) http://online.kitp.ucsb.edu/online/topomat11/kitaev/. 7

[4] D. Gaiotto, A. Kapustin, N. Seiberg, and B. Willett, "Generalized Global Symmetries," *JHEP* **1502** (2015) 172, 1412.5148. 8

[5] X.-G. Wen, "Topological order: from long-range entangled quantum matter to an unification of light and electrons," 1210.1281. 10

[6] T. Grover, Y. Zhang, and A. Vishwanath, "Entanglement entropy as a portal to the physics of quantum spin liquids," *New Journal of Physics* **15** (Feb., 2013) 025002, 1302.0899. 10

[7] D. Poulin, A. Qarry, R. Somma, and F. Verstraete, "Quantum Simulation of Time-Dependent Hamiltonians and the Convenient Illusion of Hilbert Space," *Phys.Rev.Lett.* **106** (2011) 170501, 1102.1360. 14

[8] M. B. Hastings, "An area law for one-dimensional quantum systems," *Journal of Statistical Mechanics: Theory and Experiment* **8** (Aug., 2007) 24, 0705.2024. 16

[9] I. Arad, A. Kitaev, Z. Landau, and U. Vazirani, "An area law and sub-exponential algorithm for 1D systems," *ArXiv e-prints* (Jan., 2013) 1301.1162. 16

[10] B. Swingle and J. McGreevy, "Renormalization group constructions of topological quantum liquids and beyond," *ArXiv e-prints* (July, 2014) 1407.8203. 16

[11] R. Orus, "A Practical Introduction to Tensor Networks: Matrix Product States and Projected Entangled Pair States," *Annals Phys.* **349** (2014) 117–158, 1306.2164. 16

[12] R. Orus, "Advances on Tensor Network Theory: Symmetries, Fermions, Entanglement, and Holography," *Eur.Phys.J.* **B87** (2014) 280, 1407.6552. 16

[13] J. C. Bridgeman and C. T. Chubb, "Hand-waving and Interpretive Dance: An Introductory Course on Tensor Networks," *ArXiv e-prints* (Mar., 2016) 1603.03039. 16

[14] U. Schollwöck, "The density-matrix renormalization group in the age of matrix product states," *Annals of Physics* **326** (Jan., 2011) 96–192, 1008.3477. 16

[15] G. Vidal, "Class of Quantum Many-Body States That Can Be Efficiently Simulated," *Phys. Rev. Lett.* **101** (Sep, 2008) 110501, http://link.aps.org/doi/10.1103/PhysRevLett.101.110501. 17

[16] R. N. Pfeifer, G. Evenbly, and G. Vidal, "Entanglement renormalization, scale invariance, and quantum criticality," *Phys.Rev.* **A79** (2009) 040301, 0810.0580. 17

[17] G. Evenbly and G. Vidal, "Quantum Criticality with the Multi-scale Entanglement Renormalization Ansatz," *ArXiv e-prints* (Sept., 2011) 1109.5334. 17

[18] D. Gioev and I. Klich, "Entanglement Entropy of Fermions in Any Dimension and the Widom Conjecture," *Phys. Rev. Lett.* **96** (Mar, 2006) 100503, http://link.aps.org/doi/10.1103/PhysRevLett.96.100503. 18

[19] M. M. Wolf, "Violation of the Entropic Area Law for Fermions," *Phys. Rev. Lett.* **96** (Jan, 2006) 010404, http://link.aps.org/doi/10.1103/PhysRevLett.96.010404. 18

[20] B. Swingle, "Entanglement Entropy and the Fermi Surface," *Phys.Rev.Lett.* **105** (2010) 050502, 0908.1724. 18

[21] B. Swingle, "Conformal Field Theory on the Fermi Surface," *Phys.Rev.* **B86** (2012) 035116, 1002.4635. 18

[22] Y. Zhang, T. Grover, and A. Vishwanath, "Entanglement Entropy of Critical Spin Liquids," *Physical Review Letters* **107** (Aug., 2011) 067202, 1102.0350. 18

[23] K. Van Acoleyen, M. Mariën, and F. Verstraete, "Entanglement Rates and Area Laws," *Physical Review Letters* **111** (Oct., 2013) 170501, 1304.5931. 19