19 Entanglement Entropy in Quantum Field Theory

So far we have discussed entanglement in ordinary quantum mechanics, where the Hilbert space of a finite region is finite dimensional. Now we will discuss geometric entanglement entropy in quantum field theory. Space (not spacetime) is divided into two regions, A and B, by a continuous curve:

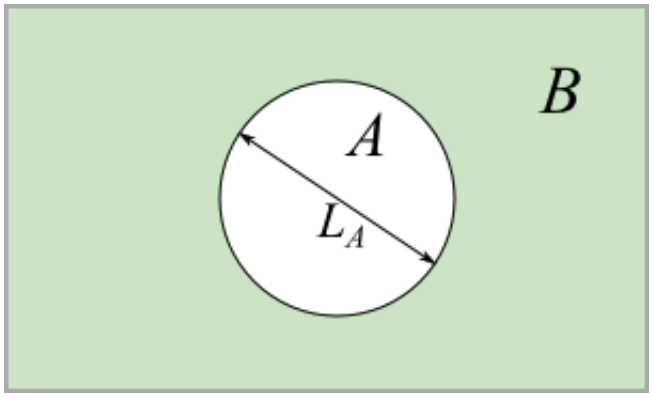

(19.1)

This picture is at a fixed time. Region A is drawn as a circle, but for now it could be any shape. (It could also be disjoint, but we will assume it is connected unless specified otherwise.)

Quantum field theory is strictly speaking *not* bipartite,

$$\mathcal{H}_{AB} \neq \mathcal{H}_A \otimes \mathcal{H}_B \ . \tag{19.2}$$

There are two things to worry about: first, in gauge theories, you cannot really localize states. The gauge constraint is applied to the full system, so by looking at any subregion, you cannot decide whether it is a physical state obeying the constraint. This issue (which also appears in ordinary quantum mechanical gauge systems) has been addressed in some nice papers just in the last year or so, and we will ignore it entirely. It turns out to not affect the discussion that follows very much.

The second issue is UV divergences. In a continuum QFT there are UV modes at arbitrarily small scales across the dividing surface ∂A, and this makes it impossible to actually split the full Hilbert space. To deal with this, we must impose a UV cutoff

by introducing the 'lattice scale' ϵ_{UV}. With a finite cutoff, the Hilbert space of a finite region is finite-dimensional, and most of the results of the previous section — in particular, positivity of relative entropy and strong subadditivity — apply to QFT. In the end we usually want to regulate the divergences somehow, but the leftover finite pieces do not immediately obey the same properties, so we need to be careful about tracking cutoff dependence throughout the problem.

19.1 Structure of the Entanglement Entropy

The divergent terms in S_A come from UV physics. In the UV, any finite energy state is the same as the vacuum state. Therefore to discuss the structure of the divergent terms we can restrict to $\rho = |0\rangle\langle 0|$, the vacuum state of the full system.

UV divergences

The divergent terms depend on the theory and on the shape of region A. In a local QFT, we expect the divergent piece to be a *local* integral over the entangling surface ∂A,

$$S_A^{(div)} \sim \int_{\partial A} d^{d-2}\sigma\sqrt{h}\, F[K_{ab}, h_{hab}] \ , \tag{19.3}$$

where F is some (theory-dependent) functional of the extrinsic curvature and induced metric on ∂A. This is for the same reason that when we do renormalization, we are only allowed to add local counterterms to the Lagrangian; non-local terms come from IR physics.

Let's organize (19.3) as an expansion in powers of K_{ab}. Since $K_{ab} \sim 1/L_A$, this is an expansion in powers of the size L_A. What sort of terms can appear? In a pure state, $S_A = S_B$, an in particular $S_A^{(div)} = S_B^{(div)}$. The extrinsic curvature is $K \sim \nabla n$ with n the unit normal; this flips sign if we consider region A vs its complement, region B. Therefore $S_A^{(div)} = S_B^{(div)}$ implies that only even powers of K_{ab} are allowed:

$$S_A^{(div)} \sim a_1 L_A^{d-2} + a_2 L_A^{d-4} + \cdots \ , \tag{19.4}$$

where a_i depend on the theory but not on L_A.

The leading term in (19.4) is a UV divergence proportional to Area(A). This makes sense: UV modes entangled across ∂A give a divergent contribution, and the number of these modes is proportional to the area.

General structure and universal terms

Now let us further assume the theory is scale invariant. In the vacuum state of a scale invariant theory, the only scales in the problem are ϵ_{UV} and L_A. Therefore, by dimensional analysis, $a_1 \sim \epsilon_{UV}^{2-d}$, $a_2 \sim \epsilon_{UV}^{4-d}$, etc. Thus, allowing also for a finite contribution, we find the general behavior of the entanglement entropy in a CFT. In odd dimensions d:

$$S_A^{CFT} \sim b_{d-2}\left(\frac{L_A}{\epsilon_{UV}}\right)^{d-2} + b_{d-4}\left(\frac{L_A}{\epsilon_{UV}}\right)^{d-4} + \cdots + b_1\frac{L_A}{\epsilon_{UV}} + (-1)^{\frac{d-1}{2}}\tilde{S} + O(\epsilon_{UV}) \ , \quad (19.5)$$

and in even dimensions:

$$\begin{aligned} S_A^{CFT} \ \sim \ & b_{d-2}\left(\frac{L_A}{\epsilon_{UV}}\right)^{d-2} + b_{d-4}\left(\frac{L_A}{\epsilon_{UV}}\right)^{d-4} + \cdots + b_2\left(\frac{L_A}{\epsilon_{UV}}\right)^{2} \\ & + (-1)^{\frac{d-2}{2}}\tilde{S}\log\frac{L_A}{\epsilon_{UV}} + \text{const} + O(\epsilon_{UV}) \ , \end{aligned} \quad (19.6)$$

The difference between even and odd comes from the fact that the "$1/\epsilon_{UV}^0$" term that would appear in even dimensions actually turns into a log divergence (just as it does in Feynman diagrams). The powers of (-1) are inserted by convention.

In the vacuum state, the b_i and $\tilde{S}$ depend on the theory, but not on L_A or ϵ_{UV}.

In a non-scale-invariant QFT, or in an excited state of a CFT, there are other scales.* So in general, $\tilde{S}$ depends on the theory, the shape, and the state ρ_{total}. Furthermore, $\tilde{S}$ is *universal* in the sense that it does not depend on ambiguities in the choice of regulator. For this reason it is sometimes called the *renormalized entanglement entropy.* †

*This is sometimes confusing in the CFT case but obviously true: even in a scale invariant theory, an excited state with lumps of stuff 1 meter apart is different from an excited state with lumps of stuff 2 meters apart!

†The fact that it is independent of regulator is clear in the even dimensional case, since it is the coefficient of a log. It is less clear for the constant term odd dimensions, since evidently shifting $\epsilon_{UV} \to \epsilon_{UV} + a$ would change the finite term. In practice it seems to be well defined in CFT for reasons I won't get into here, but I'm not sure about the non-conformal case.

Area vs volume terms

The leading UV divergence is always proportional to Area(A), in any state. In the vacuum we do not expect any extensive contribution to $\tilde{S}$, but in a random excited state, we expect

$$\tilde{S} \sim \text{Volume}(A) \ . \tag{19.7}$$

This is for the same reason that we argued for volume scaling in a random state of a lattice system. In a highly excited random state, the IR modes that contribute to $\tilde{S}$ should all be highly entangled with the outside, and the number of such modes scales with volume.

Example: 2d CFT in vacuum

As a simple example, consider a 2d CFT in the vacuum state of the full system. Space is a line, and region A is an interval of length L_A. In this case the entanglement entropy can be computed exactly (we will do this calculation later in the course) with the result

$$S_A = \frac{c}{3} \log \frac{L_A}{\epsilon_{UV}} \ . \tag{19.8}$$

Here c is the central charge of the CFT (which, remember, roughly speaking counts the degrees of freedom). This agrees with the general formula in even spacetime dimensions (19.6), with $\tilde{S} = \frac{c}{3}$.

If we instead consider a highly excited state, then we can't do the calculation in general, but in cases where it can be done the result in a typical state scales as $\tilde{S} \sim cL_A$.

19.2 Lorentz invariance

In a Lorentz-invariant QFT, the density matrix of a spatial region A must contain all of the same information as the density matrix of a spatial region A' that shares the

same causal diamond. That is, for this setup:

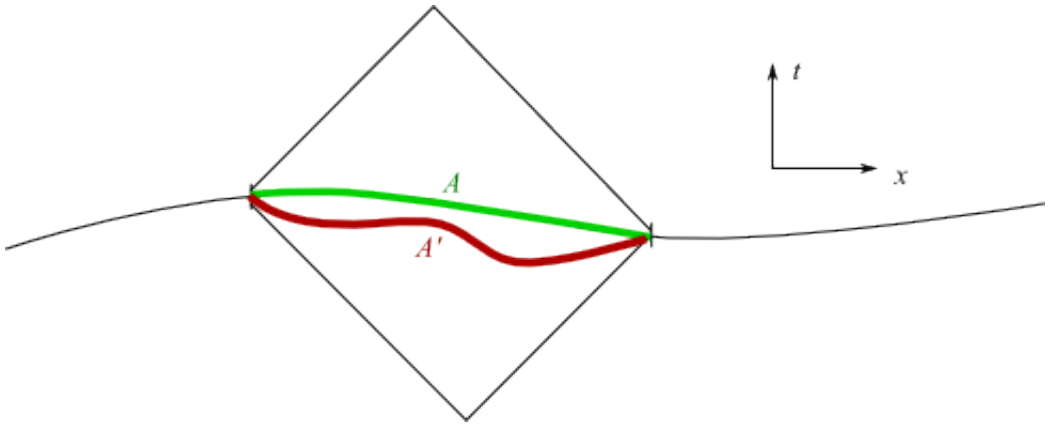

(19.9)

we must have

$$S_A = S_{A'} \ . \tag{19.10}$$

In words, this is because if we know everything about A, we can time-evolve to learn everything about A'. In formulas, it is because the reduced density matrices are related by (perhaps very complicated and nonlocal!) unitary operation,

$$\rho_{A'} = U^\dagger \rho_A U \ . \tag{19.11}$$