18 Introduction to Entanglement Entropy

The next few lectures are on entanglement entropy in quantum mechanics, in quantum field theory, and finally in quantum gravity. Here's a brief preview: Entanglement entropy is a measure of how quantum information is stored in a quantum state. With some care, it can be defined in quantum field theory, and although it is difficult to calculate, it can be used to gain insight into fundamental questions like the nature of the renormalization group. In holographic systems, entanglement entropy is encoded in geometric features of the bulk geometry.

We will start at the beginning with discrete quantum systems and work our way up to quantum gravity.

References: Harlow's lectures on quantum information in quantum gravity, available on the arxiv, may be useful. See also Nielsen and Chuang's introductory book on quantum information for derivations of various statements about matrices, traces, positivity, etc.

18.1 Definition and Basics

A *bipartite system* is a system with Hilbert space equal to the direct product of two factors,

$$\mathcal{H}_{AB} = \mathcal{H}_A \otimes \mathcal{H}_B \ . \tag{18.1}$$

Starting with a general (pure or mixed) state of the full system ρ, the reduced density matrix of a subsystem is defined by the partial trace,

$$\rho_A = \mathrm{tr}_B\, \rho \tag{18.2}$$

and the *entanglement entropy* is the von Neumann entropy of the reduced density matrix,

$$S_A \equiv -\,\mathrm{tr}\, \rho_A \log \rho_A \ . \tag{18.3}$$

Example: 2 qubit system

If each subsystem A or B is a single qubit, then the Hilbert space of the full system is spanned by

$$|00\rangle, \quad |01\rangle, \quad |10\rangle, \quad |11\rangle \ , \tag{18.4}$$

where the first bit refers to A and the second bit to B, *i.e.,* we use the shorthand

$$|ij\rangle \equiv |i\rangle_A |j\rangle_B \equiv |i\rangle_A \otimes |j\rangle_B \ . \tag{18.5}$$

Suppose the system is in the pure state

$$|\psi\rangle = \frac{1}{\sqrt{2}} \left(|00\rangle + |11\rangle\right) \ , \tag{18.6}$$

so $\rho = |\psi\rangle\langle\psi|$. As a 4x4 matrix, ρ has diagonal and off-diagonal elements. Diagonal density matrices are just classical probability distributions, but the off-diagonal elements indicate entanglement and are intrinsically quantum.

The reduced density matrix of system A is

$$\begin{aligned} \rho_A &= \mathrm{tr}_B\, \rho \\ &= \frac{1}{2}\, {}_B\langle 0| \left(|00\rangle + |11\rangle\right) \left(\langle 11| + \langle 00|\right) |0\rangle_B \\ &\quad + \frac{1}{2}\, {}_B\langle 1| \left(|00\rangle + |11\rangle\right) \left(\langle 11| + \langle 00|\right) |1\rangle_B \\ &= \frac{1}{2}\big(|0\rangle_A\, {}_A\langle 0| + |1\rangle_A\, {}_A\langle 1|\big) \\ &\propto \mathbf{1_{2x2}} \ . \end{aligned} \tag{18.7}$$

The last line says ρ_A is proportional to the identity matrix of a 2-state system. In this case we say ρ_A is *maximally mixed*, and the initial state $|\psi\rangle$ is *maximally entangled*.

The entanglement entropy of subsystem A is easy to calculate for a diagonal matrix,

$$\begin{aligned} S_A &= -\,\mathrm{tr}\, \rho_A \log \rho_A \\ &= -2 \times \frac{1}{4} \log \frac{1}{4} \\ &= \log 2 \ . \end{aligned} \tag{18.8}$$

Interpretation of entanglement entropy

In fact the 2-qubit example illustrates a useful way to put entanglement entropy into words:

Entanglement entropy counts the number of entangled bits between A and B.

If we had k qubits in system A and k qubits in system B, then in a maximally entangled state $S_A = k \log 2$. So S_A counts the number of bits, or equivalently, e^{S_A} counts the number of entangled states (since k qubits have 2^k states).

Rephrased slightly:

Given a state ρ_A with entanglement entropy S_A, the quantity e^{S_A} is the minimal number of auxiliary states that we would need to entangle with A in order to obtain ρ_A from a pure state of the enlarged system.

Schmidt decomposition

A very useful tool is the following theorem, called the Schmidt decomposition: Suppose we have a system AB in a pure state $|\psi\rangle$. Then there exist orthonormal states $|i\rangle_A$ of A and $|i\rangle_B$ of B such that

$$|\psi\rangle = \sum_i \lambda_i |i\rangle_A |\tilde{i}\rangle_B \ , \tag{18.9}$$

with λ_i real numbers in the range $[0, 1]$ satisfying

$$\sum_i \lambda_i^2 = 1 \ . \tag{18.10}$$

The number of terms in the sum is (at most) the dimension of the smaller Hilbert space $\mathcal{H}_A$ or $\mathcal{H}_B$.

Proof: See Wikipedia, or Nielsen and Chuang chapter 2.

If A is small and B is big, this is intuitive. It says we can pick a basis for $|i\rangle_A$, and each of these states will be correlated with a particular state of system B. The thermofield double is an obvious example.

Complement subsystems

An immediate consequence of the Schmidt decomposition is that a pure state of system

AB has

$$S_A = S_B \qquad \text{(pure states)} \ . \tag{18.11}$$

To see this, write the reduced density matrices in the Schmidt basis,

$$\rho_A = \sum_i \lambda_i^2 |i\rangle_{A\,A}\langle i| \ , \qquad \rho_B = \sum_i \lambda_i^2 |i\rangle_{B\,B}\langle i| \ . \tag{18.12}$$

Both density matrices have eigenvalues λ_i^2 so they have the same entropy. (18.11) does not hold for mixed states of AB.

18.2 Geometric entanglement entropy

Entanglement entropy can be defined whenever the Hilbert space splits into two factors. A very important example is when we define A as a subregion of space.

Example: N spins on a lattice in 1+1 dimensions
Let's arrange N spins in a line. Define A to be a spatial region containing k spins, and $B = A^C$ is everything else:

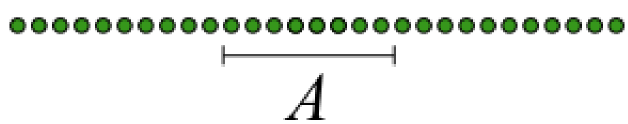

The most general state of this system is

$$|\psi\rangle = \sum_{\{s_i\}} c_{s_1\cdots s_N} |s_1\rangle|s_2\rangle\cdots|s_N\rangle \tag{18.13}$$

where $s_i = 0$ or 1 (meaning 'up' or 'down'), and the c's are complex numbers.

Scaling with system size
Let's restrict to $1 \ll |A| \ll |B|$, so that we can think of subsystem A as large and B as infinite. In a random state, *i.e.,* one in which the coefficients $c_{ij\cdots}$ are drawn from a uniform distribution, we expect any subsystem A to be almost maximally entangled with B. In the language of the Schmidt decomposition, this means that λ_i is nonzero

and $\sim 1/\sqrt{2^k}$ for a complete basis of states $|i\rangle_A$. In fact this is a theorem, see Harlow's lectures for the exact statement.

Accordingly, the entanglement entropy scales as the number of spins in region A. In 1+1d this is linear in the size of A, and more generally,

$$S_A \sim \text{Volume(A)} \qquad \text{(random state).} \tag{18.14}$$

In other words, most states in the Hilbert space of the full system have entanglement scaling with volume.

However, often we are interested in the groundstate. Ground states of a local Hamiltonian are very non-generic, and the corresponding entanglement entropies obey special scaling laws. Usually, if the system is gapped (*i.e.,* correlations die off exponentially), the ground state must obey the *area law*:

$$S_A \sim \text{Area(A)} \qquad \text{(ground state of local, gapped Hamiltonian)}\ . \tag{18.15}$$

(This is a theorem in 1+1d, and usually true in higher dimensions.)

Thus groundstates occupy a tiny, special corner of the Hilbert space. This is a corner with especially low 'complexity.' Intuitively speaking, a large degree of entnaglement is what makes quantum information exponentially more powerful than classical information; so states with lower entanglement entropy are less complex. More specifically, this actually means that you can encode a groundstate wavefunction with far fewer parameters than the 2^N complex numbers appearing in (18.13).

DMRG

In 1+1d, the area law becomes simply

$$S_A \sim \text{const} \tag{18.16}$$

independent of the system size. This special feature is responsible for a hugely important technique in quantum condensed matter called the *density matrix renormalization group* (DMRG). This technique is used to efficiently compute groundstate wavefunctions of 1+1d systems using a computer. This would not be possible for general states, since (we think) classical computers require exponential time to simulate quantum systems. But (18.16) means that, in a precise sense, groundstates of gapped 1d systems are no more complex than classical systems.

Scaling at a critical point

The area law applies to gapped systems. Near a critical point, where dof become massless and long-distance correlations are power-law instead of exponentially suppressed, the area law can be violated. In a 1+1d critical system, and therefore also in 1+1d conformal field theory, (18.16) is replaced by

$$S_A \sim \log L_A \tag{18.17}$$

where L_A is the size of region A. This is bigger than the area law, but still much lower than the volume-scaling of a random state.

18.3 Entropy Inequalities

Relative entropy

Much of the recent progress in QFT based on entanglement comes from a few inequalities obeyed by entanglement entropy. Define the *relative entropy*

$$S(\rho||\sigma) \equiv \operatorname{tr}\rho\log\rho - \operatorname{tr}\rho\log\sigma \ . \tag{18.18}$$

(Note that this is not symmetric in ρ, σ.) This obeys

$$S(\rho||\sigma) \geq 0 \tag{18.19}$$

with equality if and only if $\rho = \sigma$. The proof of this statement is straightforward, see Wikipedia. It just involves some matrix manipulations. The key ingredient is the fact that density matrices in quantum mechanics are very special: they have a positive spectral decomposition,

$$\rho = \sum_i p_i v_i v_i^* \tag{18.20}$$

where p_i is non-negative and v_i is a basis vector. This is necessary for quantum mechanics to have a sensible probabilistic interpretation and is closely related to unitarity.

The relative entropy can be viewed as a measure of how 'distinguishable' ρ and σ are. In the classical case (diagonal ρ and σ), it is error we will make in predicting the uncertainty of a random process if we think the probability distribution is σ, but actually it is ρ. Given this interpretation, positivity is obvious — clearly we will never do *better* using the wrong distribution.

Triangle inequality

Positivity or relative entropy implies the *triangle inequality*,

$$|S_A - S_B| \leq S_{AB} \ . \tag{18.21}$$

Mutual information

Define the *mutual information*,

$$I(A, B) \equiv S_A + S_B - S_{AB} \ . \tag{18.22}$$

This can be written as a relative entropy, and is therefore non-negative:

$$I(A, B) = S(\rho_{AB}||\rho_A \otimes \rho_B) \geq 0 \ . \tag{18.23}$$

Roughly, $I(A, B)$ measures the amount of information that A has about B (or vice-versa, since it is symmetric).

In a pure state of AB, the only correlations between A and B come from entanglement, so in this case $I(A, B)$ measures entanglement between A and B. However, in a mixed state, $I(A, B)$ also gets classical contributions. For example in a 2-qubit system, it is easy to check that the classical mixed state

$$\rho_{AB} \propto |00\rangle\langle 00| + |11\rangle\langle 11| \tag{18.24}$$

has non-zero mutual information.

Strong subadditivity

So far we have discussed partitioning a system into two pieces A and B, but we can partition further and find new inequalities. The *strong subadditivity* inequality (*SSA* for short) applied to a tripartite system $\mathcal{H}_{ABC} = \mathcal{H}_A \otimes \mathcal{H}_B \otimes \mathcal{H}_C$, is

$$S_{ABC} + S_B \leq S_{AB} + S_{BC} \ . \tag{18.25}$$

This is less mysterious if written in terms of the mutual information,

$$I(A, B) \leq I(A, BC) \ . \tag{18.26}$$

Although this inequality seems obvious — clearly A has more information about BC than about B alone — and is 'just' a feature of positive matrices, it is surprisingly difficult to prove. See Nielsen and Chuang for a totally unenlightening derivation.

Sometimes it is useful to express (18.25) in different notation, where A and B are two overlapping subsystems, which are not independent:

$$S_{A\cup B} + S_{A\cap B} \leq S_A + S_B \ . \tag{18.27}$$

Exercise: Positivity of classical relative entropy

Prove that the classical relative entropy is non-negative. That is, prove (18.19), assuming ρ and σ are diagonal.

Exercise: Mutual information practice

Consider a 2-qubit system. First, calculate the mutual information of the two bits in the classical mixed state

$$\rho = \tfrac{1}{2}\left(|00\rangle\langle 00| + |11\rangle\langle\, 11|\right) \ . \tag{18.28}$$

This is a clearly a state with the maximal amount of classical correlation — if we measure one bit, we know the value of the second bit.

Now, what is the maximal amount of mutual information for a *quantum* (pure or mixed) state of 2 qubits? Write an example of a state with this maximal amount of mutual information. (Quantum states with more mutual information than is possible in any classical state are sometimes called *supercorrelated*.)

Exercise: Purification and the Triangle Inequality

Use strong subadditivity to prove the following identities for a tripartite system:

$$S_A \leq S_{AB} + S_{BC} \quad (18.29)$$

$$S_A \leq S_{AB} + S_{AC} \quad (18.30)$$

$$S_{AB} \geq |S_A - S_B| \quad (18.31)$$

Hint: Purify the tripartite system that appears in strong subadditivity by adding a 4th system, D, with $ABCD$ in a pure state. This is always possible.