# Boolean Constraint Satisfaction Problems

Consider a collection of Boolean variables $s_1, ..., s_n$ , with $s_i \in \{0, 1\}$ for each i.

Using Boolean logic functions like AND, OR, and NOT we can form propositions (Boolean functions), e.g.

$$R(s_i, s_j, s_k) = (s_i \wedge s_j) \vee s_k$$

The proposition is true if R = 1, and false if R = 0. Using AND, OR, and NOT we can build up other propositions, for example implications

$$(A \Rightarrow B) = \neg A \vee B$$

We can also use identities to simplify complicated propositions e.g. $\neg(A \vee B) = \neg A \wedge \neg B$

**Theorem**: every Boolean proposition can be expressed using AND, OR, and NOT, in the following conjuctive normal form (CNF),

$$R(s_1, ..., s_n) = \wedge_{i=1}^{m} R_i(s_{i,1}, ..., s_{i,k})$$

where each $R_i(s_{i,1}, ..., s_{i,k})$ is a disjunction (an OR function) , written here as acting on k Boolean variables (“literals”) or their negations.

# Boolean Constraint Satisfaction Problems

The previous theorem tells us that general Boolean functions can be written in terms of AND, OR, NOT.

Each $R_i$ is called a "clause", and the conjunction of all the clauses is true iff each clause is true.

Given a Boolean function written in CNF, we can ask whether there exists some assignment of variables for which the value of the function is 1 ("all the clauses are true"),

$$(\exists s_1 \ldots s_n) \wedge_{i=1}^{m} R_i(s_{i,1}, \ldots, s_{i,k})$$

The expression above including the existential quantifier is called a Boolean formula, and it is "satisfied" if and only if there exists at least one assignment to the variables $s_1, \ldots, s_n$ that makes each clause true.

The computational problem which takes a Boolean formula of the form above as input, and outputs 1 if its true, and 0 if its false, is called SATISFIABILITY or SAT.

If each clause contains at most k literals, then this restricted version of the problem is called k-SAT.

# Boolean Constraint Satisfaction Problems

Given a k-SAT problem involving n Boolean variables, one approach to solving the problem would be to evaluate the function over all $2^n$ possible assignments of the variables.

Later in the course we will prove the **Cook-Levin theorem**, which states that solving k-SAT for any $k \geq 3$ is as hard as performing "classical nondeterministic polynomial time computation" ("NP-hard"). When we find out what this means it will give us evidence that solving 3-SAT is practically intractable in general.

In fact, one of the great open problems in computer science is to prove the expected result that solving 3-SAT is so hard, one can hardly improve upon the brute-force solution described above.

**Exponential time hypothesis**: there are constants $s_k > 0$ such that any classical algorithm for solving k-SAT requires time $2^{s_k n}$. The strong ETH further asserts,

$$(\exists k_0)\left(k \geq k_0 \Rightarrow \lim_{n\to\infty} s_k = 1\right)$$

The same is expected to be true for quantum algorithms, except in this case the constants $s_k$ are divided by 2 because of a general quantum algorithmic technique called amplitude amplification (which we will learn later).

# Boolean Constraint Satisfaction Problems

k-SAT is an example of a Boolean **constraint satisfaction problem**. The clauses can be viewed as constraints, and the formula is satisfiable depends iff the constraints can be simultaneously satisfied.

We can also express Boolean formulas using our notation for quantum systems by promoting the Boolean variables $s_i \in \{0,1\}$ to qubits $|s_i\rangle$. The goal is to count +1 for each constraint that is violated.

There is only one way to violate each 3-SAT clause, for example $\neg a \vee b \vee c$ is only violated if

$$(a,b,c) = (1,0,0)$$

If we think of this as a qubit state $|100\rangle$, then the projector $H = |100\rangle\langle 100|$ detects this violation.

$$\neg a \vee b \vee c \text{ is false } \Leftrightarrow \langle abc|H|abc\rangle > 0$$

This projector can also be written in terms of Pauli Z's: $|100\rangle\langle 100| = \frac{1}{8}(I + Z_a)(I - Z_b)(I - Z_c)$

# Boolean Constraint Satisfaction Problems

Similarly, consider the conjunction of two clauses $(a \vee \neg b) \wedge (b \vee c)$ .

The first clause forbids (a,b) = (0,1). The second clause forbids (b,c) = (0,0).

We can once again detect these violations with projectors,

$$H_1 = |01\rangle\langle 01|_{ab} \otimes I_c \quad , \quad H_2 = I_a \otimes |00\rangle\langle 00|_{bc}$$

Where $H_1$ acts non-trivially on qubits a,b, and $H_2$ acts nontrivially on qubits b and c. We will in general suppress the tensor products with the identity. These projectors can also be re-expressed as Pauli Z's:

$$H_1 = \frac{1}{4}(I - Z_a)(I + Z_b) \quad , \quad H_2 = \frac{1}{4}(I - Z_b)(I - Z_c)$$

The expectation $\langle abc|H_1 + H_2|abc\rangle$ counts the number of clauses that is violated by the assignment a,b,c.

# Boolean Constraint Satisfaction Problems

Putting these ideas together, we can express a general k-SAT problem as a quantum system, with a Hermitian operator that is diagonal in the computational basis, with each eigenvalue counting the number of violated constraints in the corresponding assignment.

$$\wedge_{i=1}^{m} R_i(s_{i,1}, \ldots, s_{i,k}) \mapsto \langle s_1 \ldots s_n | H | s_1 \ldots s_n \rangle \quad , \quad H = \sum_{i=1}^{m} H_i$$

where each $H_i$ acts nontrivially on k qubits, projecting onto the string that violates that clause:

$$H_i = \otimes_{j=1}^{k} \left( I - (-1)^{r_{ij}} Z_{ij} \right)$$

The formula is satisfiable if and only if the Hermitian operator H has 0 as an eigenvalue.

Computationally, we've traded a brute-force search over an exponentially large state space for finding the minimum eigenvalue of an exponentially large matrix that is diagonal in a known basis, so there is no change.

If the problem was "NP-hard" or we believed it took exponential time before, then it remains "NP-hard" and we still believe it takes exponential time.

# Noncommutative Constraint Satisfaction Problems

From a quantum point of view it is very special that all of the local constraints are written in terms of Pauli Z. The local constraints commute with one another, $[H_i, H_j] = 0$, and are all diagonal in the same basis.

For a general a set of k-local quantum constraints, $\{H_i\}_{i=1}^m$, each $H_i$ acts nontrivially on at most k qubits, but each $H_i$ is a positive semi-definite Hermitian operator (which could include any combination of Pauli operators acting on those qubits). The conjunction of all the constraints is:

$$H = \sum_{i=1}^{m} H_i$$

And the constraints are simultaneously satisfied iff H has 0 as an eigenvalue. This is a noncommutative constraint satisfaction problem. But what does it mean?

# Noncommutative Constraint Satisfaction Problems

We can find an operational meaning for noncommutative CSPs as follows. Each $H_i$ having 0 as an eigenvalue is necessary for H to have 0 eigenvalue. For each i, consider the eigendecomposition of $H_i$:

$$H_i = \sum_{j=1}^{2^k} \lambda_j^i |\psi_j^i\rangle\langle\psi_k^i|$$

Recall that H is k-local, so each terms acts on k qubits. This Hilbert space has dimension $2^k$. Both sides of the equation have implicit identities acting on the rest of the Hilbert space. The $\lambda_j^i$ are all non-negative, and some of them must be 0 in order for the global H to have 0 as an eigenvalue.

Using this decomposition of the local terms, define projectors $\Pi_i = \sum_{j=1}^{2^k} |\psi_j^i\rangle\langle\psi_j^i|$ that have the same kernels as the $H_i$, and define the new Hamiltonian:

$$H' = \sum_{i=1}^{m} \Pi_i$$

Which is equivalent as a noncommutative CSP in the sense that H has 0 as an eigenvalue iff H' does. The difference is that now each of the local constraints is a projector $\Pi_i$. Now we can think of the local terms as constraints on measurement outcomes, and we are asking whether a global state is consistent with local constraints.

# Quantum Constraint Satisfaction Problems

Returning to SAT problems, even if the formula is not satisfiable we may about the maximum number of clauses that can be simultaneously satisfied. This version of the problem is called MAX k-SAT.

Although solving MAX k-SAT seems like it should return a number (the maximum number of satisfied clauses), this is closely linked to the decision problem "can at least C clauses be satisfied?" (one can be reduced to the other by binary searching on the values of C in the range 1,…,m).

If we use our mapping from SAT to quantum systems, then MAX SAT is about the minimum eigenvalue of H,

$$\wedge_{i=1}^{m} R_i(s_{i,1}, ..., s_{i,k}) \mapsto \langle s_1 ... s_n | H | s_1 ... s_n \rangle$$

Which may differ from 0. Similarly, a general quantum constraint satisfaction problem seeks the minimum eigenvalue:

$$H = \sum_{i=1}^{m} H_i$$

If the minimum eigenvalue differs from 0, we cannot necessarily transform this into an equivalent problem in terms of local projectors. So here minimizing H is a constraint on many simultaneous expectation values.

# Quantum Constraint Satisfaction Problems

The eigenvalues of a Hermitian operator are characterized by a result known as the variational principle, or more generally (for the higher eigenvalues) the Cauchy mix-max theorem. Let the eigenvalues be ordered:

$$\lambda_0 \leq \lambda_1 \leq \ldots \leq \lambda_{N-1}$$

The variational principle states that the expectation of any non-zero state yields an upper bound on $\lambda_0$:

$$\lambda_0 = \min_{|\psi\rangle \in \mathcal{H}: |\psi\rangle \neq 0} \frac{\langle\psi|H|\psi\rangle}{\langle\psi|\psi\rangle}$$

In matrix analysis, the quantity on the RHS is called the Rayleigh quotient, $R_H(\psi) = \frac{\langle\psi|H|\psi\rangle}{\langle\psi|\psi\rangle}$.

# Quantum Constraint Satisfaction Problems

If $|\psi_0\rangle$ is the eigenstate with eigenvalue $\lambda_0$, then the next smallest eigenvalue is

$$\lambda_1 = \min_{\substack{|\psi\rangle \in \mathcal{H}: |\psi\rangle \neq 0 \\ \langle\psi|\psi_0\rangle = 0}} R_H(\psi)$$

Which means that once we've found the minimum eigenvector, then any state orthogonal to it yields an upper bound on the next largest eigenvector, and so on. The same idea works for all the eigenvalues.

The Cauchy mix-max theorem lets us restate these ideas without referring to the eigenvectors explicitly,

$$\lambda_k = \min_{\mathcal{V}:\dim(\mathcal{V})=k+1} \quad \max_{|\psi\rangle \in \mathcal{V}: |\psi\rangle \neq 0} R_H(\psi)$$

# Quantum Constraint Satisfaction Problems

How hard is it to find the minimum eigenvalue of H? To make this a meaningful computational question, we need to consider an issue of precision that did not arise in the classical case.

In the classical case, constraints are either discretely violated or not. In a system of m constraints, either 0 are violated or at least 1 is. But noncommutative CSPs open the possibility of satisfying "part" of a constraint.

What if the minimum eigenvalue of a quantum CSP is not zero, but it is doubly exponentially or even uncomputably close to zero? (note this could happen just by encoding strange numbers in the matrix elements of the $H_i$ )

Therefore a rigorously stated quantum CSP should include some reasonable limit on the precision with which one is required to estimate the minimum eigenvalue. Similarly, even if $H|\psi\rangle = 0$ for some $|\psi\rangle$ , there may be orthogonal states with eigenvalues very close to 0.

(Minus) the difference between the minimum eigenvalue and the next smallest is called the spectral gap $\Delta > 0$, and this quantity will play an important role in our discussion of quantum CSPs.

# Local Hamiltonians

Quantum CSPs are also well-motivated in the context of physics, where the local constraints describe the energy due to local interactions between quantum spins. The observable corresponding to energy is the **Hamiltonian**.

Just as we consider quantum CSPs acting on n qubits, but with each local constraint acting on at most k qubits (so that it can be efficiently described), we similarly consider **k-local Hamiltonians**:

$$H = \sum_{i=1}^{m} H_i$$

The role of local Hamiltonians in describing physical interactions comes from the fact that the universe contains many particles (*many-body physics*), but each interaction only depends on relatively few others (e.g. k of them) at a time.

**Terminology note**: even in computer science, the quantum constraint satisfaction problems we just described are much more commonly known as *local Hamiltonian problems*.

# Local Hamiltonians

For our first example of a local Hamiltonian, we will work backwards from a classical CSP.

$$(s_1 \vee \neg s_2) \wedge (s_2 \vee \neg s_3) \wedge \ldots \wedge (s_{n-1} \vee \neg s_n) \wedge (s_n \vee \neg s_1)$$

Each of these clauses forbids the state $(s_i, s_{i+1}) = (0, 1)$ (note we identify $s_{n+1} = s_1$) , so the Hamiltonian that counts the number of violated constraints in any given configuration is

$$H = \sum_{i=1}^{n} |01\rangle\langle 01|_{i,i+1}$$

From this form, we can see the constraints forbid all strings with a 0 appearing to the left of a 1. So 00…0 and 11…1 are both satisfying assignments. 11…100…0 might look ok, but it violates the term connecting qubits n and 1.

# Local Hamiltonians

Each of our clauses can be replaced with Pauli Z's: $|01\rangle\langle 01|_{i,i+1} = \frac{1}{4}(I - Z_i)(I + Z_{i+1})$

Expanding out the product and summing over the local terms,

$$H = \frac{1}{4}\sum_{i=1}^{n}(I - Z_i + Z_{i+1} - Z_iZ_{i+1}) = \frac{1}{4}\sum_{i=1}^{n}(I - Z_iZ_{i+1})$$

Where each single Z cancels a corresponding Z acting on the same qubit, but with the opposite sign, in the sum.

Just as in elementary physics where energy is only defined up to an additive constant, in a physics context it makes sense to shift and rescale a Hamiltonian in order to yield the simplest possible expression:

$$H = -\sum_{i=1}^{n} Z_iZ_{i+1}$$

This Hamiltonian describes a system of qubits in which it is energetically favorable for neighboring spins to align in the Z direction. Proposed by Lenz and analyzed by Ising in 1929, it is known as the Ising model.

# Local Hamiltonians

So far the Ising model is just a classical CSP, and in particular it is diagonal in a known basis and easy to read off the states with the minimum eigenvalue, $|00...0\rangle, |11...1\rangle$ .

To make the problem more genuinely quantum, consider a simple Hamiltonian in the X basis, $H_0 = -\sum_{i=1}^{n} X_i$, for which the state with minimum energy is $|++...+\rangle$.

The Hamiltonian $H_0$ is 1-local. Since it does not couple qubits together we say it is non-interacting; each qubit simply feels an energetic preference to be in the state $|+\rangle$, for example due to an external magnetic field.

However if we form linear combinations of these two Hamiltonians, we get a noncommuting CSP:

$$-\sum_{i=1}^{n} X_i - \sum_{i=1}^{n-1} Z_i Z_{i+1}$$

And it is no longer easy to determine the state with minimum energy (the “ground state”) by inspection.

# Local Hamiltonians

Another early example of a many-body Hamiltonian that arose from physics was the Heisenberg model of a ferromagnet. In it's simplest form, nearest-neighbor spins want to align:

Another early example of a many-body Hamiltonian that arose from physics was the Heisenberg model of a ferromagnet. In it's simplest form, nearest-neighbor spins want to align:

$$H = -\sum_{i=1}^{n} \vec{S}_i \cdot \vec{S}_{i+1} = -\sum_{i=1}^{n} (X_i X_{i+1} + Y_i Y_{i+1} + Z_i Z_{i+1})$$

$$X \otimes X = \begin{pmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{pmatrix} \quad Y \otimes Y = \begin{pmatrix} 0 & 0 & 0 & -1 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ -1 & 0 & 0 & 0 \end{pmatrix} \quad Z \otimes Z = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}$$

$$XX + YY + ZZ = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & -1 & 2 & 0 \\ 0 & 2 & -1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}$$

# Local Hamiltonians

$$\vec{S}_i \cdot \vec{S}_j = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & -1 & 2 & 0 \\ 0 & 2 & -1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \qquad \frac{1}{2}\left(I - \vec{S}_i \cdot \vec{S}_j\right) = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 1 & -1 & 0 \\ 0 & -1 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}$$

$$\frac{1}{2}\left(I - \vec{S}_i \cdot \vec{S}_j\right) = |\Psi^-\rangle\langle\Psi^{-1}| \quad , \quad |\Psi^-\rangle = \frac{1}{\sqrt{2}}\left(|01\rangle - |10\rangle\right)$$

Therefore the local terms of the Heisenberg model assign higher energy to the antisymmetric state $|\Psi^-\rangle$, and lower energy to the symmetric states $|11\rangle, |00\rangle, (|01\rangle + |10\rangle)$. The local terms do not commute, so we can immediately go from this observation to knowing the global ground state. By the way,

$$\frac{1}{2}\left(I + \vec{S}_i \cdot \vec{S}_j\right) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}$$

Is a 2-qubit unitary we have seen before, it is the SWAP gate (you can see its unitary because it's a permutation matrix).

# Local Hamiltonians

Heisenberg model: $H = -\sum_{i=1}^{n} \vec{S}_i \cdot \vec{S}_{i+1} = -\sum_{i=1}^{n} (X_i X_{i+1} + Y_i Y_{i+1} + Z_i Z_{i+1})$

It turns out that this system has a degenerate ground space: many eigenvectors corresponding to the minimum eigenvalue. These ground states maximize the sum of the total angular momentum,

$$\vec{S}_{\text{tot}} = \vec{S}_1 + \vec{S}_2 + \ldots + \vec{S}_n$$

The Hamiltonian commutes with the total spin in the Z direction, $S_z = \sum_{i=1}^{n} Z_i$ , $[H, S_z] = 0$, and so the degenerate ground states are labeled by different values of $S_z$ .

These ground states are uniform superpositions of bit strings with a given Hamming weight:

$$|\psi_m\rangle = \frac{1}{\sqrt{\binom{n}{m}}} \sum_{x \in \{0,1\}^n : |x| = m} |x\rangle$$

Obtaining this full solution, and other information such as the spectral gap, requires more advanced techniques (I use a connection to random walks). It was solved by Hans Bethe in 1931. Mathematical physics is, in the narrow sense, the rigorous study of ground states of many-body local Hamiltonians. Such exact solutions are rare.

# Local Hamiltonians

Traditional physics most often considers local Hamiltonians like the quantum Ising model that includes an additional notion of how the qubits are arranged geometrically in some Euclidean space $\mathbb{R}^D$ e.g. a lattice.

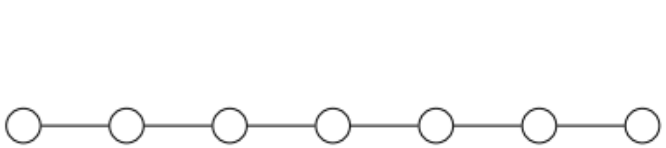

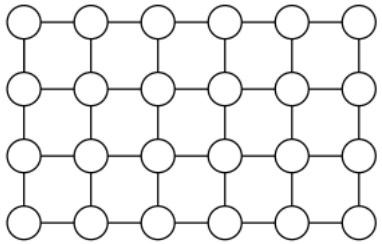

This reflects the historical development of local Hamiltonian problems, which were motivated by the study of ordered quantum matter e.g. magnets, crystals, conducting metals. For similar reasons one often considers "translation invariant" local Hamiltonians, in which H inherits the translation symmetry of a spatial lattice.

More formally, given a lattice $\Lambda$ (a set of sites) a Hamiltonian H is called "spatially local" if $H = \sum_{v \in \Lambda} H_v$ , where each $H_v$ acts nontrivially only a sites within some ball of radius $\mathcal{O}(1)$ around v.

# Local Hamiltonians

Not every k-local Hamiltonian is spatially local with respect to some lattice.

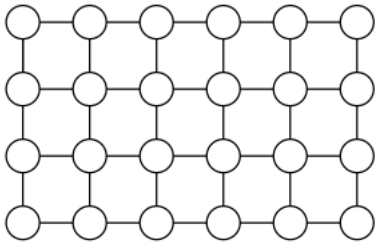

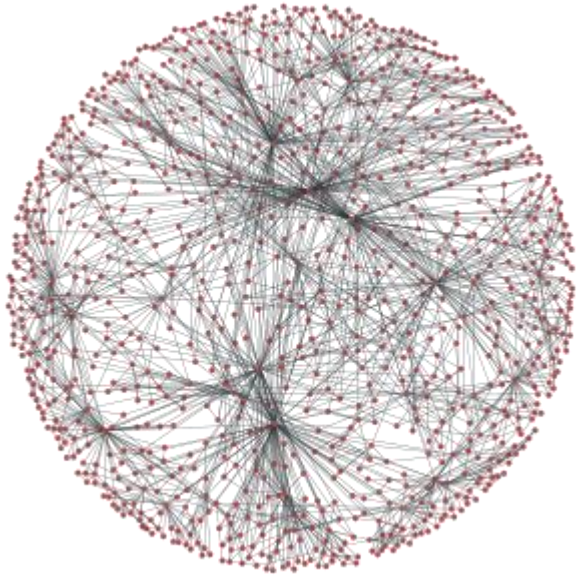

Be careful talking about “local Hamiltonians” to someone with a traditional physics background, they may think you mean spatially local (or even translation invariant) Hamiltonians.

The idea that k-local should mean “each local term acts nontrivially on at most k qubits” is due to Kitaev (a Breakthrough Prize winning physicist), and it is wise in context of quantum information (as we will soon see).

Although “k-local” is the most standard term, I advocate the term “combinatorically local” when contrast with the spatially local case is desired.

# Local Hamiltonians

From a physics point of view, by cooling down the temperature of a system we isolate it from its environment, which is a necessary condition for coherence. At very low temperatures the system will fall into its ground state.

Therefore the study of ground states, solutions to quantum CSPs, is of central importance in physics because these are the state for with the most pronounced quantum behavior.

Given a succinct description of a local H, we may sometimes be able to find the ground state (just as some 3-SAT instances are easy), but in general we know it is at least as hard as 3-SAT and so it is "NP-hard."

A more reasonable problem is the following. Given a local H and a state $|\psi\rangle$ (or many copies of the state), determine the expectation value

$$\langle\psi|H|\psi\rangle$$

This problem can be solved efficiently with a quantum computer using an algorithm called "phase estimation", but we don't know any efficient classical algorithm for doing this calculation in general (e.g. given a succinct classical description of $|\psi\rangle$).