CMSC 858L: Quantum Complexity

Instructor: Daniel Gottesman

Spring 2023

6 Clifford group

For the Clifford group, see D. Gottesman, "The Heisenberg Representation of Quantum Computers," quant-ph/9807006 and Aaronson and Gottesman, "Improved Simulation of Stabilizer Circuits," quant-ph/0406196.

Correction from earlier: If you want to define PSPACE using circuits, you need something stronger than a polynomial time Turing machine to generate the circuit. (Since the circuits for PSPACE are exponential in size, polynomial time is not enough to output them.) Instead, we say that a Turing machine can output the ith gate in the circuit in polynomial time, along with the other required circuit parameters such as its size.

OK, what about this set: $\mathcal{G} = \{H, CNOT, R_{\pi/4}\}$. If we replace $R_{\pi/4}$ by $R_{\pi/8}$ or $CNOT$ by Tof, this is universal. But what about this gate set itself? What is its computational power?

The group generated by these gates is a finite group known as the *Clifford group*. Note that these gates can generate entangled states such as a Bell state $|00\rangle + |11\rangle$ or a GHZ state $|000\rangle + |111\rangle$. The group is also of practical importance since it is all that is needed to do encoding and error correction on the large class of stabilizer quantum error-correcting codes. Nevertheless, this gate set is not just not universal, but can actually be efficiently simulated on a classical computer:

Theorem 1. *There is a polynomial time classical algorithm such that, for any quantum circuit consisting of qubits initialized in the state $|0\rangle$, gates from the Clifford group, and ending with standard basis measurements of all qubits, the algorithm calculates the conditional probability of a measurement result, conditioned on the outcome of some or all other qubits.*

This is what is known as a *strong* simulation (and an exact one, whereas one might have an approximate simulation in some cases). A *weak* simulation is where the simulation only solves the sampling problem, i.e., generates outcomes according to the correct (or approximately correct) probability distribution.

Proof. The main insight needed for this theorem is that the gates in the Clifford group are exactly those which conjugate the Pauli group into itself. The Pauli group is consists of tensor products of the Pauli matrices

$$I = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix},\ X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix},\ Y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix},\ Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \tag{1}$$

with overall phase $\pm 1, \pm i$. For instance, $HXH^\dagger = Z$, $HZH^\dagger = X$ and $CNOT(X \otimes I)CNOT^\dagger = X \otimes X$.

We then note that the initial state of all n qubits in the state $|0\rangle$ is the $+1$ eigenstate of the Paulis $M_1 = Z_1$, $M_2 = Z_2$, $\dots$, $M_n = Z_n$ (here Z_i means "Z acting on qubit i"), and is the unique state (up to global phase) with that property. We could thus equally well describe the initial state by listing the n operators $M_1, \dots, M_n$ (generators of the *stabilizer* of the state). When we perform a gate U from the Clifford group, the state changes from $|\psi\rangle$ to $U|\psi\rangle$, and if it was a $+1$ eigenstate of M before, then

$$(UMU^\dagger)U|\psi\rangle = UM|\psi\rangle = U|\psi\rangle. \tag{2}$$

That is, the state $U|\psi\rangle$ is a $+1$ eigenstate of $UMU^\dagger$. Thus, if the state of the n-qubits is a $+1$ eigenstate of $M_1, \dots, M_n$ before the gate, after the gate it is the $+1$ eigenstate of $UM_1U^\dagger, \dots, UM_nU^\dagger$ and vice-versa. If it is the unique eigenstate before the gate, it is also the unique eigenstate after the gate.

The upshot is that we can uniquely specify the state throughout the circuit by updating the stabilizer: Whenever we perform a gate U, replace M_i with $UM_iU^\dagger$. When U is a general unitary, this might be a complicated thing, but if U is in the Clifford group and M_i is in the Pauli group, then $UM_iU^\dagger$ is also in the Pauli group, and therefore can be specified using just $2n+2$ bits: 2 bits for each of the n Paulis in the tensor product and 2 more for the overall phase. The full description of the state thus requires only $O(n^2)$ bits. Updating each M_i takes a constant time, since the only bits that need to be changed are those specifying the Paulis on the qubits acted on by the gate plus the bits specifying the global phase. Thus, simulating a single Clifford group gate takes time $O(n)$.

The measurement at the end is a little trickier. Measuring qubit i in the standard basis corresponds to measuring the eigenvalue of Z_i. If Z_i or $-Z_i$ is one of our generators M_j, then this is straightforward to compute, since if $M_j = Z_i$, the outcome for measuring qubit i will always be 0 (since the $+1$ eigenstate of Z_i is $|0\rangle$) and if $M_j = -Z_i$, the outcome for measuring qubit i will always be 1 (since the state is a $+1$ eigenstate of $-Z_i$, which means it is a -1 eigenstate of Z_i, namely $|1\rangle$). Also note that it is not possible that $\pm iZ_i$ is one of the generators M_j, since the state is a $+1$ eigenstate of M_j and $\pm iZ_i$ has eigenvalues $\pm i$.

But what if Z_i is not equal to a generator? One possibility is that $\pm Z_i$ is equal to a *product* of generators

$$\pm Z_i = \prod_{j=1}^{n} M_j^{b_j}, \tag{3}$$

with each b_j a bit. If Z_i satisfies this equation, then the state is a $+1$ eigenstate of Z_i as well:

$$Z_I|\psi\rangle = \prod_{j=1}^{n} M_j^{b_j}|\psi\rangle = |\psi\rangle, \tag{4}$$

since $|\psi\rangle$ is a $+1$ eigenstate of each M_j. The measurement outcome will then be 0. SImilarly, if $-Z_i$ satisfies (3), then the state is a -1 eigenstate of Z_i and the measurement outcome will always be 1.

We can find out if (3) holds by doing linear algebra. In particular, suppose we ignore the global phase for the moment and represent each M_j by a $2n$-bit vector $(\mathbf{x}|\mathbf{z})$. If x_k is the kth bit of $\mathbf{x}$ and z_k is the kth bit of $\mathbf{z}$, then the tensor factor Pauli of M_j acting on the kth qubit is

- I if $(x_k, z_k) = (0, 0)$,
- X if $(x_k, z_k) = (1, 0)$,
- Y if $(x_k, z_k) = (1, 1)$,
- Z if $(x_k, z_k) = (0, 1)$.

Note that if $(\mathbf{x}|\mathbf{z})$ is the binary vector corresponding to M and $(\mathbf{x}'|\mathbf{z}')$ is the binary vector corresponding to M', then $(\mathbf{x}+\mathbf{x}'|\mathbf{z}+\mathbf{z}')$ is the binary vector corresponding to MM'.

This means that (3) holds iff

$$(\mathbf{0}|\mathbf{e_i}) = \sum_{j=1}^{n} b_j(\mathbf{x}_j|\mathbf{z}_j). \tag{5}$$

Here, $\mathbf{e_i}$ is the vector which is 0 except in the ith coordinate, which is 1, and $(\mathbf{x}_j|\mathbf{z}_j)$ is the binary vector corresponding to M_j. This equation can be rewritten as

$$(\mathbf{0}|\mathbf{e_i})^T = M\mathbf{b}^T, \tag{6}$$

where $\mathbf{b}$ is the row vector of the b_j's and M is the $2n \times n$ matrix with columns equal to the $(\mathbf{x}_j|\mathbf{z}_j)$ vectors.

This is a system of linear equations over the binary field and can be solved by standard techniques, such as Gaussian elimination. If it has a solution, we find the values of b_j. This procedure also tells us if (3) does *not* hold.

Note, however, that we are not quite done with this case. We have found the b_j's but we do not yet know whether the measurement outcome is 0 or 1 because we dropped the global phase for this calculation. Now we must restore it, computing $\prod M_j^{b_j}$ in the Pauli group to see if we get Z_i or $-Z_i$.

What about if (3) does not hold? Actually, there is a shortcut we can use to determine that. Note that the initial generators $M_i = Z_i$ all commute with each other under multiplication, and when we conjugate them by U that is still true:

$$(UM_iU^\dagger)(UM_jU^\dagger) = UM_iM_jU^\dagger = UM_jM_iU^\dagger = (UM_jU^\dagger)(UM_iU^\dagger). \tag{7}$$

When we take the binary vector representations of the initial $M_i = Z_i$, the vectors we get are all linearly independent. This remains true after performing Clifford group gates because the gates are invertible; if P is a product of the $UM_iU^\dagger$s, then $U^\dagger PU$ is the same product of the M_i's.

Thus, the M_i's at all times are independent, commuting Pauli operators. It turns out that we can have at most n independent commuting Pauli operators on n qubits.

Claim 1. *N commutes with every M_i iff $\pm N$ is a product of some M_i's.*

Proof of claim. Certainly, if $\pm N$ is a product of M_i's, then it commutes with all of them, since they all commute with each other.

The forward direction can again be seen as a consequence of linear algebra. Let $(\mathbf{x}|\mathbf{z})$ be the binary vector corresponding to M and let $(\mathbf{x}'|\mathbf{z}')$ be the binary vector corresponding to M'. Then we can determine by direct calculation that M and M' commute iff

$$\mathbf{x} \cdot \mathbf{z}' \oplus \mathbf{z} \cdot \mathbf{x}' = 0. \tag{8}$$

The commutation of Paulis corresponds to a *symplectic product* in the binary vector space. In particular, if M is the matrix whose columns are the binary vectors corresponding to M_i, then N with binary vector $(\mathbf{x}|\mathbf{z})$ commutes with all of the M_i's iff

$$(\mathbf{z}|\mathbf{x})M = \mathbf{0}. \tag{9}$$

(Note here that the z and x terms in the vector are switched due to the symplectic product.) This means that the vector $(\mathbf{x}|\mathbf{z})$ is again a solution to a set of linear equations. But since the M_i's are all independent, the matrix M has maximum rank n. That means that the dimension of the solution space is n (as a binary vector space). But the columns of M, the vectors corresponding to M_i, are solutions already, since the M_i's commute with each other, and there are n of them. They are linearly independent, so they span the solution space and any vector that solves (9) is a sum of the vectors corresponding to M_i. This, in turn, means that $\pm N$ is a product of the M_i's. □

So the only remaining case is when Z_i fails to commute with one or more of the M_j's. Elements of the Pauli group either commute or *anticommute*, $PQ = -QP$. Therefore, there must be some j such that $Z_iM_j = -M_jZ_i$. In this case, the measurement outcome must be a random bit.

To see this, note that the projector onto the ± 1 eigenspace of Z_i is $(I \pm Z_i)/2$. This means that the probability of getting the outcome 0 to when measuring the ith bit of the state $|\psi\rangle$ is

$$\frac{1}{2}\langle\psi|(I + Z_i)|\psi\rangle. \tag{10}$$

But if $|\psi\rangle$ is a $+1$ eigenstate of M_j and M_j anticommutes with Z_i, we have

$$\frac{1}{2}\langle\psi|(I + Z_i)|\psi\rangle = \frac{1}{2}\langle\psi|(I + Z_i)M_j|\psi\rangle = \frac{1}{2}\langle\psi|M_j(I - Z_i)|\psi\rangle = \frac{1}{2}\langle\psi|(I - Z_i)|\psi\rangle, \tag{11}$$

which is the probability of getting outcome 1. Thus, outcome 0 and outcome 1 both have probability 1/2.

If we are only measuring a single qubit, we can stop here. But that won't let us calculate conditional probabilities. To go further, we want to figure out the residual state of the remaining qubits after we measure one of them. We can do so by noting that if we measure qubit i and get outcome 0, the overall state is now $\frac{1}{\sqrt{2}}(I + Z_i)|\psi\rangle$, the projector onto the $+1$ eigenspace of Z_i, renormalized to take into account that the probability of this outcome is $1/2$. We won't bother to track the normalization from now on, since it is automatic.

We can update the stabilizer generators to take the measurement into account. Note that if M_k commutes with Z_i, then the state is still a $+1$ eigenstate of M_k:

$$M_k(I + Z_i)|\psi\rangle = (I + Z_i)M_k|\psi\rangle = (I + Z_i)|\psi\rangle. \tag{12}$$

The state is not still an eigenstate of M_j, which anticommuted with Z_i, and any other M_k that anticommute with Z_i have a similar problem. However, note that before the measurement, if the state is a $+1$ eigenstate of M_j and M_k, then it is also a $+1$ eigenstate of M_jM_k, and if M_j and M_k both anticommute with Z_i, then M_jM_k commutes with Z_i:

$$Z_i(M_jM_k) = -M_jZ_iM_k = +(M_jM_k)Z_i. \tag{13}$$

Therefore, after the measurement, the state is a $+1$ eigenstate of M_jM_k.

We therefore have the following algorithm to compute a new set $\{M_1, \ldots, M_n\}$ for which the post-measurement state is a $+1$ eigenstate:

1. Find j such that M_j anticommutes with Z_i
2. Run through all $k = 1, \ldots, n$, $k \neq j$. If M_k commutes with Z_i, leave it. If M_k anticommutes with Z_i, replace it by M_jM_k.
3. Replace M_j by Z_i if the measurement outcome was 0 and by $-Z_i$ if the measurement outcome was 1.

It is not hard to see that the resulting new set of M_i's all commute with each other and are independent.

We therefore have the following algorithm to determine conditional probabilities of measurements. Suppose we want to find the probability of measuring 0 on qubit d conditioned on having the outcomes $c_1, \ldots, c_{d-1}$ on qubits 1 through $d-1$. (This is WLOG since we can relabel the qubit numbers as needed.)

1. For qubit i running from 1 to $d-1$:
   (a) Determine if Z_i commutes with all M_j.
   (b) If Z_i and all M_j commute, solve (6) to find the expansion of $\pm Z_i$ as a product of the M_j's and then determine if the outcome of measuring Z_i should be 0 or 1. If the result matches c_i, then continue; otherwise, this condition is not possible, so halt and return that result.
   (c) If Z_i anticommutes with some M_j, update that stabilizer as above assuming that the (random) measurement outcome is c_i.
2. Determine if Z_d commutes with all M_j
3. If Z_d and all M_j commute, solve (6) to find the expansion of $\pm Z_d$ as a product of the M_j's and then determine if the outcome of measuring Z_d is 0 or 1. If the outcome is 0, return probability 1; if the outcome is 1, return probability 0.
4. If Z_i anticommutes with some M_j, return probability 1/2.

Solving the systems of linear equations by Gaussian elimination takes time $O(n^3)$, so that is the complexity of this algorithm. By tracking some additional information, we can speed this up to an algorithm taking time $O(n^2)$. □

Note that the conditional probability of getting 0 on a qubit is always 0, 1, or 1/2 (or the conditional cannot occur). This is a consequence of the special structure of the Clifford group.