CMSC 858L: Quantum Complexity

Instructor: Daniel Gottesman

Spring 2023

8 Computational Universality

The computational universality of real gates is from Bernstein and Vazirani, "Quantum Complexity Theory," SIAM Journal on Computing, vol. 26, 1411-1473 (1997), which also contains a lot about quantum Turing machines and foundational results on quantum complexity. The proof that Toffoli plus Hadamard is computationally universal is from Shi, "Both Toffoli and Controlled-NOT need little help to do universal quantum computation," quant-ph/0205115. My presentation of both results is from Aharonov, "A Simple Proof that Toffoli and Hadamard are Quantum Universal," quant-ph/0301040. For a discussion of encoded universality, see Kempe, Bacon, Lidar, and Whaley, "Theory of Decoherence-Free Fault-Tolerant Universal Quantum Computation," quant-ph/0004064.

8.1 Real gates and computational universality

Now let us consider the following set of gates: $\mathcal{G} = \{\text{all 3-qubit unitary gates with real matrix elements}\}$. Clearly this set of gates, even though it includes many-qubit gates, is not universal or even approximately universal: We can't get close to implementing any gate with complex matrix elements no matter how many real gates we multiply together. Nevertheless, this is clearly a very large subgroup of the unitary group. (It is actually $SO(2^n)$, the real orthogonal matrices.) That seems unlikely to be classically simulatable.

In fact, it is just as computationally powerful as a universal set. For instance, if we define BQP using efficiently computable gates in $\mathcal{G}$, we get the same complexity class BQP as if we define it using an approximately universal set of gates.

The most straightforward way to see this is to realize that while we can't *approximate* an arbitrary unitary operation, we can *simulate* an arbitrary unitary:

Theorem 1. *For any initial quantum state followed by any unitary operation and measurement of every qubit, there exists an exact weak simulation of the circuit using an initial quantum state with real amplitudes and a circuit consisting of gates from $\mathcal{G}$, followed by measurement of every qubit and a constant-depth classical computation.*

Proof. An n-qubit system with complex amplitudes will be simulated in an $(n+1)$-qubit system with *real* amplitudes. When the extra qubit is $|0\rangle$, that gives us the real part of the original state and when the extra qubit is $|1\rangle$, that gives us the imaginary part. In particular, the n-qubit complex state $|\psi\rangle = \sum_x \alpha_x |x\rangle$ corresponds to the $(n+1)$-qubit real state

$$|\tilde{\psi}\rangle = \sum_x (\mathrm{Re}(\alpha_x)|0\rangle + \mathrm{Im}(\alpha_x)|1\rangle) \otimes |x\rangle. \tag{1}$$

Note that the notion of a state or unitary with real coefficients is inherently a basis-dependent concept. And also note that the map $|\psi\rangle \mapsto |\tilde{\psi}\rangle$ is *not* unitary.

We can then simulate a unitary U as follows:

$$\tilde{U}|0\rangle \otimes |x\rangle = |0\rangle \otimes \mathrm{Re}(U)|x\rangle + |1\rangle \otimes \mathrm{Im}(U)|x\rangle \tag{2}$$

$$\tilde{U}|1\rangle \otimes |x\rangle = -|0\rangle \otimes \mathrm{Im}(U)|x\rangle + |1\rangle \otimes \mathrm{Re}(U)|x\rangle. \tag{3}$$

Lemma 1. *For all* $|\psi\rangle$,

$$\tilde{U}|\tilde{\psi}\rangle = \widetilde{U|\psi\rangle}. \tag{4}$$

Proof of lemma.

$$\tilde{U}|\tilde{\psi}\rangle = \sum_x \mathrm{Re}(\alpha_x)(|0\rangle \otimes \mathrm{Re}(U)|x\rangle + |1\rangle \otimes \mathrm{Im}(U)|x\rangle) + \mathrm{Im}(\alpha_x)(-|0\rangle \otimes \mathrm{Im}(U)|x\rangle + |1\rangle \otimes \mathrm{Re}(U)|x\rangle) \tag{5}$$

$$= \sum_x (\mathrm{Re}(\alpha_x)|0\rangle + \mathrm{Im}(\alpha_x)|1\rangle) \otimes \mathrm{Re}(U)|x\rangle + (\mathrm{Re}(\alpha_x)|1\rangle - \mathrm{Im}(\alpha_x)|0\rangle) \otimes \mathrm{Im}(U)|x\rangle. \tag{6}$$

Now,

$$U|\psi\rangle = (\mathrm{Re}(U) + i\mathrm{Im}(U))(\sum_x (\mathrm{Re}(\alpha_x) + i\mathrm{Im}(\alpha_x))|x\rangle) \tag{7}$$

$$= \sum_x (\mathrm{Re}(\alpha_x)\mathrm{Re}(U)|x\rangle - \mathrm{Im}(\alpha_x)\mathrm{Im}(U)|x\rangle) + i(\mathrm{Im}(\alpha_x)\mathrm{Re}(U)|x\rangle + \mathrm{Re}(\alpha_x)\mathrm{Im}(U)|x\rangle) \tag{8}$$

$$= \sum_x (\mathrm{Re}(\alpha_x) + i\mathrm{Im}(\alpha_x))\mathrm{Re}(U)|x\rangle + (-\mathrm{Im}(\alpha_x) + i\mathrm{Re}(\alpha_x))\mathrm{Im}(U)|x\rangle. \tag{9}$$

Thus,

$$\widetilde{U|\psi\rangle} = \sum_x (\mathrm{Re}(\alpha_x)|0\rangle + \mathrm{Im}(\alpha_x)|1\rangle) \otimes \mathrm{Re}(U)|x\rangle + (-\mathrm{Im}(\alpha_x)|0\rangle + \mathrm{Re}(\alpha_x)|1\rangle) \otimes \mathrm{Im}(U)|x\rangle, \tag{10}$$

which is equal to $\tilde{U}|\tilde{\psi}\rangle$, as claimed.

□

This lemma means that if we start with a state encoded in this way, we can maintain the encoding and evolve the state according to the appropriate unitary. Note that if U is a t-qubit gate, then $\tilde{U}$ will be a $(t+1)$-qubit gate since it involves the extra qubit. In particular, since $\mathcal{G}$ includes all 3-qubit real gates, it can simulate all 2-qubit complex gates, and since those are universal, circuits generated by $\mathcal{G}$ can simulate all unitaries in this encoding.

So far, we have seen how to simulate an arbitrary initial state and maintain the encoding through arbitrary unitary circuits. But what about the final measurement? If we measure every qubit of $|\tilde{\psi}\rangle$, we get the outcome $0x$ with probability $|\mathrm{Re}(\alpha_x)|^2$ and the outcome $1x$ with probability $|\mathrm{Im}(\alpha_x)|^2$. If we discard the extra qubit, we get that the probability of outcome x is

$$|\mathrm{Re}(\alpha_x)|^2 + |\mathrm{Im}(\alpha_x)|^2 = |\alpha_x|^2, \tag{11}$$

which is exactly the probability of outcome x in the original circuit. □

This is an interesting example, which shows that we can have the full power of a quantum computer without having a universal set of gates. What we have is a weaker notion of universality, *computational universality*. I don't think there is a single accepted definition of computational universality, but here is a sensible one:

Definition 1. *A set of quantum resources (e.g., a set of quantum gates) is* computationally universal *if there exists a polynomial classical algorithm which, given any quantum circuit Q with n qubits and T gates starting with standard initial states and ending with standard basis measurements, produces an algorithm using classical computation and subroutine calls using the quantum resources which produces an approximate weak simulation whose outcome is within ϵ statistical distance of the outcome distribution of Q and uses $O(\mathrm{poly}(n, T, \log 1/\epsilon))$ classical and quantum gates or other resources.*

The statistical distance between $\{p_i\}$ and $\{q_i\}$ is

$$D(\{p_i\},\{q_i\}) = \frac{1}{2}\sum_i |p_i - q_i|, \tag{12}$$

which is just the classical version of the quantum trace distance.

There are alternatives that are both weaker and stronger than this definition. For instance, this doesn't say anything about the ability to simulate situations where you want to keep a quantum state around at the end of the computation. Such a definition would need to make explicit mention of a map between the state in the original circuit and in the simulation; but it would exclude other simulations that perform some sort of transformation on the whole circuit and don't explicitly simulate the instantaneous state. We could also insist on a stricter degree of approximation, but we don't want to be too strict, since we would like regular approximate universality to be a special case of computational universality. On the other hand, rather than asking for a weak simulation, we could instead say that the computationally universal set of resources can decide any language in BQP while using only polynomial resources. I am not aware of any interesting examples which achieve the BQP definition and not the weak simulation one (although I believe you can construct somewhat artificial examples, for instance by saying that you can only measure one qubit of the final state of a quantum circuit).

8.2 Tof and H

Now let us look at the set of gates $\mathcal{G} = \{Tof, H\}$. These are real gates, so we can't hope for more than computational universality, but it turns out we do indeed have computational universality.

To show this, we use the same simulation given by the general real simulation. We can start from a somewhat standard approximately universal set $\{H, C - R_{\pi/4}\}$ and see what real gates we need to simulate that. The Hadamard is a real gate already, so $\mathrm{Im}(H) = 0$ and

$$\tilde{H}|0\rangle \otimes |x\rangle = |0\rangle \otimes H|x\rangle \tag{13}$$
$$\tilde{H}|1\rangle \otimes |x\rangle = |1\rangle \otimes H|x\rangle, \tag{14}$$

so $\tilde{H}$ is just H acting on the same qubit and doing nothing to the extra one.

The $C - R_{\pi/4}$ is the controlled-$\pi/4$ rotation, $C - R_{\pi/4} = \mathrm{diag}(1,1,1,i)$. Note that the phase in $R_{\pi/4}$ is an arbitrary choice, but in $C - R_{\pi/4}$, because of the control, the choice of phase gives different 2-qubit gates. We will choose this particular phase, which is standard. Then

$$\mathrm{Re}(C - R_{\pi/4}) = \mathrm{diag}(1,1,1,0) \tag{15}$$
$$\mathrm{Im}(C - R_{\pi/4}) = \mathrm{diag}(0,0,0,1), \tag{16}$$

which means

$$\widetilde{C - R_{\pi/4}}|000\rangle = |000\rangle \tag{17}$$
$$\widetilde{C - R_{\pi/4}}|001\rangle = |001\rangle \tag{18}$$
$$\widetilde{C - R_{\pi/4}}|010\rangle = |010\rangle \tag{19}$$
$$\widetilde{C - R_{\pi/4}}|011\rangle = |111\rangle \tag{20}$$
$$\widetilde{C - R_{\pi/4}}|100\rangle = |100\rangle \tag{21}$$
$$\widetilde{C - R_{\pi/4}}|101\rangle = |101\rangle \tag{22}$$
$$\widetilde{C - R_{\pi/4}}|110\rangle = |110\rangle \tag{23}$$
$$\widetilde{C - R_{\pi/4}}|111\rangle = -|011\rangle \tag{24}$$

We can recognize this as $CC - Z$, a twice-controlled Z gate on the three qubits (which gives a phase -1 if all three qubits are 1), followed by a Toffoli gate with the last two qubits as controls and the first one as target. The $CC - Z$ gate is equal to $(I \otimes I \otimes H)Tof(I \otimes I \otimes H)$. Therefore, to perform the real simulation of H and $C - R_{\pi/4}$, it suffices to have H and Tof, proving these gates are computationally universal.

Because Tof is universal for classical reversible computation, this result means we can think of the extra computational power of quantum computing as coming just from the Hadamard gate. But don't read too much into that: the field has lots of results like this, with quantum coming from something classical plus just one more thing.

8.3 Encoded Universality

Another important type of computational universality is *encoded universality.* Like the example of real gates, encoded universality gives a simulation of a general quantum circuit, but unlike real gates, the encoding is unitary. A system with encoded universality is genuine universal (exactly or approximately), but only on a subspace of the whole Hilbert space.

One particularly striking example of encoded universality is the *exchange interaction.* Specifically, we will look at the gate set $\mathcal{G} = \{e^{-itH_E}\}$ for arbitrary t (or at least computable t), and

$$H_E = \frac{1}{2}(I \otimes I + X \otimes X + Y \otimes Y + Z \otimes Z). \tag{25}$$

This is called the exchange interaction because it implements the SWAP gate between the two qubits. In some sense, therefore, it seems trivial since all it does is switch the order of qubits, but with the ability to run this Hamiltonian for varying amounts of time, it is possible to do non-trivial interactions which lead to computational universality. However, note that on the initial state $|00\rangle$ it does nothing, so we need a starting state other than that.

The exchange interaction permutes qubits, so in particular, it commutes with any unitary of the form $U \otimes U$ on the two qubits. It turns out that with enough qubits, there are large subspaces that are invariant under all operations $U^{\otimes n}$ on n qubits. When $n = 4$, we can use the subspace spanned by

$$|\overline{0}\rangle = \frac{1}{2}(|01\rangle - |10\rangle)(|01\rangle - |10\rangle) \tag{26}$$

$$|\overline{1}\rangle = \frac{1}{12}(2|0011\rangle + 2|1100\rangle - |0101\rangle - |1010\rangle - |1001\rangle - |0110\rangle). \tag{27}$$

(You can check by hand that these states have the claimed properties.) Because the Hamiltonian commutes with $U \otimes U \otimes U \otimes U$ for all U, the state remains invariant under such operations when we act on it with a gate from $\mathcal{G}$, so those gates keep us in this two-dimensional subspace. An analysis of the Lie algebra shows that by applying the Hamiltonian on different pairs of qubits in succession, we can generate arbitrary elements of $SU(2)$ on the subspace. (You can reduce this to a subspace of a 3-qubit system to encode on qubit, but the encoding is a bit more complicated because each logical basis state corresponds to a 2-dimensional subspace rather than a single state.)

When we have 8 qubits, which is enough to encode two logical qubits, the subspace of states invariant under $U^{\otimes n}$ is larger: It is 14-dimensional, so it includes much more than the tensor product of the two encoded qubit Hilbert spaces. The exchange Hamiltonian acting on different pairs of qubits keeps us in this 14-dimensional Hilbert space but not in the two-qubit logical Hilbert space. It generates the full $SU(14)$ unitary group on this larger Hilbert space. When we have all gates of the form e^{-itH_E}, we can exactly get all such gates and we can restrict ourselves to only using those gates within the $SU(4)$ subgroup corresponding to two-qubit logical operations. That will give us a universal set of logical gates.

In the case where we instead have only a limited set of times t that we can use (and then rely on the Solovay-Kitaev theorem to generate a dense set of gates in these subgroups), we might not get any exact two-qubit gates. However, we can approximate them to high accuracy. This acts as some additional "leakage" errors taking us out of the computational Hilbert space, but by making them small enough, we can make them negligible over the course of the computation. In this case, we can approximate an arbitrary unitary in the computational subspace.

In either case, the ability to exactly or approximately perform arbitrary elements of $SU(2^n)$ using $O(n)$ qubits lets us accurately simulate any quantum algorithm. Thus, this set of gates is also computationally universal.