Quantum speedup for unstructured search - Grover's algorithm

Jeffrey Epstein

March 28, 2022

The goal of Grover's algorithm is to perform an unstructured search, looking for a single marked element from a set of N items. It is one of a class of algorithms known as oracle algorithms, where we're given a black box that performs some kind of operation and we will ask how many times we have to use the black box in order to determine something of interest about it. An algorithm for search in the classical setting can be framed in terms of an oracle that evaluates a function

$$f_a : \{1, 2, \ldots, N\} \to \{0, 1\} \tag{1}$$

where $f_a(a) = 1$ and $f_a(x) = 0$ for $x \neq a$. Here, a is the marked element, and the goal of the search algorithm is to determine a given that we're allowed to evaluate f_a as many times as we need. What we're not allowed to do is to "look inside the box", i.e. we don't know the definition of f_a. All we can do is to give a number x between 1 and N to the black box and receive the answer $f_a(x)$, i.e. to ask "Is x the marked element?"

How many times do we need to call the classical oracle if we want to find a with high probability? This question can be formalized and a rigorous answer provided, but it seems intuitively that the best thing we can really do is just to query the oracle with different inputs until it tells us we've found the marked item, which on average will require $N/2$ queries before we find a. This reasoning turns out to be correct, and so we say that they query complexity of the algorithm is linear in N. In other words, if we double N, we also double the expected runtime of the search algorithm.

Now let's think about the quantum setting. Here our oracle will basically evaluate the same function, but will do so coherently, i.e. in a way that respects superpositions. The action of the oracle U_a (encoding the information that the marked item is a) is to flip the sign of the amplitude of $|a\rangle$ in whatever quantum state is presented to it, i.e.

$$U_a : \sum_x \alpha_x |x\rangle \longrightarrow \sum_{x \neq a} \alpha_x |x\rangle - \alpha_a |a\rangle \,. \tag{2}$$

If you're concerned that this isn't quite parallel to the classical definition, note that we can implement U_a using another oracle F_a that acts on an input register and a single auxiliary qubit as

$$F_a : |x\rangle |y\rangle \longrightarrow |x\rangle |y \oplus f_a(x)\rangle \,, \tag{3}$$

where $\oplus$ is addition mod 2. If we give this oracle an input state $|x\rangle$ and an auxiliary qubit prepared in the state $|0\rangle$, apply a Z gate to the auxiliary qubit, and then apply the oracle again (to disentangle the two registers), the effect is the same as applying U_a directly. So we can think in terms of U_a and then just remember that each application requires two uses of F_a when we're counting queries.

A natural starting point for a quantum algorithm is a uniform superposition over all inputs, which in the case of search is the same as a uniform superposition over all elements:

$$|u\rangle = \frac{1}{\sqrt{N}} \sum_{x=1}^{N} |x\rangle \,. \tag{4}$$

We'd like at the end of the algorithm to have our system in the state $|a\rangle$ so that when we measure we get the (classical) information a, i.e. we find the marked element. So at a high level, what the search algorithm should do is to use the oracle U_a to perform a rotation from $|u\rangle$ to $|a\rangle$. Now we can think entirely two-dimensionally, i.e. in the subspace spanned by these two states (remember that we don't actually know what $|a\rangle$ is, but it is some fixed state known to the oracle).

How does the oracle behave in this geometric picture? It just reflects across the line perpendicular to $|a\rangle$. We can now use the geometric fact that two consecutive reflections result in a net rotation by twice the angle between the planes of reflection (see Fig. 1). Since we've already thought about the state $|u\rangle$, and it is the second state defining our two-dimensional subspace let's consider using a reflection defined in terms of this state:

$$R_u = 2\,|u\rangle\,\langle u| - 1. \tag{5}$$

The unit vector corresponding to the line across which U_a reflects is

$$\left|a^{\perp}\right\rangle = \frac{1}{\sqrt{N-1}} \sum_{x \neq a} |x\rangle\,, \tag{6}$$

so that we can compute the angle between the two lines of reflection as

$$\theta = \arccos\left(\langle u|a^{\perp}\rangle\right) = \arccos\left(\sqrt{1 - 1/N}\right) \approx \sqrt{N}, \tag{7}$$

where the approximation is valid for large N.

But this is very exciting! It means that in the regime where N is very large, we only need to apply $R_u U_a$ about $\sqrt{N}$ times in order to rotate from our initial state $|u\rangle$ to something close to the state $|a\rangle$ corresponding to the marked element, so that we can measure and with high probability obtain the correct answer. This means that we can conclude the following:

> *Given access to the quantum oracle U_a (or F_a if you prefer), we can find the marked element with $\mathcal{O}(\sqrt{N})$ queries, which is strictly better than the required $\mathcal{O}(N)$ queries required if we only have access to the classical oracle f_a.*

It turns out, but is a bit more involved to prove, that this scaling is optimal, i.e. there is no algorithm that uses fewer than $\mathcal{O}(\sqrt{N})$ queries to the quantum oracle to find a with high probability.

There are a few things to note about this algorithm:

1. It provides a $\sqrt{N}$ speedup, *not* an exponential speedup. There are other oracle algorithms, such as Simon's algorithm, that do provably give such speedups, but solve more contrived problems.

2. If we are trying to use Grover's algorithm to solve a search problem in the real world, we need to build the oracle. In general this would require knowledge of a, so this is not really a way to do classical unstructured search. However, it might be useful if the oracle is some quantum algorithm. In that case, though, proving advantage is much trickier because in principle there might be information about the structure of the oracle that we can exploit to do something efficient classically to find a. It turns out to be much easier to prove separations between quantum and classical complexities using oracle problems of this type rather than problems that make no use of black boxes.

3. The algorithm highlights the importance of thinking of different bases as being equally "real", since we use the computational basis state $|a\rangle$ and the uniform superposition $|u\rangle$ in identical ways.

4. Somehow the speedup reflects the fact that quantum mechanics normalizes the sum of the *squares* of its amplitudes, rather than just the sum of the probabilities as in classical probability theory. This is where the $\sqrt{N}$ comes from.

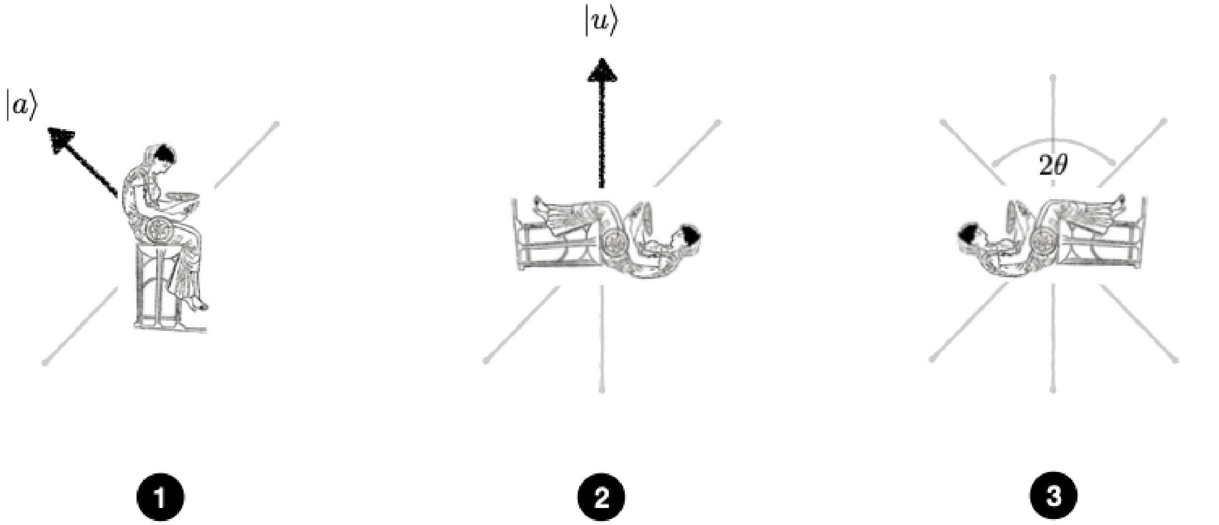

Figure 1: Two reflections performed one after the other result in a rotation by twice the angle between the planes of reflection. In Grover's algorithm, we reflect across the plane perpendicular to the marked state $|a\rangle$, and then about the plane parallel to the uniform superposition $|u\rangle$. The net effect is a rotation by $\mathcal{O}(\sqrt{N})$. The image is the one surviving depiction of the Pythia, the high priestess of the Temple of Apollo at Delphi, from the time when the oracle there was active.