Last time we discussed the HSP for finitely generated abelian groups. For a black box function f: G -> X that is constant and distinct on the cosets of H < G, classically it takes $\Omega(\sqrt{|G/H|})$ queries to identify the generators of H, while quantumly $O(\text{polylog}(|G/H|))$ queries suffice, and the number of gates other than queries is also $O(\text{poly}\log(|G/H|))$.

The idea of the algorithm is to generate an H-invariant coset state in a single query:

$$\frac{1}{\sqrt{|G|}}\sum_{g\in G}|g\rangle\otimes|0\rangle \xrightarrow[f]{} \frac{1}{\sqrt{|G|}}\sum_{g\in G}|g\rangle\otimes|f(g)\rangle$$

$$\xrightarrow[\text{measure}]{} \frac{1}{\sqrt{|H|}}\sum_{h\in H}|gh\rangle \equiv |g,H\rangle$$

By "H-invariant" we mean that, if $U(g_0)$ is right multiplication by g_0,

$$U(g_0)|g\rangle = |g g_0\rangle$$

then $U(h)|g,H\rangle = |g,H\rangle$

For abelian G, the quantum Fourier transform over G maps (a superposition of) elements of G to (a superposition of) irreducible representations of G. Because of H-invariance, when we apply the QFT to the coset state $|g,H\rangle$, we obtain a uniform (except for phases that depend on the coset) superposition of irreps of G that represent H trivially. These irreps themselves form a group (i.e. we can multiply their characters together), which I called the "dual lattice" (or dual group) $H^\perp$ in the previous lecture.

Thus, by generating a coset state, doing the QFT over G, and measuring, we sample uniformly from $H^\perp$. If G has n generators and R is an upper bound on the number $|G/H|$ of cosets, then $O(n\log(nR))$ queries suffice to find a generating set for $H^\perp$ with high prob; then H is found by an easy classical computation.

Now a natural question is, what if the group G is not abelian? It is shown in the homework problem (5.4) that the query complexity is still reasonable. In the algorithm analyzed there, the coset states are not measured individually as in the algorithm for abelian G. Rather, m coset states are generated, and then a sequence of *collective* measurements is performed on the m copies. If R is an upper bound on the number of candidates for the hidden subgroup H, then the algorithm has success probability

$$P_{success} \geq \left(1 - \frac{R}{2^{m/2}}\right)^2$$

So $m = O(\log R)$ queries suffices for constant success probability.

However, although the number of queries is efficient, the algorithm is not, because the number of collective measurements required in $O(R)$ — i.e. exponentially large.

The nonabelian HSP, then, seems to be intrinsically harder than the abelian case. In particular, the "Fourier sampling" method that works well in the abelian case is less powerful in the nonabelian case. To understand why, it is helpful to consider an example.

I'll choose an example for which no efficient algorithm is known, despite considerable effort to find one. (It turns out that a solution to the problem would have applications to cryptography, but I won't explain that.)

Let the group G be $G = D_N$, the dihedral group. This is the symmetry group of an N-sided regular polygon in the plane. The N-gon is left invariant under

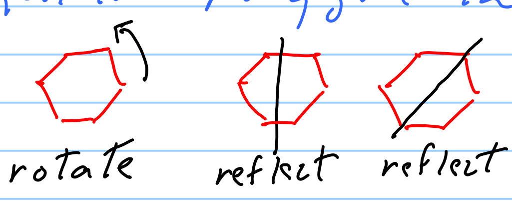

- rotation by $\theta = 2\pi/N \times$ integer
- reflection about any one of N bilateral symmetry axes.

Thus $|D_N| = 2N$ (N rotations and N reflections)

The group is generated by two (noncommuting) elements:

x = ccw rotation by $\theta = 2\pi/N$

y = reflection about x-axis

These generators satisfy 3 defining relations:

$x^N = e$ (rotation by 2π)

$y^2 = e$ (reflect twice $\Rightarrow$ identity)

$yxy = x^{-1}$ (reflection - rot by θ - reflection is equivalent to rot by $-\theta$)

- The reflection turns the z-axis $\perp$ to plane upside down, so a ccw rotation becomes cw

We are promised that the hidden subgroup H is $H = \mathbb{Z}_2$, generated by a reflection

$$H = \{e, yx^r\} \text{ where } r \in \{0,1,2,\ldots,N-1\}$$

There are N reflections, and hence N possible choices for H. The $2N$ elements of D_N can be parameterized by

$$D_N = \{y^t x^s, \; t \in \{0,1\}, \; s \in \{0,1,2,\ldots,N-1\}$$

For fixed H (i.e. r), the cosets can be labeled by s: $gH = \{x^s, yx^{r+s}\}$. Denoting $g \in D_N$ by (t,s) a coset state can be expressed as

$$|gH\rangle = \frac{1}{\sqrt{2}}\left(|0,s\rangle + |1,r+s\rangle\right)$$

Performing the $\mathbb{Z}_N$ QFT, this state is mapped to

$$|gH\rangle \xrightarrow[QFT_N]{} \frac{1}{\sqrt{2N}} \sum_{k=0}^{N-1} \left(e^{2\pi i ks/N}|0,k\rangle + e^{2\pi i k(r+s)/N}|1,k\rangle\right)$$

If we now measure k, the probability distribution governing the outcome is uniform on k. Conditioned on the outcome we obtain, up to a physically irrelevant overall phase that depends on the coset label s,

$$|\psi_r^{(k)}\rangle = \frac{1}{\sqrt{2}}\left(\quad |0\rangle + e^{2\pi i k r/N}|1\rangle\right).$$

This is a state of a single qubit, where k is known, and r is what we want to find. The unknown r is encoded in the state, but the trouble is that N is exponentially large and there are exponentially many values of r that we need to distinguish. It is hard to extract much info about r from the state $|\psi_r^{(k)}\rangle$, or even polynomially many such states $\{|\psi_r^{(k_1)}\rangle, |\psi_r^{(k_2)}\rangle, \ldots, |\psi_r^{(k_m)}\rangle\}$.

In the abelian case, Fourier sampling was more informative. What went wrong? The explanation involves some group representation theory.

A coset state is H-invariant: $U(h)|gH\rangle = |gH\rangle$ where $h \in H$. When we Fourier transform

$$\tilde{U}(h)\left(FT|gH\rangle\right) \equiv \left(FT\, U(h) (FT)^{-1}\right)\left(FT|gH\rangle\right) = FT|gH\rangle$$

Furthermore, the Fourier transform over G block diagonalizes $\tilde{U}(g)$, where the blocks are the irreducible representations (irreps) of G. When we Fourier sample (FT and then measure) we identify a particular block — i.e. the label of a particular irreducible representation of G.

All the irreducible representations of D_N are either one-dimensional or two-dimensional. Recall that a representation assigns a matrix to each group element

$$g \mapsto D^{(\nu)}(g)$$

such that $D^{(\nu)}(g_1 g_2) = D^{(\nu)}(g_1) D^{(\nu)}(g_2)$. For D_N, the representations are

$$D^{(k)}(x^s) = \begin{pmatrix} e^{2\pi i k s/N} & 0 \\ 0 & e^{-2\pi i k s/N} \end{pmatrix}$$

$$D^{(k)}(y) = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$$

we verify this is a rep by checking the defining relations

$$D^{(k)}(x^N) = D^{(k)}(y^2) = I$$

$$D^{(k)}(yxy) = D^{(k)}(x^{-1})$$

For N even, there are 2D irreps for $k = 1, 2, \ldots, \frac{N}{2} - 1$. The cases $k=0$ and $k = N/2$ each split into two 1D irreps. For N odd, there are 2D irreps for $k = 1, 2, \ldots, (N-1)/2$, and the case $k=0$ splits into two 1D irreps.

In the case where G is abelian, all irreps are 1D, and only a fraction of these irreps are H-invariant (those corresponding to elements of $H^\perp$). But <u>every</u> 2D irrep of D_N has an H invariant state.

$$D^{(k)}(yx^r) = \begin{pmatrix} 0 & \omega^* \\ \omega & 0 \end{pmatrix} \quad \text{where } \omega = e^{2\pi i k r/N} \text{ and } yx^r \in H$$

This matrix has eigenvector $\begin{pmatrix} 1 \\ \omega \end{pmatrix}$ with eigenvalue one; this is

$$|\psi_r^{(k)}\rangle = \frac{1}{\sqrt{2}}\left(|0\rangle + e^{2\pi i k r/N}|1\rangle\right),$$

the outcome of Fourier sampling when the measured irrep label is k.

So -- the label k of the irrep does not provide much useful information about r (that is, H), but we can try to get further information by measuring the H-invariant state $|\psi_r^{(k)}\rangle$. For example, suppose we measure in the X-eigenstate basis

$$|\pm\rangle = \frac{1}{\sqrt{2}}\left(|0\rangle \pm |1\rangle\right)$$

Then the probability distribution governing the measurement outcome depends on r:

$$\mathrm{Prob}_r(+, K) = \frac{1}{N}\cos^2\left(\frac{\pi K r}{N}\right)$$
$$\mathrm{Prob}_r(-, K) = \frac{1}{N}\sin^2\left(\frac{\pi K r}{N}\right)$$

That is, prob is uniform in K, and
$$P(\pm | K) = \left|\frac{1}{2}\left(1 \pm e^{2\pi i K r/N}\right)\right|^2$$

Thus, the conditional probability of the outcome K, given the outcome $+$, is

$$\mathrm{Prob}_r(K|+) = \frac{2}{N}\cos^2\left(\frac{\pi K r}{N}\right)$$

and similarly

$$\mathrm{Prob}_r(K|-) = \frac{2}{N}\sin^2\left(\frac{\pi K r}{N}\right)$$

By sampling from these distributions, we can extract a maximum likelihood estimate of r

There is good news and bad news:

- With $m = O(\log N)$ samples, max likelihood estimate returns the correct value of r with success prob at least $1 - \frac{\text{const}}{N}$ (Ettinger and Hoyer)
- But, there is no known classical algorithm that computes the most likely value of r in subexponential time.

We also can make a sharper statement. With $m = \nu \log_2 N$ uniformly chosen coset states, for $\nu < 1$, the optimal collective measurement on the m copies finds r with exponentially small success probability, while for $\nu \geq 1$, the success probability is constant (Bacon, Childs, van Dam). So logarithmic number of coset states really are needed to extract r.

Also, there <u>is</u> a quantum algorithm with $2^{O(\sqrt{\log N})}$ query <u>and</u> time complexity (Kuperberg). That is a substantial speedup relative to the $2^{O(\log N)}$ classical query complexity, but unfortunately still superpolynomial.

It is frustrating! Since the irreps of D_N are at most 2-dimensional, it may not seem to be so different from an abelian group. Yet we still don't have an efficient quantum algorithm for the D_N HSP.