Homework 2

Out: September 10, 2007 **In:** October 1, 2007

1. **(A non-abelian hidden subgroup problem)** Some hidden subgroup problems (HSPs) in non-abelian groups G can be efficiently reduced to HSPs in abelian groups, giving efficient quantum algorithms for HSP in G. In this exercise, we shall see one such example. For $r \geq 2$, consider the non-abelian group G with two generators x, y defined as follows:

$$G := \langle x, y : x^{2^r} = y^2 = 1, yx = x^{2^{r-1}+1}y \rangle.$$

Assume that the elements of G are encoded by tuples (i, j), $0 \leq i < 2^r$, $j = 0, 1$, where (i, j) denotes the group element $x^i y^j$. In the jargon of group theory, G is a particular kind of a *semidirect* product $\mathbb{Z}_{2^r} \rtimes \mathbb{Z}_2$. Define $G_0 := \langle x \rangle$. Suppose $f : G \to S$ is a function hiding a subgroup $H \leq G$. Define $H_0 := H \cap G_0$. We shall now see how one can find H efficiently by a quantum algorithm.

   (a) Show that H_0 can be found efficiently by a quantum algorithm.

   (b) Prove that H_0 is a normal subgroup of G and H/H_0 is either trivial or has order two.

   (c) Suppose $H_0 = \langle x^{2^k} \rangle$ for some $0 \leq k < r$. Show that G/H_0 is an abelian group isomorphic to $\mathbb{Z}_{2^k} \times \mathbb{Z}_2$. Hence, show that H can be found efficiently by a quantum algorithm using the HSP in $\mathbb{Z}_{2^k} \times \mathbb{Z}_2$ as a subroutine.

   (d) Show that the $H_0 = \{1\} = \langle x^{2^r} \rangle$ case can be taken care of separately.

2. **(Finding the structure of a finite abelian group)** In this exercise, we shall see that we can find a decomposition of a finite abelian group into a direct product of cyclic groups efficiently by using a quantum algorithm. A (not necessarily abelian) group $\hat{G}$ is given via a black box, that is, the elements of $\hat{G}$ are represented by elements of $\{0, 1\}^n$ for some n; thus, $|\hat{G}| \leq 2^n$. There may be bit strings that do not correspond to any group element. We assume that elements of $\hat{G}$ are uniquely encoded by bit strings and the identity element of $\hat{G}$ is represented by the all zeroes string. We are given a black box or oracle $\mathcal{M}$ for multiplying two group elements: $\mathcal{M} : |x\rangle|y\rangle|s\rangle|b_1\rangle|b_2\rangle \mapsto |x\rangle|y\rangle|s \oplus (x \cdot y)\rangle|b_1 \oplus c_x\rangle|b_2 \oplus c_y\rangle$, where $x, y, s \in \{0, 1\}^n$, $b_1, b_2 \in \{0, 1\}$, $x \cdot y$ is the bit string representing the product of x and y if both are valid group elements and the all zeroes string otherwise, and c_x, c_y are bits which are one iff x, y are invalid group elements.

   (a) Show that there is a quantum algorithm with running time poly(n) to find the inverse of a group element x, that is, we can implement the following unitary operator cleanly:

$$\mathcal{I} : |x\rangle|s\rangle|b\rangle \mapsto |x\rangle|s \oplus x^{-1}\rangle|b \oplus c_x\rangle.$$

   You may want to use ideas from solving HSP in $\mathbb{Z}$ that we saw in class, as well as clean reversible classical computation.

   (b) Now suppose we are given valid group elements $x_1, \ldots, x_k$ from $\hat{G}$ such that $G := \langle x_1, \ldots, x_k \rangle$ is abelian (observe that this is easy to check). Consider the surjective group homomorphism $h : \mathbb{Z}^k \to G$ defined by $h(z_1, \ldots, z_k) := x_1^{z_1} \cdots x_k^{z_k}$. Show how to find the kernel K of h, that is, $\vec{z} \in \mathbb{Z}^k$ such that $h(\vec{z}) = 1_G$ in time poly(n, k) by a quantum algorithm. You may want to use ideas about HSP in $\mathbb{Z}$ and in finite abelian groups for this part. Thus, $G \cong \mathbb{Z}^k/K$.

(c) Show that every $m \times k$ integer matrix can be put into a diagonal form (called *Smith normal form*) by elementary row and column operations involving only integers. This can be done by a Gaussian elimination style algorithm together with Euclid's GCD algorithm. Argue that this (classical deterministic) procedure takes $\mathrm{poly}(n, m, k)$ time, where n is an upper bound on the bit sizes of the integer entries of the matrix.

(d) Consider the $m \times k$ integer matrix $\mathcal{K}$ whose rows are the generators of K obtained above (there are $m = \mathrm{poly}(n, k)$ of them). Argue that elementary row operations on $\mathcal{K}$ continue to give generating sets for K, and elementary column operations change bases for $\mathbb{Z}^k$ (initially, we start off with the standard Dirac point mass basis for $\mathbb{Z}^k$).

(e) Explain how a cyclic decomposition of G can be obtained in deterministic $\mathrm{poly}(n, m, k)$ time from a Smith normal form of $\mathcal{K}$.

3. **(Finding non-strict periods in finite abelian groups)** In this exercise, we shall see that the standard quantum algorithm to find a strict period, that is, a hidden subgroup, in a finite abelian group also works in time $\mathrm{polylog}|G|$ if the period is only 'approximately strict'. In other words, the standard Fourier sampling based algorithm has a 'robustness property' about strict period finding. Fix $\delta > 0$. Suppose $f : G \to S$ is a function with period subgroup $F \leq G$ such that any function $h : G \to S$ that has a period group $H \geq F$ differs from f in at least $\delta|G|$ elements of G. The intuition behind this definition is that if F were a strict period for f, one would have to corrupt f in at least half the elements of G in order to increase its period group; hence viewed this way, our condition on f is an approximation of strict periodicity. Also, it is a necessary condition because if f is ϵ-close in Hamming distance to an H-periodic function h, then the oracle for f is close to the oracle for h in spectral distance and hence every quantum procedure making a small number of queries to the function oracle will be unable to distinguish between period group F versus H.

   We shall see that under our condition on f, $O(\log|G|/\delta)$ iterations of Fourier sampling allow us to find F with high probability. But first, we will need a few definitions. By a *probabilistic function* $\mu : G \to S$, we shall mean a map $x \mapsto \mu_x$ from elements x of G to probability distributions μ_x on S. For every $x \in G$, define the unit ℓ_1-norm vector $|\mu_x\rangle := \sum_{s \in S} \mu_x(s)|s\rangle$. Then, the *uniform superposition over* μ is defined as $|\mu\rangle := |G|^{-1/2} \sum_{x \in G} |x\rangle|\mu_x\rangle$. For a (deterministic) function $f : G \to S$, the uniform superposition over f boils down to $|f\rangle = |G|^{-1/2} \sum_{x \in G} |x\rangle|f(x)\rangle$. Note that $|\mu\rangle$ has unit ℓ_2-norm if μ is a deterministic function, otherwise its ℓ_2-norm is smaller. For a (deterministic) function $f : G \to S$ and a subgroup $H \leq G$, define a H-periodic probabilistic function $\mu^{f,H}$ by

$$\mu_x^{f,H}(s) := \frac{|f^{-1}(s) \cap (x + H)|}{|H|},$$

   that is, $\mu_x^{f,H}(s)$ is the proportion of elements in the coset $x + H$ where f takes the value s. When f is H-periodic, $|\mu^{f,H}\rangle = |f\rangle$. For any $F \leq G$, define $F^\perp := \{y \in G : \forall x \in F, \chi_y(x) = 1\}$. For $x \in G$ and $F \leq G$, define

$$|F^\perp(x)\rangle := \sqrt{\frac{|F|}{|G|}} \sum_{y \in F^\perp} \chi_x(y)|y\rangle.$$

Now suppose $f : G \to S$ is a (deterministic) function. The standard procedure for Fourier sampling f is given below. Recall that the quantum Fourier transform QFT_G over G is the unitary transformation $|x\rangle \to |G|^{-1/2} \sum_{y \in G} \chi_x(y)|y\rangle$.

- Start off with the state $|0\rangle_G|0\rangle_S$.
- Apply QFT_G to the first register.
- Query the oracle for f.
- Apply QFT_G to the first register.
- Measure the first register and output the result.

(a) For $x \in G$ and $H \leq G$, prove that $|x + H\rangle \overset{\mathrm{QFT}_G}{\longmapsto} |H^\perp(x)\rangle$.

(b) Fix a subgroup $H \leq G$. Show that the probability that Fourier sampling f outputs a $y \notin H^\perp$ is

$$\left\| \frac{1}{\sqrt{|G|}} \sum_{x \in G} |\{0\}^\perp(x)\rangle|f(x)\rangle - \frac{1}{\sqrt{|G||H|}} \sum_{x \in G} |H^\perp(x)\rangle|f(x)\rangle \right\|^2 .$$

(c) From the above two results, conclude that the probability that Fourier sampling f outputs a $y \notin H^\perp$ is equal to $\||f\rangle - |\mu^{f,H}\rangle\|^2$.

(d) For $H \leq G$ and a function $f : G \to S$, define a H-periodic function $f^H : G \to S$ by $f^H(x) := \mathrm{Maj}_{h \in H} f(x + h)$, where Maj is the *majority* function which returns the most frequent value taken by its input, ties being broken arbitrarily. One can view f^H as the 'correction' of f with respect to H-periodicity. Show that

$$\||f^H\rangle - |\mu^{f,H}\rangle\| \leq \||f\rangle - |\mu^{f,H}\rangle\|.$$

(e) Fix a subgroup $H \leq G$. Now suppose f is at least δ-far in Hamming distance from any H-periodic function. Show that the probability that Fourier sampling outputs a $y \notin H^\perp$ is at least $\delta/2$.

(f) Suppose f is F-periodic for some subgroup $F \leq G$. Show that Fourier sampling f will only output $y \in F^\perp$.

(g) Suppose f is F-periodic for some subgroup $F \leq G$. Also suppose f is at least δ-far in Hamming distance from any H-periodic function for any subgroup $F \leq H \leq G$. Suupose we do $k := O(\log |G|/\delta)$ iterations of Fourier sampling f obtaining output $y_1, \ldots, y_k$. Show that with probability at least $3/4$, $F^\perp = \langle y_1, \ldots, y_k \rangle$.