QUANTUM INFORMATION AND COMPUTATION
EXERCISE SHEET 3

Nilanjana Datta n.datta@damtp.cam.ac.uk (Lent 2025-2026)

(1) (Bernstein-Vazirani problem)
For n-bit strings $x = x_1 \dots x_n$ and $a = a_1 \dots a_n$ in B_n we have the sum $x \oplus a$ which is an n-bit string, and now introduce the 1-bit "dot product" $x \cdot a = x_1 a_1 \oplus x_2 a_2 \oplus \dots \oplus x_n a_n$.
For any fixed n-bit string $a = a_1 \dots a_n$ consider the function $f_a : B_n \to B_1$ given by

$$f_a(x_1, \dots, x_n) = x \cdot a \tag{1}$$

(a) Show that for any $a \neq 00\dots0$, f_a is a balanced function i.e. f_a has value 0 (respectively 1) on exactly half of its inputs x.
(b) Given a classical black box that computes f_a describe a classical deterministic algorithm that will identify the string $a = a_1 \dots a_n$ on which f_a is based. Show that any such black box classical algorithm must have query complexity at least n.

Now for any n let $H_n = H \otimes \dots \otimes H$ be the application of H to each qubit of a row of n qubits. Show that (for $x \in B_1$ and $a \in B_n$)

$$H|x\rangle = \frac{1}{\sqrt{2}} \sum_{y=0}^{1} (-1)^{xy} |y\rangle \qquad H_n |a\rangle = \frac{1}{\sqrt{2^n}} \sum_{y \in B_n} (-1)^{a \cdot y} |y\rangle$$

(c) (the Bernstein–Vazirani problem/algorithm)
For each a consider the function f_a which is a balanced function if $a \neq 00\dots0$ (as shown above). Show that the Deutsch-Jozsa algorithm will perfectly distinguish and identify the $2^n - 1$ balanced functions f_a (for $a \neq 00\dots0$) with only *one* query to the function (quantum oracle for f). Indeed, show that the n bit output of the final measurements of the algorithm gives the string a with certainty for these special balanced functions.

(2) (Classical complexity – integer exponentiation mod N)
Exponentiation of integers mod N is a basic arithmetic task (it'll be used for example in Shor's algorithm) and it is important to know that it can be done in poly(n) time where $n = \log N$ is the number of digits for integers in $\mathbb{Z}_N$.
To compute say $3^k \bmod N$ (for $k \in \mathbb{Z}_N$ and $N > 3$) we could multiply 3 together k times. Show that this is not a poly(n) time computation.
Devise an algorithm that *does* run in poly(n) time. (Hint: consider repeated squaring).
You may assume that multiplication of integers in $\mathbb{Z}_N$ may be done in $O(n^2)$ time.
Generalise to a poly time computation of $k_1^{k_2} \bmod N$ for $k_1, k_2 \in \mathbb{Z}_N$ showing that it may be computed in $O(n^3)$ time.

(3) (Simon's algorithm)
Simon's decision problem is the following:
Input: an oracle for a function $f : B_n \to B_n$,
Promise: f is either (a) a one-to-one function or (b) a two-to-one function of the following special form – there is an $\xi \in B_n$ such that $f(x) = f(y)$ iff $y = x \oplus \xi$ (i.e. ξ is the period of f when its domain is viewed as being the group $(\mathbb{Z}_2)^n$).
Problem: determine which of (a) or (b) applies (with any prescribed success probability $1 - \epsilon$ for any $\epsilon > 0$).

It can be argued (e.g. as indicated in lecture notes) that for classical computation, this requires at least $O(2^{n/4})$ queries to the oracle. In this question we will develop a quantum algorithm that that solves the problem with quantum query complexity only $O(n)$. Even more, the algorithm will determine the period ξ if (b) holds. Thus (unlike the balanced vs. constant problem) we'll

have a provable exponential separation between classical and quantum query complexities, even in the presence of bounded error.

To begin, consider $2n$ qubits with the first (resp. last) n comprising the input (resp. output) register for a quantum oracle U_f computing f i.e. $U_f |x\rangle |y\rangle = |x\rangle |y \oplus f(x)\rangle$ for n-bit strings x and y.
(a) With all qubits starting in state $|0\rangle$ apply H to each qubit of the input register, query U_f and then measure the output register (all measurements being in the computational basis). Write down the generic form of the n-qubit state $|\alpha\rangle$ of the input register, obtained after the measurement. Suppose we then measure $|\alpha\rangle$. Would the result provide any information about the period ξ?

(b) Having obtained $|\alpha\rangle$ as in (a), apply H to each qubit to obtain a state denoted $|\beta\rangle$. Show that if we measure $|\beta\rangle$ then the n-bit outcome is a uniformly random n-bit string y satisfying $\xi \cdot y = 0$ (so any such y is obtained with probability $1/2^{n-1}$).

Now we can run this algorithm repeatedly, each time independently obtaining another string y satisfying $\xi \cdot y = 0$. Recall that $B_n = (\mathbb{Z}_2)^n$ is a vector space over the field $\mathbb{Z}_2$. If $y_1, \ldots, y_s$ are s linearly independent vectors (bit strings) then their linear span contains 2^s of the 2^n vectors in B_n. Furthermore to solve systems of linear equations over B_n we can use the standard Gaussian elimination method (calculating with the algebra of the field $\mathbb{Z}_2$), which runs in poly(n) time.

(c) Show that if $(n-1)$ bit strings y are chosen uniformly randomly and independently satisfying $y \cdot \xi = 0$ then they will be linearly independent (and not include the all-zero string $00\ldots0$) with probability

$$\prod_{k=1}^{n-1} \left(1 - \frac{2^{k-1}}{2^{n-1}}\right) = \frac{1}{2} \prod_{k=1}^{n-2} \left(1 - \frac{2^{k-1}}{2^{n-1}}\right).$$

Show that this is at least $1/4$. (It may be helpful here to recall that for a and b in $[0,1]$ we have $(1-a)(1-b) \geq 1-(a+b)$).

(d) Show how the above may be used to solve Simon's problem with $O(n)$ quantum query complexity (for any desired success probability $0 < 1 - \epsilon < 1$).

(4) (Rotation in Grover's algorithm)
For the plane $\mathcal{P}(x_0)$ spanned by $|x_0\rangle$ and $|\psi_0\rangle = \frac{1}{\sqrt{N}} \sum_{x \in B_n} |x\rangle$ set up an orthonormal basis in the plane. Then using the basis, show algebraically (rather than geometrically as in lectures) that the Grover iteration operator Q is a rotation in the plane and derive the angle of rotation.

(5) (Grover's algorithm with an arbitrary starting state)
Consider Grover's algorithm for a unique good item x_0 in a search space of size $N = 2^n$. Suppose that instead of the usual uniform superposition state $|\psi_0\rangle$, we start with any state $|\eta_0\rangle$ of n qubits and conduct the algorithm just as before i.e. apply $\frac{\pi}{4}\sqrt{N}$ iterations of Q and measure.
With $|\eta_0\rangle = |\psi_0\rangle$ the final measurement gives x_0 with probability 1 up to terms of order $1/N$. If we instead begin with some other starting state $|\eta_0\rangle$, describe geometrically how $|\eta_0\rangle$ evolves in the course of the computation. Give an expression (up to terms of order $1/N$) for the probability of obtaining x_0 in the final measurement. Show that this may generally be improved by changing the number of Grover iterations.

(6) (An algebraic interpretation of Grover's algorithm)
(a) Consider the operator $-I_{|\psi_0\rangle} = 2 |\psi_0\rangle \langle\psi_0| - I$ with $|\psi_0\rangle = \frac{1}{\sqrt{N}} \sum_{x \in B_n} |x\rangle$ and $N = 2^n$. Show that

$$-I_{|\psi_0\rangle} = \frac{2}{N} \sum_{\text{all } x,y} |x\rangle \langle y| - I.$$

(b) Let $|\alpha\rangle = \sum_x a_x |x\rangle$ be any n-qubit state. The average amplitude $\overline{a}$ is defined to be $\overline{a} = (\sum_x a_x)/N$. The operation R of "inversion in the average" is defined as follows: $R|\alpha\rangle = \sum_x a'_x |x\rangle$ where $a'_x = a_x - 2(a_x - \overline{a})$ i.e. the value of each amplitude is inverted about the average. Pictorially:

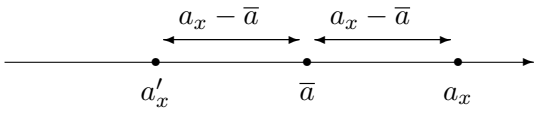

Using the formula in (a) show that $-I_{|\psi_0\rangle}|\alpha\rangle = R|\alpha\rangle$.

(c) Hence Grover's algorithm may be described as follows: start with state $|\psi_0\rangle$; then flip the sign of the x_0 amplitude; then do R, an inversion of all amplitudes in the average; then iterate the last two steps alternately. We can represent states (with real amplitudes) pictorially as a graph of the amplitudes: the x axis has the labels x and each amplitude is a (positive or negative) vertical bar. In terms of this pictorial representation, starting with $|\psi_0\rangle$, carry out one or two iterations of the "flip x_0 and then do R" operation to see how the initial amplitude distribution, uniform over all x, begins to become concentrated at x_0.

(d) Consider the definite case of $N = 4$ (so $x \in \{0, 1, 2, 3\}$) and take $x_0 = 3$ say. (In lectures we saw that for this case of "1 in 4", one Grover iteration serves to find x_0 with certainty i.e. rotating $|\psi_0\rangle$ exactly onto $|x_0\rangle$). Draw the pictorial graph representation of $|\psi_0\rangle$ and carry out one Grover iteration as a flip followed by inversion in the average. Show that as a result, the amplitude becomes exactly zero at $x \neq x_0$ and 1 at $x = x_0$.

(7) (Grover search with multiple good items; application to collision finding)
(a) Write $N = 2^n$ and let $f : B_n \to B_1$ be a function taking value 1 exactly K times, with $f(x) = 1$ iff $x \in G = \{x_1, \ldots, x_K\}$. The Grover operator is defined by $Q = -H_n I_0 H_n I_G$ where $H_n = H \otimes \ldots \otimes H$ is the Hadamard operation on each of n qubits, and for all $x \in \mathcal{B}_n$, I_0 and I_G are defined by

$$I_0|x\rangle = \begin{cases} -|x\rangle & \text{if } x = 0\ldots0 \\ |x\rangle & \text{if } x \neq 0\ldots0 \end{cases} \qquad I_G|x\rangle = \begin{cases} -|x\rangle & \text{if } x \in G \\ |x\rangle & \text{if } x \notin G. \end{cases}$$

Write $|\psi_0\rangle = \frac{1}{\sqrt{N}} \sum_{x \in B_n} |x\rangle$ and introduce $|\psi_G\rangle = \frac{1}{\sqrt{K}} \sum_{x \in G} |x\rangle$. Derive a geometrical interpretation of the action of Q in the two-dimensional subspace of n qubits spanned by $|\psi_0\rangle$ and $|\psi_G\rangle$. Using this interpretation, show that if I_G is given as a black box then an x in G may be obtained with high probability (better than a half say) with $O(\sqrt{N/K})$ uses of I_G, if N is large, and K is small compared to N.

(b) (optional) Let $g : B_n \to B_n$ be a 2-to-1 function i.e. for every y in the range of g there are precisely two strings $x \in B_n$ with $g(x) = y$. A *collision* is a pair of strings $x_1, x_2 \in B_n$ with $g(x_1) = g(x_2)$. The standard quantum oracle U_g for g is the unitary operation on $2n$ qubits defined by

$$U_g|x\rangle|y\rangle = |x\rangle|y \oplus g(x)\rangle \qquad x, y \in B_n$$

where $\oplus$ denotes bitwise addition of n-bit strings.

Suppose that we are given U_g as a black box operation. Using the result of (a), or otherwise, show that a collision may be found with high probability (better than a half say) with $O(N^{1/3})$ uses of U_g.
Hint: start by partitioning the domain of g into sets A and B of sizes $N^{1/3}$ and $(N - N^{1/3})$ and listing all the values of $g(x)$ for $x \in A$. We might find a collision there, but if we're not so lucky, what should we do next with B?

Remark: the classical query complexity for collision finding is $O(\sqrt{N})$. The $O(N^{1/3})$ upper bound on its quantum query complexity established in this question can also be shown to be optimal.