Quantum Computing (5334QUCO8Y), Final exam

Ronald de Wolf

Monday June 13, 2022
10:00–13:00, REC A0.03

The exam is "open book," meaning you can use any kind of paper you want but no electronic devices beyond what is needed to download and read the exam questions, and (at the end) to scan and upload your solutions. Please answer in English. Use a black or blue pen, not a pencil. Write clearly and explicitly, and explain your answers. For a multipart-question, you may assume answers for earlier parts of the question to answer later parts, even if you don't know the earlier answers. The total number of points adds up to 9; your exam grade will be your number of points +1. An exam grade of at least 5 is a necessary condition for passing the course. Your final grade will be 60% exam + 40% homework, rounded to the nearest integer.

1. **(1 point)** Suppose you have a 2-qubit basis state $|b,0\rangle$, with unknown bit $b \in \{0,1\}$, and you'd like to copy the classical bit b into the second qubit (overwriting the $|0\rangle$). How can you do this using only Hadamard gates and controlled-Z gates?

2. **(2.5 points)** This exercise is about efficiently finding the gradient $\nabla f(z)$ of a function $f : \mathbb{R}^d \to \mathbb{R}$ at a point $z \in \mathbb{R}^d$. The gradient is the d-dimensional real vector of the d partial derivatives $\partial f/\partial x_i$, evaluated at the point z.

   (a) Let $f(x) = a + bx$ be a linear function from $\mathbb{R}$ to $\mathbb{R}$, where the real number $b \in [0,1)$ can be written with n bits of precision. Suppose we have a unitary O_f that maps $|x,0\rangle \to |x,f(x)\rangle$ (assume we have enough qubits to write down x and $f(x)$). Give a quantum algorithm to compute b using one application of O_f and one application of O_f^{-1}, and some unitaries that do not depend on f.
   Hint: Start with a uniform superposition over all $x \in \{0,1\}^n$ and end with an inverse QFT. You're allowed to use a unitary like $|c\rangle \mapsto e^{2\pi ic}|c\rangle$ since it does not depend on f.

   (b) Let $f(x_1,\ldots,x_d) = a + b_1x_1 + \cdots + b_dx_d$ be a linear function from $\mathbb{R}^d$ to $\mathbb{R}$, where $a, b_1, \ldots, b_d \in \mathbb{R}$. Show that the gradient $\nabla f(z)$ is equal to $(b_1,\ldots,b_d)$ for every $z \in \mathbb{R}^d$.

   (c) Assume that for the function f in (b), each coefficient b_k is $\in [0,1)$ and can be written with n bits of precision. Suppose we have a unitary O_f that maps $|x_1,\ldots,x_d,0\rangle \to |x_1,\ldots,x_d,f(x_1,\ldots,x_d)\rangle$. Give a quantum algorithm that computes the gradient $\nabla f(z)$ using one application of O_f and O_f^{-1}, and some unitaries that do not depend on f.

3. **(1.5 points)** Consider the intersection problem from communication complexity: Alice has input $x \in \{0,1\}^n$, Bob has input $y \in \{0,1\}^n$, and they want to find (with success probability $\geq 2/3$) an i such that $x_i = y_i = 1$, if such an i exists. We know that using $r = O(\sqrt{n})$ messages between Alice and Bob, they can solve the intersection problem with $O(\sqrt{n}\log n)$ qubits of communication (see Section 15.4). We also know that with only $r = 1$ message (i.e., one-way communication) $\Theta(n)$ qubits of communication are necessary and sufficient (see Exercise 15.5). Now suppose we limit them to some $r \in \{1, \ldots, \sqrt{n}\}$ messages. This r is known to Alice and Bob. Show how they can solve the intersection problem with $O((n/r)\log n)$ qubits of communication.

4. **(2 points)** Suppose you have a classical description of an n-qubit Hamiltonian H that is the sum of $m = n^2$ 2-local terms. Assume the eigenvalues of the Hermitian matrix H lie in $[0,1)$, and can all be written exactly with $2\log n$ bits of precision. You would like to exactly determine the smallest eigenvalue $\lambda_{\min}$ of H, corresponding to unknown n-qubit eigenstate $|\psi_{\min}\rangle$. You're given (as a quantum state) an n-qubit state $|\psi\rangle$ that has a significant overlap with $|\psi_{\min}\rangle$: $|\langle\psi|\psi_{\min}\rangle|^2 \geq 0.7$. Give a polynomial-size quantum circuit that, with probability $\geq 2/3$, outputs $\lambda_{\min}$ exactly.
   NB: You don't need to write down the circuit to the last detail; a clear description of the different parts of the circuit (possibly with some reference to details in the lecture notes) suffices.

5. **(2 points)** This question is about the quantum complexity of inverting a permutation, which is an important problem in cryptography. Let N be a power of 2 and $S = \{0, \ldots, N-1\}$. Let $x \in S^N$ correspond to a permutation on S, meaning that each $j \in S$ occurs exactly once as an entry of x (so the map $i \mapsto x_i$ is a permutation). Suppose we can query x, i.e., we have a unitary O_x that maps $|i,j\rangle \to |i, x_i + j \bmod N\rangle$ for all $i,j \in S$, and we can also apply O_x^{-1}.

   (a) Show how we can find the unique index $i \in S$ for which $x_i = 0$, with success probability $\geq 2/3$, using $O(\sqrt{N})$ queries to O_x and O_x^{-1}.

   (b) The adversary lower bound of Section 11.3 still works with the following modifications: (1) the x's and y's are not binary strings, but strings over a larger alphabet, such as S, and (2) let $\ell_{x,i}$ be the number of $y \in Y$ such that $(x,y) \in R$ and $x_i \neq y_i$; $\ell_{y,i}$ be the number of $x \in X$ such that $(x,y) \in R$ and $x_i \neq y_i$; and $\ell_{\max} = \max\{\ell_{x,i}\ell_{y,i} : (x,y) \in R, i \in \{0,\ldots,N-1\}, x_i \neq y_i\}$.
   In this case the quantum query lower bound is $\Omega(\sqrt{m_0 m_1/\ell_{\max}})$. You may assume this without proof.
   Use this strengthened adversary bound to show a lower bound of $\Omega(\sqrt{N})$ quantum queries for computing the task of part (a).
   Hint: Distinguish the inputs depending on whether 0 sits at an odd or even location i in the string x.

Intended solutions

1. This copying of a classical bit is exactly a CNOT gate, which (since $X = HZH$) we can implement as $(I \otimes H)U(I \otimes H)$ where U is the controlled-Z gate.

2. (a) Let $N = 2^n$. Use Hadamard gates on $|0^n\rangle$ to create a uniform superposition over all $x \in \{0,1\}^n = \{0, \ldots, N-1\}$. Compute $f(x)$ in a second register using O_f to obtain state

$$\frac{1}{\sqrt{2^n}} \sum_{x=0}^{N-1} |x\rangle|f(x)\rangle.$$

Now apply unitary $|c\rangle \mapsto e^{2\pi ic}|c\rangle$ to the second register, and apply O_f^{-1} to put the second register back to $|0\rangle$. Thus we obtain the state

$$\frac{1}{\sqrt{2^n}} \sum_{x=0}^{N-1} |x\rangle e^{2\pi i f(x)}|0\rangle = e^{2\pi ia} \left(\frac{1}{\sqrt{2^n}} \sum_{x=0}^{N-1} e^{2\pi ibx}|x\rangle \right) |0\rangle.$$

Applying inverse QFT to the first register followed by a measurement in the computational basis gives us b.

(b) f is linear, so the partial derivative of f w.r.t. the variable x_k is just the coefficient b_k, independent of z. Hence $\nabla f(z) = (b_1, \ldots, b_d)$ for every $z \in \mathbb{R}^d$

(c) We basically do (a) separately for all d coordinates, using a superposition over $x_j \in \{0,1\}^n$ for each j. This results (after the second query) in state

$$\begin{aligned} &\frac{1}{\sqrt{2^{dn}}} \sum_{x_1,\ldots,x_d \in \{0,1\}^n} |x_1, \ldots, x_d\rangle e^{2\pi i f(x_1,\ldots,x_d)}|0\rangle \\ &= e^{2\pi ia} \left(\frac{1}{\sqrt{2^n}} \sum_{x_1=0}^{N-1} e^{2\pi i b_1 x_1}|x_1\rangle \right) \otimes \cdots \otimes \left(\frac{1}{\sqrt{2^n}} \sum_{x_d=0}^{N-1} e^{2\pi i b_d x_d}|x_d\rangle \right) \otimes |0\rangle. \end{aligned}$$

Separate inverse QFTs for each of the d n-qubit registers, followed by measurements, give us $b_1, \ldots, b_d$.

3. Alice and Bob divide the n indices into $\lceil n/r^2 \rceil$ subsets of r^2 elements each (the last subset could have fewer than r^2 elements if r^2 does not divide n). They run the protocol of Section 15.4 in parallel for each of the $\lceil n/r^2 \rceil$ subsets, and see if one of them finds an intersection. This uses $O(\sqrt{r^2}) = O(r)$ messages, each consisting of $\lceil n/r^2 \rceil$ blocks of $O(\log n)$ qubits, so each message consists of $O(\lceil n/r^2 \rceil \log n)$ qubits. The total communication is the number of messages times the communication-per-message, which adds up to $O(r) \cdot O((n/r^2) \log n) = O((n/r) \log n)$ qubits. By tweaking the constants in the division into subsets we can make sure the number of messages is exactly r instead of $O(r)$ (though I would already give full points if you end with $O(r)$ messages).

4. Using Hamiltonian simulation (Chapter 9) we can implement the unitary $U = e^{2\pi i H}$ with a polynomial-sized circuit up to very small (even exponentially small in n) error ε. We can also implement powers U^t with a polynomial-sized circuit as long as t itself is at most polynomially large. Let A be a unitary that does phase estimation with $2\log n$ bits of precision, using U. The circuit for A still has polynomial size, because the largest power of U that we need to implement (in a controlled manner) for phase estimation with $2\log n$ bits of precision, is $2^{2\log(n)-1} = O(n^2)$. If we apply A to the state $|0^{2\log n}\rangle|\psi_{\min}\rangle$, then the resulting state would be $|\phi\rangle = |n^2\lambda_{\min}\rangle|\psi_{\min}\rangle$ (up to very small error which is due to the fact that we only implement the U^t's up to very small error). The multiplication with $n^2 = 2^{2\log(n)}$ is to convert $\lambda_{\min}$ (which lies in $[0,1)$) to a binary string of $2\log n$ bits, since phase estimation gives us the bits of $\lambda_{\min}$. Then the first register would tell us $\lambda_{\min}$ exactly. If instead we start with the state $|0^{2\log n}\rangle|\psi\rangle$, then the resulting state $|\phi'\rangle$ satisfies $|\langle\phi'|\phi\rangle|^2 = |\langle\psi|\psi_{\min}\rangle|^2 \geq 0.7$ because A is unitary and hence preserves inner product. Therefore a measurement of the first register of $|\phi'\rangle$ tells us $\lambda_{\min}$ exactly, with probability $\geq 2/3$. The difference between 0.7 and 2/3 amply accounts for the tiny errors made in Hamiltonian simulation.

5. (a) Let $z \in \{0,1\}^N$ be a binary string defined by $z_i = 1$ iff $x_i = 0$. Note that because x is a permutation, z has exactly one 1-bit, which is at the location i where $x_i = 0$. We can implement a phase-query to z using 1 query O_x, a unitary on the query's target register that puts a -1 in front of $|0\ldots0\rangle$, and one O_x^{-1} (to set the query's target register back to $|0\ldots0\rangle$). Now use Grover's algorithm to find the location of the unique 1-bit in z. This uses $T = O(\sqrt{N})$ queries to z and hence $T = O(\sqrt{N})$ applications of O_x and O_x^{-1}.

   (b) Let the set X consist of those permutations $x \in S^N$ where the 0 sits at an even location (i.e., $x_i = 0$ for even i), and Y be the set of permutations y where 0 sits at an odd location. Define relation $R \subseteq X \times Y$ as follows: start with an $x \in X$, suppose 0 sits at location i (which is even) and j be an arbitrary odd location; obtain y from x by switching x_i and x_j. Do this for each choice of odd j, so x is in relation with $N/2$ different y's. Now we can calculate the parameters of the strengthened adversary bound for the Boolean-valued function that determines whether 0 sits at an odd or even location in x. Clearly $m_0 = m_1 = N/2$. If $(x,y) \in R$ and $x_i \neq y_i$ then either $x_i = 0$ (in which case $\ell_{x,i} = N/2$ and $\ell_{y,i} = 1$) or $y_i = 0$ (in which case $\ell_{x,i} = 1$ and $\ell_{y,i} = N/2$). This implies $\ell_{\max} = N/2$. Then the strengthened adversary bound implies we need $\Omega(\sqrt{m_0 m_1/\ell_{\max}}) = \Omega(\sqrt{N})$ quantum queries to decide whether 0 sits at an odd or even location i in x. Hence the same lower bound applies to the at-least-as-hard problem of *finding* that location i.