Quantum Computing (5334QUCO8Y), Exam

Ronald de Wolf

Monday June 12, 2023, 10:00–13:00, Science Park USC Sporthal 1

The exam is "open book", meaning you can use any kind of paper you want but no electronic devices (except after 1pm for scanning and uploading your solutions). 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 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 is your number of points +1. An exam grade ≥ 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 $N = 2^n$ and $x \in \{0,1\}^N$ is an unknown bitstring. You are given one copy of the $(n+1)$-qubit state $\displaystyle\frac{1}{\sqrt{N}}\sum_{i=0}^{N-1}|i\rangle|x_i\rangle$. Show how you can convert this into state $\displaystyle\frac{1}{\sqrt{N}}\sum_{i=0}^{N-1}(-1)^{x_i}|i\rangle|1\rangle$ with success probability 1/2, such that you know when you succeeded.

2. **(1.5 points)** This question is about parallelizing search. Let $p \geq 1$ be a fixed integer. Suppose you have an input $x \in \{0,1\}^N$ and you have a special kind of oracle Q_x that answers p binary queries to x in parallel:

$$Q_x : |i_1, b_1,\, i_2, b_2, \ldots, i_p, b_p\rangle \mapsto |i_1, b_1 \oplus x_{i_1},\, i_2, b_2 \oplus x_{i_2}, \ldots, i_p, b_p \oplus x_{i_p}\rangle,$$

   where the i_j's are in $\{0, \ldots, N-1\}$ and the b_j's are bits.

   Show how you can find a solution to the search problem (i.e., an $i \in \{0, \ldots, N-1\}$ such that $x_i = 1$, if such an i exists) using $O(\sqrt{N/p})$ applications of Q_x. You may assume for simplicity that N/p is a power of 2. A precise higher-level description suffices, no need to draw a circuit.

3. **(2 points)** A *stochastic process* is a recursion of the form $x_{t+1} = Ax_t + b$, where A is an $N \times N$ matrix with real entries and b is an N-dimensional real vector. Given an initial vector x_0, the process induces a time-series $x_0, x_1, x_2, \ldots$ A vector $x^* \in \mathbb{R}^N$ is called a "stable state" of this process if it doesn't change under this recursion (i.e., $x_{t+1} = x_t$ for all t if $x_0 = x^*$).

   (a) Show that the stable state can be written as $x^* = (I-A)^{-1}b$, assuming $I-A$ is invertible.

   (b) Assume $N = 2^n$, A is Hermitian and sparse, $I-A$ is well-conditioned (which in particular implies that $I-A$ is invertible), and $|b\rangle$ can be efficiently prepared. Show how you can efficiently compute a state that's close to the n-qubit quantum state $|x^*\rangle = \frac{1}{\|x^*\|}\sum_{i=0}^{N-1} x_i^*|i\rangle$

corresponding to the stable state x^*. A precise description with references to the lecture notes suffices; I'm being deliberately a bit vague about words like "sparse", "well-conditioned", "efficient", "close" (you can be too in your answer) to focus on the ideas.

(c) Suppose we have two different stochastic processes: $x_{t+1} = Ax_t + b$ and $y_{t+1} = By_t + c$, where A, B are $N \times N$ matrices and $b, c \in \mathbb{R}^N$, with the same assumptions as in part (b). We are promised that their stable states x^* and y^* are either equal or have an inner product that's close to 0. Show how you can efficiently distinguish these two situations.
Hint: Remember the SWAP-test.

4. **(2 points)**

(a) Let $|\psi\rangle$ be an EPR-pair, $\frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$. Show that

$$\mathrm{Tr}((X \otimes X)|\psi\rangle\langle\psi|) = \mathrm{Tr}((Z \otimes Z)|\psi\rangle\langle\psi|) = 1 \text{ and } \mathrm{Tr}((X \otimes Z)|\psi\rangle\langle\psi|) = \mathrm{Tr}((Z \otimes X)|\psi\rangle\langle\psi|) = 0$$

(b) Show that $\frac{1}{\sqrt{2}}(X+Z)$ and $\frac{1}{\sqrt{2}}(X-Z)$ are ± 1-valued observables (i.e., Hermitian matrices with eigenvalues $+1$ and -1).

(c) Consider the CHSH game from Section 17.2: Alice and Bob each receive a uniformly-distributed input bit (x and y respectively), and they each produce an output bit (a and b respectively). They win the game if the condition $a \oplus b = x \cdot y$ holds.
Give a protocol for CHSH (different from the one in Section 17.2) that uses one EPR-pair between Alice and Bob, with winning probability $\frac{1}{2} + \frac{1}{2\sqrt{2}}$, by specifying ± 1-valued observables for Alice and Bob that depend on their respective input (i.e., Alice's observable depends on x, and Bob's observable depends on y).

5. **(2.5 points)** "Private information retrieval" is the following cryptographic problem. Alice has a string $x \in \{0,1\}^n$, and Bob has an index $i \in \{0, \ldots, n-1\}$. Bob wants to learn the bit x_i (with success probability 1). If Bob didn't mind telling Alice what i is, then this information retrieval is easy: Bob sends i to Alice and she sends back x_i, costing only $\log(n) + 1$ bits of communication. However, now suppose Bob doesn't want to give Alice any information about his i, but he still wants to learn x_i. It's fine if Bob learns more than x_i, but Alice should learn nothing about i (hence the adjective "private").

(a) Show that n qubits of communication between Alice and Bob are *sufficient* to achieve this private information retrieval.

(b) Let $|\psi\rangle_{AB}$ be a bipartite quantum state of the form $\frac{1}{\sqrt{2^n}} \sum_{x \in \{0,1\}^n} |x\rangle_A |\phi_x\rangle_{AB} |x\rangle_B$, where the $|\phi_x\rangle_{AB}$ are arbitrary (normalized) states shared between Alice and Bob. Show that n qubits of communication between Alice and Bob are *necessary* to create $|\psi\rangle_{AB}$ if they start from an unentangled state.

(c) Show that n qubits of communication between Alice and Bob are *necessary* to achieve private information retrieval.
Hint: Consider a private information retrieval protocol where Alice and Bob's communication and local operations on initial state $|x\rangle_A|i\rangle_B$ correspond to a unitary U that requires q qubits of communication to implement in total (the unitary comes with a specification which qubits are Alice's and which are Bob's at the end to reflect the communication; you may just assume that such a U exists). Analyze the Schmidt decompositions of the n different states $|\psi_i\rangle_{AB} = \frac{1}{\sqrt{2^n}} \sum_{x \in \{0,1\}^n} |x\rangle_A U(|x\rangle_A |i\rangle_B)$, use this to create some bipartite state of the form of (b) by local operations on Bob's side, and then use (b) to conclude $q \geq n$.

Intended solutions

1. Apply a Hadamard gate to the last qubit. This turns the state into

$$\frac{1}{\sqrt{N}}\sum_{i=0}^{N-1}|i\rangle\frac{1}{\sqrt{2}}(|0\rangle+(-1)^{x_i}|1\rangle)=\frac{1}{\sqrt{2N}}\sum_{i=0}^{N-1}|i\rangle|0\rangle+\frac{1}{\sqrt{2N}}\sum_{i=0}^{N-1}(-1)^{x_i}|i\rangle|1\rangle.$$

Measure the last qubit. You get outcome 1 with probability 1/2 (and you of course know when this happened), and then the state becomes the desired state $\frac{1}{\sqrt{N}}\sum_{i=0}^{N-1}(-1)^{x_i}|i\rangle|1\rangle$.

2. Think of the N-bit string x as consisting of p separate pieces of N/p bits each. Run p Grovers in parallel, one for each of the separate pieces, using $O(\sqrt{N/p})$ queries for each piece. Each block of p parallel queries (one to each of the pieces) corresponds to one application of Q_x, with the p target qubits set to $|-\rangle$ in order to get the answers in the phase. The p parallel runs of Grover yield p potential solutions at the end, which can all be checked (to see which ones of them are actual solutions) using one more application of Q_x. If the N-bit string has at least one solution, then the run of Grover on the N/p-bit piece containing that solution will have probability $\geq 2/3$ of finding a solution.

3. (a) The equation $x^*=Ax^*+b$ is equivalent to $(I-A)x^*=b$, hence $x^*=(I-A)^{-1}b$.

   (b) Let $A'=I-A$. Then A' is Hermitian, sparse, and well-conditioned. A' has at most one additional nonzero entry in each row and column compared to A, namely on the diagonal, so we can efficiently turn the sparse-access oracles for A into those for A'. So all the conditions for applying the HHL algorithm are in place, hence we can find a quantum state that is close to the state $|x^*\rangle$ corresponding to the solution x^* of the linear system $A'x=b$.

   (c) Use part (b) twice, once to generate a state very close to $|x^*\rangle$ and once to generate a state very close to $|y^*\rangle$. Then use the SWAP-test (Section 16.6) to test if these two states are approximately equal or approximately orthogonal. The SWAP-test will yield measurement outcome 0 with probability ≈ 1 if the two states are equal, and with probability $\approx 1/2$ if the two states are almost orthogonal. You can repeat this algorithm a few times to reduce the error probability to some small constant.

4. (a) We have $(X\otimes X)|\psi\rangle=(Z\otimes Z)|\psi\rangle=|\psi\rangle$, which implies the first part (because $\mathrm{Tr}(M|\psi\rangle\langle\psi|)=\langle\psi|M|\psi\rangle$ due to the cyclicity of the trace). We have $(X\otimes Z)|\psi\rangle=\frac{1}{\sqrt{2}}(|10\rangle-|01\rangle)$ and $(Z\otimes X)|\psi\rangle=\frac{1}{\sqrt{2}}(|01\rangle-|10\rangle)$, which are both orthogonal to $|\psi\rangle$, implying the second part.

   (b) The two matrices are both Hermitian, so their eigenvalues are real. Hence to conclude that their eigenvalues are in $\{-1,+1\}$, it suffices to show that these matrices both square to identity. We have $(\frac{1}{\sqrt{2}}(X+Z))^2=\frac{1}{2}(X^2+Z^2+XZ+ZX)$, which is I because $X^2=Z^2=I$ while $XZ=-ZX$. Similarly $(\frac{1}{\sqrt{2}}(X-Z))^2=\frac{1}{2}(X^2+Z^2-XZ-ZX)=I$.

   (c) Alice and Bob start with one EPR-pair $|\psi\rangle$. Alice measures ± 1-valued observable $A_0=X$ if $x=0$ and she measures observable $A_1=Z$ if $x=1$; if her measurement outcome is $+1$ then Alice outputs $a=0$, and if her measurement outcome is -1 then she outputs $a=1$. Bob measures observable $B_0=\frac{1}{\sqrt{2}}(X+Z)$ if $y=0$ and $B_0=\frac{1}{\sqrt{2}}(X-Z)$ if

$y = 1$, and similarly converts the measurement outcome to $b \in \{0, 1\}$ (part (b) implies that Bob's observables are ± 1-valued as well). Using part (a) and linearity of the trace, we calculate the probability that $a \oplus b = 0$, minus the probability that $a \oplus b = 1$, as

$$\mathrm{Tr}((A_x \otimes B_y)|\psi\rangle\langle\psi|) = \begin{cases} \frac{1}{\sqrt{2}} & \text{if } x \cdot y = 0 \\ -\frac{1}{\sqrt{2}} & \text{if } x \cdot y = 1 \end{cases}$$

Hence the probability that $a \oplus b = 0$ is $\frac{1}{2} + \frac{1}{2\sqrt{2}}$ if $x \cdot y = 0$, and that probability is $\frac{1}{2} - \frac{1}{2\sqrt{2}}$ if $x \cdot y = 1$ (i.e., if $x = y = 1$). Thus the probability of a winning output-pair ab is $\frac{1}{2} + \frac{1}{2\sqrt{2}}$ for all 4 possible input-pairs xy.

Comment: This probability $\frac{1}{2} + \frac{1}{2\sqrt{2}}$ *happens to be equal to* $\cos(\pi/8)^2 \approx 0.85$*, the same winning probability as the protocol in Section 17.2. This is optimal for CHSH because of the Tsirelson bound (Exercise 17.6).*

5. (a) Alice just sends the whole x to Bob (which takes n bits of communication), Bob sends nothing to Alice. Then Bob knows every bit of x including x_i, while Alice has learned nothing about i.

   (b) The initial unentangled state has Schmidt rank 1, the final state $|\psi\rangle_{AB}$ has Schmidt rank at least 2^n because each $|\phi_x\rangle$ has Schmidt rank at least 1 and there are 2^n orthonormal $|x\rangle$'s. It is easy to see that one qubit of communication can at most double the Schmidt rank of a bipartite state: if Alice sends Bob one qubit (i.e., one qubit changes ownership) then the number of states in Bob's Schmidt basis will at most double. Hence at least n qubits of communication are needed to go from a bipartite state of Schmidt rank 1 to a bipartite state of Schmidt rank $\geq 2^n$.

   (c) Fix any quantum communication protocol for information retrieval where Alice learns nothing about Bob's input i. Define the n states $|\psi_i\rangle_{AB}$ as in the hint. These states have to be the same on Alice's side (i.e., if you trace out Bob's qubits from $|\psi_i\rangle_{AB}$ then you get a mixed state ρ_A that doesn't depend on i), otherwise she could get nonzero information about i by measuring her part of the state at the end of the protocol. Hence there exist Schmidt decompositions $|\psi_i\rangle_{AB} = \sum_k \lambda_k |a_k\rangle_A |b_k^i\rangle_B$ that only differ in Bob's orthonormal basis $\{|b_k^i\rangle\}_k$, which can depend on i. Bob can learn x_1 from $|\psi_1\rangle_{AB}$ with probability 1, so without disturbing the state. Then he can locally change $|\psi_1\rangle_{AB} \mapsto |\psi_2\rangle_{AB}$ by applying the unitary map $|b_k^1\rangle \mapsto |b_k^2\rangle$ to his part of $|\psi_1\rangle_{AB}$, which costs no communication. From $|\psi_2\rangle_{AB}$ Bob can learn x_2, then locally change to $|\psi_3\rangle_{AB}$ etc., eventually recovering x completely. Since this actually happens in superposition over all x, the bipartite state $|\psi\rangle_{AB}$ after Bob has recovered x, will be of the form of (b). But (b) says that at least n qubits of communication are needed to produce $|\psi\rangle_{AB}$ starting from an unentangled state. Therefore the number of qubits of communication needed to implement U must be at least n.