Quantum Computing (5334QUCO8Y), Exam

Ronald de Wolf

Monday Jan 27, 2025, 14:00–17:00
room REC D4.02 Brug, UvA Roeterseiland

The exam is "open book": you can bring any kind of paper you want, but no electronic devices. 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 the earlier parts in order to answer later parts, even if you didn't provide answers to the earlier parts. You are allowed to refer to things explained in the lecture notes, or things that were part of homework exercises that you did, without re-explaining their details in your answer (just state clearly what the invoked construction achieves).

The total number of points adds up to 9; your exam grade is number of points +1. An exam grade ≥ 5 is a necessary condition for passing the course. Your final grade is 60% exam + 40% homework, rounded to the nearest half-integer (except 5.5).

1. **(1 point)** We know that symmetrization cannot increase the degree of an N-variate polynomial. Can it *decrease* it? If yes, give an example polynomial; if no, argue why.

2. **(2 points)** This question is about a variant of quantum phase estimation. Section 4.6 assumed controlled-U's for this, some of which were applied sequentially, but in this exercise we are only allowed to apply uncontrolled U's, and only in parallel (= simultaneous, not one after another). Suppose $U = \begin{pmatrix} 1 & 0 \\ 0 & e^{2\pi i\phi} \end{pmatrix}$ with unknown $\phi = \{0, 1/4, 1/2, 3/4\}$, so 2 bits suffice to exactly represent the number ϕ.

   (a) Show how to prepare a qubit $\frac{1}{\sqrt{2}}(|0\rangle + e^{2\pi i\phi}|1\rangle)$ using one application of U and some elementary gate(s).

   (b) Show how to prepare a qubit $\frac{1}{\sqrt{2}}(|0\rangle + e^{2\pi i2\phi}|1\rangle)$ using two parallel applications of U and some elementary gate(s). You're allowed to use an auxiliary qubit.
   Hint: Set up an EPR-pair.

   (c) Give a circuit that acts on 3 qubits, that uses $U^{\otimes 3}$ and no other applications of U, and that returns ϕ with probability 1.

3. **(2 points)** Suppose we are given k copies of an unknown m-qubit state $|\psi\rangle$, and we want to learn (by means of some measurement on the km qubits) the 2^m-dimensional vector of amplitudes of $|\psi\rangle$ up to Euclidean distance $\leq 1/10$, with success probability $\geq 2/3$. Show that this requires $k \geq \Omega(2^m/m)$, i.e. a measurement that works for every $|\psi\rangle$ needs to measure at least some constant times $2^m/m$ many copies of $|\psi\rangle$ to be able to produce such an approximating vector.

   Hint: Consider the 2^n different fingerprint states $|\phi_x\rangle$ on $m = \log n + O(1)$ qubits from Section 16.6, for all 2^n possible $x \in \{0,1\}^n$, and how much information one can get from the km qubits of k copies of $|\phi_x\rangle$. Argue that if you know the state $|\phi_x\rangle$ up to Euclidean distance $\leq 1/10$, then you can find the n-bit string x.

4. **(1.5 points)** Suppose we have a qubit whose density matrix is $\rho = \alpha_0 I + \alpha_1 X + \alpha_2 Y + \alpha_3 Z$, where $\alpha_0, \alpha_1, \alpha_2, \alpha_3$ are real coefficients and I, X, Y, Z are the Pauli matrices.

   (a) Show that $\alpha_0 = 1/2$.

   (b) *Depolarizing noise* (of strength $p \in [0,1]$) acts on a qubit as follows: with probability $1-p$ nothing happens to the qubit, and with probability p the qubit is replaced by the "completely mixed state" of a qubit, whose density matrix is $I/2$.

   Show that depolarizing noise on the above qubit doesn't change the coefficient α_0, but shrinks each of $\alpha_1, \alpha_2, \alpha_3$ by a factor of $1-p$.

5. **(2.5 points)**

   (a) Let $H \in \{0,1\}^{2^n \times 2^n}$ be a symmetric Boolean matrix whose entries we can query in the usual way (we have a unitary O_H such that $O_H|j,k,b\rangle = |j,k,b \oplus H_{jk}\rangle$ for all $j, k, \in \{0,1\}^n$ and $b \in \{0,1\}$). We are promised that there is an unknown $j \neq 0^n$ such that $H_{0^n,j} = H_{j,0^n} = 1$, and all other entries of H are 0. Prove that finding this j (with success probability $\geq 2/3$) requires $\Omega(\sqrt{2^n})$ quantum queries.

   (b) For the same H as in (a), show that the state $e^{iH}|0^n\rangle$ is $\cos(1)|0^n\rangle + i\sin(1)|j\rangle$.

   Hint: Use Taylor series $e^{iH} = \sum_{k\geq 0}(iH)^k/k!$, $\cos(x) = \sum_{\text{even}\, k\geq 0}(ix)^k/k!$, $i\sin(x) = \sum_{\text{odd}\, k\geq 0}(ix)^k/k!$

   (c) Consider the Hamiltonian simulation problem for an n-qubit Hamiltonian H, which is to implement (as a circuit, including queries to entries of H) a unitary $\widetilde{U}$ that is ε-close to e^{iHt}. Now suppose that H is s-sparse but (in contrast to p.77) we do not have access to the $O_{H,loc}$ oracle, but only to the O_H oracle for accessing the entries of H. Show an $\Omega(\sqrt{2^n})$ quantum query lower bound in this setting for Hamiltonian simulation for the case of $s = 1$, $t = 1$, and $\varepsilon = 1/100$.

   Hint: $(\sin(1) - 1/100)^2 > 2/3$.

   Comment: This shows that the efficient access to the locations of the few non-zero entries of H that is provided by $O_{H,loc}$ (on p.77), is essential for efficient Hamiltonian simulation for a sparse H; with only entry-wise access, Hamiltonian simulation requires exponentially large circuits, even for sparsity $s = 1$.

Model solutions

1. Yes. For example $p(x_0, x_1) = x_0 - x_1$ symmetrizes to the 0 polynomial (so the degree goes from 1 to 0).

2. (a) Start with $|0\rangle$, apply H, and then U.

   (b) Start with $|00\rangle$, apply $H \otimes I$ and then CNOT to create an EPR-pair. Apply $U \otimes U$ to get the right phase on $|11\rangle$. Apply a CNOT to set the last qubit back to 0.

   (c) Use (a) and (b) to create the 2-qubit product state $F_4|\phi\rangle$ (and a now-irrelevant 3rd qubit in state $|0\rangle$). Applying the inverse QFT circuit F_4^{-1} gives us the 2-bit basis state $|4\phi\rangle$, whose 2 bits are the bits of ϕ.

3. Consider m-qubit quantum fingerprints $|\phi_x\rangle$ and $|\phi_y\rangle$ for distinct $x, y \in \{0,1\}^n$ (as defined halfway p.138, with $m = \log n + O(1)$). Their distance is large, because their inner product is small:
$$\| \, |\phi_x\rangle - |\phi_y\rangle \, \|^2 = 2 - 2\langle\phi_x|\phi_y\rangle \geq 1.96.$$
If vector $v \in \mathbb{C}^{2^m}$ is 1/10-close to $|\phi_x\rangle$ then it cannot also be 1/10-close to $|\phi_y\rangle$, because otherwise we would have $\| \, |\phi_x\rangle - |\phi_y\rangle \, \| \leq \| \, |\phi_x\rangle - v \, \| + \| \, v - |\phi_y\rangle \, \| \leq 2/10$ by triangle inequality, contradicting that the distance between $|\phi_x\rangle$ and $|\phi_y\rangle$ must be large. Accordingly, if we can obtain a vector v that is 1/10-close to an unknown quantum fingerprint $|\phi_x\rangle$ (by doing a measurement on the km-qubit state $|\phi_x\rangle^{\otimes k}$) then that uniquely determines $|\phi_x\rangle$, which means we can compute x from v. Think of $x \in \{0,1\}^n$ as uniformly random, then learning x gives us $n = 2^{m-O(1)} = \Omega(2^m)$ bits of information. If we can obtain such a v with success probability $\geq 2/3$ rather than with success probability 1, then we are still obtaining $\Omega(n)$ bits of information (about x) from the km qubits of the k copies of $|\phi_x\rangle$. By Holevo's theorem, we must have $km \geq \Omega(n) = \Omega(2^m)$ and hence $k \geq \Omega(n/m) = \Omega(2^m/m)$.

4. (a) $1 = \mathrm{Tr}(\rho) = \alpha_0\mathrm{Tr}(I) + \alpha_1\mathrm{Tr}(X) + \alpha_2\mathrm{Tr}(Y) + \alpha_3\mathrm{Tr}(Z) = 2\alpha_0$, hence $\alpha_0 = 1/2$.

   (b) The mixed state resulting from applying depolarizing noise to a qubit ρ is
$$(1-p)\rho + pI/2 = (\underbrace{(1-p)\alpha_0 + p/2}_{1/2})I + (1-p)\alpha_1 X + (1-p)\alpha_2 Y + (1-p)\alpha_3 Z.$$

5. (a) Consider an input $x \in \{0,1\}^N$ to the search problem (from Chapter 7) with $N = 2^n - 1$. We can define a symmetric $2^n \times 2^n$ matrix H by setting its 0th row and column to $(0, x)$, and all other entries to 0. We can now implement a query to entries of this H by means of one query to x. If we can find the j that the exercise talks about using T queries to entries of H, then we can solve the search problem on x with the same success probability, using T queries to x. Since we know the latter takes $\Omega(\sqrt{N})$ queries (this was proved multiple times in Chapter 11), we must have $T = \Omega(\sqrt{N}) = \Omega(\sqrt{2^n - 1}) = \Omega(\sqrt{2^n})$.

(b) Note that $H|0^n\rangle = |j\rangle$ and $H|j\rangle = |0^n\rangle$, so $H^k|0^n\rangle = |0^n\rangle$ for all even integers k, and $H^k|0^n\rangle = |j\rangle$ for all odd k. Using the three Taylor series from the hint, we have

$$\begin{aligned} e^{iH}|0^n\rangle &= \sum_{k\geq 0} \frac{(iH)^k}{k!}|0^n\rangle \\ &= \sum_{\text{even}\, k\geq 0} \frac{i^k}{k!} H^k|0^n\rangle + \sum_{\text{odd}\, k\geq 0} \frac{i^k}{k!} H^k|0^n\rangle = \cos(1)|0^n\rangle + i\sin(1)|j\rangle \end{aligned}$$

(c) Consider a matrix H as in part (a) of the question, which is an n-qubit Hamiltonian that is s-sparse for $s = 1$. Let $\widetilde{U}$ be a circuit (including queries to entries of H, elementary gates that are independent of H, possibly even some auxiliary qubits that start and end in $|0\rangle$) that does Hamiltonian simulation for this H for time $t = 1$ and $\varepsilon = 1/100$. In other words, $\widetilde{U}$ is close to e^{iH}, within operator-norm error $1/100$. Apply $\widetilde{U}$ to $|0^n\rangle$. The resulting state has distance $\leq 1/100$ to the state of (b), which in particular means that the magnitude of the amplitude on basis state $|j\rangle$ is at least $\sin(1) - 1/100$. Measuring the state in the computational basis gives us j with probability at least $(\sin(1) - 1/100)^2 > 2/3$. By part (a) this requires at least $\Omega(\sqrt{2^n})$ queries to entries of H, so the circuit $\widetilde{U}$ has to contain at least that many queries.