Quantum Proofs, Semester A 2024

Homework # 2 Solutions **January 21st, 2024**

Please read the document and check your understanding of the answer. If you feel that your solution was correct, but I mistakenly did not award you all points, please talk to me. If my sketch is not detailed enough and you would like to see a full solution, please ask me as well.

Problems:

1. (2 points) **Product test**

   Consider the following quantum algorithm, that takes as input two $2n$-qubit quantum states $|\psi\rangle$ and $|\psi'\rangle$ as input:

   (a) Initialize an ancilla qubit in state $|0\rangle$;

   (b) Apply a Hadamard on the ancilla;

   (c) Controlled on the ancilla being in state $|1\rangle$, exchange the first n qubits of $|\psi\rangle$ with the first n qubits of $|\psi'\rangle$;

   (d) Apply a Hadamard on the ancilla;

   (e) Measure the ancilla in the computational basis. Accept if the outcome is 0.

   To be clear, on an input of the special form $|\psi\rangle = |\phi_1\rangle \otimes |\phi_2\rangle$, $|\psi'\rangle = |\phi'_1\rangle \otimes |\phi'_2\rangle$ where each $|\phi_i\rangle, |\phi'_i\rangle$ is on n qubits for $i \in \{1, 2\}$, this algorithm generates the following sequence of states:

$$\begin{aligned}
|\phi_1\rangle|\phi_2\rangle|\phi'_1\rangle|\phi'_2\rangle &\mapsto |0\rangle|\phi_1\rangle|\phi_2\rangle|\phi'_1\rangle|\phi'_2\rangle \\
&\mapsto |+\rangle|\phi_1\rangle|\phi_2\rangle|\phi'_1\rangle|\phi'_2\rangle \\
&\mapsto \frac{1}{\sqrt{2}}\big(|0\rangle|\phi_1\rangle|\phi_2\rangle|\phi'_1\rangle|\phi'_2\rangle + |1\rangle|\phi'_1\rangle|\phi_2\rangle|\phi_1\rangle|\phi'_2\rangle\big) \\
&\mapsto \frac{1}{\sqrt{2}}\big(|+\rangle|\phi_1\rangle|\phi_2\rangle|\phi'_1\rangle|\phi'_2\rangle + |-\rangle|\phi'_1\rangle|\phi_2\rangle|\phi_1\rangle|\phi'_2\rangle\big) \ ,
\end{aligned}$$

   followed by a measurement of the first qubit in the standard basis.

   (a) For input states in product form as above, compute the probability that this algorithm accepts.

   A direct calculation, continuing on the derivation above, gives a probability of $\frac{1}{2}(1 + |\langle\phi_1|\phi'_1\rangle|^2)$

   (b) Now suppose that the input is two copies of the *same* arbitrary pure state $|\psi\rangle = |\psi'\rangle$. Compute the probability that the algorithm accepts, as a function of the reduced density matrix ρ of $|\psi\rangle$ on the first n qubits. *[Hint: You should know that $|\psi\rangle$ can be written in the form $|\psi\rangle = \sum_i \sqrt{\lambda_i}|u_i\rangle|v_i\rangle$, for orthonormal*

families $\{|u_i\rangle\}$ and $\{|v_i\rangle\}$ — this is called the "Schmidt decomposition". Then, $\rho = \sum_i \lambda_i |u_i\rangle\langle u_i|$. Use the Schmidt decomposition to analyze the algorithm.]

The calculation from the previous question generalizes to give a result $\frac{1}{2}(1 + \mathrm{Tr}(\rho^2))$.

(c) Deduce that, if the algorithm accepts $|\psi\rangle$ and $|\psi'\rangle = |\psi\rangle$ with probability at least $1 - \varepsilon$, then $|\psi\rangle$ is within trace distance at most η of a product state — give an upper bound on η as a function of ε.

There were multiple ways to answer this question. If $|\psi\rangle = \sum_i \sqrt{\lambda_i}|u_i\rangle|v_i\rangle$, there are two natural candidates for the closest product state: $|u_1\rangle|v_1\rangle$, where λ_1 is the largest Schmidt coefficient, or $(\sum_i \sqrt{\lambda_i}|u_i\rangle)(\sum_i \sqrt{\lambda_i}|v_i\rangle)$. The first state has a squared overlap with $|\psi\rangle$ of λ_1, and the second gives $(\sum_i \lambda_i^{3/2})^2$. Using $\sum_i \lambda_i = 1$ we get $\lambda_1 \geq \sum_i \lambda_i^2 \geq 1 - 2\varepsilon$ by the previous question. The second state will only have a worse overlap, since by the Cauchy-Schwarz inequality

$$\Big(\sum_i \lambda_i^{3/2}\Big)^2 \leq \Big(\sum_i \lambda_i\Big)\Big(\sum_i \lambda_i^2\Big) = \sum_i \lambda_i^2 \,.$$

Translating back to the trace norm,

$$\big\||\psi\rangle\langle\psi| - |u_1\rangle\langle u_1| \otimes |v_1\rangle\langle v_1|\big\|_1^2 = 1 - |\langle\psi| \cdot |u_1\rangle \otimes |v_1\rangle|^2 \leq 2\varepsilon \,.$$

(Even though the second choice of a product state, or other choices, is somewhat worse, I accepted them as answer.)

2. (2 points) **Measuring a Hermitian operator**

Let $q \geq 1$ and H a Hermitian matrix on the space $\mathcal{H} = (\mathbb{C}^2)^{\otimes q}$ of q qubits such that $\|H\| \leq 1$. Describe a measurement procedure on $\mathcal{H}$ that returns an outcome $b \in [-1, 1]$ such that for any density ρ on q qubits, on expectation, $\mathbf{E}[b] = \mathrm{Tr}(H\rho)$. The "measurement procedure" is only allowed to perform *arbitrary unitary* operations, ancilla creation in state $|0\rangle$, and measurements in the *standard* basis.

Let $H = \sum_i \lambda_i |u_i\rangle\langle u_i|$ be the singular value decomposition. (This can be computed in time $O(2^{3q})$ given an explicit matrix representation of H.) A short answer is to apply the unitary $U = \sum_i |i\rangle\langle u_i|$ to the q qubits on which H acts, measure in the computational basis to obtain an i and return λ_i. If we want to make this entire procedure into a quantum circuit we can also apply

$$U' = \sum_i |i\rangle\langle u_i| \otimes \big(|+_i\rangle\langle 0| + |-_i\rangle\langle 1|\big) \,,$$

where

$$|\pm_i\rangle = \sqrt{\frac{1 + \lambda_i}{2}}|0\rangle \pm \sqrt{\frac{1 - \lambda_i}{2}}|1\rangle \,,$$

the first tensor factor acts on the q qubits acted on by H, and the second factor acts on an additional ancilla qubit initialized in the $|0\rangle$ state. (The second part, $|-i\rangle\langle 1|$,

doesn't do anything when the ancilla is $|0\rangle$; it is there only so that I can claim that U' is indeed a unitary.) We then measure the ancilla qubit in the standard basis to obtain an outcome $o \in \{0,1\}$ and return $b = (-1)^o$. This has the same effect.

3. (2 points) **Basic properties of the Feynman-Kitaev construction**

   Everyone gave a good answer, so I will not include a solution.

4. (5 points) **QMA with log-depth verifier**

   (a) Let V be a QMA verification circuit acting on q proof qubits, m ancilla qubits, and with T gates in total, such that the completeness and soundness parameters of V are $1-2^{-p(n)}$ and $2^{-p(n)}$ respectively. Show that there is a verification circuit V', acting on q proof qubits and $O((q+m)T)$ ancilla qubits and with $O((q+m)T)$ gates, such that V' has the same completeness and soundness parameters as V (i.e. V' performs the same computation as V) and furthermore each qubit of V' (proof or ancilla) takes part in a *constant* number of gates.

   Everyone had the right idea, which is to use abundant swap gates to perform at most one step of the original computation on each wire.

   (A nice observation is that you can imagine replacing the swap steps by teleportation. This means that any quantum computation can be implemented in "constant adaptive depth" in the following manner: start with the input to the computation and a lot of ancillas initialized as EPR pairs. At each step, apply one gate from the original circuit, and then immediately after perform a teleportation measurement on the output qubit(s) and half of one (or two) EPR pairs. Apply the reqiured correction on the target EPR pair(s) and proceed. Observe that on each EPR pair we apply (a) a correction, (b) a unitary, (c) a teleportation measurement; then we move to another EPR pair. The combination (a)(b)(c) can be thought of as a single two-qubit measurement. And so the whole computation reduces (other than the preparation of the EPR pairs) to a sequence of two-qubit measurements, each on *non-overlapping* pairs of qubits, and chosen adaptively as a function of previous measurement outcomes. This is very related to the model of "measurement-based computation" (which achieves the same, but now with only *single-qubit* measurements! Albeit with a more complex, but computation-independent, initial state).

   (b) Show how to appropriately modify the Feynman-Kitaev construction (with unary clock) so that when it is applied to V', it returns a local Hamiltonian such that furthermore each qubit is acted on by at most a *constant* number of local terms. Here, you do not need to prove soundness of your modified construction; only state what modification you suggest making to satisfy the "constant degree" constraint. *[Hint: be mindful of the clock qubits!]*

Here again almost everyone had a good idea, which amounts to delaying the ancilla checks to a time when they are needed, so that they are not all performed at the same time 0.

(c) Show that for any promise problem $L = (L_y, L_n)$ in QMA, there is a polynomial-time computable family of verifiers $z \mapsto V_z$ for L such that for each z, V_z is specified by a quantum circuit of depth $O(\log |z|)$.[1] Here the constant implicit in $O(\cdot)$ may depend on L, but not on z. *[Hint: use the previous questions. It only remains to design a logarithmic-depth* QMA *verifier for verifying a local Hamiltonian as in part (b). Describe how this verifier proceeds, and argue that it has the required depth. You are not asked to formally show soundness of the verifier.]*

This question required more thought; yet here again most of you found their way to a satisfactory solution, which is great.

A first observation is that individual terms of the local Hamiltonian obtained by combining steps (a) and (b) of the question can be grouped in (at most) a logarithmic number of "layers", such that terms in the same layer never overlap. (This can be shown e.g. by greedy counting, or using "known facts" on coloring of bounded-degree graphs.) From there, there were different options on how to proceed:

- One option is to sequentially measure each group, storing each measurement outcome, for each individual term, in an ancilla qubit. All terms in a single group can then be measured at the same time, because they act on disjoint qubits. Once this is completed, we need to compute the OR of all m ancilla qubits, which can again be done in logarithmic depth using a tree-like construction.
- Another option is to initialize a register of log(#layers)-many qubits in a random (all-$|+\rangle$) state, and then control on that register to decide which layer to measure. (So we measure a layer at random, as opposed to a single layer.) One caveat here is that all terms in the layer need to look at the same random bits at the same time; this can be solved by copying the random bits (in the standard basis, using CNOTs organized in a tree) before using them.

While the "standard method" which we saw in class measures a single term uniformly at random, it is clear that measuring an entire layer, chosen at random, will only improve the soundness of the verification procedure. That is is also possible to measure the layers in a fixed sequential order is harder to show, but still true — although it induces a higher loss in soundness; but I was not requiring you to verify this so I accepted the solution as well.

A final caveat common to both methods is that the resulting completeness-soundness gap is only 1/poly. Luckily this can be amplified in parallel, using

[1] The depth of a circuit is the smallest number of layers that its gates can be organized in. Multiple gates are allowed to operate in parallel if they act on disjoint sets of wires.

again a log-depth circuit for aggregating the results of each of the parallel repetitions.