Quantum Complexity Theory *Instructor: Anand Natarajan*

Problem set 1

Due: Tues, Nov 3, 2026 *TA: Tina Zhang*

Instructions

You may work with collaborators and consult textbooks or other references, but please list your collaborators and cite any references you use.

AI guidance: We will not police your use of AI but recommend that you use it judiciously to enhance your learning. Use of AI to look up mathematical facts (e.g. standard combinatorial identities or inequalities) or to review prerequisite material is acceptable and encouraged.

Q1: Polynomials, moments, and Simon's problem (7 points)

In class, we used the polynomial method to show that a problem that's classically intractable is still quantumly intractable (albeit slightly less so). In this problem we'll see how the polynomial perspective sheds light on an instance of exponential quantum advantage.

(a) (2 points) Let μ_1, μ_2 be two distributions over $\{0,1\}^N$ and let d be an integer with $0 \leq d \leq N$. Say that μ_1 and μ_2 are *d-wise identical* if for all sets of d indices $\{i_1, \ldots, i_d\} \subseteq [N]$ and for every $z \in \{0,1\}^d$, it holds that

$$\Pr_{\mu_1}[(x_{i_1}, \ldots, x_{i_d}) = z] = \Pr_{\mu_2}[(x_{i_1}, \ldots, x_{i_d}) = z].$$

Show that no T-query quantum algorithm can distinguish samples from μ_1 from samples from μ_2 if $T \leq \lfloor d/2 \rfloor$.

(b) (2 points) Now, say that μ_1 is d-wise $(1+\varepsilon)$-bounded by μ_2 if for all sets of d indices $\{i_1, \ldots, i_d\} \subseteq [N]$ and for every $z \in \{0,1\}^d$, it holds that

$$\Pr_{\mu_1}[(x_{i_1}, \ldots, x_{i_d}) = z] \leq (1+\varepsilon) \cdot \Pr_{\mu_2}[(x_{i_1}, \ldots, x_{i_d}) = z].$$

Show that if $T \leq d$, then for any T-query *classical* deterministic algorithm $\mathcal{A}$ outputting a single bit,

$$|\Pr_{x \sim \mu_1}[\mathcal{A}(x) = 1] - \Pr_{x \sim \mu_2}[\mathcal{A}(x) = 1]| \leq \varepsilon.$$

This quantity is called the *distinguishing advantage* of $\mathcal{A}$.

(c) (3 points) We will now give an example that shows that the result of the previous part does not extend to quantum algorithms. Let $N = n \cdot 2^n$, so that a string $x \in \{0,1\}^N$ can be interpreted as the truth table of a function $f_x : \{0,1\}^n \to \{0,1\}^n$. Let μ_1 be the uniform distribution over truth tables of functions with Simon's property (recall that f has Simon's property if there exists $s \in \{0,1\}^n \setminus \{0^n\}$ such that $f(u) = f(v)$ iff $u = v$ or $u = v \oplus s$). Let μ_2 be the uniform distribution over truth tables of bijections. Use the previous part to show that for any classical algorithm $\mathcal{A}$ making T queries, its distinguishing advantage for μ_1 and μ_2 is at most $O(T^2/2^n)$.

These two distributions are of course distinguished quantumly by Simon's algorithm, which can get a constant distinguishing advantage using $\mathrm{poly}(n)$ queries.

Q2: Circuit-to-Hamiltonian for simple quantum circuits (7 points)

Here we will look at circuit-to-Hamiltonian mappings for simple classes of quantum circuits, and try to understand why they do not generalize to general quantum circuits. This helps us appreciate the power of the Feynman-Kitaev mapping. All gates are assumed to act on at most two qubits.

(a) (1 point) Suppose $U = U_T \dots U_1$ is a quantum circuit of depth d on n qubits, that is $U = V_d \dots V_2 V_1$, where V_i is a layer of gates on n qubits in which no two gates overlap. Consider the state $U|0\rangle^n$. Show that this state is the unique ground state of the Hamiltonian $H_U = \sum_i U|1\rangle\langle 1|_i U^\dagger$.

(b) (1 point) In the above construction, what is the locality of the Hamiltonian H_U? What is the spectral gap (i.e. the difference between the two smallest eigenvalues) of H_U? Why is H_U not a good option for QMA-completeness of the local Hamiltonian problem?

(c) (3 points) Next, consider a quantum circuit $W = W_T \dots W_1$ on a d dimensional quantum system $\mathcal{H}_S$. Consider the quantum state $W|0\rangle$, where $|0\rangle$ is a fixed state in $\mathcal{H}_S$. Introduce $T+1$ copies of $\mathcal{H}_S$, labelled $S_0, S_1 \dots S_T$, and consider the Hamiltonian

$$H_W = \sum_{i=0}^{T-1} (\mathbb{1} \otimes (W_{i+1})_{S_{i+1}})(\mathbb{1} - \Pi_{sym,S_i,S_{i+1}})(\mathbb{1} \otimes (W_{i+1})^\dagger_{S_{i+1}}),$$

where

$$\Pi_{sym,S_i,S_{i+1}} = \frac{1}{2}(I_{S_i} \otimes I_{S_{i+1}} + SWAP_{S_i S_{i+1}})$$

is the projector onto the symmetric subspace of systems S_i and S_{i+1}. Prove that the ground space of H_W is

$$\text{span}\{|\psi\rangle_{S_0} \otimes W_1|\psi\rangle_{S_1} \otimes W_2 W_1|\psi\rangle_{S_2} \otimes \dots (W_T \dots W_1)|\psi\rangle_{S_T}, \forall \psi \in \mathcal{H}_S\}.$$

(d) (1 point) Write down the ground space of $(\mathbb{1} - |0\rangle\langle 0|)_{S_0} + H_W$.

(e) (1 point) If S consists of n qubits ($d = 2^n$) and W_i are 2-qubit gates, what is the locality of H_W? Why is $(\mathbb{1} - |0\rangle\langle 0|)_{S_0} + H_W$ not a good option for QMA-completeness of the local Hamiltonian problem?

Q3: Secret sharing (7 points)

Threshold secret-sharing is a primitive commonly studied in cryptography that allows a secret to be distributed among n parties such that, for some threshold $t \leq n$ decided in advance, any $t-1$ parties or fewer have effectively no information about the secret, but any set of t or more parties can pool together their information and recover the secret. The following definition is based on that of Beimel [Bei11].

Definition 1 ((t,n)-threshold secret-sharing scheme)**.** Let $\{p_1, \dots, p_n\}$ be a set of parties, let $1 \leq t \leq n$, and let K be a finite set of secrets with $|K| \geq 2$.

A distribution scheme $\Sigma = \langle \Pi, \mu \rangle$ consists of a probability distribution μ on a finite set R of random strings and a mapping

$$\Pi : K \times R \longrightarrow K_1 \times \dots \times K_n,$$

where K_j is the domain of shares of party p_j. To distribute $k \in K$, the dealer samples r according to μ and gives party p_j the jth component of $\Pi(k, r)$. For a set of parties A, let $\Pi(k, r)_A$ denote the restriction of $\Pi(k, r)$ to its A-entries.

The scheme is a (t, n)-threshold secret-sharing scheme if the following requirements hold:

1. **Correctness.** For every set $B = \{p_{i_1}, \ldots, p_{i_{|B|}}\}$ with $|B| \geq t$, there exists a reconstruction function
$$\mathrm{Recon}_B : K_{i_1} \times \cdots \times K_{i_{|B|}} \longrightarrow K$$
such that, for every $k \in K$,
$$\Pr[\mathrm{Recon}_B(\Pi(k, r)_B) = k] = 1.$$

2. **Perfect privacy.** For every set T of parties with $|T| < t$, every two secrets $a, b \in K$, and every vector of shares $\langle s_j \rangle_{p_j \in T}$ with $s_j \in K_j$,
$$\Pr\left[\Pi(a, r)_T = \langle s_j \rangle_{p_j \in T}\right] = \Pr\left[\Pi(b, r)_T = \langle s_j \rangle_{p_j \in T}\right].$$

All probabilities are over the choice of r according to μ.

(a) (3 points) Propose a threshold secret-sharing scheme based on the Reed-Solomon code, and prove that it satisfies correctness and perfect privacy. (Hint: Work over a finite field $\mathbb{F}_q$ with $q > n$, and take the secret space to be $K = \mathbb{F}_q$. Try setting the secret to be the constant term in a polynomial.)

Hint: we recommend looking up properties of Vandermonde matrices on Wikipedia.

(b) (3 points) Threshold secret-sharing can also be done when the proposed secret that needs to be reconstructed is a quantum state, provided the shares are also quantum. Consider the following definition of quantum threshold secret-sharing:

Definition 2 (Perfect (t, n)-threshold quantum secret sharing)**.** Let $P = \{p_1, \ldots, p_n\}$ be a set of parties and $1 \leq t \leq n$. Let $\mathcal{H}$, $\dim(\mathcal{H}) \geq 2$ be a finite-dimensional Hilbert space containing the secret, and let $\mathcal{H}_j$ be the finite-dimensional Hilbert space containing party p_j's share.

A secret-sharing scheme is a quantum channel
$$\Pi : \mathcal{D}(\mathcal{H}) \longrightarrow \mathcal{D}(\mathcal{H}_1 \otimes \cdots \otimes \mathcal{H}_n).$$

For $A \subseteq P$, write
$$\mathcal{H}_A = \bigotimes_{p_j \in A} \mathcal{H}_j, \qquad \Pi(\rho)_A = \mathrm{Tr}_{P \setminus A}\big(\Pi(\rho)\big),$$
where the partial trace discards the shares outside A.

The scheme is a perfect (t, n)-threshold quantum secret-sharing scheme if:

(a) **Correctness.** For every $B \subseteq P$ with $|B| \geq t$, there exists a quantum channel
$$\mathrm{Recon}_B : \mathcal{D}(\mathcal{H}_B) \longrightarrow \mathcal{D}(\mathcal{H})$$
such that, for every $\rho \in \mathcal{D}(\mathcal{H})$,
$$\mathrm{Recon}_B\big(\Pi(\rho)_B\big) = \rho.$$

(b) **Perfect privacy.** For every $T \subseteq P$ with $|T| < t$ and every $\rho, \sigma \in \mathcal{D}(\mathcal{H})$,

$$\Pi(\rho)_T = \Pi(\sigma)_T.$$

Let $t \leq n < 2t$. Suppose a $[[2t-1, 1, t]]$ quantum code $\mathcal{C}$ is given. Show how to construct, from *any* such $\mathcal{C}$, a perfect (t, n)-threshold quantum secret-sharing scheme for $\mathcal{H} = \mathbb{C}^2$, and prove correctness and perfect privacy.

Hint 1: first construct a $(t, 2t-1)$ scheme, then discard $2t-1-n$ shares. Hint 2: Use local indistinguishability to prove that your scheme works.

(Relatedly, observe that classically, a high-distance code alone does not suffice to construct a secret sharing scheme, as can be seen by considering the repetition code.)

(c) (1 point) Intuitively, why do we need the additional restriction $n < 2t$ in the quantum case compared with the classical case?

References

[Bei11] Amos Beimel. Secret-sharing schemes: A survey. In *Coding and Cryptology (IWCC 2011)*, volume 6639 of *Lecture Notes in Computer Science*, pages 11–46. Springer, 2011.