Consistency of Local Density Matrices *Instructor: Anand Natarajan, Anurag Anshu*

Lecture 23: Consistency of Local Density Matrices

November 25, 2025 *Scribe: Andrew Huang*

1 Consistency of Local Density Matrices

Definition 1.1 (k-Consistency of local density matrices (k-CLDM) [Liu06])**.** Let $n \in \mathbb{N}$. The input to the k-consistency of local density matrices problem consists of $((C_1, \rho_1), \ldots, (C_m, \rho_m))$ where $C_i \subseteq [n]$ and $|C_i| \leq k$; and ρ_i is a density matrix on $|C_i|$ qubits and each matrix entry of ρ_i has precision $\mathsf{poly}(n)$. Given two parameters α and β, assuming that one of the following conditions is true, we have to decide which of them holds.

- **Yes.** There exists some n-qubit state τ such that for every $i \in [m]$, $\mathrm{TV}(\mathsf{Tr}_{\overline{C_i}}(\tau), \rho_i) \leq \alpha$.
- **No.** For every n-qubit state τ, there exists some $i \in [m]$ such that $\mathrm{TV}(\mathsf{Tr}_{\overline{C_i}}(\tau), \rho_i) \geq \beta$.

Classically, marginal problems are NP-hard: consider the 3-coloring problem, which consists of a graph $G = (V, E)$ which is either 3-colorable or not. We can map G to a marginal problem as such:

- For each edge $(i, j) \in E$, set $\rho_{i,j}(c_1, c_2) = \begin{cases} 1/6 & \text{if } c_1 \neq c_2, \\ 0 & \text{otherwise.} \end{cases}$

We observe that if G is 3-colorable, then there is a simple probability distribution which matches the marginals $\{\rho_{i,j}\}_{(i,j)\in E}$: we can take any valid 3-coloring over G and apply a random permutation over the 3 colors. We observe that each edge in G will be assigned a random pair of distinct colors. On the other hand, if G is not 3-colorable, then any distribution of colorings over G must contain an edge which is assigned the same color with nonzero probability. Thus G is 3-colorable if and only if the marginal problem has a solution, which shows that the classical version of CLDM is NP-hard.

Lemma 1.2 ([Liu06, BG22])**.** *k-$\mathsf{CLDM} \in \mathsf{QMA}$ for $k = O(\log n)$ and $\frac{\beta}{4^k} - \alpha \geq \frac{1}{\mathsf{poly}(n)}$.*

Proof. The proof is relatively simple: let $((C_1, \rho_1), \ldots, (C_m, \rho_m))$ be an instance for CLDM. The verifier V will more or less conduct state tomography. V expects $\Omega(n/\varepsilon^2)$ copies of a state, where $\varepsilon = \frac{\beta}{4^k} - \alpha$. V samples a random $i \in [m]$ and Pauli $P \in \mathcal{P}_{C_i}$ before measuring $P_{C_i} \otimes I_{\overline{C_i}}$ on each of the copies of the state and accepting if the average value matches $\mathrm{Tr}(P\rho_i)$ within $\alpha + \varepsilon/2$.

In the case where the proof is actually a product state, the completeness and soundness follow nearly immediately. To ensure the proof looks like a product state, we can appeal to a quantum de Finetti theorem: the verifier picks a random subset of registers and discards the registers which were not selected before measuring P on the rest of the registers. □

2 QMA-Hardness of CLDM

We saw that k-CLDM is in QMA for small k, and in [Liu06], it was shown that CLDM is QMA-hard under Turing reductions. This means that $\mathsf{QMA} \subseteq \mathsf{P}^{\mathsf{CLDM}}$. It was conjectured that CLDM was also QMA-hard under *Karp* reductions, but this problem remained open until Broadbent and Grilo resolved this question in the positive [BG22]. We now give a rough outline of their proof.

The primary issue with showing QMA-hardness of CLDM is that the normal Feynman-Kitaev construction produces a history state which depends heavily on the witness used. The natural question of course, is if there is a way for the verifier to compute a local view of a witness *without* knowing the witness itself, and yet still be able to check validity.

Here, we note that the first property (of local indistinguishability) is precisely what is satisfied by the encoded states of a quantum error-correcting code! This observation was first made by [GSY19]. In order to use an error-correcting code, of course, we need to be able to perform logical operations on encoded states. We will use an QECC that allows us to compute all desired logical operations using only transversal gates and encoded magic states. In particular, we will use the 3-fold Steane code which is a $[[7^3, 1, 3^3]]$-QECC that admits transversal Clifford gates and only requires magic states to compute T gates. We first show that the Steane code is locally simulatable, which means we can compute the reduced density matrices of codewords as well as intermediate steps of transversal Clifford gates and T-gadgets.

We will use the following lemma which states a fairly standard property of QECCs that roughly captures the idea that the local reduced density matrix of a codeword is independent of the state being encoded.

Lemma 2.1. *Let $\mathcal{C}$ be an $[[N, 1, D]]$-QECC, A be a quantum register (possibly entangled with some environment register E), and A' be the register that is the output of encoding a quantum system in register A under $\mathcal{C}$. Let S be a subset of size at most $(D-1)/2$ of the qubits in A' and $\overline{S} := A' \setminus S$. Then there exists some state τ_S such that for all states $|psi\rangle_{\mathsf{AE}}$,*

$$\mathsf{Tr}_{\overline{S}}((\mathsf{Enc} \otimes I_{\mathsf{E}})(|\psi\rangle\langle\psi|_{\mathsf{AE}})) = \tau_S \otimes \mathsf{Tr}_{\mathsf{A}}(|\psi\rangle\langle\psi|_{\mathsf{AE}}).$$

By taking a trivial system E, we get the following corollary.

Corollary 2.2. *Let $\mathcal{C}$ be an $[[N, 1, D]]$-QECC. For any S such that $|S| \leq (D-1)/2$, there exists some density matrix τ_S such that for all states $|\psi\rangle$, $\tau_S = \mathsf{Tr}_{\overline{S}}(\mathsf{Enc}(\psi))$.*

Observe that Corollary 2.2 implies that the simulator can easily compute the reduced density matrix of any codeword by simply computing $\mathsf{Tr}_{\overline{S}}(\mathsf{Enc}(|0\rangle\langle 0|))$ in time $\mathsf{poly}(2^N)$.

Transversal gates are not too difficult to deal with either; consider applying a logical gate G by applying gates $U_1, \ldots, U_N$, where U_i applies the physical gate G on i'th physical qubit per logical qubit being acted on. Our goal is to be able to compute a state

$$\rho(G, t, S) = \mathsf{Tr}_{\overline{S}}((U_t \ldots U_1)\mathsf{Enc}(\sigma)(U_1^\dagger \ldots U_t^\dagger))$$

which is independent of the state σ being encoded (for S where $|S| \leq (D-1)/4$ and any $0 \leq t \leq N$). Note that we can decompose $\rho(G, t, S)$ into its Pauli basis and compute the coefficient associated with each Pauli on S. That is, we can write $\rho(G, t, S) = \frac{1}{2^{|S|}} \sum_{P \in \mathcal{P}_{|S|}} \mathrm{Tr}\left(\rho(G, t, S)P\right) \cdot P$. Since each Pauli P acts on at most $|S|$ qubits and U_i are Clifford, we can compute the Pauli P' such that $P'U_1^\dagger \ldots U_t^\dagger = U_1^\dagger \ldots U_t^\dagger P$, and note that P' acts on at most $2|S|$ qubits (since U_i each act on at most

2 qubits which are distinct from each other). Our goal is to compute for each $P \in \mathcal{P}_{|S|}$ the state $\mathrm{Tr}\left(\rho(G,t,S)P\right) = \mathrm{Tr}\left(\mathsf{Enc}(\sigma)P'\right)$. Note that since $2|S| <= (D-1)/2$, we can compute the reduced state of $\mathsf{Enc}(\sigma)$ on the qubits for which P' acts trivially before computing $\mathrm{Tr}\left(\mathsf{Enc}(\sigma)P'\right)$. Thus we can compute $\mathrm{Tr}\left(\rho(G,t,S)P\right)$ iteratively before computing $\rho(G,t,S)$ in time $\mathsf{poly}(2^N)\cdot\mathsf{poly}(2^{|S|}) = \mathsf{poly}(2^N)$. One can easily extend this to controlled applications of transversal Clifford gates when the control bit is known.

Dealing with the T gate requires more care, and so we will only briefly describe how to simulate such computations. First, note that we need to implement the unitary version of the T-gadget which first applies two logical CNOTs before decoding one qubit, applying controlled X and P gates (transversally), and then re-encoding. As argued earlier, we can already simulate the application of the two logical CNOTs, so it remains to see why we can deal with the operations where some qubits are physical. Roughly speaking, if we are in the decoding or re-encoding stage, the qubits not involved in the decoding/re-encoding are codewords or can be simulated already by Lemma 2.1, and it will suffice to compute the state of $1/2\sum_{b\in\{0,1\}}\mathsf{Enc}(|b\rangle\langle b|)$ while it is being encoded or decoded. On the other hand, when we are in the stage of applying of X or P, the intermediate state looks somewhat like the application of a classically-controlled transversal Clifford gate with a known control qubit which is chosen uniformly at random, which is easy to compute as argued earlier.

With this in hand, we are nearly done, and we can (more or less) show the following lemma.

Lemma 2.3. *For any problem $A = (A_{yes}, A_{no}) \in \mathsf{QMA}$, there is a uniform family of verification algorithms $V_x = U_T \ldots U_1$ for A that acts on a witness of size $p(|x|)$ and $q(|x|)$ auxiliary qubits such that there exists a $\mathsf{poly}(|x|)$-time deterministic algorithm Sim_V that on input $x \in A$ and subset $S \subseteq [T+p+q]$ such that $|S| \leq 5$, $\mathsf{Sim}_V(x,S)$ outputs the classical description of an $|S|$-qubit density matrix $\rho(x,S)$ with the following properties:*

1. *If x is a YES instance, then there exists a witness ψ that makes V_x accept with probability at least $1-\mathsf{negl}(n)$ such that $\|\rho(x,S) - \mathsf{Tr}_{\overline{S}}(\Phi_{V_x,\psi})\|_1 \leq \mathsf{negl}(n)$, where*

$$\Phi_{V_x,\psi} = \frac{1}{T+1}\sum_{t,t'\in[T+1]} |1^t 0^{T-t}\rangle\langle 1^{t'}0^{T-t'}| \otimes U_t\cdots U_1(\psi\otimes|0\rangle\langle 0|^{\otimes q})U_1^\dagger\cdots U_{t'}^\dagger$$

   is the history state of V_x on the witness ψ.

2. *If x is a NO instance, then all states make V_x accept with probability at most $\mathsf{negl}(n)$.*

3. *Let H_i be one term from the circuit-to-local Hamiltonian construction from V_x and S_i be the set of qubits on which H_i acts non-trivially. Then for every $x \in A$, $\mathrm{Tr}\left(H_i\rho(x,S_i)\right) = 0$.*

Proof. The full proof of this lemma is quite technical, so we will define the appropriate family of circuits V_x and describe how to locally simulate *snapshots* of the computation of V_x, i.e. the reduced state of $U_t\cdots U_1(\psi\otimes|0\rangle\langle 0|^{\otimes q})U_1^\dagger\cdots U_t^\dagger$ for any $t \in [T+1]$. Extending this result to the full history state simulation requires some work, but is not very interesting, so we omit the proof for brevity.

We know that since A is in QMA, there is a uniform family of QMA verifiers V_x' with completeness $1-\mathsf{negl}(n)$ and soundness $\mathsf{negl}(n)$. The basic idea is to run a fault-tolerant/encoded version of V_x' so we can leverage the simulatability of QECCs. We consider the QMA verifier V_x which expects an encoded witness and applies the logical version of its gates using transversal gates and encoded

ancilla and magic states. V_x will first check that its received input is in fact the encoding of some state before generating/encoding the required resource states; it will then perform the logical version of V'_x before decoding the output logical qubit and measuring the physical decoded qubit. It is straightforward to see that V_x preserves completeness and soundness, so it remains to show simulatability for YES instances.

As argued earlier, the stage of computation corresponding to the logical version of V'_x is simulatable, as is the generation and encoding of ancilla/magic states, since these are fixed and do not depend on the witness. We only need to consider the simulation of the decoding and encoding check stages of V_x. At first, one might be concerned that the decoding step involves operations on the physical level, and thus may be dangerous. However, this is not a concern: note that for a YES instance, the output qubit of V'_x should be negligibly close to $|1\rangle$, and thus the state right before the decoding stage must be negligibly close to $\mathsf{Enc}(|1\rangle)$. Thus, we can simulate what these qubits should look like, and the non-output qubits remain encoded so they can also be simulated.

Dealing with the encoding check is much trickier, and in fact is not quite possible with the current construction of V'_x. Here we employ a simple trick: we ask the prover to apply a quantum one-time pad to his witness along with the one-time pad values before encoding this whole state. Our V'_x verifier will decrypt the witness according to the one-time pad before behaving as before. Clearly this modification preserves completeness and soundness, but we now note that when we perform the encoding check, the underlying state (should) consist of the encoding of a fully one-time-padded state along with uniformly random pad values. Each qubit is either part of the padded witness or the one-time pads, but note that both of these states on their own are the fully mixed state, whose encoding can be trivially computed/simulated! With this trick we are able to simulate the encoding check of V_x, completing the analysis of all stages of the circuit as desired. □

We can now show that CLDM is QMA-hard. Consider any promise problem $A = (A_{yes}, A_{no}) \in \mathsf{QMA}$. By Lemma 2.3, there exists a family of verifier circuits $\{V_x\}_x$ with the related properties. We now use the simulator Sim_V to compute the CLDM instance

$$\{(S, \rho(x, S) : S \subseteq [T + p + q], |S| \leq 5\}.$$

If x is a YES instance, then there exists some witness ψ such that $\|\rho(x, S) - \mathsf{Tr}_{\overline{S}}(\Phi_{V_x,\psi})\| \leq \mathsf{negl}(n)$, and so the associated history state Φ is consistent by construction with all the density matrices in the CLDM instance.

If x is a NO instance, then note that every state makes V_x accept with probability at most $\mathsf{negl}(n)$. Now suppose there is some state τ such that for all subsets $S \subseteq [T + p + q]$ where $|S| \leq 5$, $\|\rho(x, S) - \mathsf{Tr}_{\overline{S}}(\tau)\| \leq \frac{1}{T^5}$. Taking $H := H_{V_x} = \sum_i H_i$ to be the Hamiltonian that results from the circuit-to-Hamiltonian reduction applied to V_x where H_i is a 5-local term that acts nontrivially on qubits S_i, we observe that

$$\mathrm{Tr}\left(H\tau\right) = \sum_i \mathrm{Tr}\left(H_i \mathsf{Tr}_{\overline{S_i}}(\tau)\right) \leq \sum_i \left(\mathrm{Tr}\left(H_i \rho(x, S_i)\right) + \frac{1}{T^5}\right) = \sum_i \frac{1}{T^5} = O\left(\frac{1}{T^4}\right),$$

contradicting the fact that H has ground energy $\Omega\left(\frac{1}{T^3}\right)$. Thus for all states τ, there exists some subset $S \subseteq [T + p + q]$ where $|S| \leq 5$ and $\|\rho(x, S) - \mathsf{Tr}_{\overline{S}}(\tau)\| > \frac{1}{T^5}$, concluding the claim.

References

[BG22] Anne Broadbent and Alex Bredariol Grilo. Qma-hardness of consistency of local density matrices with applications to quantum zero-knowledge. *SIAM Journal on Computing*, 51(4):1400–1450, 2022.

[GSY19] Alex B Grilo, William Slofstra, and Henry Yuen. Perfect zero knowledge for quantum multiprover interactive proofs. In *2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 611–635. IEEE, 2019.

[Liu06] Yi-Kai Liu. Consistency of local density matrices is qma-complete. In *International Workshop on Approximation Algorithms for Combinatorial Optimization*, pages 438–449. Springer, 2006.