Hamiltonian Property Testing

Andreas Bluhm— Univ. Grenoble Alpes, CNRS, Grenoble INP, LIG

joint work with Matthias C. Caro and Aadil Oufkir

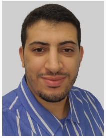

ICMP Strasbourg, July 5, 2024

Introduction

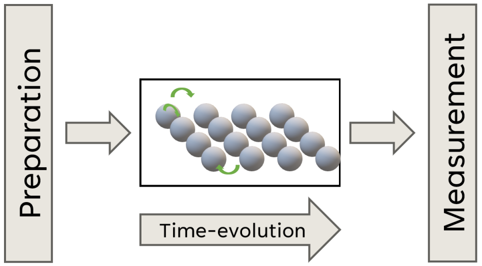

- Quantum systems are governed by their Hamiltonians
- Often, we want to learn the Hamiltonian from access to its time evolution
- What happens if we only want to test a property (e.g. whether it is local)? Is this an easier problem?

Setup

- Consider a system of n qubits, dimension 2^n
- Pauli expansion on n-qubits: $H = \sum_{P \in \mathbb{P}_n} \alpha_P P$, where $\mathbb{P}_n = \{I, X, Y, Z\}^{\otimes n}$

Learning:

- For learning, you want your algorithm to output an estimator $\hat{H}$ such that $\left|\left|\left|\hat{H} - H\right|\right|\right| \leq \varepsilon$ with probability at least $2/3$
- Often, learning algorithms assume that the Hamiltonian they want to learn is local

Locality:

- We call the Hamiltonian H k-local (k-body) if $\alpha_P = 0$ holds for all $P \in \mathbb{P}_n$ with $|P| > k$. Here, $|P|$ denotes the number of non-identity tensor factors in P
- Example: $X \otimes I \otimes I \otimes X$ is 2-local, $X \otimes Y \otimes Z \otimes X$ is 4-local

Problem statement

Definition (Hamiltonian locality testing)

Given a locality parameter $1 \leq k \leq n$, a norm $|||\cdot|||$, and an accuracy parameter $\varepsilon \in (0,1)$, the Hamiltonian k-locality testing problem, denoted as $\mathcal{T}^{\mathrm{loc}}_{|||\cdot|||}(\varepsilon)$, is the following task: Given access to the time evolution according to an unknown Hamiltonian H, decide, with success probability $\geq 2/3$, whether

(i) H is k-local, or

(ii) H is ε-far from being k-local, i.e., $\left|\left|\left| H - \tilde{H} \right|\right|\right| \geq \varepsilon$ for all k-local Hamiltonians $\tilde{H}$.

If H satisfies neither (i) nor (ii), then any output of the tester is considered valid.

This is a promise problem. Instead of $2/3$ we could take any constant probability larger than $1/2$.

Different types of algorithms

Incoherent adaptive strategy

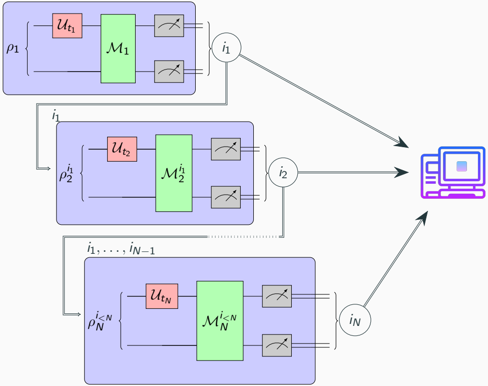

$U_t(\cdot) = e^{-itH} \cdot e^{itH}$, $\mathcal{M}$ measurement

Coherent strategy

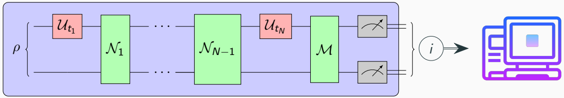

Hardness of Hamiltonian locality testing w.r.t. the operator norm

Theorem

For $k \leq \tilde{\mathcal{O}}(n)$, any ancilla-free, incoherent, adaptive quantum algorithm that solves the k-locality testing problem $\mathcal{T}^{\mathrm{loc}}_{\|\cdot\|_\infty}(\varepsilon)$, even only under the additional promise that the unknown Hamiltonian H satisfies $\mathrm{tr}[H] = 0$ and $\|H\|_\infty \leq 1$, has to make at least $N \geq \tilde{\Omega}(2^n)$ queries to the unknown Hamiltonian and has to use an expected total evolution time of at least $\mathbb{E}[T] \geq \tilde{\Omega}\left(\frac{2^n}{\varepsilon}\right)$. Even any coherent quantum algorithm achieving the same has to make at least $N \geq \Omega\left(2^{n/2}\right)$ many queries and has to use a total evolution time of at least $T \geq \Omega\left(\frac{2^{n/2}}{\varepsilon}\right)$.

This result actually rules out efficient property testing in any Schatten p-norm

Efficient Hamiltonian locality testing w.r.t. normalized Frobenius norm

- Locality testing with respect to p-norms is hard
- What if we use instead the norm $|||\cdot||| = 2^{-n/2}\|\cdot\|_2$, where $\|A\|_2 = (\mathrm{tr}[A^*A])^{1/2}$ is the Frobenius norm?
- This represents the average case setting, whereas $\|\cdot\|_\infty$ corresponds the worse case

Theorem

Let $k \leq \tilde{\mathcal{O}}(n)$. When promised that the unknown Hamiltonian H satisfies $\mathrm{tr}[H] = 0$ and $\|H\|_\infty \leq 1$, there is an ancilla-free, incoherent, non-adaptive quantum algorithm that solves the Hamiltonian k-locality testing problem $\mathcal{T}^{\mathrm{loc}}_{\frac{1}{\sqrt{2^n}}\|\cdot\|_2}(\varepsilon)$ using $\mathcal{O}\left(\varepsilon^{-4}\right)$ many queries to the unknown Hamiltonian, a total evolution time of $\mathcal{O}\left(\varepsilon^{-3}\right)$, and a classical post-processing time of $\mathcal{O}\left(\frac{n^{k+3}}{\varepsilon^4}\right)$. Moreover, the testing algorithm uses only stabilizer states as inputs and stabilizer basis measurements at the output.

Description of the algorithm

- We construct $d+1$ stabilizer bases $\mathcal{B}_i = \{|\phi_{i,j}\rangle\}_{j\in\{1,\ldots,d\}}$ from maximal Abelian subgroups of the Pauli group. A QC can prepare and measure them efficiently.

Efficient algorithm for locality testing (polynomial runtime)

1. Choose $(i,j) \in [d+1] \times [d]$ uniformly at random and prepare the state $|\phi_{i,j}\rangle$
2. Let it evolve under the unknown Hamiltonian H for time $t = \mathcal{O}(\varepsilon)$
3. Perform a measurement in the basis $\mathcal{B}_i$ and observe outcome ℓ
4. Repeat this procedure $N = \mathcal{O}(\varepsilon^{-4})$ times
5. If at least one of the N rounds produces an output ℓ such that

$$|\langle\phi_{i,\ell}| P |\phi_{i,j}\rangle| = 0 \quad \text{for all } P \in \mathbb{P}_n \text{ with } |P| \leq k\,,$$

then we conclude that H is ε-far from k-local. Otherwise, we claim that H is k-local.

The checks in the last step can be efficiently performed on a classical computer

Proof idea

- Consider a simpler version with commuting Hamiltonians consisting of terms from $\{I, X\}^{\otimes n}$, $|0\rangle\langle 0|$ as input state, and computational basis measurement
- Note that $|j\rangle = X^j |0\rangle$ for any $j \in \{0,1\}^n$ and that $U_t = \mathrm{e}^{\mathrm{i}tH} \approx I + \mathrm{i}tH$ for short t
- For any n-bit string j with weight $|j| > 0$, it holds that $|\langle j| U_t |0\rangle|^2 \approx t^2 |\alpha_{X^j}|^2$
- If H is indeed k-local, then $\alpha_{X^j} = 0$ holds whenever $|j| > k$, and we find that

$$\sum_{j:|j|>k} |\langle j| U_t |0\rangle|^2 \approx 0\,,$$

  i.e., we make approximately no error
- Conversely, if H is ε-far from any k-local Hamiltonian, then $\sum_{j:|j|>k} |\alpha_{X^j}|^2 \geq \varepsilon^2$
- Thus, $\sum_{j:|j|>k} |\langle j| U_t |0\rangle|^2 \gtrapprox t^2 \varepsilon^2$
- Repeating $\mathcal{O}(t^{-2}\varepsilon^{-2})$ times makes success probability constant
- To make the proof precise, we need to deal with higher order terms and non-commutative Hamiltonians

Hardness of Hamiltonian learning w.r.t. normalized Frobenius norm

Theorem

Any (even coherent) quantum algorithm with a constant number of auxiliary qubits that, when given time evolution access to an arbitrary n-qubit Hamiltonian H, promised to satisfy $\mathrm{tr}[H] = 0$ and $\|H\|_\infty \leq 1$, with success probability $\geq 2/3$, outputs (the classical description of) a Hamiltonian $\hat{H}$ such that $\frac{1}{\sqrt{2^n}}\left\|H - \hat{H}\right\|_2 \leq \varepsilon$ has to make at least $\tilde{\Omega}\left(2^{2n}\right)$ many queries to H.

Any non-adaptive incoherent quantum algorithm achieving the same without auxiliary qubits has to use a total evolution time of at least $\tilde{\Omega}\left(\frac{2^{2n}}{\varepsilon}\right)$.

Proof idea

- We want to prove that learning w.r.t. the norm. Frobenius norm is hard ($\tilde{\Omega}(\frac{2^{2n}}{\varepsilon})$ total evolution time).
- Strategy: Identify a distinguishing problem (probabilistic argument) that any successful general Hamiltonian learner can solve
- Lower bounds for that distinguishing task through information-theoretic arguments
- Construct $M = \exp(\Omega(4^n))$ unitaries U_x such that the Hamiltonians $H_x = \varepsilon U_x O U_x^\dagger$ are pairwise ε-far apart w.r.t. $\frac{1}{\sqrt{2^n}}\|\cdot\|_2$, $O = \mathrm{diag}(+1, \ldots, +1, -1, \ldots, -1)$
- Fano's inequality: mutual information lower bound $\mathcal{I}(X : Y) \geq \Omega(\log M) \geq \Omega(4^n)$, where $X \sim \mathrm{Uniform}([M])$ and Y outcomes observed by the learner
- Remaining work: Upper bounds on $\mathcal{I}(X : Y)$; uses Weingarten calculus

Additional results and open questions

Some things which did not make it into the talk:

- We can actually test any property, i.e., whether H has only terms in some subset $\mathcal{S} \subset \mathbb{P}_n$ or is at least ε-far from it. If $\mathcal{S}$ is too big, we need ancillas
- We can actually test M properties at the same time with only $\log M$ overhead
- We can also do tolerant property testing (checking whether H is ε_1-close to having $\mathcal{S}$ of at least ε_2-far from any such Hamiltonian, for $\varepsilon_1 < \varepsilon_2$)

Some open questions:

- Are our bounds optimal? Is the scaling $N = \mathcal{O}(\epsilon^{-4})$ necessary?
- What about other distance measures, such as Wasserstein distances?
- What about other access models, e.g., learning from Gibbs states?

Summary

We have considered the task of locality testing, i.e., testing whether a Hamiltonian is k-local or ε-far from any such Hamiltonian

We have found a setting in which learning is hard, but in which we can give an efficient algorithm for locality testing, thereby separating the two tasks

For more details, see

arXiv:2403.02968