



Frustration Free Stoquastic Local Hamiltonian with Sub-Constant Gap is in **NP**

Eshan Chattopadhyay*

Oren Renard†

Nicholas Spooner‡

September 14, 2026

Abstract

We continue the study of the Stoquastic Local Hamiltonian problem, a physically motivated restriction of the **QMA**-complete Local Hamiltonian problem [KSV02]. For the β -gapped, frustration-free case, Bravyi, Bessen, and Terhal [BBT06] showed that the problem is **MA**-complete when $\beta = \frac{1}{\text{poly}(n)}$. Aharonov and Grilo [AG19] derandomized this algorithm and proved membership in **NP** for constant gap $\beta = \Omega(1)$.

We present an improved algorithm and analysis, establishing membership in **NP** even when $\beta = \Omega\left(\frac{1}{\log \log n}\right)$. We complement our result with an explicit example demonstrating why the analysis does not extend directly to $\beta = o(1/\log \log n)$.

Contents

1	Introduction	2
1.1	Our results	3
1.2	Technical Overview	3
1.3	Discussions	7
2	Preliminaries	7
2.1	Definitions	7
2.2	The Associated Graph and Its Analysis	8
3	Our Algorithm	9
3.1	Setup	10
3.2	Uniformly Random Sequence	12
3.3	The Chopped Sequence	12
4	An Instance For Which Exponentially Many Overlapping Sequences Are δ-valid	15
4.1	Construction	15
4.2	Validity of the Instance	15
4.3	Doubly Exponentially Many Sequentially Frustrated Sequences	17

*Cornell University. Email: eshan@cs.cornell.edu.

†Cornell University. Email: oren@cs.cornell.edu.

‡Cornell University. Email: nspooner@cornell.edu.

1 Introduction

The k -Local Hamiltonian problem is the canonical complete problem for **QMA** [KSV02], the quantum analogue of **NP**, and is itself the quantum analogue of 3SAT. An instance is a tuple (H, α, β) where $H = \frac{1}{m} \sum_i H_i$ is a normalized sum of Hermitian operators, each H_i acting nontrivially on at most k qubits. The goal is to decide whether the ground energy of the system, namely its minimum eigenvalue $\lambda_{\min}(H)$, is at most α or at least β , promised that one of these holds. The quantity $\beta - \alpha$ is called the *promise gap*, and the problem is **QMA**-complete for any inverse-polynomial gap $\beta - \alpha \geq 1/\text{poly}(n)$ [KSV02]. The locality parameter k plays an important role, and in this work, we focus on arbitrary constant k , though our results extend naturally to any locality.

Due to the intractability of the Local Hamiltonian problem in full generality, a vast line of research has studied restricted variants, seeking better upper bounds (e.g., geometrically local instances [OT08; AGIK09], commuting Hamiltonians [BV05; Has13], translationally invariant systems [BCO17; CM16], and much more; see [GHLS14] for a survey).

In this work, we consider two standard restrictions. The first is *frustration-free* instances, where the ground state (the minimal eigenvector) simultaneously satisfies every local term H_i (equivalently, $\alpha = 0$). This is a well-studied condition known to preserve hardness in many settings [BBT06; CM16; AGIK09].

The second, and more central to this work, is the restriction to *Stoquastic* Hamiltonians [BBT06; BT10], a natural “sign-free” variant of the Local Hamiltonian problem. Since each Hermitian matrix H_i can have complex entries, it permits destructive interference between amplitudes, a key source of quantum hardness and the main obstacle to classical simulation. Stoquastic Hamiltonians eliminate this by restricting each local term H_i to be real-valued, with non-positive off-diagonal entries. In this case, one can show there exists a *ground state* with real, non-negative amplitudes [BBT06], rendering it tantamount to a *classical* probability distribution amenable to classical Markov chain methods [BBT06; AG19; BH17; PM17; BDOT08; KT19; AGL25; REG26].

[BBT06] define the complexity class **StoqMA** (α, β) , a variant of **QMA** in which the verifier is fully classical except for a final Hadamard-basis measurement, and show that the Stoquastic Local Hamiltonian problem is complete for this class. It follows from the definition of the class that

$$\mathbf{MA} \subseteq \mathbf{StoqMA}(\alpha, \beta) \subseteq \mathbf{QMA}.$$

Intuitively, it seems that **StoqMA** lies closer to **MA**, since the **StoqMA** verifier is “barely quantum”. Indeed it is conjectured that **StoqMA** = **MA** [BBT06; AGL25], though this remains open.¹

Bravyi, Bessen, and Terhal [BBT06] showed that the *frustration-free projective* Stoquastic Local Hamiltonian problem (in which every local term is a projector and $\alpha = 0$) is **MA**-complete at inverse-polynomial soundness $\beta = 1/\text{poly}(n)$. The result was later extended to the *uniform projective* case [AG19].² Most notably, these yielded the first **MA**-complete problem(s), offering a surprising quantum lens through which to study **MA**. We focus on this setting and denote the problem over q -dimensional qudits by $\text{PROJUSLH}_q(k, d, \alpha, \beta)$, where α and β are the YES- and NO-case energy thresholds.

Theorem 1.1 ([BBT06; AG19]). *There exist constant integers q, k, d such that $\text{PROJUSLH}_q(k, d, 0, \frac{1}{\text{poly}(n)})$ is **MA**-complete.*

Remarkably, Aharonov and Grilo [AG19] showed that at constant soundness $\beta > 0$, the problem can be decided in non-deterministic polynomial time, namely $\text{PROJUSLH}_q(k, d, 0, \beta = \Omega(1)) \in \mathbf{NP}$. Remarkably, this implies that a stoquastic PCP theorem, i.e., gap amplification for $\text{PROJUSLH}(k, 0, \cdot)$, would collapse **MA** = **NP**, a major open problem in derandomization [NW94; IW97; IKW02; KM02].

¹It is known that **StoqMA** $\subseteq$ **AM** [BDOT08].

²A projector is called uniform if it can be decomposed into sum of orthogonal uniform subset state projectors. See Section 2 for a formal definition.

1.1 Our results

We present an improved algorithm for PROJUSLH that pushes the tractability threshold to a sub-constant soundness while remaining within **NP**, yielding strict improvement over [AG19]. Formally, although our results extend naturally to arbitrary parameter choices, we focus on constant parameters $k, q, d = O(1)$.

Our main contribution is an alternative analysis framework that yields an improved **NP** algorithm for subconstant soundness $\beta = o(1)$. Our framework uses the probabilistic method (though the algorithm itself is deterministic), which yields a stronger, simpler analysis.

Theorem 1.2. *For any $k, d, q \in \mathbb{N}$ there exists η such that $\text{PROJUSLH}_q(k, d, 0, \eta \cdot \frac{1}{\log \log n}) \in \mathbf{NP}$.*

Combining the above with Theorem 1.1, we conclude that a gap amplification reduction for PROJUSLH_q from inverse-poly up to $\beta = \Omega(\frac{1}{\log \log n})$ would suffice to derandomize Merlin-Arthur games.

Corollary 1.3. *There exists a constant $d \in \mathbb{N}$ such that, if there exist $k', d', q' \in \mathbb{N}$ such that there is a reduction from $\text{PROJUSLH}_{q=2}(k = 6, d, 0, \frac{1}{\text{poly}(n)})$ to $\text{PROJUSLH}_{q'}(k', d', 0, \eta' \cdot \frac{1}{\log \log n})$ for large enough constant $\eta' = \eta'(k', d', q')$, then $\mathbf{MA} = \mathbf{NP}$.*

Finally, we give evidence that the soundness bound in Theorem 1.2 is tight for our current algorithmic approach. At a high level, the verifier searches for a short sequence of local terms that leads from a proposed witness string to a contradiction. Its running time depends not only on the length of such sequences, but also on the number of candidates that appear promising according to the progress measure used in our analysis. In Section 4, we construct explicit frustrated instances admitting exponentially many such candidates. Consequently, this progress measure alone cannot prune the search sufficiently to extend Theorem 1.2 to $\beta = o(1/\log \log n)$; such an extension would require a new approach.

1.2 Technical Overview

A critical component in the analysis of stoquastic Hamiltonians is the graph-based framework of [BBT06]. We use its simplified undirected form that was presented by [AG19].³

For this overview, an instance $\mathcal{P} = \{P_i\}_{i \in [m]}$ of m many k -local Projector Uniform Stoquastic terms over n qubits consists of projectors $P \in \mathcal{P}$ that are (a) entrywise real and non-negative, and (b) k -local. To simplify the exposition, we ignore the “uniform” restriction here, though it remains essential for the proof (see the full definition at Definition 2.3).

We embed $\mathcal{P}$ into an exponentially large graph $G = (\{0, 1\}^n, E)$, with vertices corresponding to computational basis elements $x \in \{0, 1\}^n$ and edges defined by

$$(x, y) \in E \iff \langle x | \frac{1}{m} \sum_{i \in [m]} P_i | y \rangle > 0.$$

Bravyi, Bessen, and Terhal [BBT06] introduced the notion of BAD strings $x \in \{0, 1\}^n$, that are defined as those satisfying $\langle x | P_i | x \rangle = 0$ for some term P_i . Subsequently, they showed that these strings sharply characterize the perfect-completeness gap problem, where YES instances are frustration-free while NO instances are ε -frustrated on average:

- (YES) if $|\theta\rangle$ is a ground state that is supported on the set $\text{Supp}(|\theta\rangle) = \{x \mid \langle \theta | x \rangle \neq 0\}$, then the connected component of $\text{Supp}(|\theta\rangle)$ in G consists entirely of GOOD strings;
- (NO) every connected component of G contains a BAD string, and a random walk from any vertex hits a BAD string with high probability.

³The undirected and directed variants correspond to the uniform and non-uniform variants of Stoquastic Local Hamiltonian, respectively. Both variants are **StoqMA**-complete and therefore fully characterize the problem.

This yields a natural **MA** verifier: it expects a witness $x \in \text{Supp}(|\theta\rangle)$, walks randomly on G , and rejects upon encountering a BAD string.

Aharonov and Grilo [AG19] subsequently derandomized this verifier. Their key insight is that, in NO instances which are ε -frustrated for constant $\varepsilon > 0$, one can deterministically locate a BAD string because *frustration implies expansion*. To this end, they relied on two main lemmas in the analysis of such NO instances:

First, they showed that every state $|\varphi\rangle$ is “ $\varepsilon/2$ -frustrated” by at least an $\varepsilon/2$ fraction of the terms. This follows from a simple averaging argument: there exist many terms $P_i \in \mathcal{P}$ that satisfy

$$\|P_i |\varphi\rangle\|_2^2 \leq 1 - \varepsilon/2. \quad (1)$$

Second, a short algebraic calculation shows that any such P_i *expands* the support of $|\varphi\rangle$, when $|\varphi\rangle$ is a “uniform subset state”.⁴ At a high level, because the terms are stoquastic projectors, i.e., entrywise non-negative, frustration implies growth of the support size of entry-wise non-negative states. Since Equation (1) gives step frustration at least $\varepsilon/2$, the frustration–expansion lemma yields

$$|\text{Supp}(P_i |\varphi\rangle)| \geq (1 + \varepsilon/4) |\text{Supp}(|\varphi\rangle)|. \quad (2)$$

Then [AG19] organizes the expanding projectors into $\ell = O(1/\varepsilon^2)$ layers, each of which contains a sequence of terms that act on non-overlapping qubits. Then, starting from the state $|\varphi_0\rangle = |x\rangle$ and applying the terms sequentially, as long as no BAD string is encountered, each term (or layer) expands the support; since the support has size at most 2^n , a BAD string must eventually be reached (which always exists for a frustrated instance).

Tracing backward from a term witnessing this BAD string, i.e. P , gives a lightcone (denoted C_P), whose size is bounded by $|C_P| \approx k^\ell$. Thus the BAD string lies within radius $|C_P|$ of the witness x in the graph G , and their verifier finds it by exploring the entire surrounding ball in the graph G . Their final argument is therefore to bound the number of vertices in this ball:

$$|\text{Ball}_G(x, |C_P|)| \leq n^{|C_P|} \approx n^{k^\ell}. \quad (3)$$

The issue is that plugging $\ell \approx 1/\varepsilon^2$ and enforcing polynomial runtime requires Equation (3) to be polynomial, forcing $\varepsilon = \Omega(1)$ to be constant.

Improvement to $\varepsilon \geq \frac{1}{\sqrt{\log \log n}}$. As a first step, we observe that enumerating the entire ball around x is wasteful. Instead of enumerating the entire ball, we can restrict the walk to terms that form a *lightcone*, of which there are far fewer:

$$|\text{CONES}_{P,\ell}| \leq 2^{kd^\ell}. \quad (4)$$

This bound uses two separate properties of the layered lightcones from [AG19]: terms within each layer act on disjoint sets of qubits, while the lightcone construction retains from each layer only terms that overlap a retained term in the next layer closer to the root. Since all terms are k -local, the number of candidates overlapping with any fixed term is bounded by kd , and iterating this count across ℓ layers yields the bound above.

Enumerating $|\text{CONES}_{P,\ell}|$ in polynomial time would enforce Equation (4) to be at most $\text{poly}(n)$, and plugging $\ell = O(1/\varepsilon^2)$ into Equation (4) yields $\varepsilon \geq \frac{1}{(\log \log n)^{1/2}}$ and therefore an immediate improvement over the earlier **NP** algorithm for soundness $\varepsilon = \Omega(1)$.

⁴We say $|\varphi\rangle$ is a subset state if $|\varphi\rangle = \frac{1}{\sqrt{|S|}} \sum_{x \in S} |x\rangle$ for some $S \subseteq \{0, 1\}^n$.

Improvement to $\varepsilon \geq \frac{1}{\log \log n}$. We next present how to improve the algorithm to support quadratically smaller ε . The reason for the square root loss comes from the number of layers in the argument above: the deterministic expansion argument of [AG19] requires $\ell \approx 1/\varepsilon^2$ layers. We avoid this loss by replacing the carefully chosen (but otherwise “worst-case”) layers with a sufficiently long random sequence of terms. (Note that this step only appears in the analysis; the **NP** algorithm remains deterministic.) We then show that—similarly to [AG19]—although the sequence itself may be long, only a short overlapping subsequence is relevant for reaching a BAD string.

Fix an arbitrary witness $x \in \{0, 1\}^n$. Then observe that by definition, since every state $|\varphi\rangle$ has energy at most $1 - \varepsilon$, a randomly chosen term P_i must have at least ε frustration with respect to $|\varphi\rangle$. That is, denoting $\gamma_i = 1 - \|P_i |\varphi\rangle\|_2^2$, we have that

$$\mathbb{E}_{i \in_R [m]} [\gamma_i] \geq \varepsilon.$$

This gives rise to a very natural approach: choose independently and uniformly at random T terms, for some integer T to be determined:

$$P_{i_1}, \dots, P_{i_T} \in_R \mathcal{P}.$$

Starting from $S_0 = \{x\}$, we apply the terms sequentially. We denote the t -th step support and frustration as:

$$\begin{aligned} S_t &:= \text{Supp}(P_{i_t} |S_{t-1}\rangle), \\ \gamma_t &:= 1 - \|P_{i_t} |S_{t-1}\rangle\|_2^2. \end{aligned}$$

Since the instance is ε -frustrated and P_{i_t} is chosen uniformly, we have that by definition, the expected frustration at any step is $\mathbb{E}_{P_{i_t}} [\gamma_t] \geq \varepsilon$. Moreover, as long as we have not encountered a BAD string, the frustration/expansion lemma (Equation (2)) implies that the sets $S_1, \dots, S_T$ are gradually expanding:

$$|S_t| \geq \left(1 + \frac{\gamma_t}{2}\right) |S_{t-1}|.$$

Thus, if no BAD string were reached, one can show that the expected support after T steps would satisfy

$$\mathbb{E}[|S_T|] \geq \exp\left(\frac{1}{4} \sum_{t=1}^T \mathbb{E}[\gamma_t]\right) \geq \exp\left(\frac{\varepsilon T}{4}\right). \quad (5)$$

Therefore, choosing $T = 16n/\varepsilon$, this lower bound is at least 2^{2n} , which exceeds the 2^n strings in the entire space. Hence some sequence of terms must reach a BAD string.

However, enumerating all sequences of length T would be far too expensive. The key point is that most terms in such a sequence are irrelevant for producing the final BAD string. To see this, for every step t , consider the (reversed order) sequence

$$P_{i_t}, P_{i_{t-1}}, \dots, P_{i_1}.$$

Now beginning from the first term P_{i_t} , iteratively retain only terms that overlap some term already retained. Denote the resulting sequence as

$$P'_{i'_t} = P_{i_t}, P'_{i'_{t-1}}, \dots, P'_{i'_1},$$

and denote the length of this subsequence as D_t .

This is precisely the backward lightcone of the root term within the random sequence. In contrast to the structured lightcone that emerges from the analysis of [AG19], our lightcone’s only structure requirement is that each term overlaps a previously chosen term (but not necessarily in layers).

It is not too hard to show that the expected length of the retained overlapping sequence is at most exponentially large in $1/\varepsilon$:

$$\mathbb{E}[D_t] \approx 2^{O(1/\varepsilon)}.$$

Therefore, setting $\ell_\varepsilon = \frac{2}{\varepsilon} \mathbb{E}[D_t]$, it follows by Markov's inequality that with high probability the retained overlapping sequence has length at most ℓ_ε :

$$\Pr[D_t > \ell_\varepsilon] \leq \frac{\varepsilon}{2}.$$

We then show that even after *chopping* the random sequence whenever $D_t > \ell_\varepsilon$, namely removing some terms P_{it} , the expected frustration is still relatively high, i.e. $\geq \varepsilon/2$. The preceding expansion argument (Equation (5)) therefore continues to work, with only a constant-factor loss. In particular, the chopped sequence must still reach a BAD string. On the other hand, every term that survives the chopping has a bounded (backward) lightcone of size at most

$$\ell_\varepsilon \approx \exp(1/\varepsilon),$$

in contrast to our previous argument bounding it by $\approx \exp(1/\varepsilon^2)$. This allows us to choose $\varepsilon > 1/\log \log n$ while still enumerating all overlapping sequences of length ℓ_ε in polynomial time.

Tightness of the algorithm. One may hope to improve the algorithm by *pruning* overlapping sequences that do not make progress. In particular, we could enumerate only sequences in which every applied term has noticeable frustration. We show that this does not help: our last contribution is demonstrating an explicit frustrated instance, which has exponentially many overlapping sequences, that all appear equally promising according to this criterion (namely, all of them are sequentially frustrated).

Let us describe the example at a high level. The instance is arranged on a complete k -ary tree over the nodes $V = [n]$, where the qudits live over the alphabet $\Sigma = \{0, \dots, q-1\}$ for $q \geq 3$. We associate a single term $P^{(v)}$ with each node $v \in V$, that acts as follows:

- it acts only on the qudit v and its children $v[1], \dots, v[k]$.
- $P^{(v)}$ projects the v -th qudit to the subset state $|A\rangle$, for some $A \subseteq \Sigma$.
- $P^{(v)}$ projects each child register into the fully uniform state $|\Sigma\rangle$.

By choosing $|A| \approx (1 - \delta)q$, one can show that these two requirements are slightly incompatible, making the instance $\Omega(\delta/k)$ -frustrated, and therefore a NO instance.

Now start from the witness $x = 0^n$. For any rooted subtree, list its terms top down from the distinguished root, as in the rooted overlapping-sequence convention. The corresponding operator product applies them in reverse, bottom-up order. Immediately before applying $P^{(v)}$, register v is therefore still in state $|0\rangle$, and hence

$$\|P^{(v)}|\psi\rangle\|_2^2 \leq 1 - \delta.$$

Thus, every term frustrates the state when it is applied, and every such rooted subtree gives a sequentially δ -frustrated overlapping sequence.

The problem is that there are many possible subtrees. At every reached node, we may independently choose which of its k children to continue exploring. Among sequences containing at most L applied terms, the construction gives at least $2^{\Omega(L)}$ different candidates. Taking $L = \ell_\varepsilon \approx \exp(1/\varepsilon)$ therefore yields

$$2^{\Omega(L)} = \exp(\exp(1/\varepsilon)).$$

Although one of them is guaranteed to reach a BAD string, local frustration gives no way to distinguish it from the others. Hence a pruning strategy based on local frustration cannot improve the $\varepsilon = \Theta(1/\log \log n)$ threshold of our algorithm.

1.3 Discussions

Geometric constraints. Our techniques combine naturally with the geometric constraints of a D -dimensional lattice, yielding an algorithm for exponentially smaller soundness parameters, up to $\beta = \frac{1}{(\log n)^{O(1/D)}}$.

However, the ground energy of any (general) Local Hamiltonian problem on a D -dimensional lattice can be approximated deterministically in polynomial time (namely, in $\mathbf{P}$) to additive accuracy $\frac{1}{\text{poly}(\log n)}$. Indeed, as noted by Hastings and Terhal (see [AGIK09]), one can decompose the Hamiltonian into blocks over $O(\log n)$ qudits, find the ground-state energy within each block in exponential time $2^{O(\log n)} = \text{poly}(n)$, and then deduce an $\approx (\log n)^{-1/D}$ additive approximation for the original instance. See [AAV13] for additional details.

It is a curious coincidence that the deterministic block-decomposition technique and the frustration/expansion framework of [AG19] both reach $\frac{1}{\text{poly}(\log n)}$ approximation scales.

On the importance of uniformity-preserving reductions. Our results build on Aharonov and Grilo [AG19] and are therefore restricted to stoquastic Local Hamiltonian instances satisfying the *uniformity* condition. As observed there, extending the frustration–expansion framework to non-uniform weights is challenging because frustration may depend on global properties of the Hamiltonian. Our improvements thus further motivate understanding the distinction between the uniform and non-uniform variants.

2 Preliminaries

2.1 Definitions

We follow definitions and notations from [AG19].

Let $\Sigma = [q]$ be an alphabet for $q \in \mathbb{N}$. For any subset $S \subseteq \Sigma^n$, the Uniform Subset State $|S\rangle$ is defined as

$$|S\rangle \stackrel{\text{def}}{=} \frac{1}{\sqrt{|S|}} \sum_{x \in S} |x\rangle.$$

The support of a quantum state $|\varphi\rangle$ is defined as

$$\text{Supp}(|\varphi\rangle) \stackrel{\text{def}}{=} \{x \mid \langle x|\varphi\rangle \neq 0\}.$$

The Uniform Support Subset State of a quantum state $|\varphi\rangle$, denoted as $|\widehat{\varphi}\rangle$, is defined as the uniform subset state over the support of the state, namely

$$|\widehat{\varphi}\rangle \stackrel{\text{def}}{=} |\text{Supp}(|\varphi\rangle)\rangle.$$

Definition 2.1. A linear operator $P \in \mathbb{C}^{q^n \times q^n}$ is k -local Uniform Stoquastic Projector if it satisfies the following:

1. *Stoquastic Projector:* P is entrywise real and non-negative, namely $P \in (\mathbb{R}^{\geq 0})^{q^n \times q^n}$.
2. *k -local:* P acts non-trivially on at most k qubits, namely $P = P'_S \otimes I_{[n] \setminus S}$ for some $|S| \leq k$.
3. *Uniform:* the non-trivial part can be written as $P' = \sum_{j \in [r]} |T_j\rangle\langle T_j|$, where each $|T_j\rangle$ is a uniform subset state, and the subset states $|T_j\rangle$ are orthonormal $\langle T_i | T_j \rangle = \delta_{ij}$.

For example, if P is a Uniform Stoquastic Projector acting non-trivially on exactly the first k qubits, there exist an integer $r \leq q^k$ and sets $T_j \subseteq [q]^k$ such that

$$P = \sum_{j \in [r]} \sum_{z \in [q]^{n-k}} (|T_j\rangle\langle T_j|)_{[k]} \otimes (|z\rangle\langle z|)_{[k+1,n]}.$$

The support of a k -local linear operator $P \in \mathbb{C}^{q^n \times q^n}$, denoted by $\text{Supp}(P)$, is the set of qudits $S \subseteq [n]$ on which it acts non-trivially.

Definition 2.2. We define $\text{ProjUSLH}_q(k, d)$ as the set of all collections $\mathcal{P} = \{P_i\}_{i \in [m]}$ such that each P_i is a k -local Uniform Stoquastic Projector over n qudits with alphabet $\Sigma = [q]$, and each qudit participates in at most $d = O(1)$ terms.

Definition 2.3. The Projection Uniform Stoquastic Local Hamiltonian promise problem is defined as

$$\text{PROJUSLH}_q(k, d, \alpha, \beta) \stackrel{\text{def}}{=} \text{YES}_{q,k,d,\alpha} \cup \text{NO}_{q,k,d,\beta}.$$

We omit the subscripts and abbreviate YES, NO when the parameters are clear from context.

An instance $\mathcal{P} \in \text{ProjUSLH}_q(k, d)$ belongs to one of the sets if:

1. $\mathcal{P} \in \text{YES}$: there exists a “ground state” $|\psi\rangle$ such that $\langle \psi | \frac{1}{m} \sum_{i \in [m]} P_i | \psi \rangle \geq 1 - \alpha$;
2. $\mathcal{P} \in \text{NO}$: for all states $|\psi\rangle$ it holds that $\langle \psi | \frac{1}{m} \sum_{i \in [m]} P_i | \psi \rangle \leq 1 - \beta$.

Any instance not belonging to either YES or NO is outside the promise.

Our results work with an arbitrary alphabet size q , but for simplicity one should think of a constant-size alphabet, e.g., binary $q = 2$. We omit q from the notation when it is clear from context.

2.2 The Associated Graph and Its Analysis

Following Bravyi, Bessen, and Terhal [BBT06] and Aharonov and Grilo [AG19], we associate each instance of PROJUSLH with a graph. Given $\mathcal{P} = \{P_i\}_{i \in [m]} \in \text{ProjUSLH}_q(k, d)$, define the undirected graph $G = (\Sigma^n, E)$ as follows:

$$(x, y) \in E \iff \exists i \in [m] \text{ such that } \langle x | P_i | y \rangle > 0.$$

Definition 2.4 ([BBT06]). A string $x \in \Sigma^n$ is called P_i -BAD if $\langle x | P_i | x \rangle = 0$. Similarly, x is BAD if it is P_i -BAD for some $P_i \in \mathcal{P}$. A string $x \in \Sigma^n$ is GOOD if it is not BAD.

From here on we focus on analyzing instances $\mathcal{P}$ of the gap problem $\text{ProjUSLH}_q(k, d, 0, \varepsilon)$, namely those that admit either perfect completeness or non-trivial soundness $\varepsilon > 0$. We introduce the high-level analysis of such graphs.

In the YES case, we rely on earlier work:

Lemma 2.1 ([BBT06, Appendix A.2.1]). Consider a YES instance $\mathcal{P}$ of $\text{ProjUSLH}_q(k, d, 0, \varepsilon)$, for arbitrary $\varepsilon > 0$.

Let $|\theta\rangle$ be the ground state of $\mathcal{P}$. Then the connected component of every $w \in \text{Supp}(|\theta\rangle)$ in the graph of $\mathcal{P}$ contains no BAD strings.

In the NO case, [AG19] demonstrated the frustration vs. expansion lemma, which plays a critical role in finding BAD strings.

Lemma 2.2 ([AG19]). Let $P \in \mathbb{R}^{|\Sigma|^n \times |\Sigma|^n}$ be a uniform stoquastic projector. Let $S \subset \Sigma^n$ be an arbitrary set that contains no P -BAD strings.

Denote the frustration $\gamma = 1 - \|P|S\rangle\|_2^2$. Then frustration implies expansion:

$$|\text{Supp}(P|S\rangle)| \geq (1 + \gamma/2)|S|.$$

3 Our Algorithm

In this section, we present a probabilistic analysis as an alternative to the deterministic one of [AG19]. To this end, instead of relying on the original notion of *structured* lightcones from [AG19], we introduce a new notion of *unstructured* lightcones, namely of overlapping sequences.

Definition 3.1. Let $\mathcal{P} \in \text{ProjUSLH}_q(k, d)$. Let $(Q_1, \dots, Q_t)$ be a sequence of terms in $\mathcal{P}$. The sequence is called *overlapping* if every term Q_j for $j \geq 2$ overlaps some earlier term $Q_{j'}$ for $j' < j$.

Theorem 3.1. $\text{PROJUSLH}_q(k, d, 0, \eta \cdot \frac{1}{\log \log n}) \in \text{NP}$ for some constant $\eta = \eta(k, d, q)$.

Algorithm 1:

Input: An instance $\mathcal{P} = \{P_i\}_{i \in [m]}$, and $\varepsilon > 0$.

Witness: $x \in \Sigma^n$.

```

for all root terms  $P \in \mathcal{P}$ , and sequence length  $\ell \leq O(q^{32kd/\varepsilon})$  do
  for all overlapping sequences  $(P, P_{i_1}, \dots, P_{i_\ell})$  do
    for all reachable strings  $y \in \text{Supp}(P_{i_1} \cdots P_{i_\ell} | x)$  do
      if  $y$  is  $P$ -BAD then
        return reject;
return accept;

```

Proof. The verifier is described in algorithm 1.

Correctness. Consider first a YES instance. Let $|\theta\rangle$ be a frustration-free ground state, and suppose that the witness x belongs to $\text{Supp}(|\theta\rangle)$. By Lemma 2.1, the connected component of x consists entirely of GOOD strings. Every string examined by the verifier is reachable from x , and is therefore GOOD. Hence, the verifier accepts.

Now consider a NO instance. By Lemma 3.7, for every $x \in \Sigma^n$, there exist a term $P \in \mathcal{P}$, and an overlapping sequence $(P, P_{i_1}, \dots, P_{i_\ell})$ such that $\text{Supp}(P_{i_1} \cdots P_{i_\ell} | x)$ contains a P -BAD string y , where

$$\ell \leq \ell_\varepsilon \leq q^{32kd/\varepsilon}.$$

The verifier enumerates every such P , length ℓ , and overlapping sequence, and therefore eventually examines this particular sequence and finds y . Therefore the verifier always rejects.

Time complexity. We bound the cost of each stage of the verifier.

1. Fix $\ell \leq \ell_\varepsilon$ and $P \in \mathcal{P}$. After the root and $r - 1$ applied terms have been selected, there are r retained terms in total. The next term must overlap at least one of them, so there are at most kdr choices. Thus, the number of overlapping sequences with ℓ applied terms rooted at P is at most

$$\leq \prod_{r=1}^{\ell} kdr \leq (k d \ell)^\ell \leq \exp(O(\ell_\varepsilon \log(k d \ell_\varepsilon))) \leq \exp\left(\exp\left(O\left(\frac{k d \log q}{\varepsilon}\right)\right)\right).$$

2. Fix an overlapping sequence $(P, P_{i_1}, \dots, P_{i_\ell})$. Each applied projector acts on at most k symbols and can increase the current support size by a factor of at most q^k . Consequently, the final support set can be enumerated in time

$$|\text{Supp}(P_{i_1} \cdots P_{i_\ell} | x)| \leq q^{k\ell}.$$

3. For each enumerated string y , checking whether y is P -BAD requires only the k symbols on which P acts, and hence takes time $O(k)$.

There are $|\mathcal{P}| = \text{poly}(n)$ choices for P and at most ℓ_ε choices for ℓ . Therefore, the total running time is at most

$$\begin{aligned}
&\leq |\mathcal{P}| \cdot \ell_\varepsilon \cdot \exp\left(\exp\left(O\left(\frac{kd \log q}{\varepsilon}\right)\right)\right) \cdot q^{O(k\ell_\varepsilon)} \cdot O(k) \\
&\leq \text{poly}(n) \cdot q^{O(kd/\varepsilon)} \cdot \exp\left(\exp\left(O\left(\frac{kd \log q}{\varepsilon}\right)\right)\right) \cdot q^{k_\varepsilon^2 q^{O(kd/\varepsilon)}} \\
&\leq \exp\left(\exp\left(O\left(\frac{kd \cdot \ln q}{\varepsilon}\right)\right)\right).
\end{aligned}$$

Assuming $k, d, q = O(1)$ are all constants, it follows that there exists some constant $\eta = \eta(k, d, q)$ such that the above is bounded by $\text{poly}(n)$ for

$$\varepsilon \geq \frac{\eta}{\log \log n}.$$

□

3.1 Setup

Our main structural lemma, Lemma 3.7, shows that if the instance is ε -frustrated, then from every string there is a bounded-length overlapping sequence reaching a BAD string.

To this end, we use a probabilistic argument, showing that a sufficiently long randomly chosen sequence of terms $P_{i_1}, \dots, P_{i_T}$ satisfies the following:

- at some time t , the set S_{t-1} contains a P_{i_t} -BAD string;
- there exists a relatively short overlapping subsequence that leads x to a BAD string, say w .

Since all the terms act locally, being P -BAD depends only on the few qudits on which P acts. Using a standard lightcone argument, we can trace backward the terms that influenced these qudits throughout our sequence. Only these terms are eventually relevant for modifying x into a P -BAD string. It follows that we should bound the length of the longest *overlapping subsequence* in order to bound the graph distance between x and w .

Definition 3.2. Let $\vec{Q} = (Q_1, \dots, Q_t)$ be a sequence of terms in $\mathcal{P}$. The overlapping subsequence of $\vec{Q}$, denoted by $\text{OverlapSubseq}(\vec{Q})$, is the longest ordered subsequence $Q_{i_1}, \dots, Q_{i_t}$ such that $i_1 = 1$ and for every $2 \leq j \leq t$, Q_{i_j} overlaps Q_{i_k} for some $k < j$.

Formally, it can be constructed as follows:

Algorithm 2: OverlapSubseq

Input: A sequence of terms $\vec{Q} = (Q_1, \dots, Q_t)$.
Output: An overlapping subsequence $\vec{W} \subseteq \vec{Q}$.
 $\vec{W} \leftarrow (Q_1);$
for $j = 2, \dots, t$ **do**
 if $\exists Q \in \vec{W}$ such that $\text{Supp}(Q_j) \cap \text{Supp}(Q) \neq \emptyset$ **then**
 append Q_j to $\vec{W};$
return $\vec{W};$

Without loss of generality, we assume that $m \geq n$.⁵ Throughout, we consider the parameters

$$T \stackrel{\text{def}}{=} \left\lceil \frac{16m \ln q}{\varepsilon} \right\rceil \quad \text{and} \quad \ell_\varepsilon \stackrel{\text{def}}{=} \left\lceil \frac{2}{\varepsilon} q^{16kd/\varepsilon} \right\rceil. \quad (6)$$

Fix an ε -frustrated instance $\mathcal{P}$. Throughout this section, we use the following notations. Let $x \in \Sigma^n$ be arbitrary. Consider T -many indices, which we will later show how to choose:

$$i_1, \dots, i_T \in [m].$$

Fixing these choices, we define the following quantities.

- For every $t \leq T$, denote by D_t the length of the overlapping subsequence of the reverse sequence $(P_{i_t}, \dots, P_{i_1})$, namely

$$D_t := |\text{OverlapSubseq}(P_{i_t}, P_{i_{t-1}}, \dots, P_{i_1})|.$$

- Starting from $S_0 \stackrel{\text{def}}{=} \{x\}$, for every $t \leq T$ define the set $S_t \subseteq \Sigma^n$ as

$$S_t \stackrel{\text{def}}{=} \text{Supp}(P_{i_t} | S_{t-1}).$$

We identify every set S_t with its uniform subset state $|S_t\rangle$.

- For every $t \leq T$, define the t -th step frustration as

$$\gamma_t \stackrel{\text{def}}{=} 1 - \|P_{i_t} | S_{t-1}\|_2^2.$$

Our goal is to show that, for some choice of indices $i_1, \dots, i_T$, the sets $S_1, \dots, S_T$ steadily increase in size. In turn, S_T would be too large unless a BAD string appears earlier. To this end, let $\mathcal{I}$ be an arbitrary distribution on $[m]^T$, and suppose the indices are sampled from it:

$$(i_1, \dots, i_T) \sim \mathcal{I}.$$

We now relate the expected size of the final set, S_T , in this random process, to the expected frustration at every intermediate step.

Lemma 3.2. *Let $\mathcal{I} \subseteq [m]^T$ be an arbitrary distribution, and let $(i_1, \dots, i_T) \sim \mathcal{I}$.*

Assume that, with probability one, for every t the set S_{t-1} contains no P_{i_t} -BAD string. Then, for every $x \in \Sigma^n$,

$$\mathbb{E}_{(i_1, \dots, i_T) \sim \mathcal{I}} [|S_T|] \geq \exp \left(\frac{1}{4} \sum_{t=1}^T \mathbb{E}_{(i_1, \dots, i_T) \sim \mathcal{I}} [\gamma_t] \right).$$

Proof. By the hypothesis, for every t , the set S_{t-1} contains no P_{i_t} -BAD strings, and hence we can apply Lemma 2.2. Using the inequality $1 + \delta \geq e^{\delta/2}$ for all $\delta \in [0, 1]$, we conclude

$$|S_t| \geq \left(1 + \frac{\gamma_t}{2}\right) |S_{t-1}| \geq e^{\gamma_t/4} |S_{t-1}|.$$

Applying this argument inductively, taking expectations and applying Jensen's inequality, we obtain

$$\mathbb{E}_{(i_1, \dots, i_T) \sim \mathcal{I}} [|S_T|] \geq \mathbb{E}_{(i_1, \dots, i_T) \sim \mathcal{I}} \exp \left(\frac{1}{4} \sum_{t=1}^T \gamma_t \right) \geq \exp \left(\frac{1}{4} \sum_{t=1}^T \mathbb{E}_{(i_1, \dots, i_T) \sim \mathcal{I}} [\gamma_t] \right).$$

□

⁵After removing unused qudits, we have $n \leq km$. We can therefore duplicate all terms at most k times to obtain $m \geq n$. This leaves the averaged Hamiltonian, and hence its frustration, unchanged, while increasing the degree by only a constant factor.

The role of the preceding lemma is to force a BAD string at an intermediate time whenever the chosen distribution makes $\sum_t \mathbb{E}[\gamma_t]$ sufficiently large. Indeed, if every fixation of indices and time t had S_{t-1} containing no P_{i_t} -BAD string, then the lemma would give $\mathbb{E}[|S_T|] > |\Sigma^n|$, contradicting the trivial bound $S_T \subseteq \Sigma^n$. Hence, for the distributions considered below, some indices and time t satisfy that S_{t-1} contains a P_{i_t} -BAD string. It remains to bound the distance from x to such a string, which reduces to bounding D_t .

In Sections 3.2 and 3.3 we show two bounds for two distributions $\mathcal{I}$.

3.2 Uniformly Random Sequence

In this subsection we analyze the above process when $\mathcal{I}$ is the uniform distribution on $[m]^T$. First, observe that the expected frustration of a randomly chosen term is at least the frustration of $\mathcal{P}$:

Lemma 3.3. *For every $t \leq T$, if $i_t \in_R [m]$ is selected uniformly at random, we have that*

$$\mathbb{E}_{i_t \in_R [m]} [\gamma_t] \geq \varepsilon.$$

Proof. Fix $i_1, \dots, i_{t-1}$ arbitrarily, and let $|S'_{t-1}\rangle$ be the corresponding fixed state. Then, since $\mathcal{P}$ is ε -frustrated,

$$\mathbb{E}_{i_t \in_R [m]} [\gamma_t | i_1, \dots, i_{t-1}] = \frac{1}{m} \sum_{i_t \in [m]} (1 - \|P_{i_t} |S'_{t-1}\rangle\|^2) = 1 - \langle S'_{t-1} | \frac{1}{m} \sum_{i_t \in [m]} P_{i_t} |S'_{t-1}\rangle \geq \varepsilon.$$

The lemma follows since the above holds for every fixed choice of $i_1, \dots, i_{t-1}$. $\square$

We observe that Lemmas 3.2 and 3.3 already imply that some choice of indices must encounter a P_{i_t} -BAD string in S_{t-1} at some time t . Indeed, if this never occurred, then for $T \geq \frac{8n}{\varepsilon} \ln q$,

$$\mathbb{E}_{i_1, \dots, i_T \in_R [m]} [|S_T|] \geq \exp \left(\frac{1}{4} \sum_{t=1}^T \mathbb{E}_{i_1, \dots, i_T \in_R [m]} [\gamma_t] \right) \geq \exp \left(\frac{T\varepsilon}{4} \right) \geq q^{2n}.$$

This contradicts the trivial bound $|S_T| \leq q^n$.

However, the issue is that such a sequence does not guarantee any *bound* on the maximal path from x to the BAD string, besides the trivial bound $D_t \leq T$. We address this issue next.

3.3 The Chopped Sequence

In this subsection, we still pick the indices uniformly at random, but we remove some of them. Informally, we consider the “chopped” random sequence of $(P_{i_1}, \dots, P_{i_T})$, where we remove terms P_{i_t} that can lead to long overlapping subsequences. We then analyze this random process and show that eventually the modified random sequence both (a) expands sufficiently and contains a BAD string, and, more importantly, (b) gives a bound on the maximal distance between x and a BAD string.

Formally, define the new “chopped” random sequence as follows. The indices are chosen uniformly at random:

$$(i_1, \dots, i_T) \in_R [m]^T.$$

We now define the sequence $\widetilde{P}_{i_1}, \dots, \widetilde{P}_{i_T}$, which is related to the random sequence $P_{i_1}, \dots, P_{i_T}$. For every t we define

$$\widetilde{P}_{i_t} \stackrel{\text{def}}{=} \begin{cases} P_{i_t}, & D_t \leq \ell_\varepsilon, \\ I, & D_t > \ell_\varepsilon. \end{cases}$$

Correspondingly, we define the support states as $\tilde{S}_0 = \{x\}$, and for all $t \geq 1$,

$$\tilde{S}_t \stackrel{\text{def}}{=} \text{Supp}(\tilde{P}_{i_t} | \tilde{S}_{t-1}).$$

The corresponding frustration of the t -th step in the chopped random process is defined as

$$\tilde{\gamma}_t \stackrel{\text{def}}{=} 1 - \left\| \tilde{P}_{i_t} | \tilde{S}_{t-1} \right\|_2^2.$$

We now relate the expected frustration of the chopped sequence to the chopping probability.

Lemma 3.4. *For every $t \leq T$,*

$$\mathbb{E}[\tilde{\gamma}_t] \geq \varepsilon - \Pr[D_t > \ell_\varepsilon].$$

Proof. Define the “hybrid” frustration random variable, of the original term P_{i_t} with respect to the current state of the chopped process:

$$\bar{\gamma}_t := 1 - \left\| P_{i_t} | \tilde{S}_{t-1} \right\|_2^2.$$

Note that if $\tilde{P}_{i_t} = I$, then by definition there is no frustration:

$$\tilde{\gamma}_t = 1 - \left\| I \cdot | \tilde{S}_{t-1} \right\|_2^2 = 0,$$

and since that happens whenever $D_t > \ell_\varepsilon$, we can write

$$\tilde{\gamma}_t = \bar{\gamma}_t \mathbf{1}_{D_t \leq \ell_\varepsilon}.$$

Then, since the frustration is bounded by $\bar{\gamma}_t \leq 1$,

$$\begin{aligned} \mathbb{E}[\tilde{\gamma}_t] &= \mathbb{E}[\bar{\gamma}_t \mathbf{1}_{D_t \leq \ell_\varepsilon}] \\ &= \mathbb{E}[\bar{\gamma}_t] - \mathbb{E}[\bar{\gamma}_t \mathbf{1}_{D_t > \ell_\varepsilon}] \\ &\geq \varepsilon - 1 \cdot \Pr[D_t > \ell_\varepsilon], \end{aligned}$$

where the proof that $\mathbb{E}[\bar{\gamma}_t] \geq \varepsilon$ follows similarly to Lemma 3.3. $\square$

Therefore, if we can bound the probability of having a long overlapping subsequence by, say, $\varepsilon/2$, we can guarantee both (a) the expected chopped frustration is still high, $\mathbb{E}[\tilde{\gamma}_t] \geq \varepsilon/2$, which is necessary for the gradual expansion of the sets $\tilde{S}_t$, and (b) a bound on the maximal path between x and a BAD string.

We now show that, for every fixed step, its overlapping subsequence is relatively short with high probability:

Lemma 3.5. *For every $t \leq T$,*

$$\Pr[D_t > \ell_\varepsilon] \leq \frac{\varepsilon}{2}.$$

Proof. Fix t , and consider the construction of $\text{OverlapSubseq}(P_{i_t}, \dots, P_{i_1})$ at step $j \in [t-1, 1]$. Suppose that r terms have already been retained. Therefore at most kdr terms of $\mathcal{P}$ overlap at least one of them, and since i_j is uniform and independent of the previously fixed indices $i_t, \dots, i_{j+1}$, it follows that the current term is retained with probability at most

$$\Pr_{i_j} [P_{i_j} \text{ is retained} \mid i_t, \dots, i_{j+1}] \leq \frac{kdr}{m}.$$

Thus, the expected number of retained terms grows by a factor of at most $(1 + \frac{kd}{m})$ at every step. Starting from the single root term, and using $t - 1 \leq T - 1 < \frac{16m \ln q}{\varepsilon}$, we obtain

$$\mathbb{E}[D_t] = \mathbb{E}[|\text{OverlapSubseq}(P_{i_t}, \dots, P_{i_1})|] \leq \left(1 + \frac{kd}{m}\right)^{t-1} \leq e^{kd(t-1)/m} \leq q^{16kd/\varepsilon}.$$

Therefore, by Markov's inequality,

$$\Pr[D_t > \ell_\varepsilon] \leq \frac{\mathbb{E}[D_t]}{\ell_\varepsilon} \leq \frac{\varepsilon}{2}.$$

□

Plugging the above bounds into Lemma 3.2, we obtain:

Lemma 3.6. *Let $x \in \Sigma^n$ be arbitrary. Suppose that, with probability one, for every t the set $\tilde{S}_{t-1}$ contains no $\tilde{P}_{i_t}$ -BAD string. Then,*

$$\mathbb{E}[|\tilde{S}_T|] \geq q^{2n}.$$

Proof. Combining Lemmas 3.4 and 3.5, we have that $\mathbb{E}[\tilde{\gamma}_t] \geq \varepsilon/2$. Let $\tilde{\mathcal{I}}$ denote the distribution on chopped sequences induced by the uniformly random indices. Applying Lemma 3.2 to $\tilde{\mathcal{I}}$ and the states $\tilde{S}_t$ gives

$$\mathbb{E}[|\tilde{S}_T|] \geq \exp\left(\frac{1}{4} \sum_{t=1}^T \mathbb{E}[\tilde{\gamma}_t]\right) \geq \exp\left(\frac{\varepsilon T}{8}\right) \geq q^{2n}.$$

□

At a high level, the preceding expectation bound forces some realization of the chopped process to encounter a BAD string. Moreover, by design, the chopped process induces a relatively short path from any string to a BAD one, allowing us to conclude:

Lemma 3.7. *For every $x \in \Sigma^n$, there exist a term $P \in \mathcal{P}$ and a rooted overlapping sequence $(P, Q_1, \dots, Q_\ell)$, of length $\ell < \ell_\varepsilon$, such that $\text{Supp}(Q_1 \cdots Q_\ell | x)$ contains a P -BAD string.*

Proof. Fix $x \in \Sigma^n$. Suppose, toward a contradiction, that no fixing of the indices encounters a $\tilde{P}_{i_t}$ -BAD string in $\tilde{S}_{t-1}$ at any time t . Then the hypothesis of Lemma 3.6 holds for every choice of indices, and hence

$$\mathbb{E}[|\tilde{S}_T|] \geq q^{2n}.$$

This contradicts the trivial bound $|\tilde{S}_T| \leq q^n$. Therefore, for some $i_1, \dots, i_T$ and some time t , the set $\tilde{S}_{t-1}$ contains a $\tilde{P}_{i_t}$ -BAD string. Necessarily $\tilde{P}_{i_t} = P_{i_t}$, since no string is I -BAD. By the definition of the chopped sequence, this implies $D_t \leq \ell_\varepsilon$.

Finally, form the backward lightcone of the *chopped* sequence $(\tilde{P}_{i_t}, \dots, \tilde{P}_{i_1})$, discarding identity operators. Denote this lightcone, with its root listed first, by $(P_{i_t}, Q_1, \dots, Q_\ell)$. Replacing terms by the identity can only remove dependencies, so this lightcone is contained in $\text{OverlapSubseq}(P_{i_t}, \dots, P_{i_1})$ and contains at most D_t terms (including its root P_{i_t}).

We now show that the non-root terms of this lightcone still lead from x to a P_{i_t} -BAD string. To this end, choose a P_{i_t} -BAD string $y \in \tilde{S}_{t-1}$. Since the projectors are entrywise non-negative,

$$0 < \langle y | \tilde{P}_{i_{t-1}} \cdots \tilde{P}_{i_1} | x \rangle. \quad (7)$$

Let $z \stackrel{\text{def}}{=} y|_{\text{Supp}(P_{i_t})}$ be the local BAD pattern witnessed by y . Consider any term of the chopped sequence $\tilde{P}_{i_j}$ that does not appear among $Q_1, \dots, Q_\ell$. By construction of the backward lightcone, this term shares no qudit with P_{i_t} or with any retained term Q_r applied after it. Thus, any change caused by this term cannot propagate to the final symbols on $\text{Supp}(P_{i_t})$. Removing all such terms from the product in Equation (7) therefore gives some $y' \in \Sigma^n$ such that

$$0 < \langle y' | Q_1 \cdots Q_\ell | x \rangle, \quad y'|_{\text{Supp}(P_{i_t})} = z.$$

Hence $y' \in \text{Supp}(Q_1 \cdots Q_\ell | x)$ is also P_{i_t} -BAD, which has distance $\ell \leq D_t - 1 \leq \ell_\varepsilon$ from x , as required. □

4 An Instance For Which Exponentially Many Overlapping Sequences Are δ -valid

In this section, we show for every $\delta > 0$, an explicit $O(\delta)$ frustrated instance $\mathcal{P} \in \text{ProjUSLH}_q(k, d)$ which exhibits the following phenomenon: a large subset of the rooted overlapping sequences are sequentially δ -frustrated with respect to some single witness $x \in \Sigma^n$. Thus, restricting the enumeration only to such a subset does not improve the worst-case dependence on the sequence length.

To have some sense of the parameter, the reader can think of any $\delta \leq 1 - \sqrt{2/3}$, for which we may take constant-size alphabet dimension $q = 3$ (and the set $A = \{0, 1\}$ downstream). In particular, for all sufficiently large n , the construction applies with constant alphabet size to any sub-constant frustration $\delta = o(1)$.

4.1 Construction

Let $\delta \in (0, 1)$ be arbitrary. Fix an integer $k \geq 2$. Let the alphabet size be $q := \lceil \frac{2}{(1-\delta)^2} \rceil$, and set $\Sigma = \{0, \dots, q-1\}$. We fix the set

$$A \stackrel{\text{def}}{=} \{0, \dots, \lfloor (1-\delta)q \rfloor - 1\} \subseteq \Sigma.$$

We begin by describing the interaction graph. Let $T \geq 3$ be an integer TBD. Consider the complete depth- T k -ary tree with root $\emptyset$, whose vertex set is defined as

$$V := \{\emptyset\} \bigcup_{t=1}^{T-1} [k]^t.$$

Internal nodes at depth $t < T-1$ are denoted by strings $v \in [k]^t$, and their children are denoted by $v1, \dots, vk$. Each of them is naturally associated with their parent, which is defined as

$$\text{parent}(vj) \stackrel{\text{def}}{=} v.$$

We associate one qudit over Σ per node, so $[n] = V$, and therefore we choose T such that

$$n = \frac{k^T - 1}{k - 1}.$$

Next, we describe the Hamiltonian's terms. The instance is defined as $\mathcal{P} = \{P^{(v)}\}_{v \in [n]}$ where every term is defined as

$$P^{(v)} \stackrel{\text{def}}{=} |T_v\rangle\langle T_v|, \quad \text{where} \quad |T_v\rangle \stackrel{\text{def}}{=} \begin{cases} |A\rangle_v \otimes_{j=1}^k |\Sigma\rangle_{vj}, & v \text{ internal,} \\ |A\rangle_v, & v \text{ leaf.} \end{cases} \quad (8)$$

Figure 1 illustrates the construction.

4.2 Validity of the Instance

First, observe that the instance is valid:

Lemma 4.1. $\mathcal{P} \in \text{ProjUSLH}_q(k+1, d=2)$.

Proof. By construction, each $P^{(v)}$ is either 1-local or $(k+1)$ -local. Additionally, observe that each non-root qudit is acted upon only by $P^{(v)}$ and $P^{(\text{parent}(v))}$, so each qudit is acted upon by at most $d = 2$ terms. Finally, each $|T_v\rangle$ is a singleton uniform subset state, corresponding to either the set $A \times (\Sigma)^k$ or A . $\square$




Figure 1: The instance $\mathcal{P}$ for $k = \ell = 3$. Circles denote qudits, labeled by their register name v ; each block groups the $k + 1$ qudits acted on by a single term $P^{(v)}$, with the “parent” register v on top and the “child” registers $v1, \dots, vk$ beneath. Every register v is shared by exactly two terms, $P^{(v)}$ and $P^{(\text{parent}(v))}$. The line and shaded cone mark this overlap, connecting the register’s appearance as the parent’s i -th child slot to its appearance as the child block’s top register. The two ends are styled differently to mark the asymmetric action on the shared register: the parent applies $|\Sigma\rangle\langle\Sigma|_v$ (coiled end) and the child applies $|A\rangle\langle A|_v$ (straight end).

Second, we show that the instance is frustrated, thereby a NO case.

Lemma 4.2. $\mathcal{P}$ is $\delta/8k$ -frustrated.

Proof. Let $|\psi\rangle$ be an arbitrary state, and consider its energy: Recall that $|V| = n$, and hence

$$\langle\psi|\frac{1}{n}\sum_{v\in V}P^{(v)}|\psi\rangle = \frac{1}{n}\sum_{v\in V}\|P^{(v)}|\psi\rangle\|^2 \leq \frac{1}{n}\sum_{v\in V}\|P^{(v)}|\psi\rangle\|. \quad (9)$$

Fix a pair of parent-child registers $(v, v1)$, and we show how to bound $\|P^{(v)}|\psi\rangle\| + \|P^{(v1)}|\psi\rangle\|$. Later we’ll show how to pair all registers in V .

Each pair of terms $P^{(v)}, P^{(v1)}$ acts simultaneously only on register $v1$. Since $|\phi\rangle\langle\phi| \preceq I$ for any unit-state $|\phi\rangle$, we have that each projector is smaller in the Löwner partial order than the following projectors:

$$\begin{aligned} P^{(v)} &\preceq |\Sigma\rangle\langle\Sigma|_{v1} \otimes I_{v2, \dots, vk} = |\Sigma\rangle\langle\Sigma|_{v1} \\ P^{(v1)} &\preceq |A\rangle\langle A|_{v1} \otimes I_{v11, \dots, v1k} = |A\rangle\langle A|_{v1}. \end{aligned} \quad (10)$$

Therefore, using the inequality $\sqrt{x} + \sqrt{y} \leq \sqrt{2(x+y)}$ we conclude that

$$\begin{aligned} \|P^{(v)}|\psi\rangle\| + \|P^{(v1)}|\psi\rangle\| &= \sqrt{\langle\psi|P^{(v)}|\psi\rangle} + \sqrt{\langle\psi|P^{(v1)}|\psi\rangle} \\ &\leq \sqrt{2(\langle\psi|P^{(v)}|\psi\rangle + \langle\psi|P^{(v1)}|\psi\rangle)} \\ &\leq \sqrt{2\langle\psi|(|\Sigma\rangle\langle\Sigma|_{v1} + |A\rangle\langle A|_{v1})|\psi\rangle} \\ &\stackrel{\text{def}}{=} \sqrt{2\langle\psi|M^{(v1)}|\psi\rangle}. \end{aligned}$$

In Lemma 4.3, we show that the operator $M^{(v_1)}$ has maximal eigenvalue $\lambda_{\max}(M^{(v_1)}) \leq 2 - \delta/2$. Using the inequality $\sqrt{1-x} \leq 1 - x/2$ for $x \in [0, 1]$, we upper bound the above by

$$\left\| P^{(v)} |\psi\rangle \right\| + \left\| P^{(v_1)} |\psi\rangle \right\| \leq \sqrt{4-\delta} = 2\sqrt{1-\delta/4} \leq 2 - \delta/4.$$

Now, let $W \subseteq V$ be the set of $(T-2)$ -depth nodes, namely the internals whose children are all leaves, so $|W| = k^{T-2}$. Pair each $v \in W$ with its first child v_1 . Therefore, using $T \geq 3$ and $k \geq 2$,

$$\frac{|W|}{n} = k^{T-2} \frac{k-1}{k^T-1} \geq \frac{k-1}{k^2} \geq \frac{1}{2k}.$$

Finally, we use the preceding pairwise bound and bound the remaining $n-2|W|$ projection norms trivially by 1. Plugging it into Equation (9), we conclude that

$$\langle \psi | \frac{1}{n} \sum_{v \in V} P^{(v)} | \psi \rangle \leq \frac{n-2|W|}{n} \cdot 1 + \frac{|W|}{n} \cdot (2 - \delta/4) = 1 - \frac{|W|\delta/4}{n} \leq 1 - \frac{\delta}{8k}.$$

□

Lemma 4.3. *For every internal node $u \in V$, define the operator*

$$M^{(u)} \stackrel{\text{def}}{=} |A\rangle\langle A|_u + |\Sigma\rangle\langle \Sigma|_u.$$

Then $\lambda_{\max}(M^{(u)}) \leq 2 - \delta/2$.

Proof. First, observe that the relative size of $A \subseteq \Sigma$ is bounded by

$$\mu \stackrel{\text{def}}{=} \langle A | \Sigma \rangle = \sqrt{|A|/|\Sigma|} \leq \sqrt{1-\delta} \leq 1 - \delta/2.$$

Now observe that

$$\begin{aligned} M^{(u)}(|A\rangle + |\Sigma\rangle) &= (1 + \mu)(|A\rangle + |\Sigma\rangle), \\ M^{(u)}(|A\rangle - |\Sigma\rangle) &= (1 - \mu)(|A\rangle - |\Sigma\rangle). \end{aligned}$$

On the other hand, $M^{(u)}|\chi\rangle = 0$ for every state $|\chi\rangle$ orthogonal to $\text{span}\{|A\rangle, |\Sigma\rangle\}$. Hence $M^{(u)}$'s eigenvalues are $\{1 \pm \mu, 0\}$, and since $\mu > 0$, the maximal eigenvalue is $\lambda_{\max}(M^{(u)}) \leq 1 + \mu \leq 2 - \delta/2$. □

4.3 Doubly Exponentially Many Sequentially Frustrated Sequences

Definition 4.1. *Let $x \in \Sigma^n$. A sequence of terms $(Q_1, \dots, Q_t)$ is sequentially δ -frustrated with respect to x if, for every $j \in [t]$:*

$$\left\| Q_j \cdot \widehat{(Q_{j-1} \cdots Q_1 | x)} \right\|_2^2 \leq 1 - \delta.$$

Fix a term $P^{(u)} \in \mathcal{P}$ and a length bound $\ell \geq 1$.

Definition 4.2. *A sequence $(P^{(v_1)}, \dots, P^{(v_t)})$ is nice below $P^{(u)}$ if its nodes form an ordered top-down subtree below u . Equivalently, for every $j \in [t]$, the parent of v_j is either u , or some $v_{j'}$ with $j' < j$.*

The set of such sequences having at most ℓ applied terms is denoted by

$$\text{NiceOverlapSeq}_{P^{(u)}, \ell} = \left\{ (P^{(v_1)}, \dots, P^{(v_t)}) \mid 0 \leq t \leq \ell \text{ and the sequence is nice below } P^{(u)} \right\}.$$

We first observe that the number of nice sequences is large.

Lemma 4.4. *Let $\ell_\varepsilon \geq 2$. For every node v with at least $\ell_\varepsilon - 1$ levels below it,*

$$\left| \text{NiceOverlapSeq}_{P(v), \ell_\varepsilon} \right| \geq 2^{\ell_\varepsilon/(2k)}.$$

Proof. Let h be maximal such that the complete depth- h subtree rooted at v contains at most ℓ_ε nodes, namely

$$1 + k + \dots + k^h \leq \ell_\varepsilon.$$

Every subset of the k^h leaves determines a distinct subtree of size at most ℓ_ε , and hence a distinct nice overlapping sequence. By the maximality of h we have $\ell_\varepsilon \leq 2k^{h+1}$, and hence $k^h > \ell_\varepsilon/(2k)$. Therefore, the number of nice overlapping sequences of length at most ℓ_ε is at least

$$2^{k^h} \geq 2^{\ell_\varepsilon/(2k)}.$$

□

Next, we prove that all nice overlapping sequences are δ -frustrated with respect to some witness $x \in \Sigma^n$.

Lemma 4.5. *For every $P \in \mathcal{P}$ and length bound $\ell > 0$, if $(Q_1, \dots, Q_t) \in \text{NiceOverlapSeq}_{P, \ell}$, then the reversed sequence $(Q_t, \dots, Q_1)$ is sequentially δ -frustrated with respect to $x = 0^n$.*

Proof. Fix $P \in \mathcal{P}$, a length bound $\ell > 0$, and a nice sequence $(Q_1, \dots, Q_t) \in \text{NiceOverlapSeq}_{P, \ell}$. Fix an arbitrary $j \in [t - 1]$, and let

$$|\Psi\rangle := Q_{j+1} \dots Q_t |x\rangle.$$

We show that

$$\left\| Q_j |\widehat{\Psi}\rangle \right\|_2^2 \leq 1 - \delta. \quad (11)$$

Suppose $P^{(v)} = Q_j$. The key observation is that, immediately before applying $P^{(v)}$, register v has never been acted upon. Indeed, the only terms acting on register v are $P^{(v)}$ and, if v is non-root, $P^{(\text{parent}(v))}$. Since the sequence is listed top down, $P^{(\text{parent}(v))}$ appears before $P^{(v)}$ and is therefore applied later; if $P^{(\text{parent}(v))} = P$, then it is not part of the sequence. Hence none of the terms $Q_{j+1}, \dots, Q_t$ acts on register v , and therefore

$$|\widehat{\Psi}\rangle_v = |0\rangle_v.$$

On register v , the projector $P^{(v)}$ acts as $(P^{(v)})_v = |A\rangle\langle A|_v$. Recall that $P^{(v)} \preceq |A\rangle\langle A|_v \otimes I$ (Equation (10)). Therefore,

$$\begin{aligned} \left\| P^{(v)} |\widehat{\Psi}\rangle \right\|_2^2 &= \langle \widehat{\Psi} | P^{(v)} | \widehat{\Psi} \rangle \\ &\leq \langle \widehat{\Psi} | (|A\rangle\langle A|_v \otimes I) | \widehat{\Psi} \rangle \\ &\leq |\langle A | 0 \rangle|^2 = \frac{1}{|A|}. \end{aligned}$$

Then by our choice $|A| = \lfloor (1 - \delta)q \rfloor$ and $q \geq 2/(1 - \delta)^2$, we have

$$|A| \geq (1 - \delta)q - 1 \geq \frac{2}{1 - \delta} - 1 \geq \frac{1}{1 - \delta}.$$

and hence

$$\left\| P^{(v)} |\widehat{\Psi}\rangle \right\|_2^2 \leq \frac{1}{|A|} \leq 1 - \delta.$$

Since $P^{(v)}$ was arbitrary, every term application is sequentially δ -frustrated, proving Lemma 4.5. □

Acknowledgments

We thank John Bostanci and Uma Girish for collaboration in the early stages of this project, and Alex Grilo and Dorit Aharonov for helpful discussions. This work began during a visit to Henry Yuen’s group, and we thank him for his hospitality.

AI Disclosure. This research was conducted by the human authors with assistance from AI tools. Claude and ChatGPT were used as research and writing assistants for brainstorming and proof exploration. ChatGPT was used during the development of the improvement from the initial $\beta = \Omega(1/\sqrt{\log \log n})$ algorithm to the final $\beta = \Omega(1/\log \log n)$ result, and Claude assisted with exploratory ideas related to Section 4. The authors verified the correctness and originality of all content including references.

References

- [AAV13] Dorit Aharonov, Itai Arad, and Thomas Vidick. “Guest column: the quantum PCP conjecture”. In: *Acm sigact news* 44.2 (2013), pp. 47–79.
- [AG19] Dorit Aharonov and Alex Bredariol Grilo. “Stoquastic PCP vs. randomness”. In: *2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE. 2019, pp. 1000–1023.
- [AGIK09] Dorit Aharonov, Daniel Gottesman, Sandy Irani, and Julia Kempe. “The power of quantum systems on a line”. In: *Communications in mathematical physics* 287.1 (2009), pp. 41–65.
- [AGL25] Dorit Aharonov, Alex B Grilo, and Yupan Liu. “StoqMA vs. MA: the power of error reduction”. In: *Quantum* 9 (2025), p. 1853.
- [BBT06] Sergey Bravyi, Arvid J. Bessen, and Barbara M. Terhal. *Merlin-Arthur Games and Stoquastic Complexity*. 2006. arXiv: [quant-ph/0611021](https://arxiv.org/abs/quant-ph/0611021) [quant-ph]. URL: <https://arxiv.org/abs/quant-ph/0611021>.
- [BCO17] Johannes Bausch, Toby Cubitt, and Maris Ozols. “The complexity of translationally invariant spin chains with low local dimension”. In: *Annales Henri Poincaré* 18.11 (2017), pp. 3449–3513.
- [BDOT08] Sergey Bravyi, David P DiVincenzo, Roberto Oliveira, and Barbara M Terhal. “The complexity of stoquastic local Hamiltonian problems”. In: *Quantum Information & Computation* 8.5 (2008), pp. 361–385.
- [BH17] Sergey Bravyi and Matthew Hastings. “On complexity of the quantum Ising model”. In: *Communications in Mathematical Physics* 349.1 (2017), pp. 1–45.
- [BT10] Sergey Bravyi and Barbara Terhal. “Complexity of Stoquastic Frustration-Free Hamiltonians”. In: *SIAM Journal on Computing* 39.4 (2010), pp. 1462–1485. DOI: [10.1137/08072689X](https://doi.org/10.1137/08072689X). eprint: <https://doi.org/10.1137/08072689X>. URL: <https://doi.org/10.1137/08072689X>.
- [BV05] Sergey Bravyi and Mikhail Vyalyi. “Commutative version of the k-local Hamiltonian problem and common eigenspace problem”. In: *Quantum Information & Computation* 5.3 (2005), pp. 187–215.
- [CM16] Toby Cubitt and Ashley Montanaro. “Complexity classification of local Hamiltonian problems”. In: *SIAM Journal on Computing* 45.2 (2016), pp. 268–316.
- [GHLS14] Sevag Gharibian, Yichen Huang, Zeph Landau, and Seung Woo Shin. “Quantum hamiltonian complexity”. In: *Foundations and Trends® in Theoretical Computer Science* 10.3 (2014), pp. 159–282.

- [Has13] Matthew B Hastings. “Trivial low energy states for commuting Hamiltonians, and the quantum PCP conjecture”. In: *Quantum Information & Computation* 13.5-6 (2013), pp. 393–429.
- [IKW02] Russell Impagliazzo, Valentine Kabanets, and Avi Wigderson. “In search of an easy witness: Exponential time vs. probabilistic polynomial time”. In: *Journal of Computer and System Sciences* 65.4 (2002), pp. 672–694.
- [IW97] Russell Impagliazzo and Avi Wigderson. “P= BPP if E requires exponential circuits: Derandomizing the XOR lemma”. In: *Proceedings of the twenty-ninth annual ACM symposium on Theory of computing*. 1997, pp. 220–229.
- [KM02] Adam R. Klivans and Dieter van Melkebeek. “Graph nonisomorphism has subexponential size proofs unless the polynomial-time hierarchy collapses”. In: *SIAM Journal on Computing* 31.5 (2002), pp. 1501–1526.
- [KSV02] Alexei Yu Kitaev, Alexander Shen, and Mikhail N Vyalyi. *Classical and quantum computation*. 47. American Mathematical Soc., 2002.
- [KT19] Joel Klassen and Barbara M Terhal. “Two-local qubit Hamiltonians: when are they stoquastic?” In: *Quantum* 3 (2019), p. 139.
- [NW94] Noam Nisan and Avi Wigderson. “Hardness vs. randomness”. In: *Journal of Computer and System Sciences* 49.2 (1994), pp. 149–167.
- [OT08] Roberto Oliveira and Barbara M Terhal. “The complexity of quantum spin systems on a two-dimensional square lattice”. In: *Quantum Information & Computation* 8.10 (2008), pp. 900–924.
- [PM17] Stephen Piddock and Ashley Montanaro. “The complexity of antiferromagnetic interactions and 2D lattices”. In: *Quantum Information & Computation* 17.7-8 (2017), pp. 636–672.
- [REG26] Asad Raza, Jens Eisert, and Alex B Grilo. “Complexity of geometrically local stoquastic Hamiltonians”. In: *Quantum* 10 (2026), p. 2004.