Electronic Colloquium on Computational Complexity, Report No. 10 (2019)

Stoquastic PCP vs. Randomness

Dorit Aharonov[1] and Alex B. Grilo[2]

[1]Hebrew University of Jerusalem, Jerusalem, Israel
[2]CWI and QuSoft, Amsterdam, The Netherlands

Abstract

The derandomization of MA, the probabilistic version of NP, is a long standing open question. In this work, we connect this problem to a variant of another major problem: the quantum PCP conjecture. Our connection goes through the surprising quantum characterization of MA by Bravyi and Terhal. They proved the MA-completeness of the problem of deciding whether the groundenergy of a uniform stoquastic local Hamiltonian is zero or inverse polynomial. We show that the gapped version of this problem, i.e. deciding if a given uniform stoquastic local Hamiltonian is frustration-free or has energy at least some constant ϵ, is in NP. Thus, if there exists a gap-amplification procedure for uniform stoquastic Local Hamiltonians (in analogy to the gap amplification procedure for constraint satisfaction problems in the original PCP theorem), then MA = NP (and vice versa). Furthermore, if this gap amplification procedure exhibits some additional (natural) properties, then P = RP. We feel this work opens up a rich set of new directions to explore, which might lead to progress on both quantum PCP and derandomization. As a small side result, we also show that deciding if *commuting* stoquastic Hamiltonian is frustration free is in NP.

1 Introduction

It is a long standing open question, whether the randomized version of NP, called MA (for Merlin and Arthur) can be derandomized, namely, whether MA = NP. In MA, a randomized polynomial-time verification algorithm has to decide if an input x is a positive or negative instance, with the help of an untrusted polynomially long proof y. If x is indeed a yes-instance, there must exists some proof y that makes the verification algorithm accept with probability 1. On the other hand, if x is a no-instance, the verification algorithm should reject with high probability for all possible y's.

The derandomization of MA is implied by widely believed conjectures such as NEXP does not have polynomial size circuits [IKW02]. It is also implied by the stronger derandomization conjecture that P = BPP [GZ11], which itself is implied by the existence of one-way functions [HILL99] as well as by commonly conjectured circuit lower bounds [BFNW93, NW94, IW97, STV01, KI04]. The (somewhat counter-intuitive at first) connection between lower-bound on computation, and derandomization (which can be viewed as an upper-bound result), was coined the intriguing name "Hardness versus Randomness" [NW94]. Our work follows this path, and provides a result of a somewhat similar flavor: we connect the derandomization of MA (as well as that of RP) with a different hardness problem in computational complexity – that of quantum PCP – hence the title of this paper.

ISSN 1433-8092

Our starting point is a seminal and beautiful paper of Bravyi and Terhal [BT09], where they prove the MA-completeness of a quantum-related problem. Until their work, natural MA-complete problems essentially were not known[1] so it is very surprising that the first problem of this kind is *quantum*-defined. To explain this characterization of MA, and how we use it to make the connection to quantum PCP, let us make a detour to quantum Hamiltonian complexity.

1.1 Hamiltonian complexity and stoquastic Hamiltonians

The power of quantum proofs [KSV02, AN02] had been the major area of study in the recent decade extending the notions of NP and MA; this direction had led to enormous progress of our understanding of the complexity of quantum states and the reductions between Hamiltonians (see [Osb12, GHLS15]). A central player here is the class QMA, in which a polynomial time quantum verification algorithm receives a *quantum* proof $|\psi\rangle$ for some classical input x. The verification algorithm should accept x with high probability if x is a positive instance, otherwise, no matter what the quantum proof is, x should be rejected with high probability. In addition to being a natural generalization of classical proof systems, the relevance of QMA was evidenced by Kitaev, who showed that estimating the groundenergy of a local Hamiltonian, a central problem in physics, is complete for QMA [KSV02]. Kitaev's theorem is the quantum analog of the seminal Cook-Levin theorem [Coo71, Lev73], and it makes a very strong connection between a major question in condensed matter physics (namely, groundstates of local Hamiltonians), and a major problem in Theoretical Computer Science, (namely, optimal solutions for constraint satisfaction problems). In fact, the connection is even deeper since what is shown is that the latter is simply a special case of the former.

More concretely[2], the evolution of a physical system is described by a Hamiltonian, which mathematically is represented by a self-adjoint operator. In Nature, Hamiltonians can be usually decomposed as a sum of terms which correspond to interactions between just a small (constant) number of particles. Looking at this problem through Theoretical Computer Science lens, Kitaev defined the k-Local Hamiltonian problem, whose input is a Hamiltonian H on an n particle system that can be decomposed as a sum of m local terms, each of them acting non-trivially on at most k out of the n particles (with no geometrical restrictions on the interacting particles in each term). We then ask if there is a state whose energy is smaller that some parameter α (or mathematically, the smallest eigenvalue of H is at most α), or all states have energy larger than β (or, all eigenvalues are at least β). The hardness of the Local Hamiltonian problem depends on the promise gap $\beta - \alpha$ and Kitaev showed that the problem is QMA-complete for some inverse polynomial promise gap [KSV02].

Bravyi, DiVincenzo, Oliveira and Terhal [BDOT08] studied the problem of Local Hamiltonian when the local terms have more structure. They asked how hard the Local Hamiltonian problem is, when the off-diagonal elements of the local terms are non-positive, a property that they named "stoquastic". This property implies a lot of structure on lowest-energy states[3], and in physics it is associated with the lack of what is known as the "sign problem"; it is considered a far easier

[1] For PromiseMA, it is a folklore that one can prove complete problems by extending NP-complete problems (see, e.g. [SW]): we define an exponential family of 3SAT formulas (given as input succinctly) and we have to decide if there is an assignment that satisfies all of the formulas, or for every assignment, a random formula in the family will not be satisfied with good probability.

[2] See Section 3 for a detailed definition.

[3] See Lemma 3.2 for more details.

case than general Hamiltonians, since one can often use Markov Chain Monte Carlo experiments to study it[4]. In [BBT06], this structure is used to show that the stoquastic Local Hamiltonian problem is StoqMA-complete, where StoqMA is a complexity class defined by them and sits between MA and QMA.

Importantly for this paper, Bravyi and Terhal [BT09] then showed that the stoquastic Local Hamiltonian problem is MA-complete if we pick $\alpha = 0$ and $\beta \geq \frac{1}{poly(n)}$, or in other words, if we want to decide whether the Hamiltonian is *frustration-free*[5] or the frustration is at least inverse polynomial. This problem is, to the best of our knowledge, the first MA-complete problem which is not an extension of NP-complete problems into the randomized setting[6] (see also [Bra14]). We notice that in the MA-complete problem of [BT09], the groundspaces of the local terms are in fact all spanned by *subset-states*, i.e., states which are the *uniform* superposition of a subset of strings. We call these Hamiltonians *uniform stoquastic Hamiltonians*.

This paper is concerned with the *gapped* version of the above defined uniform stoquastic Local Hamiltonian problem. Gapped versions of NP-hard problems have played a crucial role in the topic of probabilistically checkable proofs (PCPs) and have revolutionized classical Theoretical Computer Science over the past three decades. Before we define the gapped version of the uniform stoquastic Hamiltonian problem, let us introduce PCPs in more detail.

1.2 PCP theorem

The "mother" of all NP-complete problems is 3SAT. An instance to this problem is a boolean formula ϕ in the form $\phi(x) = C_1 \wedge C_2 \wedge ... \wedge C_m$, where $C_i = (y_{i,1} \vee y_{i,2} \vee y_{i,3})$ is a clause and $y_{i,j} \in \{x_1, ..., x_n, \overline{x_1}, ...\overline{x_n}\}$. We ask if there exists an assignment $x \in \{0,1\}^n$ such that ϕ is satisfied.

The problem MAX3SAT$_\delta$, parametrized by some function $\delta(n)$, is a generalization of 3SAT. In this problem we have to distinguish between the cases where ϕ is satisfiable or for every assignment of the input variables, at least a $\delta(n)$ fraction of the clauses is not satisfied.

By picking $\delta(n) = \frac{1}{m}$, then MAX3SAT$_\delta$ becomes equivalent to 3SAT, and therefore it is NP-complete. In PCPs, we are interested in versions of the problem for δ significantly larger. The celebrated PCP theorem [ALM+98, AS98] states the remarkable result that there exists some constant ε independent of n, such that the problem MAX3SAT$_\varepsilon$ is NP-complete. This problem with constant ε is called the *gapped* version of the problem, and the PCP theorem proves that the gapped version of this problem is as hard as the original one.

In her celebrated alternative proof to the PCP theorem, Dinur [Din07] used an explicit *gap amplification* procedure. More concretely, the input to this procedure is an instance ϕ of 3SAT (or equivalently of MAX3SAT$_{\frac{1}{m}}$) and reduces it to an instance ϕ' of MAX3SAT$_\varepsilon$, for some constant ε. In this reduction, ϕ is satisfiable iff ϕ' is satisfiable. The key point is that if ϕ is not satisfiable, then every assignment to the variables of ϕ' leaves at least ε fraction of the clauses of ϕ' unsatisfied, amplifying the gap between the two cases.

The PCP theorem is considered one of the crown jewels of Computational Complexity Theory, and it has far-reaching applications such as inapproximability results (e.g. [Hås01]), verifiable delegation of computation (e.g. [GKR15]), program obfuscation (e.g. [BISW17]) and can even be used in the construction of cryptocurrencies (e.g. [BSCG+14]).

[4] See [BDOT08] for more background and references on this.

[5] A Hamiltonian is frustration-free if there is a state with no energy-penalties from any term.

[6] As commented on in Footnote 1

Recently, there has been a lot of attention regarding the question of whether the quantum version of the PCP theorem holds [AALV09, AAV13]. The quantum PCP conjecture can be stated as follows: the Local Hamiltonian problem is QMA-complete even when the promise gap is ε, for some constant ε independent of the other parameters of the problem[7]. The interest in providing a quantum generalization to the PCP theorem stems also from its implication to our understanding of multipartite entanglement: it is related to the question of stability of entanglement at "room-temperature" (see [AAV13]). Despite much work on this direction [AALV09, Ara11, Has13b, BH13, FH14, AE15, EH17, NVY18], progress on the gapped version of the quantum PCP conjecture had so far been limited. We mention that the game version of quantum PCP theorem was recently proven [NV18], however it is not known to be equivalent to the Hamiltonian version.

What about PCPs for MA? Unfortunately, there is not a lot of research around PCPs for randomized complexity classes. To the best of our knowledge, the only work on this area proves a PCP-like theorem for AM [Dru11], another randomized generalization of NP. In AM, the randomness used by the verifier is public, and therefore known by the Prover. In this case, an instance of an AM-complete problem consists of a collection of boolean formulas $\{\phi_r\}$, and we want to decide if every formula in this family is satisfiable or with high probability a uniformly random formula in this family is not satisfiable. Drucker proved that there exists some constant ε such that the following problem is AM-complete: decide if every formula in the family is satisfiable or with high probability every assignment to the variables of a uniformly random formula leaves an ε fraction of the clauses unsatisfied. It is unclear how to carry this result over to MA.

1.3 Our result

In this work, we connect the derandomization of MA to the question of quantum PCP in the restricted setting of stoquastic Hamiltonians, which can be viewed as a variant of PCP for MA. To this end, we prove that the gapped version of the uniform stoquastic Hamiltonian problem (under certain restrictions) is in NP. Let us first define the problem at hand with slightly more detail, state the result, and then clarify the above sentence.

Informal Definition: (The Gapped, Uniform, Stoquastic, Frustration-Free, Local Hamiltonian problem.) For some constant ε, independent of all other parameters of the problem, the input to the gapped uniform stoquastic frustration-free k-Local Hamiltonian problem is a set of m uniform stoquastic Hamiltonian terms $H_1, \ldots, H_m$ where each H_i acts on k qudits out of the n qudit system and $\|H_i\| \leq 1$. We also have that at most d terms act non-trivially on each qudit, for some constant d. For $H = \frac{1}{m}\sum_{j=1}^{m} H_j$, we have to decide which of one of the following two conditions hold, given the promise that one is true.
Yes. There exists a quantum state $|\psi\rangle$ such that $\langle\psi|H|\psi\rangle = 0$.
No. For all quantum states $|\psi\rangle$ it holds that $\langle\psi|H|\psi\rangle \geq \varepsilon$.

Thus, we are given a uniform stoquastic k-local Hamiltonian with bounded degree, and it is promised that this input Hamiltonian is either frustration-free or constantly frustrated, i.e., any state has energy at least ε, for some constant ε. Our main result is that this problem is in NP.

[7] Other versions of quantum PCPs have also been considered [AAV13, NV18], but they are not (directly) related to this work.

Theorem 1.1 (**main: uniform stoquastic frustration-free Local Hamiltonian is in NP**). *For any constant $\varepsilon > 0$, the problem of deciding whether a uniform stoquastic Hamiltonian H is frustration-free or ε-frustrated, is in* NP.

We note that the same problem, except with inverse polynomial gap, is MA-complete. Hence, first of all, this provides a new tighter upper-bound on the problem, when the (average) promise gap is constant. This is interesting first of all for the study of stoquastic Hamiltonians and the hardness of deciding their ground energies and groundstates, per se.

We now study the implications from the complexity-theoretical point of view. In this context, our result implies that a PCP-like theorem for MA (or, more precisely, for the stoquastic Hamiltonian characterization of MA, under the uniformity and bounded degree conditions as in Definition 3.13), would imply that MA = NP.

Corollary 1.2 (uniform stoquastic PCP theorem implies derandomization of MA). *If there exists a constant ε such that solving the gapped uniform frustration free stoquastic Hamiltonian problem with ε gap is* MA*-hard, then* MA = NP.

From an optimistic perspective, by the above Corollary, our result opens a way towards proving that MA = NP via quantum arguments, in particular by proving specific types of quantum PCPs or quantum gap amplification procedures. This path could of course be very hard, but under the belief that MA = NP, such a gap amplification procedure is in fact *known* to exist. Again, we can look at this result in analogy with the famous Hardness vs. Randomness results [NW94]: our result shows that the problem of stoquastic quantum PCP, under the above listed restrictions, is *equivalent* to proving derandomization of MA. Taking the opposite point of view, and assuming the less commonly believed assumption that MA is strictly larger than NP, our work proves that no PCP exists for stoquastic local Hamiltonians, or, loosely speaking, there is no PCP for MA.

The results regarding the class MA [BT09] can be "scaled down" to the class co-RP (one-sided-error probabilistic polynomial time computations) by slightly modifying the uniform stoquastic Local Hamiltonian problem. In this case, a PCP-like theorem for MA, with some additional natural requirements, would also imply that P = co-RP, and since P is closed under complement, P = RP. We discuss this in more detail in Appendix A.

Corollary 1.3 (uniform stoquastic PCP theorem implies derandomization of RP). *If, in addition to the existence of a gap amplification procedure as in Corollary 1.2, there is a polynomial time algorithm that maps a witness of the original problem into a witness of the gapped problem, then* P = RP.

As a small side result, we show that the *commuting* version of stoquastic Hamiltonian problem is in NP (for any promise gap.)

Theorem 1.4 (commmuting stoquastic Local Hamiltonian problem is in NP). *The problem of deciding if a commuting stoquastic Hamiltonian H is frustration-free is in* NP.

1.4 Proof overview and main ideas

We prove Theorem 1.1 by derandomizing the verification algorithm used by Bravyi and Terhal [BT09] in their proof of the containment in MA of the inverse-polynomial version of the stoquastic Hamiltonian problem. The derandomization becomes possible when the gap is constant,

namely, when we know that the Hamiltonian is either frustration free or there is a large amount of frustration.

We briefly explain now the main ideas behind the randomized verification procedure of [BT09], using a random walk; then we overview our approach to derandomize it.

Bravyi and Terhal started by defining an (exponential-size) graph whose vertices are all possible n-bit strings. The edges are defined based on the stoquastic Hamiltonian: two strings x, y are adjacent in the graph, iff they are connected by some H_i, one of the local terms of the Hamiltonian, namely, if x and y appear together in some groundstate of H_i. The paper considers the following random-walk on the graph: starting from a given n-bit string, pick one of the terms uniformly at random, and go to any of the (constantly many) strings connected to the current string by that term, uniformly at random. This is called a *step*. In the non-uniform case, there are weights involved and the random-walk becomes more complicated, but here we focus only on the uniform case. Bravyi and Terhal also define the notion of a bad string, which is a string that does not appear in the support of any of the groundstates of some local term.

Bravyi and Terhal then showed that if the stoquastic Hamiltonian is frustration-free, then the connected component of any string in the support of some groundstate of the Hamiltonian, does not contain bad strings. In particular, any walk on the above defined graph, starting from some string in a groundstate, does not reach bad strings. On the other hand, if the Hamiltonian is at least $\frac{1}{p(n)}$ frustrated, for some polynomial p, then there exists some polynomial q such that a $q(n)$-step random walk starting from any initial string reaches a bad string with high probability. The MA verification algorithm then proceeds by the Prover sending some x, which is supposed to lie in the support of some groundstate of the Hamiltonian and the verifier performs a $q(n)$-step random walk starting from x, as above. The algorithm rejects if a bad string is encountered in the random-walk.

Our main technical result is showing that if the Hamiltonian is ε frustrated, for some constant ε independent of n, then from any initial string it is possible to reach a bad string in r steps, where r is a constant that only depends on ε, k and d. Therefore, we can define an NP verification algorithm which, given some initial string x, tries all possible r-size paths, and this can be performed in polynomial time since r is constant. We describe now the main ideas on how to prove that for highly frustrated Hamiltonians, such a short path always exists.

Our proof is based on the following two key ideas. First, we notice that if we start with any initial quantum state which is a uniform superposition of good strings, then in case the Hamiltonian is highly frustrated, there must be a term H_i which has large energy with respect to that string (in fact, there must be many, but for now we focus on one). When we apply on the state the projection $\tilde{P}_i$ onto the groundspace of that frustrated term H_i, then it is not very difficult to see that the number of strings in the support of the new state, after this projection, will be larger by a constant factor. Moreover, the value of this expansion factor is directly related to the amount of frustration of H_i with respect to the state we started with. In other words, the more frustrated the term is, the larger the expansion of the set of strings would be, due to projection with respect to that term. We call this phenomenon "one term expansion"; it is proven in Lemma 4.3.

Now, the idea is to start with one good string given to the verifier by the prover, and expand it to an increasingly larger set of good strings by such projections. Our goal is to perform such expansions by a "circuit of parallel non-overlapping[8] projections", as in Figure 1a. We would like to argue that if the frustration of the Hamiltonian is high for any state, as we are assuming now, then there is a constant fraction of the m projections, given by one layer in the circuit, which are

[8]Two terms are non-overlapping if the sets of qudits on which they act are disjoint

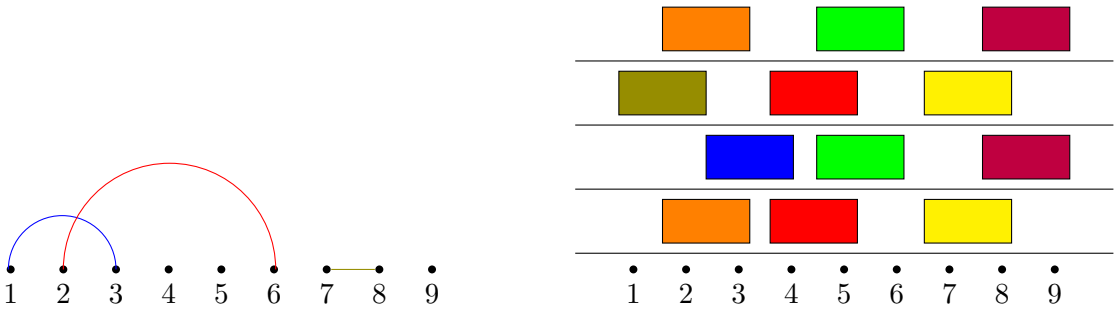

(a) (b)

Figure 1: In Figure 1a, we show an example of non-overlapping 2-local terms, where each term corresponds to an edge with different color. In Figure 1b, we depict a constant-depth circuit where each layer contains non-overlapping 2-local terms (the terms act on neighbor particles in order to simplify the drawing).

all at least constantly frustrated. By the single term expansion argument, each such term would contribute a constant multiplicative factor to the number of good strings in our set, and thus the size of the set of good strings accumulates an exponential factor due to each *layer* in the circuit. If this is true, then it must be the case that after at most constantly many layers, the argument breaks down (namely, a bad string is found) since otherwise the number of strings would just be larger than the number of all possible strings. The implication is that after constantly many layers, a bad string is reached.

Unfortunately, there is a problem in applying the above line of thought directly. The problem is that the amount of expansion of two different terms might be strongly correlated. Let us see an example of such correlation.

Example 1. Let $S = \{0000, 0011, 1100, 1111\}$ and let $P_{1,4} = P_{2,3} = |\Phi^+\rangle\langle\Phi^+| + |\Psi^+\rangle\langle\Psi^+|$, where $P_{i,j}$ acts on qubits i and j (and are implicitly tensored with identity on the other qubits), and $|\Phi^+\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$ $|\Psi^+\rangle = \frac{1}{\sqrt{2}}(|01\rangle + |10\rangle)$ are two of the Bell states. We notice that

$$\langle S|P_{1,4}|S\rangle = \frac{1}{2}.$$

and the same holds for $P_{2,3}$.

However, if we take the support of $P_{1,4}|S\rangle$, $S' = \{0000, 0110, 0011, 0101, 1100, 1010, 1111, 1001\}$, it follows that $\langle S'|P_{2,3}|S'\rangle = 0$, so $P_{2,3}$ has no frustration after we correct the frustration of $P_{1,4}$.

This example means that we cannot use the above argument as stated, for many non-overlapping terms applied in parallel: even though there are indeed a linear number of non-overlapping terms which are all frustrated, we cannot simply multiply the expansions due to each of them.

We overcome this difficulty by resorting to some "online" version of the claim: it turns out that by using an adaptive argument, a constant fraction of terms can be found which will all contribute independent multiplicative factors to the increase in size of the set of good strings. This means that each layer in the circuit of non-overlapping parallel projections does contribute an exponential increase in the number of strings.

We summarize the first part of the proof: in the case of frustration, assume we start with a subset-state of good strings $|S\rangle$, and let $L|S\rangle$ be the state which we arrive at, after applying all projections in the sequence L which we have found above, and taking the subset-state of all strings we have reached. We can show that $L|S\rangle$ contains $\left(1+\frac{\varepsilon}{4}\right)^{\frac{\varepsilon n}{2kd}}$ more strings than S. Then, we repeat this constantly many times. More concretely, set $S_0 = \{x\}$, for some initial string x. The above argument shows that either $|S_0\rangle$ contains a bad string (i.e. x is a bad string), or there is a set of terms L_1 such that the set S_1 with the strings in the support of $L_1|S_0\rangle$ has exponentially more strings than S_0. We now repeat this process starting with the state $|S_1\rangle$ instead of $|S_0\rangle$, and so on, until we reach a bad string. Since the number of strings in the set increases exponentially at every step, there exists some constant ℓ, that depends only on ε, k and d, such that S_ℓ (which we prove to be the strings in the support of $L_\ell...L_1|x\rangle$) contains a bad string. This shows that a constant depth circuit of non-overlapping projections, applied to an input string, leads to a bad string. We depict this constant depth circuit in Figure 1b. The proof of the claim that within constantly many layers a bad string is reached, is given in Lemma 4.9.

We notice that such a constant depth circuit implies that a bad string can be found within a constant number of rounds, where each round consists of a set of local steps, each changing a different local part of the string. However the brute-force search of such a path is intractable, since the number of steps in each round might be polynomial, and thus the number of possible paths is exponential.

The next part of the proof is where we show how to find a bad string efficiently, given that one is reached in the above constant depth projection circuit. We call this "the light cone argument'. To retrieve a constant size path from $L_\ell...L_1|x\rangle$, the key point is noticing that badness of a string is a *local* property, namely, if a string is bad, we can point at at least one local term which it is bad for; let us refer to this term as the frustrated term (this is a meaningful name since if a state contains that bad string, then that term will indeed be frustrated). The crux of the matter is that the fact that badness of a string is local, implies also that projections on one set of qubits does not affect the badness of terms restricted to the complementary set of qudits (see Claim 3.5). This implies, by a simple argument, that even if we remove all of the terms in $L_1, ..., L_\ell$ that are not in the *lightcone* of the frustrated term, we will still achieve a bad string. This is because if $L_\ell...L_1|x\rangle$ contains a bad string, which is bad for some term, then any projection in $L_\ell...L_1$ which is not in the light cone of that term, cannot influence its badness. We depict such argument in Figure 2.

Using the light-cone argument, we can deduce that instead of applying the layers $L_\ell...L_1$, we can apply just the terms in these layers which are contained in the lightcone; we denote them by $L_1^\triangle$, ..., $L_\ell^\triangle$, to arrive at the conclusion that also the state $L_\ell^\triangle...L_1^\triangle|x\rangle$ contains a bad string. Since the lightcone operators are of constant size, every string in $L_i^\triangle...L_1^\triangle|x\rangle$ can be reached from a string in $L_{i-1}^\triangle...L_1^\triangle|x\rangle$ in a constant number of steps, and by induction we deduce that there is a short path from x to a bad string.

1.5 Discussion and open problems

In this section we discuss the relation with related work, the implications of our result and state some open problems.

PCP for AM vs. PCP for MA. PCP theorem for AM proved by Drucker [Dru11] relies strongly on the classical PCP theorem: since the randomness is public, both Prover and Verifier

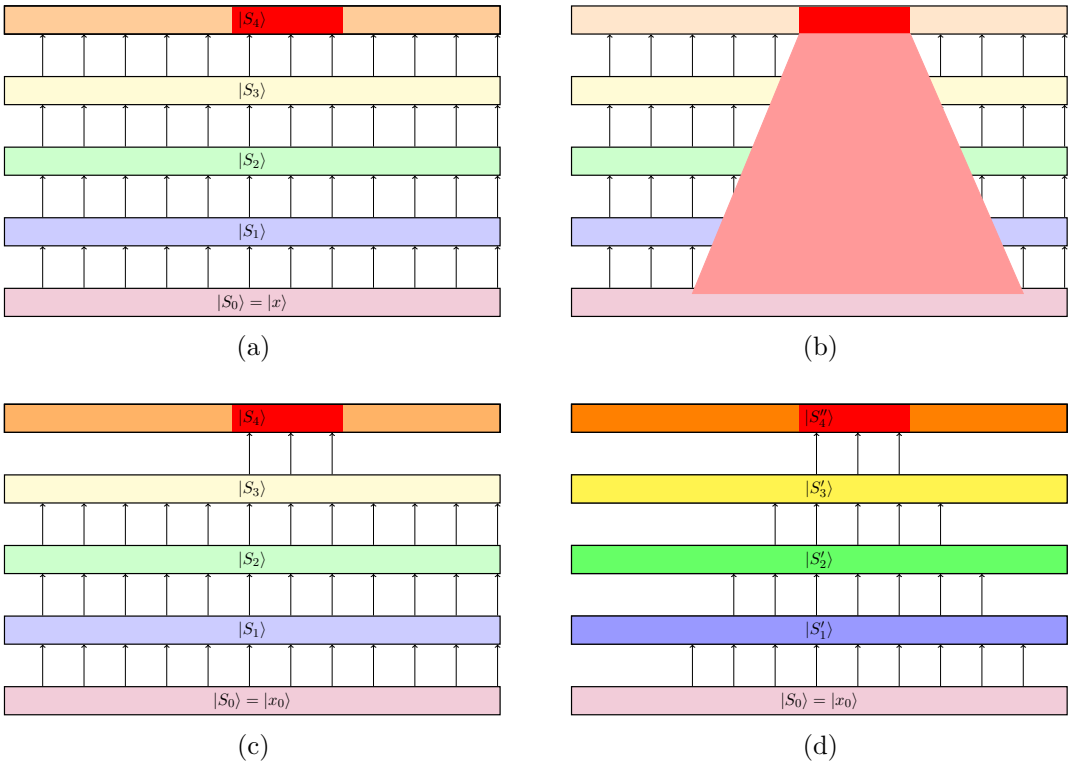

Figure 2: Example of application of the light-cone argument in order to find a constant-step path from the initial string to a bad string. The red rectangle marks the qudits of frustrated by a bad string. In Figure 2a, we have a constant number of layers (in the sense of Figure 1b) that reach a state with a bad string. In Figure 2b, we show the light-cone from the frustrated term. In Figure 2c, we remove the projections outside of the light-cone on the last layer. Notice that the state on the last layer has changed, which is depicted by the change of transparency, but it also contains a bad string. In Figure 2d, we show the projections that are left after removing all the terms outside of the light-cone. Each layer contains different states from the original ones, again depicted by the change of transparency, but the last layer still contains a bad string.

agree on the same boolean formula, and then they can apply the original PCP theorem with this formula. In the case of MA, such argument does not hold since the Prover does not know the formula that will be tested by the Verifier.

PCP for QMA vs. PCP for MA. We notice the asymmetry between the quantum and the uniform stoquastic PCP conjecture. It is widely believed that NP = MA, and therefore the uniform stoquastic Local Hamiltonian problem is believed to be both in NP, which we prove in this work, and MA-hard, which is the uniform stoquastic PCP conjecture. In this sense, the result that we prove is somehow "expected", and it does not change our belief that the stoquastic PCP conjecture does hold (as it is implied by BPP = P). However, in the fully quantum setting, we believe that NP $\neq$ QMA, and if the constant promise gap Local Hamiltonian problem is in NP, then this is a strong indication that the quantum PCP conjecture is false[9]. Our result implies also that if one expects to prove the quantum PCP conjecture without causing the extra "side-effect" of proving NP = MA, the gap-amplification process should not maintain uniform-stoquasticity.

Detectability lemma. Our setting resembles that of the Detectability lemma (DL) [AALV09], a useful tool in quantum Hamiltonian complexity [AALV11, ALV12, GH16, ALVV17]. Like in our setting, in the DL setting (see the formulation of [AAV16]) one considers a given local Hamiltonian, where each term is associated with a local projector on its groundspace. Starting with some state, one considers applying the local projectors one by one (in an arbitrary order). Under our assumption by which every qudit participates in at most constantly many local terms, this can be viewed as applying all local projections, organized in a constant depth circuit made of local projections - very much like in our setting. If the state we start with is the groundstate, and there is no-frustration in the Hamiltonian, then the norm of the state after all these projections of course remains one; this is the easy part. The DL says that if the Hamiltonian is frustrated, then the norm of the state after all these projections will have shrunk by at least some factor; the key in the lemma is to upper bound that factor. When the Hamiltonian is highly frustrated, the DL says that the factor will be a constant strictly less than 1. This is a strong statement; could it possibly be used to deduce our result? A closer look reveals important differences between the two questions. While our goal is to argue containment in NP, and thus we need an efficient classical witness (which we take as an n-dit string), the DL requires full knowledge of the quantum state on which the projections are applied, in order to deduce the behavior of the norm. We thus do not know how to make any usage of the DL in our setting, though interestingly, there seem to be some conceptual connections; In particular, like in the current paper, the proof of the DL relies on arguing that the correlations between the different projections do not matter. It would be interesting if more can be said in this direction.

Uniform vs. non-uniform case. It is very natural to try to extend our result to general (i.e. non-uniform) stoquastic frustration-free Hamiltonians; removing the uniformity restriction seems conceptually important, even though the uniform case is already MA hard. Unfortunately we do not know how to do this, and this remains for future work. The main difference between the uniform and the general case is that in the uniform case, the only source of frustration is the existence of

[9]Indeed, proving that constant promise gap Local Hamiltonian problem is in NP is often considered as disproving the quantum PCP conjecture.

bad strings. When we go to the non-uniform case, then the frustration might also appear due to amplitude inconsistency. Let us see a simple example of this.

Example 2. Let us consider a one-qubit system, and the Hamiltonian consists in the sum of two terms whose groundstate projectors are

$$P_1 = \frac{1}{2}\left(|0\rangle + |1\rangle\right)\left(\langle 0| + \langle 1|\right) \text{ and } P_2 = \left(\sqrt{1-\varepsilon}|0\rangle + \sqrt{\varepsilon}|1\rangle\right)\left(\sqrt{1-\varepsilon}\langle 0| + \sqrt{\varepsilon}\langle 1|\right),$$

for a small ε.

We notice that the Hamiltonian is frustrated but there are no bad strings! The source of the frustration is the fact that the first terms requires the amplitude of the groundstate to be the same, but on the other hand the second term pushes them far apart.

Bravyi and Terhal deal with this problem in their random walk by assigning weights to the edges, where the weights depend on the Hamiltonian term connecting the two strings. Then, frustration implies that the weight of different paths between two pairs of strings in the support of the groundstate, have different weights (a weight of a path is the product of the weights of the edges). Bravyi and Terhal then add extra tests to find these inconsistent paths. At every step of the random walk, the verifier rejects if the weight of the path is larger than one (for this the verifier should start the random walk from the string with maximal amplitude in the groundstate). They prove that if the verifier is provided a string whose amplitude was not maximal or that inconsistent paths exist, then with high probability a random walk finds a path whose weight is larger than one, leading to rejection.

Interestingly, our proof goes through almost all the way, also for the non-uniform case: we can even prove that inconsistency is achieved within constantly many layers in the projection circuit, exactly as in the uniform case. The only problem preventing us from extending the proof to the non-uniform case is the light-cone lemma, which does not seem to hold when attempting to detect inconsistencies rather than reachability of bad strings. In other words, we do not know how to find the inconsistency efficiently in such a constant depth projection circuit, even though we know it exists! A different way to say it is that as far as we can tell, the random-walk proposed by Bravyi and Terhal deals with such cases in a non-local way. The situation seems to be related to the hardness of sampling from constant depth quantum circuits [TD04, BGK18, CSV18, Gal18], but we do not know that the obstacle cannot be overcome.

Such an extension would of course result in a stronger statement, saying that stoquastic PCPs implies MA = NP, without the requirement on uniformity; and of course that the gapped stoquastic Hamiltonian problem is in NP, without the uniformity restriction. We view this seemingly technical problem as very interesting, as it might point at some interesting difference between uniform and non-uniform stoquastic Hamiltonians.

Perfect completeness. Our result strongly relies on the fact that MA can be defined with perfect completeness. It is an important problem to clarify whether any of the insights emerging from this work can be useful in other settings. Bravyi [Bra14] defined a variant of the Local Hamiltonian problem, called *guided stoquastic Hamiltonian problem* and proved it to be MA-complete; the proof does not require perfect-completeness. Could a gapped version of this problem be proven to be in NP? Another set of problems for which we can ask the same question are those related to StoqMA. As mentioned previously, the stoquastic Local Hamiltonian problem where we have to decide if the

groundstate energy is below some threshold α or above another threshold β is StoqMA-complete for inverse polynomial $\beta - \alpha$. In fact, the StoqMA-completeness holds even for a restricted class of stoquastic Hamiltonians: the Transverse-field Ising mode [BH17]. We leave as an open question if any these problems with constant promise gap is also in NP. Notice that this, together with some PCP theorem along the lines of Corollary 1.3 would in fact imply P = BPP.

Adiabatic evolution of Hamiltonians. Bravyi and Terhal used their random walk for stoquastic Hamiltonians to prove that the adiabatic evolution of frustration-free stoquastic Hamiltonians with inverse polynomial spectral-gap can be performed in randomized polynomial time. A major open problem remains to extend their result to the general case, in which the frustration-free assumption is relaxed. This would lead to a classical simulation of adiabatic optimization, e.g., of D-Wave type algorithms [FGGS00, BT09, Has13a, CCD15]. We leave as an open question whether our techniques can have any implications in that context.

Organization of the paper We start with some preliminaries in Section 2. We discuss stoquastic Hamiltonians and the proof of MA-completeness in Section 3. Our main result is proven in Section 4. We finish by proving that commuting stoquastic Hamiltonians are in NP in Section 5.

Acknowledgments

The authors thank Ayal Green for the helpful discussion at early stages of this work. DA is also grateful for a very short discussion with Noam Nisan which was as insightful as much as it was short, as well as for very useful remarks from Avi Wigderson. We also thank Henry Yuen for the help on improving the clarity of the paper. AG is supported by ERC Consolidator Grant 615307-QPROGRESS. DA's research on this project was supported by ERC grant 280157 and ISF grant 1721/17. Part of this work was done when AG was a member of IRIF, Université Paris Diderot, Paris, France, where he was supported by ERC QCC. The authors thank also the French-Israeli Laboratory on Foundations of Computer Science (FILOFOCS) and Simons collaboration grant number 385590 that allowed AG to visit Hebrew University of Jerusalem.

2 Preliminaries and Notations

2.1 Complexity classes, NP and MA

A (promise) problem $A = (A_{yes}, A_{no})$ consists of two non-intersecting sets $A_{yes}, A_{no} \subseteq \{0,1\}^*$. We define now the main complexity classes that are considered in this work. We start by formally defining the well-known class NP.

Definition 2.1 (NP)**.** A problem $A = (A_{\text{yes}}, A_{\text{no}})$ is in NP if and only if there exist a polynomial p and a deterministic algorithm D, where D takes as input a string $x \in \Sigma^*$ and a $p(|x|)$-bit witness y and decides on acceptance or rejection of x such that:
Completeness. If $x \in A_{\text{yes}}$, then there exists a witness y such that D accepts (x, y).
Soundness. If $x \in A_{\text{no}}$, then for any witness y, D rejects (x, y).

We can then generalize this notion, by giving the verification algorithm the power of flip random coins, leading to the complexity class MA.

Definition 2.2 (MA)**.** A problem $A = (A_{\text{yes}}, A_{\text{no}})$ is in MA if and only if there exist a polynomial p and a probabilistic algorithm R, where R takes as input a string $x \in \Sigma^*$ and a $p(|x|)$-bit witness y and decides on acceptance or rejection of x such that:
Completeness. If $x \in A_{\text{yes}}$, then there exists a witness y such that R accepts (x, y) with probability 1.
Soundness. If $x \in A_{\text{no}}$, then for any witness y, R accepts (x, y) with probability at most $\frac{1}{3}$.

The usual definition of MA requires yes-instances to be accepted with probability at least $\frac{2}{3}$, but it has been shown that there is no change in the computational power if we require the verification algorithm to always accept yes-instances [ZF87, GZ11].

2.2 Quantum states

We review now the concepts and notation of Quantum Computation that are used in this work. We refer to Ref. [NC00] for a detailed introduction of these topics.

Let $\Sigma = \{0, ..., q-1\}$ be some alphabet. A qudit of dimension q is associated with the Hilbert space $\mathbb{C}^\Sigma$, whose canonical (also called computational) basis is $\{|i\rangle\}_{i\in\Sigma}$. A pure quantum state of n qudits of dimension q is a unit vector in the Hilbert space $\left(\mathbb{C}^\Sigma\right)^{\otimes n}$, where $\otimes$ is the Kroeneker (or tensor) product. The basis for such Hilbert space is $\{|i\rangle\}_{i\in\Sigma^n}$. For some quantum state $|\psi\rangle$, we denote $\langle\psi|$ as its conjugate transpose. The inner product between two vectors $|\psi\rangle$ and $|\phi\rangle$ is denoted by $\langle\psi|\phi\rangle$ and their outer product as $|\psi\rangle\langle\phi|$. For a vector $|\psi\rangle \in \mathbb{C}^{|\Sigma|^n}$, its 2-norm is defined as $\||\psi\rangle\| := \left(\sum_{i\in\Sigma^n} |\langle\psi|i\rangle|^2\right)^{\frac{1}{2}}$.

We now introduce some notation which is somewhat less commonly used and more specific for this paper: the support of $|\psi\rangle$, $supp(|\psi\rangle) = \{i \in \Sigma^n : \langle\psi|i\rangle \neq 0\}$, is the set strings with non-zero amplitude. We call quantum state $|\psi\rangle$ *non-negative* if $\langle i|\psi\rangle \geq 0$ for all $i \in \Sigma^n$. For any $S \subseteq \Sigma^n$, we define the state $|S\rangle := \frac{1}{\sqrt{S}} \sum_{i\in S} |i\rangle$ as the subset-state corresponding to the set S [Wat00]. For a non-negative state $|\psi\rangle$, we define $|\widehat{\psi}\rangle := |supp(|\psi\rangle)\rangle$ as the subset-state induced by the strings in the support of $|\psi\rangle$. We say that this is the subset state corresponding to the state $|\psi\rangle$. Analogously, for some linear operator P the state $\widehat{P|\psi\rangle}$ means the subset-state corresponding to the state $P|\psi\rangle$.

2.3 Hamiltonians, Groundstates, Energies, Frustration

Definition 2.3 (Hamiltonian)**.** A *Hamiltonian* on n qudits is a Hermitian operator on $\mathbb{C}^{|\Sigma|^n}$, namely, a complex Hermitian matrix of dimension $|\Sigma|^n \times |\Sigma|^n$. A Hamiltonian on n qudits is called k-*Local* if it can be written as $H = \sum_{i=1}^m \tilde{H}_i$, where each $\tilde{H}_i$ can be written in the form $\tilde{H}_i = H_i \otimes I$, where H_i acts on at most k out of the n qudits.

Hamiltonians describe the evolution of physical systems, using Schrodinger's equation. Their eigenvalues correspond to the energy of the system; more generally, the energy of a state $|\psi\rangle$ with respect to a Hamiltonian $H = \frac{1}{m}\sum_{i=1}^m H_i$ is given by $\langle\psi|H|\psi\rangle$. Notice that we use the term energy even though we average by the number of terms, so this is the average energy *per term*; this is different from the usual usage of the term energy, or energy density, in the physics literature, where one usually considers the average energy *per particle*. This normalization is more convenient in the context of PCPs [Din07]. We also consider the energy of the state with respect to a specific term H_i, which is $\langle\psi|H_i|\psi\rangle$. The minimal energy is the smallest eigenvalue of the Hamiltonian, and an eigenstate which has this energy is called a *groundstate*.

Definition 2.4 (Groundstate, groundspace, frustration and frustration-free)**.** A *groundstate* of a Hamiltonian $H = \frac{1}{m}\sum_{i=1}^{m} H_i$ is an eigenvector associated with its minimum eigenvalue, which is called the *groundstate energy*. The *groundspace* of a Hamiltonian is the subspace spanned by its groundstates. H is called ε-frustrated if for every state $|\psi\rangle$, $\frac{1}{m}\sum_i \langle\psi|H_i|\psi\rangle \geq \varepsilon$. Finally, H is called *frustration-free* if there exists some $|\psi\rangle$ such that for every i, the local term H_i is positive definite and $\langle\psi|H_i|\psi\rangle = 0$.

Throughout this paper we use the following notation for a local Hamiltonian $H = \frac{1}{m}\sum_{i=1}^{m} \tilde{H}_i$. We set P_i to be the *local* projection on the groundspace of H_i; while $\tilde{P}_i = P_i \otimes I$ corresponds to the projection on the groundspace of $\tilde{H}_i$.

We prove now a useful lower-bound on the number of frustrated terms of a highly frustrated Hamiltonian.

Claim 2.5 (Lower bound on frustrated terms)**.** *Let $H = \frac{1}{m}\sum_{i=1}^{m} \tilde{H}_i$ be a Local Hamiltonian that is ε-frustrated. Then for every state $|\psi\rangle$, there exist at least $\frac{\varepsilon m}{2}$ terms that are at least $\frac{\varepsilon}{2}$ frustrated.*

Proof. We prove this by contradiction. Let $F = \{i : \langle\psi|\tilde{H}_i|\psi\rangle \geq \frac{\varepsilon}{2}\}$. We assume then that $|F| < \frac{\varepsilon m}{2}$. Then the energy of the state is

$$\begin{aligned}
&\langle\psi|H|\psi\rangle \\
&= \frac{1}{m}\left(\sum_{i\in F}\langle\psi|\tilde{H}_j|\psi\rangle + \sum_{i:\notin F}\langle\psi|\tilde{H}_j|\psi\rangle\right) \\
&\leq \frac{|F|}{m} + (m - |F|)\frac{\varepsilon}{2m} \\
&< \frac{\varepsilon}{2} + \frac{\varepsilon}{2} \\
&= \varepsilon,
\end{aligned}$$

where in the first inequality we used the fact that the norm of all terms is at most 1, and that the terms outside of F contribute at most $\frac{\varepsilon}{2m}$ to the above sum by definition of the set F. In the second inequality we used our assumption that $|F| < \frac{\varepsilon m}{2}$. We then have that H is not ε-frustrated, which is a contradiction. □

3 Background: the Stoquastic Hamiltonian problem

In this section, we define stoquastic Hamiltonians, prove certain basic properties, as well as state their relation to the complexity class MA.

3.1 Stoquastic Hamiltonians

In this work, we deal with a special type of Hamiltonians, which are called stoquastic.

Definition 3.1 (Stoquastic Hamiltonian [BDOT08])**.** A k-Local Hamiltonian $H = \sum_{i=1}^{m} \tilde{H}_i$ is called *stoquastic* in the computational basis if for all i, the off-diagonal elements of H_i (the local terms) in this basis are non-positive[10].

[10]Klassen and Terhal [KT18] have a different nomenclature. They call a matrix Z-symmetric if the off-diagonal elements of the local terms are non-positive and they call a Hamiltonian stoquastic if all local terms can be made Z-symmetric by local rotations.

As remarked in [MLH18], every Hamiltonian is stoquastic in the basis that diagonalizes it. However, the length of such description might be exponential in the number of qubits, since it may be impossible to write it as a sum of local terms. Some recent works [MLH18, KT18], provide evidence that deciding if a given local Hamiltonian can be made stoquastic by local basis change is computationally hard. Therefore, in our definition we assume that the stoquastic Hamiltonian is *given* in the basis where each of the *local* terms is stoquastic, i.e., has non-positive off-diagonal elements.

A property of a stoquastic local Hamiltonian is that the groundspace of the local terms can be decomposed in a sum of orthogonal non-negative rank-1 projectors.

Lemma 3.2 (Groundspace of stoquastic Hamiltonians, Proposition 4.1 of [BT09])**.** *Let H be a stoquastic Hamiltonian and let P be the projector onto its groundspace. It follows that*

$$P = \sum_j |\phi_j\rangle\langle\phi_j|, \tag{1}$$

where for all j, $|\phi_j\rangle$ is non-negative and for $j \neq j'$, $\langle\phi_{j'}|\phi_j\rangle = 0$.

Proof. We start by showing that if all of the entries of P are non-negative, then the statement holds. Let x, y, z be some strings such that $\langle x|P|y\rangle > 0$ and $\langle y|P|z\rangle > 0$. Then

$$\langle x|P|z\rangle = \langle x|P^2|z\rangle = \sum_w \langle x|P|w\rangle\langle w|P|z\rangle > \langle x|P|y\rangle\langle y|P|z\rangle > 0,$$

where in the first inequality we use the fact that P has only non-negative entries. Therefore, we can partition the string in equivalent classes $T_1, ...T_t$ regarding the property $\langle x|P|y\rangle > 0$.

It follows that the subspace spanned by the strings in T_i is P-invariant and therefore P is block-diagonal with respect to the direct sum of such subspaces. Using the Perron-Frobenius theorem for each of the blocks, we have that its largest eigenvalue is non-degenerate, and in this case the block is rank-one, since P is a projector with eigenvalues 1 and 0. Since all the entries of P are non-negative, then each one of these rank-one blocks correspond to a non-negative state.

This finishes the proof of the case where P has non-negative entries. We show now that this property holds for stoquastic Hamiltonians.

We have that the groundspace projector P of the Hamiltonian consists of the Gibbs state for temperature tending to 0, i.e., $P = \lim_{\beta\to\infty} q\frac{e^{-\beta H}}{Tr(e^{-\beta H})}$, where $q > 0$ is the dimension of the groundspace (This is a well known easy fact, see for example Proposition 4.1 in [BT09]). Thus, it suffices to prove that $e^{-\beta H}$ is a matrix of non-negative entries (we already know that the trace is non-negative by the fact that the eigenvalues of $e^{-\beta H}$ are positive).

Let s be some value such that $-\beta H + sI$ has only non-negative entries. Write

$$e^{-\beta H} = e^{-\beta H + sI - sI} = e^{-\beta H + sI}e^{-sI},$$

where the last equality holds because $-sI$ and $(-\beta H + sI)$ commute.

Note that the Taylor expansion of e^A is

$$e^A = \sum_{i=0}^{\infty} \frac{A^k}{k!}.$$

Thus we have that the entries of $e^{-\beta H + sI}$ are non-negative. Write $e^{-s} = p > 0$, and we have that all entries of $e^{-\beta H} = pe^{-\beta H + sI}$ are non-negative as well. □

Definition 3.3. (Stoquastic Projector) Given a projection matrix P acting on k qudits, if there exists a set of orthogonal k-qudit non-negative states $\{|\Phi_j\rangle\}_j$, such that

$$P = \sum_j |\Phi_j\rangle\langle\Phi_j|,$$

we say P is a stoquastic projector, and we refer to this unique decomposition as a sum of projections to non-negative states as the *non-negative decomposition* of P.

Remark 3.4. Notice that Lemma 3.2 implies that the projection on the groundspace of a stoquastic Hamiltonian is a stoquastic projector.

A crucial point in the paper is the fact that when applying a local stoquastic projector P on some set of qudits Q, we do not introduce new strings in the reduced density matrix of the set of qudits outside of Q. Notice that since we are considering non-unitary operators, namely projections, then even though they are local, such projections *can* in fact have the effect of *removing* strings away from the density matrices of qubits which they do not touch; the point of this claim is that they cannot *add* new strings away from where they act.

Claim 3.5 (Local action of projectors)**.** *Let P be a stoquastic projector on a subset Q of $k \leq n$ qudits. Consider the projection $\tilde{P}$ on n qudits derived from P by $\tilde{P} = P_Q \otimes I_{\overline{Q}}$. Then $\tilde{P}$ is also a stoquastic projector, it can be written as the following non-negative decomposition:*

$$\tilde{P} = \sum_{z\in\Sigma^{n-k},j} |\phi_j\rangle\langle\phi_j|_Q \otimes |z\rangle\langle z|_{\overline{Q}}, \tag{2}$$

where $\{|\phi_j\rangle\}_j$ is the non-negative decomposition of P, and moreover, for any non-negative state $|\psi\rangle$, we have:

$$\{x_{\overline{Q}} : x \in supp(\tilde{P}|\psi\rangle)\} \subseteq \{x_{\overline{Q}} : x \in supp(|\psi\rangle)\}.$$

Proof. For the first part of the claim, the fact that $\tilde{P}_i = P_i \otimes I$ implies that it can be written in the desired form, and this implies that it is a stoquastic projector.

For the moreover part, write $|\psi\rangle = \sum_{x\in\Sigma^n} \alpha_x|x\rangle$. We have:

$$\tilde{P}|\psi\rangle = \sum_{x\in supp(\psi)} \sum_j \sum_{z\in\Sigma^{n-k}} \alpha_x|\phi_j\rangle\langle\phi_j|x_Q\rangle|z\rangle\langle z|_{\overline{Q}}|x_{\overline{Q}}\rangle_{\overline{Q}} = \sum_{x\in supp(\psi)} \sum_j \alpha_x\langle\phi_j|x_Q\rangle|\phi_j\rangle_Q|x_{\overline{Q}}\rangle_{\overline{Q}}.$$

Let $y \in supp(\tilde{P}|\psi\rangle)$; so $\langle y|\tilde{P}|\psi\rangle \neq 0$. By the above expression there must be $x \in supp(|\psi\rangle)$ such that $x_{\overline{Q}} = y_{\overline{Q}}$. □

Remark 3.6 (Notation of local and global projectors, and their non-negative decompositions)**.** As can be seen in Claim 3.5, we use the tilde to denote the global projector, i.e., $\tilde{P} = P \otimes I_{\overline{Q}}$. We also extend (in a slightly different way) this notation to the rank-1 projectors and we denote $|\tilde{\phi}_{j,z}\rangle := |\phi_j\rangle_Q|z\rangle_{\overline{Q}}$, for $z \in \Sigma^{n-k}$.

Remark 3.7 (Uniqueness of groundstate containing a string)**.** Consider a stoquastic projector P and its global version, $\tilde{P} = \sum_{j,z} |\tilde{\phi}_{j,z}\rangle\langle\tilde{\phi}_{j,z}|$(this can be done by Claim 3.5 and we use the notation of Remark 3.6). Then for every n-dit string $x \in \Sigma^n$, there exists at most one pair of values j^*, z^* such that $\langle\tilde{\phi}_{j^*,z^*}|x\rangle > 0$. Clearly, a similar uniqueness statement holds for the local stoquastic projector P and its non-negative decomposition, with respect to x being a k-dit string.

We now define a particular type of strings, called *bad* strings, which play a crucial role in our result. Consider a string x such that $\langle x|\tilde{P}_i|x\rangle = 0$; we notice that in this case x cannot belong to any groundstate of $\tilde{P}_i$. This leads to the following definition:

Definition 3.8 (Bad strings [BT09])**.** Given a stoquastic projector P on a set Q of k out of n qudits, we say that a string $x \in \Sigma^n$ is *bad* for P (or P-bad) if $\langle x|\tilde{P}|x\rangle = 0$ for $\tilde{P} = P \otimes I_{\overline{Q}}$. Equivalently, x is bad for P, if $\langle x_Q|P|x_Q\rangle = 0$. This means that x_Q is not in the support of any of the states in the non-negative decomposition, as in Lemma 3.2, of P. If a string is not bad for P, we say it is P-good. Given a local Hamiltonian $H = \frac{1}{m}\sum_i \tilde{H}_i$, and the corresponding stoquastic projectors P_i, We say that x is *bad* for H, or H-bad, if there exists some $i \in [m]$ such that x is bad for P_i.

One other property that we use is that starting with a non-negative state $|\psi\rangle$, and applying $\tilde{P}_i$, the projector onto the groundspace of some local term H_i, maintains all the H_i-good strings in $|\psi\rangle$.

Lemma 3.9 (Strings added by Local Stoquastic Projectors)**.** *Let $|\psi\rangle$ be a non-negative n-qudit state. Consider $P = \sum_j |\phi_j\rangle\langle\phi_j|$ a stoquastic projector and its non-negative decomposition, acting on the subset Q of k qudits out of these n qudits. Then all P-good strings of $|\psi\rangle$ are also in the support of $\tilde{P}|\psi\rangle$, where $\tilde{P} = P \otimes I_{\overline{Q}}$. Moreover, it follows that*

$$supp(\tilde{P}|\psi\rangle) = \bigcup_{j,z:\langle\tilde{\phi}_{j,z}|\psi\rangle>0} supp(|\tilde{\phi}_{j,z}\rangle).$$

Proof. Let S be the support of $|\psi\rangle$ and let also $|\psi\rangle = \sum_{x\in S}\alpha_x|x\rangle$. Let also G and B be the sets of P-good and P-bad strings of n dits, respectively. We have that

$$\tilde{P}|\psi\rangle = \sum_{x\in S\cap G}\alpha_x\tilde{P}|x\rangle + \sum_{x\in S\cap B}\alpha_x\tilde{P}|x\rangle = \sum_{x\in S\cap G}\alpha_x\tilde{P}|x\rangle. \tag{3}$$

We now use Equation (2) from Claim 3.5, to apply the projector $\tilde{P}$. We have

$$\tilde{P}|\psi\rangle = \sum_{x\in S\cap G, j, z\in\Sigma^{n-k}} \alpha_x \left(|\phi_j\rangle\langle\phi_j|_Q \otimes |z\rangle\langle z|_{\overline{Q}}\right)|x_Q\rangle_Q|x_{\overline{Q}}\rangle_{\overline{Q}} = \sum_{x\in S\cap G,j}\alpha_x\langle x_Q|\phi_j\rangle|\phi_j\rangle_Q|x_{\overline{Q}}\rangle_{\overline{Q}} \tag{4}$$

If x is a P-good string, then there exists a (unique) $|\phi_j\rangle$ such that $\langle\phi_j|x_Q\rangle > 0$. Using also that $x \in S$, we see that the amplitude of x in the above expression $\tilde{P}|\psi\rangle$ is non-zero.

For the moreover part, notice that Equation (4) can be written as

$$\tilde{P}|\psi\rangle = \sum_{x\in S\cap G,j,z}\alpha_x\langle x|\tilde{\phi}_{j,z}\rangle|\tilde{\phi}_{j,z}\rangle \tag{5}$$

and the statement holds directly.

□

Corollary 3.10 (Composition of stoquastic projectors)**.** *Let $\tilde{P}_1$ and $\tilde{P}_2$ be two n-qudit stoquastic projectors (which may or may not be global versions of local projections) and let $|\psi\rangle$ be a non-negative state. Then $\widehat{\tilde{P}_1\tilde{P}_2|\psi\rangle} = \widehat{\tilde{P}_1\tilde{P}_2}|\psi\rangle = |supp(\tilde{P}_1\tilde{P}_2|\psi\rangle)\rangle$. Moreover, if $\tilde{P}_1$ and $\tilde{P}_2$ commute, then the above states are also equal to $\widehat{\tilde{P}_2\tilde{P}_1|\psi\rangle} = \widehat{\tilde{P}_2\tilde{P}_1}|\psi\rangle$.*

Proof. First, we claim that for two non-negative states $|\psi\rangle$ and $|\psi'\rangle$, if $supp(|\psi\rangle) = supp(|\psi'\rangle)$, then $supp(\tilde{P}|\psi\rangle) = supp(\tilde{P}|\psi'\rangle)$, for any stoquastic projector $\tilde{P}$. This can be argued as follows. Let y be a string in the support of $\tilde{P}|\psi\rangle$, i.e. $\langle y|\tilde{P}|\psi\rangle \neq 0$, and $|\tilde{\phi}_{j,z}\rangle$ be the unique state in the non-negative decomposition of $\tilde{P}$ that contains y, as stated in Remark 3.7. Since $|\psi\rangle$ and $|\tilde{\phi}_{j,z}\rangle$ are both non-negative states, we have that $\langle y|\tilde{P}|\psi\rangle = \langle y|\tilde{\phi}_{j,z}\rangle\langle\tilde{\phi}_{j,z}|\psi\rangle > 0$ and in particular $\langle\tilde{\phi}_{j,z}|\psi\rangle > 0$. Notice that since $|\psi'\rangle$ is also a non-negative state and its support is equal to the support of $|\psi\rangle$, we also have that $\langle\tilde{\phi}_{j,z}|\psi'\rangle > 0$ and thus $\langle y|\tilde{P}|\psi'\rangle = \langle y|\tilde{\phi}_{j,z}\rangle\langle\tilde{\phi}_{j,z}|\psi'\rangle > 0$, i.e. y is in the support of $\tilde{P}|\psi'\rangle$. The converse follows from a similar argument.

In particular, by definition, $supp(\tilde{P}_2|\psi\rangle) = supp(\widehat{\tilde{P}_2|\psi\rangle})$; and thus applying $\tilde{P}_1$ on both states, we get $supp(\tilde{P}_1\tilde{P}_2|\psi\rangle) = supp(\tilde{P}_1\widehat{\tilde{P}_2|\psi\rangle})$. By definition, we also have $supp(\tilde{P}_1\tilde{P}_2|\psi\rangle) = supp(\widehat{\tilde{P}_1\tilde{P}_2|\psi\rangle})$, which proves the first part of the Corollary.

For the moreover part, we can apply the previous argument to show $\widehat{\tilde{P}_2\tilde{P}_1|\psi\rangle} = \widehat{\tilde{P}_2\widehat{\tilde{P}_1|\psi\rangle}} = supp(\widehat{\tilde{P}_2\tilde{P}_1|\psi\rangle})$. If $\tilde{P}_1$ and $\tilde{P}_2$ commute, we have $supp(\widehat{\tilde{P}_1\tilde{P}_2|\psi\rangle}) = supp(\widehat{\tilde{P}_2\tilde{P}_1|\psi\rangle})$, and the result follows. □

3.2 Uniform Stoquastic Hamiltonians

In this work, we focus on a restricted class of stoquastic Hamiltonian which we call uniform stoquastic Hamiltonian.

Definition 3.11 (uniform stoquastic Local Hamiltonian)**.** A stoquastic Local Hamiltonian $H = \frac{1}{m}\sum_{i=1}^{m}\tilde{H}_i$ is called uniform if the states of the unique non-negative decompositions of each local stoquastic projector (P_i) are subset-states.

Following Claim 3.5 and remark 3.6, for uniform stoquastic Local Hamiltonians, the groundspace projector of $\tilde{H}_i$ is $\tilde{P}_i = (P_i)_Q \otimes I_{\overline{Q}} = \sum_{j,x}|T_{i,j}\rangle\langle T_{i,j}| \otimes |x\rangle\langle x|$, with $T_{i,j} \subseteq \Sigma^k$ and $T_{i,j} \cap T_{i,j'} = \emptyset$ for $j \neq j'$. We also denote $|\tilde{T}_{i,j,x}\rangle := |T_{i,j}\rangle|x\rangle$ (and $\tilde{T}_{i,j,x}$ as the corresponding set of n-bit strings).

We provide now a lemma which we do not strictly use in the proof, regarding stoquastic frustration-free Local Hamiltonians; that we can always assume that the groundstate of the entire Hamiltonian is a subset-state. Though the claim itself is not used, it is helpful to conceptually hold it in mind, when reading the proof.

Lemma 3.12 (The structure of groundstates of uniform stoquastic Hamiltonian)**.** *Let H be a uniform stoquastic frustration-free Local Hamiltonian. Let H_i be a local term of H. Then if $|\psi\rangle$ is a groundstate of H, it can be written in the form $|\psi\rangle = \sum_{j,z}\alpha_{i,j,z}|\tilde{T}_{i,j,z}\rangle$, for some choice of coefficients $\alpha_{i,j,z} \in \mathbb{C}$. Moreover, the subset-state $|S\rangle$, for $S = \bigcup_{j,z:\alpha_{j,z}\neq 0}\tilde{T}_{i,j,z}$, is also a groundstate of H.*

Proof. The first claim just follows from the fact that H is frustration free so any groundstate must be spanned by groundstates of a fixed term H_i, namely

$$|\psi\rangle = \sum_{j,z}\alpha_{i,j,z}|\tilde{T}_{i,j,z}\rangle. \tag{6}$$

We show now that $|S\rangle$ is also a ground-state of H. Let us consider some term $H_{i'}$ and the decomposition of $|\psi\rangle$ from Equation (6) in respect to its non-negative decomposition. It follows

that $\bigcup_{j,z:\alpha_{i',j,z}\neq 0} \tilde{T}_{i',j,z} = S$. This implies that

$$|S\rangle = \sum_{j,z:\alpha_{i',j,z}\neq 0} \frac{\sqrt{\tilde{T}_{i',j,z}}}{\sqrt{S}} |\tilde{T}_{i',j,z}\rangle,$$

and therefore $|S\rangle$ is in the groundspace of $H_{i'}$, for any $i' \in [m]$. □

We formally define now the frustration-free Uniform Stoquastic k-Local Hamiltonian problem.

Definition 3.13 (uniform stoquastic frustration-free k-Local Hamiltonian problem)**.** The *uniform stoquastic frustration-free k-Local Hamiltonian* problem, where $k \in \mathbb{N}^*$ is called the locality and $\varepsilon : \mathbb{N} \to [0,1]$ is a non-decreasing function, is the following promise problem. Let n be the number of qudits of a quantum system. The input is a set of $m(n)$ uniform stoquastic Hamiltonians $H_1, \ldots, H_{m(n)}$ where m is a polynomial, $\forall i \in m(n) : 0 \leq H_i \leq I$ and each H_i acts on k qudits out of the n qudit system. We also assume that there are at most d terms that act non-trivially on each qudit, for some constant d, and that $m \geq n$. For $H = \frac{1}{m(n)} \sum_{i=1}^{m(n)} H_i$, one of the following two conditions hold.
Yes. There exists a n-qudit quantum state $|\psi\rangle$ such that $\langle\psi|H|\psi\rangle = 0$
No. For all n-qudit quantum states $|\psi\rangle$ it holds that $\langle\psi|H|\psi\rangle \geq \varepsilon(n)$.

3.3 MA-completeness

Bravyi and Terhal [BT09] showed that there exists some polynomial $p(n) = O(n^2)$ such that the frustration-free uniform stoquastic 6-Local Hamiltonian problem with $\varepsilon(n) = \frac{1}{p(n)}$ is MA-hard. They also proved that for every polynomial p' and every constant k, the frustration-free uniform stoquastic k-Local Hamiltonian problem with $\varepsilon(n) = \frac{1}{p'(n)}$ is in MA.

Let us start with the direction which is of less technical interest to us, and thus we will not need to go into details. The MA-hardness is proved by analyzing the quantum Cook-Levin theorem [KSV02, AN02] when considering an MA verifier. A verification circuit for MA can be described as a quantum circuit consisting only of the (classical) gates from the universal (classical) gateset Toffoli and NOT, operating on input qubits in the state $|0\rangle$ (the NOT gates can then fix them to the right input) and ancillas which are either in the state $|0\rangle$ used as workspace, or in the state $|+\rangle$, used as random bits. At the end of the circuit, the first qubit is measured in the computational basis and the input is accepted iff the output is 1. It is not difficult to check that for such gateset and ancillas, the stoquastic Local Hamiltonian resulting from the circuit-to-Hamiltonian construction of the quantum Cook-Levin theorem (which forces both the correct propagation as well as the correct input state, as well as the output qubit accepting), is a uniform stoquastic Hamiltonian; in particular all the entries are in $\{0, \pm 1, -\frac{1}{2}\}$. We can also assume that each qubit is used in at most d gates, for some $d \geq 3$. This is true because all the computation done by the verifier is classical and therefore the information can be copied to fresh ancilla bits (initialized on $|0\rangle$) with a CNOT operation. Notice then that each qubit takes place on at most 3 steps: as the target of the CNOT, in some actual computation, and as the source of the next CNOT.

We now explain the other direction, namely Bravyi and Terhal's approach for showing that the stoquastic Hamiltonian problem is in MA. We actually show a simplified version of their result, since we are only interested in *uniform* stoquastic Hamiltonian. Notice that by our above description, this problem is sufficient to achieve MA-hardness.

Following [BT09] We now define the graph on which the random walk will take place; this graph is based on a given uniform stoquastic Hamiltonian.

Definition 3.14 (Graph from uniform stoquastic Hamiltonian)**.** Let $H = \frac{1}{m}\sum_i H_i$ be a uniform stoquastic Hamiltonian on n qudits of dimension $|\Sigma|$. We define the undirected graph $G(H) = (|\Sigma|^n, E)$ where $(x, y) \in E$ iff there exists a local term H_i with corresponding groundstate projector P_i such that

$$\langle x|P_i|y\rangle > 0. \tag{7}$$

From Remark 3.7 and claim 3.5, we have that for a fixed i, the neighbor strings form an equivalence class and in each class the strings differ only in the positions where H_i acts non-trivially. We also remark that given some string x, one can compute in polynomial time if x is bad for H, by just inspecting the groundspace of each local term.

The random walk starts from some initial string x_0 sent by the prover. If x_0 is bad for H, then the algorithm rejects. Otherwise, a term H_i is picked uniformly at random and a string x_1 is picked uniformly at random from $|\tilde{T}_{i,j,z}\rangle$, which is the unique rank-one subset-state from $\tilde{P}_i$ such that $x_0 \in \tilde{T}_{i,j,z}$ (see Remark 3.7). The random walk proceeds by repeating this process with x_1. We describe the random walk proposed by BT (simplified for the uniform case) in Figure 3.

---

1. Let x_0 be the initial string.
2. Repeat for T steps
   (a) If x_t is bad for H, reject
   (b) Pick $i \in [m]$ uniformly at random
   (c) Pick x_{t+1} uniformly at random from the strings in the unique $\tilde{T}_{i,j,z}$ that contains x_t
3. Accept

Figure 3: BT Random Walk

We state now the lemmas proved in [BT09].

Lemma 3.15 (Completeness, adapted from Section 6.1 of [BT09])**.** *If H is frustration-free, then there exists some string x such that there are no bad-strings in the connected component of x.*

The proof goes by showing that if H is frustration-free, then for any string x in some groundstate of H, the uniform superposition of the connected component of x is a groundstate of H. In this case, since all strings in the connected component of x are good (this is by definition of the connected component), the verifier will accept.

Lemma 3.16 (Soundness, adapted from Section 6.2 of [BT09])**.** *For every polynomial p, there exists some polynomial q such that if H is at least $\frac{1}{p(n)}$-frustrated, then for every string x, for $T = q(n)$, the random walk from Figure 3 rejects with constant probability.*

The intuition of the proof is that since the Hamiltonian is frustrated, one can upper bound the expansion on any set of good-strings by $1 - \frac{1}{p(n)}$, otherwise the Hamiltonian would not be $\frac{1}{p(n)}$-frustrated. In this case, there exists some polynomial q such that a random walk with $q(n)$ steps escape of any set of good strings with high probability.

4 Uniform Gapped stoquastic Hamiltonians are in NP

Our main technical result in this work is showing that if a stoquastic uniform Hamiltonian is ε-frustrated for some constant ε, then every string x is constantly-close to a bad string.

Lemma 4.1 (Short path to a bad string)**.** *If the stoquastic uniform k-Local Hamiltonian H is ε-frustrated, then for every string x, there is a bad string y such that the distance between x and y in $G(H)$ is at most $k^{\frac{2kd}{\varepsilon}\log_{(1+\frac{\varepsilon}{4})}|\Sigma|}$.*

Using this lemma, we can prove our main result:

Theorem 1.1 (restated). For any constant $\varepsilon > 0$, the problem of deciding whether a uniform stoquastic Hamiltonian H is frustration-free or ε-frustrated, is in NP.

Proof. The NP witness for the problem consists in some initial string x that is promised to be in the support of the groundstate of H. The verification proceeds by running over all possible $k^{\frac{2kd}{\varepsilon}\log_{(1+\frac{\varepsilon}{4})}|\Sigma|}$-step paths from x. Since for each one of the m terms there are constantly many possible steps, the number of possibilities for one step is polynomial, and so the number constantly-long paths is also polynomial. Therefore, such enumeration can be performed efficiently. For each path, we check if one of the strings it reaches is bad - again this can be done in polynomial time since badness is with respect to the local terms (see Remark 4.2 for the precision issues). The verifier rejects if any of the paths reached a bad string, otherwise it accepts.

Let x be the string sent by the Prover. If H is frustration-free, then by Lemma 3.15 all strings in the connected component of x are good. On the other hand, if H is ε-frustrated, then by Lemma 4.1, there exists a $k^{\frac{2kd}{\varepsilon}\log_{(1+\frac{\varepsilon}{4})}|\Sigma|}$-step path from x to some string y that is bad for H, and such path will be found by the brute-force search. □

Remark 4.2 (Deciding on badness of a string with respect to a local uniform stoquastic term)**.** We note that while in the non-uniform case the question of whether a string is bad for a local term or not, may depend on precision issues, this is not a problem when considering uniform stoquastic Hamiltonians. In the uniform case, the set of strings comprising the subset states in the non-negative decomposition of every projector, as in Equation (1), can be calculated *exactly* given the matrix description of the local Hamiltonian term (even if we need to apply approximations when computing the groundstates). This is because the locality of the Hamiltonian, together with uniformity, imply that if a string is in the support of one of the groundstates, its weight must be $\frac{1}{\sqrt{q}}$ for some positive integer q smaller than some constant.

The remainder of this section is dedicated to proving Lemma 4.1. Section 4.1 gives the one term expansion argument, Section 4.2 provides the proof that a constant number of layers consisting of parallel non-overlapping projections suffices to reach a bad string; and Section 4.3 provides the light-cone argument to show that if a bad string is reached within constantly many layers, then in fact we there is a bad string within *constantly* many steps from the initial string. This then allows searching for such a string by brute-force. Finally, Section 4.4 just puts all the pieces together to finish the proof of Lemma 4.1.

4.1 Expansion

We start by showing that if subset-state $|S\rangle$ does not contain any bad string but a term P is highly frustrated by $|S\rangle$, then the support of $\tilde{P}|S\rangle$ is larger than that of $|S\rangle$ by a constant factor.

Lemma 4.3 (One term expansion)**.** *Let $\tilde{P} = \sum_{j,z} |\tilde{T}_{j,z}\rangle\langle\tilde{T}_{j,z}|$ be a uniform stoquastic projector on the set Q of k out of n qudits. Let $S \subseteq \Sigma^n$ be such that $|S\rangle$ does not contain P-bad-strings, and $\left\|\tilde{P}|S\rangle\right\|^2 \leq 1-\delta$. It follows that the size of the support of $\tilde{P}|S\rangle$ is at least $(1+\frac{\delta}{2})|S|$.*

Proof. Since S does not contain bad strings, we start by noticing that from Lemma 3.9, S is contained in the support of $\tilde{P}|S\rangle$, and $\tilde{P}$ only adds the neighbors of strings in $|S\rangle$.

We have that

$$1-\delta \geq \|\tilde{P}|S\rangle\|^2 = \langle S|\tilde{P}|S\rangle = \sum_{j,z}\langle S|\tilde{T}_{j,z}\rangle\langle\tilde{T}_{j,z}|S\rangle = \sum_{j,z}|\langle S|\tilde{T}_{j,z}\rangle|^2. \tag{8}$$

Let $\mathbf{T} = \bigcup_{j,z:S\cap\tilde{T}_{j,z}\neq\emptyset}\tilde{T}_{j,z}$. It follows that

$$\begin{aligned}
\sum_{j,z}|\langle S|\tilde{T}_{j,z}\rangle|^2 &= \sum_{j,z}\frac{|\tilde{T}_{j,z}\cap S|^2}{|\tilde{T}_{j,z}||S|} \\
&= \sum_{j,z}\frac{(|\tilde{T}_{j,z}| - |\tilde{T}_{j,z}\setminus S|)^2}{|\tilde{T}_{j,z}||S|} \\
&= \sum_{j,z:\tilde{T}_{j,z}\cap S\neq\emptyset}\frac{|\tilde{T}_{j,z}|^2 - 2|\tilde{T}_{j,z}||\tilde{T}_{j,z}\setminus S| + |\tilde{T}_{j,z}\setminus S|^2}{|\tilde{T}_{j,z}||S|} \\
&\geq \frac{|\mathbf{T}| - 2|\mathbf{T}\setminus S|}{|S|} \\
&\geq 1 - \frac{2|\mathbf{T}\setminus S|}{|S|}.
\end{aligned} \tag{9}$$

where in first inequality we remove some non-negative terms and use the fact that $\tilde{T}_{j,z}$ and $\tilde{T}_{j',z}$ are disjoint for $j \neq j'$ (Remark 3.7) and in the second inequality we use the fact that $S \subseteq \mathbf{T}$ since there are no bad strings in S.

By putting together Equations (8) and (9), and noticing that $\mathbf{T} = supp(\tilde{P}|S\rangle)$ from Lemma 3.9, we have that

$$|supp(\tilde{P}|S\rangle)| = |\mathbf{T}| = |S| + |\mathbf{T}\setminus S| \geq |S| + \frac{\delta}{2}|S| = \left(1+\frac{\delta}{2}\right)|S|. \qquad \square$$

4.2 Bad string in a constant number of layers

We prove in this section that with a constant number of "layers", it is possible to reach a state with a bad string.

We first want find a linear number of non-overlapping terms that are (roughly) simultaneously frustrated by some subset-state $|S\rangle$. Let us first define what we mean by non-overlapping terms.

Definition 4.4 (Non-overlapping projectors)**.** A sequence of local projectors $L = (Q_1, ..., Q_\ell)$ is non-overlapping if for any $i \neq j$, Q_i and Q_j act on disjoint sets of qudits.

Now, we need to be more careful in order to define the notion of "simultaneous". Recall that if H is ε-frustrated, by Claim 2.5, there must exist at least $\frac{\varepsilon m}{2}$ terms that are at least $\frac{\varepsilon}{2}$-frustrated. However, as we explained in Example 1, their frustration may be correlated due to entanglement, and when we "correct" the frustration of one term, we could also be correcting the frustration of other terms,

Because of this, we need to choose the sequence of projectors more carefully. We are looking for projectors which are frustrated, even after applying the previous projectors in the sequence.

Definition 4.5 (Sequentially frustrated terms)**.** A sequence of projectors $L = (Q_1, ..., Q_\ell)$ is *sequentially* δ-frustrated by some subset-state $|S_0\rangle$, if for all $i = 1, ..., \ell$, $\left\| \tilde{Q}_i \widehat{\tilde{Q}_{i-1} ... \tilde{Q}_1 |S_0\rangle} \right\|^2 \leq 1 - \delta$.

We show now that we can find a linear-size sequence that is non-overlapping and sequentially highly frustrated, in a greedy way. At iteration i, we fix a projector Q_i such that $\tilde{Q}_i$ is highly frustrated by $\widehat{\tilde{Q}_{i-1} ... \tilde{Q}_1 |S_0\rangle}$ and Q_i does not overlap with any Q_j for $j < i$. More concretely, we choose Q_i arbitrarily from the intersection of the following sets:

- F_i, the set of terms that are at least $\frac{\varepsilon}{2}$-**F**rustrated by $\widehat{\tilde{Q}_{i-1} ... \tilde{Q}_1 |S_0\rangle}$
- A_i, the set of **A**vailable terms, i.e. the terms that do not overlap with Q_j, for $j < i$.

We describe such an algorithm in Figure 4 and analyze its correctness in Lemma 4.6.

Let $H = \frac{1}{m} \sum_{j=1}^{m} H_j$ be a stoquastic k-Local Hamiltonian with frustration at least ε and some subset-state $|S_0\rangle$. For each term H_j, we denote by P_j the projector onto its groundspace.

1. Let $i = 0$, $A_0 = \{P_1, ..., P_m\}$ and $F_0 = \{P_j : \left\| \tilde{P}_j |S_0\rangle \right\|^2 \leq 1 - \frac{\varepsilon}{2}\}$
2. While $A_i \cap F_i \neq \emptyset$
   (a) Pick any $P_j \in A_i \cap F_i$ and set $Q_i = P_j$
   (b) Let $A_{i+1} = A_i \setminus \{P_j : P_j \text{ overlaps with } Q_i\}$ and $F_{i+1} = \{P_j : \left\| \tilde{P}_j \widehat{\tilde{Q}_i ... \tilde{Q}_0 |S_0\rangle} \right\|^2 \leq 1 - \frac{\varepsilon}{2}\}$
   (c) Let $i = i + 1$
3. Output $L = (Q_0, ..., Q_{i-1})$

Figure 4: Algorithm for finding non-overlapping frustrated terms

Lemma 4.6 (Linear number of sequential non-overlapping frustrated terms)**.** *Let $H = \frac{1}{m} \sum_{i=1}^{m} H_i$ be a ε-frustrated uniform stoquastic k-Local Hamiltonian, $S_0 \subseteq \Sigma^n$, and $L = (Q_0, ..., Q_{i^*-1})$ be the output of Figure 4. Then i) L is non-overlapping, ii) L is sequentially $\frac{\varepsilon}{2}$-frustrated by $|S_0\rangle$, and iii) $i^* \geq \frac{\varepsilon n}{2kd}$.*

Proof. Properties *i)* and *ii)* follow by construction: $Q_i \in A_i$, and thus it does not overlap with Q_j for $j < i$; and $Q_i \in F_i$, therefore $\left\| \tilde{Q}_i \widehat{\tilde{Q}_{i-1} .. \tilde{Q}_0 |S_0\rangle} \right\|^2 \leq 1 - \frac{\varepsilon}{2}$.

We prove now property *iii*). Notice that $|A_{i+1}| - |A_i| \leq kd$, since the only difference between these two sets are the overlapping terms of Q_i, Q_i acts on at most k qudits, and there are at most d other terms that overlap with Q_i due to a specific qudit. Therefore, we have that

$$|A_{i^*}| \geq m - i^* kd. \tag{10}$$

Notice that for every i, we have that $|F_i| \geq \frac{\varepsilon m}{2}$ by Claim 2.5. We also have that if $|A_i| + |F_i| > m$, then $A_i \cap F_i \neq \emptyset$ by the pigeonhole principle. Therefore, $A_{i^*} \cap F_{i^*} = \emptyset$ implies that

$$\left(1 - \frac{\varepsilon}{2}\right) m \geq |A_{i^*}| \tag{11}$$

Putting Equations (10) and (11) together we have

$$m - i^* kd \leq |A_{i^*}| \leq \left(1 - \frac{\varepsilon}{2}\right) m$$

and therefore it follows that

$$i^* \geq \frac{\varepsilon m}{2kd} \geq \frac{\varepsilon n}{2kd},$$

where we use the fact that $m \geq n$. □

Definition 4.7 (A layer acting on $|S_0\rangle$)**.** We denote $L|S_0\rangle$ to be $\tilde{Q}_{i^*-1}...\tilde{Q}_0|S_0\rangle$. We notice that this state is equal to $\tilde{Q}_{\sigma(i^*-1)}...\tilde{Q}_{\sigma(0)}|S_0\rangle$ for every permutation σ on the indices $0, ..., i^* - 1$, since by property i) of Lemma 4.6 the terms in L are non overlapping, and thus commuting. Thus the resulting state does not depend on the order of application of these projections, hence the notion of *applying a layer of non-overlapping terms on a state* is well defined.

By combining Lemmas 4.3 and 4.6 we can now prove that if we apply all the projections output by Figure 4, namely all projections in L, the resulting state has exponentially more strings than the original one.

Corollary 4.8 (Multiple terms expansion)**.** *Let $H = \frac{1}{m}\sum_{i=1}^{m} H_i$ be an ε-frustrated uniform stoquastic k-Local Hamiltonian and $|S\rangle$ be a subset-state such that $|S\rangle$ does not contain any bad string for H. Then there is a sequence L of non-overlapping terms of H such that the number of strings in the support of $L|S\rangle$ is at least $(1 + \frac{\varepsilon}{4})^{\frac{\varepsilon n}{2kd}}$ times the number of strings in $|S\rangle$.*

Proof. Let $L = (Q_1, ..., Q_{i^*})$ be the output of Figure 4. Let us argue now that we can use Lemma 4.3 for each one of the $i^* \geq \frac{\varepsilon n}{2kd}$ terms in L sequentially, which would imply the statement.

First, let us claim that for all j, $\tilde{Q}_j...\tilde{Q}_1|S\rangle$ does not contain $Q_{j'}$-bad strings, for $j' > j$. This follows by induction; for $j = 0$ this is true by assumption. Also from Lemma 4.6, the terms in L are non-overlapping. Hence, by Claim 3.5 when we apply $\tilde{Q}_j$ no $Q_{j'}$-bad strings appear, for all $j' \neq j$. Secondly, directly from Lemma 4.6 we have that $\left\| \widehat{\tilde{Q}_{j+1}\tilde{Q}_j...\tilde{Q}_1|S\rangle} \right\| \leq 1 - \frac{\varepsilon}{2}$.

This means that the conditions of Lemma 4.3 are satisfied, and it follows that the number of strings in the support of the state $\widehat{\tilde{Q}_{j+1}\tilde{Q}_j...\tilde{Q}_1|S\rangle}$ is larger by a factor of $(1 + \frac{\varepsilon}{4})$ than that of the state $\widehat{\tilde{Q}_j...\tilde{Q}_1|S\rangle}$. By Corollary 3.10 we deduce that the support of $\tilde{Q}_{j+1}...\tilde{Q}_1|S\rangle$ is bigger by the same factor, than that of $\tilde{Q}_j...\tilde{Q}_1|S\rangle$. □

We now observe that this expansion is too large to be applied for many layers, since the number of strings reached will just exceed the number of n-bit strings.

Lemma 4.9 (Bad string in constant number of layers)**.** *Let $\ell^* = \lceil \frac{2kd}{\varepsilon} \log_{\left(1+\frac{\varepsilon}{4}\right)} |\Sigma| \rceil$. Consider a uniform[11] stoquastic k-Local Hamiltonian $H = \frac{1}{m}\sum_{i=1}^{m} H_i$ which is ε-frustrated. Then for every good string x, there exists some $\ell < \ell^*$ and a sequence $L_1, ...L_\ell$, where each L_i consists of a set of non-overlapping local projectors (corresponding to the projections on the local terms of H), such that $L_\ell ... L_1|x\rangle$ contains a bad string.*

Proof. Let L_1 be the output of Figure 4 for the initial state $|x\rangle$, and recursively, for $\ell \leq \ell^*$, let L_ℓ be the output of Figure 4 for the state $\widehat{L_\ell ... L_1|x\rangle}$. Let S_ℓ be the set of strings in the support of the state $\widehat{L_\ell ... L_1|x\rangle}$.

Now, if for some $\ell < \ell^*$, S_ℓ contains a bad string for H, we are done. Otherwise, for all $\ell < \ell^*$, the state $\widehat{L_{\ell-1} ... L_1|x\rangle}$ contains no bad string for H.

From Corollary 4.8, we have that for each $1 \leq \ell \leq \ell^* - 1$,

$$|supp(L_\ell(\widehat{L_{\ell-1} ... L_1|x\rangle}))| \geq \left(1 + \frac{\varepsilon}{4}\right)^{\frac{\varepsilon n}{2kd}} |S_{\ell-1}|$$

Moreover, by Corollary 3.10, we have

$$|S_\ell| = |supp(L_\ell(\widehat{L_{\ell-1} ... L_1|x\rangle}))|.$$

Thus, by a trivial induction, we arrive at

$$|S_{\ell^*}| \geq \left(\left(1 + \frac{\varepsilon}{4}\right)^{\frac{\varepsilon n}{2kd}}\right)^{\frac{2kd}{\varepsilon}\log_{(1+\frac{\varepsilon}{4})}|\Sigma|} = \left(\left(1 + \frac{\varepsilon}{4}\right)^{\log_{(1+\frac{\varepsilon}{4})}|\Sigma|}\right)^{n} = |\Sigma|^n.$$

Thus S_{ℓ^*} contains all possible strings. However if H is frustrated, bad strings exist and thus there must be a bad string in $S_{\ell^*} = supp(L_{\ell^*} ... L_1|x\rangle)$. □

4.3 Finding the bad string

In this section, we prove that if for some constant ℓ we have that $L_\ell ... L_1|x\rangle$ contains a bad string, and each L_i consists of non-overlapping terms, then there exists a constant path (namely a sequence of constantly many local steps) from x to a bad string.

We start by showing how to retrieve (possibly polynomial-size) paths from strings in some non-negative state $|\psi\rangle$ to string in some state $L|\psi\rangle$, for a non-overlapping set of projections L.

Lemma 4.10 (From non-overlapping projections to paths)**.** *Let $|\psi\rangle$ be a non-negative state and L be an arbitrary set of non-overlapping stoquastic projectors. Then for every string y in $L|\psi\rangle$, there exists a string x in $|\psi\rangle$ such that there is a $|L|$-step path between x and y in $G(H)$.*

Proof. Let $L = \{Q_1, ..., Q_g\}$. Since L is a set of non-overlapping stoquastic projectors, if $y \in supp(L|\psi\rangle)$, then there must exist some x in $supp(|\psi\rangle)$ such that

$$0 < \langle x|L|y\rangle.$$

[11] In fact, uniformity is not needed for this claim, but this requires more work.

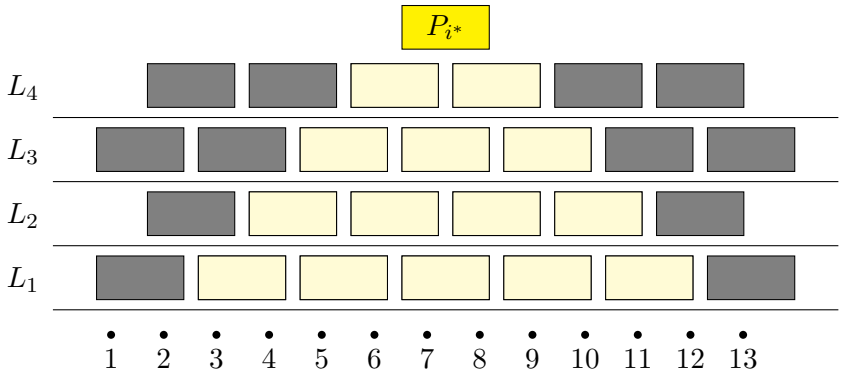

Figure 5: The layers $L_1, ..., L_4$ reach a bad string for some term P_{i^*} that does not necessarily belong to any layer, and then we consider the light-cone of P_{i*} in these layers. Notice that the light-cone is defined only within the layers, each consisting of non-overlapping terms, and *not* by considering *all* projectors which touch any qudit in P_{i^*}, then all projectors which touch those, etc.)

This is because writing $\langle\psi| = \sum_x \psi_x \langle x|$ we have (using $L^\dagger = L$ due to the fact that L consists of non-overlapping projections):

$$0 < \langle\psi|L|y\rangle = \sum_{x \in supp(|\psi\rangle)} \psi_x \langle x|L|y\rangle$$

and the coefficients ψ_x are all non-negative. We have

$$0 < \langle x|L|y\rangle = \langle x|\tilde{Q}_1...\tilde{Q}_g|y\rangle = \sum_{w_1,...,w_{g-1}} \langle x|\tilde{Q}_1|w_1\rangle\langle w_1|\tilde{Q}_2...|w_{g-1}\rangle\langle w_{g-1}|\tilde{Q}_g|y\rangle, \tag{12}$$

where we can write the $\tilde{Q}_i$ in any order since they are non-overlapping, from Definition 4.7; the w's above run over all n-dit strings.

Since every $\tilde{Q}_i$ has only non-negative entries, for every pair of strings z and z', $\langle z|\tilde{Q}_i|z'\rangle \geq 0$. Then Equation (12) holds iff there exists values $w_1^*, ..., w_{g-1}^*$ such that

$$0 < \langle w_0^*|\tilde{Q}_1|w_1^*\rangle\langle w_1^*|\tilde{Q}_2...|w_{g-1}^*\rangle\langle w_{g-1}^*|\tilde{Q}_g|w_g^*\rangle,$$

where we have set $x := w_0^*$ and $y := w_g^*$. It follows that for every $i \in [g]$, $\langle w_{i-1}^*|\tilde{Q}_i|w_{i+1}^*\rangle > 0$, and then from Equation (7), w_i^* and w_{i+1}^* are neighbors in $G(H)$. Therefore, there is a $|L|$-step path from x to y. □

Finally, we show how to find a short path between the initial string a bad string. The intuition of the proof is depicted in Figure 2.

Lemma 4.11 (The light-cone argument)**.** *For some initial string x, let $L_1, ..., L_\ell$ each be a sequence of non-overlapping projectors such that $L_\ell...L_1|x\rangle$ contains in its support a string w^* which is bad for H. Then, there is a $O(k^\ell)$-steps path from x to a bad string for H (which could be different than w^*).*

Proof. Since w^* is bad for H, we have that for some $i^* \in [m]$, $\langle w^*|\tilde{P}_{i^*}|w^*\rangle = 0$. Let $L_\ell^\triangle \subseteq L_\ell$ be the projectors of L_ℓ that touch some qudit of P_{i^*}. We define recursively $L_{j-1}^\triangle$ as the set of projectors in L_{j-1} that overlap with some projector in $\bigcup_{j' \geq j} L_{j'}^\triangle$. These are the *layers* of what we call the *light-cone* of P_{i*} and we depict it in Figure 5. Let us also define the complement of $L_j^\triangle$ in L_j to be L_j'.

For convenience, set $L_{\ell+1} = L_{\ell+1}^\triangle = \{P_{i^*}\}$. Let D_j be the set of qudits touched by the terms in $\bigcup_{j' \geq j} L_{j'}^\triangle$. We prove that $|D_j| \leq k^{\ell-j+2}$. We prove this by a downward induction from $j = \ell+1$ to $j = 1$. The basis step is true since P_{i^*} is k-local, and so $|D_{\ell+1}| = k$. Now, assume this is true for the j'th layer. D_{j-1} is defined by adding to D_j the qudits touched by the next layer of the lightcone, $L_{j-1}^\triangle$. By definition of the lightcone, these are all qudits touched by terms in L_{j-1} that overlap D_j. Since these terms are non-overlapping and k-local, this can at most multiply the number of qudits already in D_j by a factor of k. We have that the set D_1 of qudits within the entire lightcone originating from P_{i^*} contains at most $k^{\ell+1}$ qudits.

The terms in L_j commute (as they are non-overlapping), and thus $L_j = L_j' L_j^\triangle$. It follows that

$$L_\ell ....L_1|x\rangle = L_\ell' L_\ell^\triangle L_{\ell-1}' L_{\ell-1}^\triangle ...L_1' L_1^\triangle |x\rangle = L_\ell' L_{\ell-1}' ...L_1' L_\ell^\triangle ...L_1^\triangle |x\rangle, \tag{13}$$

where in the second equality, we use in fact an iterative argument (that is common in light-cone reasoning, see, e.g., [AALV09, BGK18]): we notice that the projectors in L_j' commute with the projectors in $L_{j'}^\triangle$ for all $j' \geq j$, and thus they can be commuted one by one across the lightcone operators, to the left. The fact that they commute follows by definition: level j of the lightcone, $L_j^\triangle$, contains *all* terms in L_j that overlap with any term $L_{j'}^\triangle$, for $j' > j'$; thus the remaining terms in L_j, namely L_j', do not overlap the upper layers of the lightcone, and thus commute with them.

From Equation 13 we deduce that we can first apply on x all terms in the lightcone, and delay all terms outside of the lightcone to later. From this, we can show that $L_\ell^\triangle ...L_1^\triangle |x\rangle$ also contains a bad string for $\tilde{P}_{i^*}$, and this will complete the proof. To do this, let Q be the set of positions where the term P_{i^*} acts non-trivially. We claim that the application of the terms outside of the lightcone, which do not touch Q, couldn't have added a string which is bad with respect to P_{i^*}, unless such a string was there before. This can be deduced from Claim 3.5, which when applied iteratively gives that

$$\{y_Q : y \in supp(L_\ell' L_{\ell-1}' ...L_1' L_\ell^\triangle ...L_1^\triangle |x\rangle)\} \subseteq \{y_Q : y \in supp(L_\ell^\triangle ...L_1^\triangle |x\rangle)\}. \tag{14}$$

Since $L_\ell' L_{\ell-1}' ...L_1' L_\ell^\triangle ...L_1^\triangle |x\rangle$ contains a bad string w^* for $\tilde{P}_{i*}$, from Equation (14) we have that $L_\ell^\triangle ...L_1^\triangle |x\rangle$ contains a string w' such that $w_Q^* = w_Q'$, thus w' is also bad for $\tilde{P}_{i^*}$.

Finally, we can use Lemma 4.10 together with a (highly wasteful) bound on $|L_j^\triangle| \leq |D_j| \leq k^{\ell-j+2}$ recursively for $L_\ell^\triangle ...L_1^\triangle |x\rangle$: for any string y in $L_j^\triangle ...L_1^\triangle |x\rangle$, there exists a string y' in $L_{j-1}^\triangle ...L_1^\triangle |x\rangle$, such that there is a $|L_j^\triangle| \leq k^{\ell-j+2}$-step path from y' to y in $G(H)$. Hence, there is a path of size $\sum_{j'=1}^{\ell} k^{\ell-j'+2} = \sum_{j'=2}^{\ell+1} k^{j'} = O(k^\ell)$ from x to w' in $G(H)$. □

4.4 Proof of Lemma 4.1

We can finally prove Lemma 4.1 by composing Lemmas 4.9 and 4.11.

Lemma 4.1 (restated). If the stoquastic uniform k-Local Hamiltonian H is ε-frustrated, then for every string x, there is a bad string y such that the distance between x and y in $G(H)$ is at most $k^{\frac{2kd}{\varepsilon}\log_{(1+\frac{\varepsilon}{4})}|\Sigma|}$.

Proof. Let $\ell^* = \frac{2kd}{\varepsilon}\log_{(1+\frac{\varepsilon}{4})}|\Sigma|$. By Lemma 4.9, there exists a sequence of sets $L_1, ...L_\ell$, for some $\ell \leq \ell^*$, such that each L_i consists of non-overlapping projectors, and $L_\ell ... L_1|x\rangle$ contains a bad string. We finish the proof by using $L_1, ..., L_\ell$ in Lemma 4.11 for some $\ell \leq \frac{2kd}{\varepsilon}\log_{(1+\frac{\varepsilon}{4})}|\Sigma|$. □

5 Commuting Stoquastic Hamiltonians are in NP

In this section, we give a simple proof that deciding if a commuting Stoquastic Hamiltonian is frustration-free is in NP.

Theorem 1.4 (restated). The problem of deciding if a commuting stoquastic Hamiltonian H is frustration-free is in NP.

We notice that given that we are not assuming anything on the gap, one needs to be somewhat careful with the assumptions on how the input is given; we assume here that the local terms of the commuting stoquastic Hamiltonian H are provided by giving the matrix elements, each with $poly(m)$ bits, and the terms mutually commute exactly.

Proof. Let $H_1, ..., H_m$ be the terms in the Hamiltonian, and let $\tilde{P}_1, ..., \tilde{P}_m$ be the corresponding projectors onto their groundspaces. We show that H is frustration-free iff there exist a string x that is good for H.

The first direction ($\Rightarrow$) is trivial: any string in the 0-energy groundstate is good for H.

We prove now the converse ($\Leftarrow$). Let x be a string that is good for H. Let $|\phi\rangle = \tilde{P}_1\tilde{P}_2...\tilde{P}_m|x\rangle$. Since $\tilde{P}_1, ..., \tilde{P}_m$ commute, we have that

$$|\phi\rangle = \tilde{P}_i|\phi\rangle,$$

which means that either $|\phi\rangle = 0$ or it is a $+1$-eigenstate of every $\tilde{P}_i$, and therefore it has energy 0 with respect to H. We show now that $|\phi\rangle \neq 0$. The string x is in the support of some groundstate of every term H_i, therefore for every state $|\alpha\rangle = \sum_y \alpha_y|y\rangle$ with $\alpha_y \in \mathbb{R}^+$ and $\alpha_x > 0$, $\tilde{P}_i|\alpha\rangle = \sum_y \alpha'_y|y\rangle$, with $\alpha'_y \in \mathbb{R}^+$ and $\alpha'_x > 0$. Therefore we have that $\tilde{P}_1\tilde{P}_2...\tilde{P}_m|x\rangle \neq 0$.

Finally, to show that the problem is in NP: the proof is supposed to be a string with largest amplitude in some groundstate of the Hamiltonian. The verification algorithm checks if this string is indeed good for all local terms, as follows. First, for each k-local term H_i, the verifier computes P_i to within constant approximation. More precisely, it computes each matrix elements of P_i to within $\frac{1}{4|\Sigma|^k}$. This can be done efficiently since we are working with local terms where each matrix element of H_i is specified by polynomially many bits. Let the approximated projector be P'_i. The verifier then checks if x is bad for P_i: to do this, it first restricts x to Q, the set of k qudits on which H_i acts, and then checks if $\langle x_Q|P'_i|x_Q\rangle \leq \frac{1}{2|\Sigma|^k}$. If yes, then it rejects. If the verifier does not reject for any of the i's, then it accepts.

To finish the proof, we need to show that the above procedure cannot lead to an error in the verifier's decision of whether x is bad or good for H. We start by arguing that in the frustration-free case, if x is the string with largest amplitude in some groundstate $|\psi\rangle$, then it passes the

test. This follows since we can write $\tilde{P}_i$ using its non-negative decomposition (see Remark 3.6): $\tilde{P}_i = \sum_j |\tilde{\phi}_{i,j,z}\rangle\langle\tilde{\phi}_{i,j,z}|$. We then write $|\psi\rangle = \sum_{j,z} \alpha_{i,j,z}|\tilde{\phi}_{i,j,z}\rangle$; there is a unique $|\tilde{\phi}_{i,j^*,z^*}\rangle$ that contains x. Then x must be the string with largest amplitude in $|\tilde{\phi}_{i,j^*,z^*}\rangle$ (since a string y in $|\tilde{\phi}_{i,j^*,z^*}\rangle$ has amplitude $\alpha_{i,j,z}\langle y|\tilde{\phi}_{i,j,z}\rangle$ in $|\psi\rangle$ and x maximizes this value). Since $|\tilde{\phi}_{i,j^*,z^*}\rangle$ contains at most $|\Sigma|^k$ strings, the amplitude of x in it is at least $|\Sigma|^{-\frac{k}{2}}$. Hence $\langle x|\tilde{P}_i|x\rangle = \langle x|\tilde{\phi}_{i,j^*,z^*}\rangle\langle\tilde{\phi}_{i,j^*,z^*}|x\rangle \geq \frac{1}{|\Sigma|^k}$. By our described procedure, in this case the verifier accepts. For soundness, if some string x is bad for H_i, then for some i $\langle x|P_i|x\rangle = 0$ and thus by our bound on the error due to approximation of P_i by P_i', we know $\langle x|P_i'|x\rangle \leq \frac{1}{4|\Sigma|^k}$, and the verifier will reject. □

References

[AALV09] Dorit Aharonov, Itai Arad, Zeph Landau, and Umesh V. Vazirani. The detectability lemma and quantum gap amplification. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009*, pages 417–426, 2009.

[AALV11] Dorit Aharonov, Itai Arad, Zeph Landau, and Umesh V. Vazirani. The 1d area law and the complexity of quantum states: A combinatorial approach. In *IEEE 52nd Annual Symposium on Foundations of Computer Science, FOCS 2011*, pages 324–333, 2011.

[AAV13] Dorit Aharonov, Itai Arad, and Thomas Vidick. Guest column: the quantum PCP conjecture. *SIGACT News*, 44(2):47–79, 2013.

[AAV16] Anurag Anshu, Itai Arad, and Thomas Vidick. Simple proof of the detectability lemma and spectral gap amplification. *Phys. Rev. B*, 93:205142, May 2016.

[AE15] Dorit Aharonov and Lior Eldar. The Commuting Local Hamiltonian Problem on Locally Expanding Graphs is Approximable in NP. *Quantum Information Processing*, 14(1):83–101, 2015.

[ALM+98] Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy. Proof Verification and the Hardness of Approximation Problems. *Journal of the ACM*, 45(3):501–555, 1998.

[ALV12] Itai Arad, Zeph Landau, and Umesh Vazirani. Improved one-dimensional area law for frustration-free systems. *Phys. Rev. B*, 85:195145, May 2012.

[ALVV17] Itai Arad, Zeph Landau, Umesh Vazirani, and Thomas Vidick. Rigorous RG Algorithms and Area Laws for Low Energy Eigenstates in 1D. *Communications in Mathematical Physics*, 356(1):65–105, Nov 2017.

[AN02] Dorit Aharonov and Tomer Naveh. Quantum NP - A Survey, 2002.

[Ara11] Itai Arad. A note about a partial no-go theorem for quantum PCP. *Quantum Information & Computation*, 11(11-12):1019–1027, 2011.

[AS98] Sanjeev Arora and S Safra. Probabilistic Checking of Proofs: A New Characterization of NP. *Journal of the ACM*, 45(1):70–122, 1998.

[BBT06] Sergey Bravyi, Arvid J. Bessen, and Barbara M. Terhal. Merlin-arthur games and stoquastic complexity. *arXiv preprint arXiv:0611021*, 2006.

[BDOT08] Sergey Bravyi, David P. Divincenzo, Roberto Oliveira, and Barbara M. Terhal. The complexity of stoquastic local hamiltonian problems. *Quantum Info. Comput.*, 8(5):361–385, May 2008.

[BFNW93] László Babai, Lance Fortnow, Noam Nisan, and Avi Wigderson. Bpp has subexponential time simulations unless exptime has publishable proofs. *Comput. Complex.*, 3(4):307–318, October 1993.

[BGK18] Sergey Bravyi, David Gosset, and Robert Knig. Quantum advantage with shallow circuits. *Science*, 362(6412):308–311, 2018.

[BH13] Fernando G. S. L. Brandão and Aram W. Harrow. Product-state approximations to quantum ground states. In *Proceedings of the 45th Annual ACM Symposium on Theory of Computing (STOC '13)*, pages 871–880, 2013.

[BH17] Sergey Bravyi and Matthew Hastings. On complexity of the quantum ising model. *Communications in Mathematical Physics*, 349(1):1–45, Jan 2017.

[BISW17] Dan Boneh, Yuval Ishai, Amit Sahai, and David J. Wu. Lattice-based snargs and their application to more efficient obfuscation. In *Advances in Cryptology - EUROCRYPT 2017 - 36th Annual International Conference on the Theory and Applications of Cryptographic Techniques*, 2017.

[Bra14] Sergey Bravyi. Monte Carlo simulation of stoquastic Hamiltonians. *arXiv preprint arXiv:1402.2295*, 2014.

[BSCG+14] Eli Ben-Sasson, Alessandro Chiesa, Christina Garman, Matthew Green, Ian Miers, Eran Tromer, and Madars Virza. Zerocash: Decentralized anonymous payments from bitcoin. In *IEEE Symposium on Security and Privacy*, pages 459–474. IEEE Computer Society, 2014.

[BT09] Sergey Bravyi and Barbara M. Terhal. Complexity of stoquastic frustration-free hamiltonians. *SIAM J. Comput.*, 39(4):1462–1485, 2009.

[CCD15] Cristian S. Calude, Elena Calude, and Michael J. Dinneen. Guest column: Adiabatic quantum computing challenges. *SIGACT News*, 46(1):40–61, March 2015.

[Coo71] Stephen A Cook. The complexity of theorem proving procedures. In *Proceedings of the Third Annual ACM Symposium*, pages 151–158. ACM, 1971.

[CSV18] Matthew Coudron, Jalex Stark, and Thomas Vidick. Trading locality for time: certifiable randomness from low-depth circuits. *arXiv preprint arXiv:1810.04233*, 2018.

[Din07] Irit Dinur. The PCP Theorem by Gap Amplification. *Journal of the ACM*, 54(3), 2007.

[Dru11] Andrew Drucker. A pcp characterization of am. In *Proceedings of the 38th International Colloquim Conference on Automata, Languages and Programming - Volume Part I*, ICALP'11, 2011.

[EH17] Lior Eldar and Aram W Harrow. Local hamiltonians whose ground states are hard to approximate. In *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017*, pages 427–438, 2017.

[FGGS00] Edward Farhi, Jeffrey Goldstone, Sam Gutmann, and Michael Sipser. Quantum computation by adiabatic evolution. *arXiv preprint arXiv:0001106*, 2000.

[FH14] Michael H. Freedman and Matthew B. Hastings. Quantum systems on non-k-hyperfinite complexes: a generalization of classical statistical mechanics on expander graphs. *Quantum Information & Computation*, 14(1-2):144–180, 2014.

[Gal18] Franćois Le Gall. Average-Case Quantum Advantage with Shallow Circuits. *arXiv preprint arXiv:1810.12792*, 2018.

[GH16] David Gosset and Yichen Huang. Correlation length versus gap in frustration-free systems. *Phys. Rev. Lett.*, 116:097202, Mar 2016.

[GHLS15] Sevag Gharibian, Yichen Huang, Zeph Landau, and Seung Woo Shin. Quantum hamiltonian complexity. *Foundations and Trends in Theoretical Computer Science*, 10(3):159–282, 2015.

[GKR15] Shafi Goldwasser, Yael Tauman Kalai, and Guy N. Rothblum. Delegating computation: Interactive proofs for muggles. *J. ACM*, 62(4):27:1–27:64, 2015.

[GZ11] Oded Goldreich and David Zuckerman. Another proof that BPP $\subseteq$ PH (and more). In *Studies in Complexity and Cryptography*, pages 40–53. Springer-Verlag, 2011.

[Hås01] Johan Håstad. Some optimal inapproximability results. *J. ACM*, 48(4), July 2001.

[Has13a] Matthew B. Hastings. Obstructions to classically simulating the quantum adiabatic algorithm. *Quantum Information & Computation*, 13(11-12):1038–1076, 2013.

[Has13b] Matthew B. Hastings. Trivial low energy states for commuting hamiltonians, and the quantum pcp conjecture. *Quantum Info. Comput.*, 13(5-6):393–429, May 2013.

[HILL99] Johan Håstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. A pseudo-random generator from any one-way function. *SIAM J. Comput.*, 28(4), March 1999.

[IKW02] Russell Impagliazzo, Valentine Kabanets, and Avi Wigderson. In search of an easy witness: exponential time vs. probabilistic polynomial time. *J. Comput. Syst. Sci.*, 65(4):672–694, 2002.

[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*, STOC '97, pages 220–229. ACM, 1997.

[KI04] Valentine Kabanets and Russell Impagliazzo. Derandomizing polynomial identity tests means proving circuit lower bounds. *Computational Complexity*, 13(1):1–46, Dec 2004.

[KSV02] Alexei Kitaev, A Shen, and M N Vyalyi. *Classical and quantum computation*. Graduate studies in mathematics. American mathematical society, Providence (R.I.), 2002.

[KT18] Joel Klassen and Barbara M. Terhal. Two-local qubit Hamiltonians: when are they stoquastic? *arXiv preprint arXiv:1806.05405*, 2018.

[Lev73] Leonid A Levin. Universal sequential search problems. *Problems of Information Transmission*, 9(3):265–266, 1973.

[MLH18] Milad Marvian, Daniel A Lidar, and Itay Hen. On the Computational Complexity of Curing the Sign Problem. *arXiv preprint arXiv:1802.03408*, 2018.

[NC00] Michael A Nielsen and Isaac L Chuang. *Quantum computation and quantum information*. Cambridge Series on Information and the Natural Sciences. Cambridge University Press, 2000.

[NV18] Anand Natarajan and Thomas Vidick. Low-degree testing for quantum states, and a quantum entangled games PCP for QMA. In *59th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2018*, pages 731–742, 2018.

[NVY18] Chinmay Nirkhe, Umesh Vazirani, and Henry Yuen. Approximate low-weight check codes and circuit lower bounds for noisy ground states. In *45th International Colloquium on Automata, Languages, and Programming, ICALP 2018*, pages 91:1–91:11, 2018.

[NW94] Noam Nisan and Avi Wigderson. Hardness vs randomness. *J. Comput. Syst. Sci.*, 49(2):149–167, 1994.

[Osb12] Tobias J Osborne. Hamiltonian complexity. *Reports on Progress in Physics*, 75(2):022001, 2012.

[STV01] Madhu Sudan, Luca Trevisan, and Salil Vadhan. Pseudorandom generators without the xor lemma. *Journal of Computer and System Sciences*, 62(2):236 – 266, 2001.

[SW] Peter Shor and Ryan Williams. Mathoverflow: Complete problems for randomized complexity classes. https://mathoverflow.net/questions/34469/complete-problems-for-randomized-complexity-classes. Accessed: 2019-01-14.

[TD04] Barbara M. Terhal and David P. DiVincenzo. Adaptive quantum computation, constant depth quantum circuits and arthur-merlin games. *Quantum Information & Computation*, 4:134–145, 2004.

[Wat00] John Watrous. Succinct quantum proofs for properties of finite groups. In *41st Annual Symposium on Foundations of Computer Science, FOCS 2000, 12-14 November 2000, Redondo Beach, California, USA*, pages 537–546, 2000.

[ZF87] Stathis Zachos and Martin Furer. Probabilistic quantifiers vs. distrustful adversaries. In *Seventh Conference on Foundations of Software Technology and Theoretical Computer Science*, pages 443–455, 1987.

A Complexity class co-RP and its complete problem

The complexity class co-RP is the "perfect-complete" version of BPP, i.e., errors are only allowed for negative instances.

Definition A.1 (co-RP)**.** A problem $A = (A_{\text{yes}}, A_{\text{no}})$ is in co-RP if and only if there exists a probabilistic polynomial-time algorithm R, where R takes as input a string $x \in \Sigma^*$ and decides on acceptance or rejection of x such that:
Completeness. If $x \in A_{\text{yes}}$, then R accepts x with probability 1.
Soundness. If $x \in A_{\text{no}}$, then R accepts x with probability at most $\frac{1}{2}$.

The only difference of Definitions 2.2 and A.1 is that in MA there is some witness y, unknown by the verification algorithm. We can see it as if the witness in co-RP is trivially the empty string. In this case, it is not surprising that we can define a version of Definition 3.13, where we fix some string in the groundstate.

Definition A.2 (pinned uniform stoquastic frustration-free k-Local Hamiltonian problem)**.** The *pinned uniform stoquastic frustration-free k-Local Hamiltonian* problem, where $k \in \mathbb{N}^*$ is called the locality and $\varepsilon : \mathbb{N} \to [0, 1]$ is a non-decreasing function, is the following promise problem. Let n be the number of qudits of a quantum system. The input is a set of $m(n)$ uniform stoquastic Hamiltonians $H_1, \ldots, H_{m(n)}$ where m is a polynomial, $\forall i \in m(n) : 0 \leq H_i \leq I$ and each H_i acts on k qudits out of the n qudit system. We also assume that there are at most d terms that act non-trivially on each qudit, for some constant d. For $H = \frac{1}{m(n)} \sum_{j=1}^{m(n)} H_j$, one of the following two conditions hold.
Yes. There exists a n-qudit quantum state state $|\psi\rangle$ such that $\langle\psi|0\rangle > 0$ and $\langle\psi|H|\psi\rangle = 0$.
No. For all n-qudit quantum states $|\psi\rangle$, it holds that $\langle\psi|H|\psi\rangle \geq \varepsilon(n)$.

Theorem A.3. *The pinned uniform stoquastic frustration-free k-Local Hamiltonian problem is in co-RP for every constant k and $\varepsilon(n) = \frac{1}{p(n)}$ where p is some polynomial. Also, for some $p'(n) = O(n^2)$, the pinned uniform stoquastic frustration-free 6-Local Hamiltonian problem is co-RP complete.*

Proof. The inclusion in co-RP comes directly from the random-walk proposed by Bravyi and Terhal [BT09], starting from the fixed all-zeros string.

For the hardness part, as in [BT09], we can analyze the reduction of the quantum Cook-Levin theorem for a co-RP circuit. In the yes-instance, we have that the all-zeros string must be in the groundstate, since it is a valid initial configuration, whereas for no-instances, all states have inverse-polynomial frustration. □

ECCC ISSN 1433-8092
https://eccc.weizmann.ac.il