

Bounds on the Quantum Satisfiability Threshold

Sergey Bravyi¹ Christopher Moore² Alexander Russell³

¹IBM Watson Research Center, Yorktown Heights NY 10598, USA

²Center for Quantum Information and Control, University of New Mexico
and the Santa Fe Institute

³University of Connecticut

sbravyi@us.ibm.com moore@cs.unm.edu acr@cse.uconn.edu

Abstract: Quantum k -SAT is the problem of deciding whether there is a n -qubit state which is perpendicular to a set of vectors, each of which lies in the Hilbert space of k qubits. Equivalently, the problem is to decide whether a particular type of local Hamiltonian has a ground state with zero energy. We consider random quantum k -SAT formulas with n variables and $m = \alpha n$ clauses, and ask at what value of α these formulas cease to be satisfiable. We show that the threshold for random quantum 3-SAT is at most 3.594. For comparison, convincing arguments from statistical physics suggest that the classical 3-SAT threshold is $\alpha_c \approx 4.267$. For larger k , we show that the quantum threshold is a constant factor smaller than the classical one. Our bounds work by determining the generic rank of the satisfying subspace for certain gadgets, and then using the technique of differential equations to analyze various algorithms that partition the hypergraph into a collection of these gadgets.

Keywords: quantum computing; random structures; phase transitions; disordered Hamiltonians

1 Introduction

In quantum k -SAT [1], each clause corresponds to a projection operator on the Hilbert space of n qubits,

$$C = (\mathbb{1}_k - |v\rangle\langle v|) \otimes \mathbb{1}_{n-k}.$$

Here $|v\rangle$ is a vector in the 2^k -dimensional Hilbert space of some k -tuple of qubits, $\mathbb{1}_k$ is the identity on that Hilbert space, and $\mathbb{1}_{n-k}$ is the identity on the remaining qubits. A formula is a set of clauses $\phi = \{C_1, \dots, C_m\}$. We say that ϕ is *satisfiable* if there is a state $|\psi\rangle$ which is perpendicular to all the forbidden vectors $|v\rangle$: in other words, if

$$\langle \psi | C_i | \psi \rangle = 1 \text{ for all } i.$$

We call the subspace of such states $|\psi\rangle$ the *satisfying subspace* V_{sat} .

In addition to being the quantum analogue of a canonical NP-complete problem, quantum k -SAT is an illustrative case of a k -local Hamiltonian. In that case, V_{sat} is the subspace spanned by eigenstates of a Hamiltonian $H = \sum_i (I - C_i)$ with zero energy. It was shown in [1] that the decision problem of whether a particular quantum k -SAT formula is satisfiable is in P for $k = 2$, and is QMA₁-complete for $k \geq 4$, where QMA₁ is the subclass of QMA where the probability of acceptance for a yes-instance is 1.

We are also interested in the problem of determining the rank $R_{\text{sat}} = \dim V_{\text{sat}}$ of the satisfying sub-

space, or equivalently the degeneracy of the zero-energy ground states. Determining R_{sat} is a natural quantum analogue of a classical counting problem, namely determining the number of satisfying assignments of a k -SAT formula. Classically, even for $k = 2$ this problem is #P-complete under Turing reductions [2]. In the quantum case, it is not obvious that finding R_{sat} is even in #P, since the satisfying states may be arbitrarily entangled and may have no succinct description. Indeed, it seems to us that one can define a natural quantum version of #P as the class of problems consisting of finding the rank of an eigenspace of a k -local Hamiltonian, although we do not pursue this further here.

In the classical setting, a lively collaboration between computer scientists, statistical physicists, and mathematicians has grown up around the behavior of random k -SAT formulas. These are constructed in the following way. In order to construct a random formula $\phi(n, m)$ with n variables and m clauses, we first construct a random k -uniform hypergraph with n vertices and m edges, by choosing m times, uniformly and with replacement, from the $\binom{n}{k}$ possible k -tuples of vertices. Then, for each edge, we choose uniformly from the 2^k possible combinations of signs for those k literals.

We are particularly interested in the sparse case, where $m = \alpha n$ for some constant α . There is a conjectured phase transition, where these formulas go from

satisfiable to unsatisfiable when α exceeds a critical threshold:

Conjecture 1. *For each $k \geq 3$, there is a constant α_c such that*

$$\lim_{n \rightarrow \infty} \Pr[\phi(n, \alpha n) \text{ is satisfiable}] = \begin{cases} 1 & \text{if } \alpha < \alpha_c \\ 0 & \text{if } \alpha > \alpha_c \end{cases}$$

In the absence of a proof of this conjecture, one can prove statements of the form that ϕ is unsatisfiable with high probability if $\alpha > \alpha^*$, or satisfiable with high probability if $\alpha < \alpha^\ddagger$. Then, assuming that a phase transition exists, α^* and α^\ddagger are upper and lower bounds on the threshold α_c . The state of the art for classical 3-SAT is [3–5]

$$3.52 \leq \alpha_c \leq 4.490,$$

although overwhelmingly convincing arguments from physics (e.g. [6]) indicate that

$$\alpha_c \approx 4.267.$$

2 Summary of Results

Following Laumann et al. [7] we define a random quantum k -SAT formula $\phi(n, m)$ as a random hypergraph, where for each edge we choose the forbidden vector $|v\rangle$ uniformly from the vectors of norm 1 in the Hilbert space $\mathbb{C}_2^{\otimes k}$ of those k qubits. We can then conjecture an analogous phase transition at a critical density α_c^q . Laumann et al. [7, 8] showed that

$$\alpha_{dc} \leq \alpha_c^q \leq \alpha_c,$$

where α_{dc} is the density below which the hypergraph admits a dimer covering, i.e., a one-to-one mapping from clauses to variables. It seems that this coincides with the density at which the 2-core has as many vertices as edges, so that e.g. $\alpha_{dc} \approx 0.92$ for $k = 3$, although we do not know of a rigorous proof of this.

In this paper, we show that the quantum 3-SAT threshold is bounded by

$$\alpha_c^q \leq 3.594. \tag{1}$$

Note that this upper bound is well below the accepted value of the classical 3-SAT threshold. We also show that for all $k \geq 4$,

$$\alpha_c^q \leq 2^k b, \tag{2}$$

where $b \approx 0.573$. Since the classical threshold grows as $2^k \ln 2$ [9], this shows that the ratio α_c^q/α_c is strictly less than 1. We note that, using our approach, Laumann et al. [8] recently improved this to $b = 1/2$.

In order to prove these results, we exploit the observation of [7] that, once the hypergraph G is fixed, R_{sat} takes a generic value $R_{\text{sat}}^{\text{gen}}$ with probability 1. One way to see this is to note that with probability 1, the components of the clause vectors are algebraically independent transcendentals. Then any subdeterminant of the matrix of forbidden vectors is zero if and only if it is zero when these components are replaced by indeterminates. Moreover, for any particular choice of the clause vectors $|v\rangle$ we have $R_{\text{sat}} \geq R_{\text{sat}}^{\text{gen}}$, since this choice can only result in linear dependences among the forbidden vectors and thus increase the rank of the satisfying subspace.

Our bounds work by partitioning random hypergraphs into gadgets for which we can compute $R_{\text{sat}}^{\text{gen}}$ exactly. In order to show that certain partitions exist, we use the technique of differential equations to analyze simple greedy algorithms. To our knowledge, this is the first time that differential equations have been used to prove *upper* bounds on satisfiability thresholds.

3 Expected Gadget Projectors

Our first lemma generalizes a result of [7], which showed that adding a k -clause reduces $R_{\text{sat}}^{\text{gen}}$ by a factor of $1 - 2^{-k}$. Our argument is somewhat simpler.

Lemma 1. *Let G and H be hypergraphs with n and $t \leq n$ vertices respectively. Let $G \cup H$ denote the hypergraph resulting from adding a copy of H , on some subset of G 's vertices, to G . Then*

$$R_{\text{sat}}^{\text{gen}}(G \cup H) \leq 2^{-t} R_{\text{sat}}^{\text{gen}}(H) R_{\text{sat}}^{\text{gen}}(G).$$

Proof. Let Π_H be the projection operator onto the satisfying subspace of H , viewed as a subspace of $\mathbb{C}_2^{\otimes n}$. Now consider its expectation $\mathbb{E} \Pi_H$, taken over the choice of clause vectors $|v\rangle$. Since each $|v\rangle$ is chosen uniformly from the sphere in $\mathbb{C}_2^{\otimes k}$, and since the uniform measure is invariant under any rotation of a single qubit, $\mathbb{E} \Pi_H$ commutes with any one-qubit unitary operator affecting a vertex in H . Since Π_H acts as the identity on the other $n - t$ vertices, it commutes with one-qubit operators on them as well.

Thus Π_H commutes with any Pauli operator. Since these form a basis for the full matrix algebra acting on $\mathbb{C}_2^{\otimes n}$, it follows that $\mathbb{E} \Pi_H$ must be a scalar. Since

$$\text{rk } \Pi_H = 2^{n-t} R_{\text{sat}}^{\text{gen}}(H)$$

holds with probability 1 (where the factor of 2^{n-t} comes from being able to set the other qubits of G arbitrarily), it also holds in expectation. Thus

$$\text{tr } \mathbb{E} \Pi_H = \mathbb{E} \text{tr } \Pi_H = 2^{n-t} R_{\text{sat}}^{\text{gen}}(H),$$

and therefore

$$\mathbb{E} \Pi_H = 2^{-t} R_{\text{sat}}^{\text{gen}}(H) \mathbf{1}.$$

Now suppose the clause vectors of G are in general position. For any choice of clause vectors on H we have

$$R_{\text{sat}}^{\text{gen}}(G \cup H) \leq \text{rk} \Pi_{G \cup H} \leq \text{tr}(\Pi_G \Pi_H \Pi_G).$$

(This follows from the fact that ABA is positive whenever A and B are projection operators.) This is also true in expectation over the clause vectors of H , so

$$\begin{aligned} R_{\text{sat}}^{\text{gen}}(G \cup H) &\leq \mathbb{E} \text{tr}(\Pi_G \Pi_H \Pi_G) \\ &= 2^{-t} R_{\text{sat}}^{\text{gen}}(H) \text{tr} \Pi_G \\ &= 2^{-t} R_{\text{sat}}^{\text{gen}}(H) R_{\text{sat}}^{\text{gen}}(G), \end{aligned}$$

completing the proof.

4 Two Handy Gadgets

We shall apply Lemma 1 by choosing H from two families of hypergraphs that we shall refer to as *sunflowers* and *nosegays*. A (d, k) -*sunflower* is the k -uniform hypergraph consisting of d clauses (edges), each pair sharing a common “center” vertex, see Fig. 1 for an example.

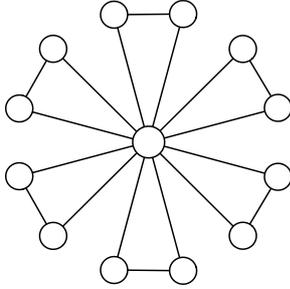


Figure 1: The $(6, 3)$ -sunflower.

A special feature of a sunflower making possible computation of its generic rank is that all qubits except for the central one are the “hanging qubits”, i.e., they participate only in a single clause. This property allows one to construct a unitary transformation mapping an instance of quantum k -SAT defined on a (d, k) -sunflower to a “direct sum” of quantum 2-SAT formulas ϕ_1, \dots, ϕ_d , where ϕ_j is defined on a star graph with j legs. Computing the generic rank of quantum 2-SAT on star graphs using the results of [7] we get

Lemma 2. *Let $S(d, k)$ denote the generic rank of the (d, k) -sunflower. Then*

$$S(d, k) = 2(2^{k-1} - 1)^d \left(\frac{d}{2^k - 2} + 1 \right).$$

Proof. Decompose the Hilbert space of the (d, k) -sunflower as

$$\mathcal{H} = \mathcal{H}_0 \otimes \mathcal{H}_1 \otimes \dots \otimes \mathcal{H}_d,$$

where

$$\mathcal{H}_0 = \mathbb{C}^2, \quad \mathcal{H}_1 = \dots = \mathcal{H}_d = (\mathbb{C}^2)^{\otimes k-1}.$$

Here \mathcal{H}_0 describes the central qubit and \mathcal{H}_j describes the other $k - 1$ qubits on the j th petal for $1 \leq j \leq d$. Let C_j denote the clause on the j th petal, and let $|v_j\rangle \in \mathcal{H}_0 \otimes \mathcal{H}_j$ be its forbidden state. Clearly $|v_j\rangle$ has at most two non-zero Schmidt coefficients, so we can always choose a unitary operator U_j acting on \mathcal{H}_j such that

$$|v_j\rangle = (I_0 \otimes U_j) |u_j\rangle \otimes |0^{\otimes k-2}\rangle. \quad (3)$$

Here $|u_j\rangle$ is some entangled state between the central qubit and the first qubit of the j th petal. Let us refer to the remaining $k - 2$ qubits of the j th petal that are projected onto $|0\rangle$ as ancillas. Clearly U_j does not change the rank, so without loss of generality we can assume that $U_j = \mathbf{1}$ for all j .

If we ignore the ancillas, then we get an instance of quantum 2-SAT on a star graph with d edges which has generic rank $d + 2$, see [7]. Clearly ancillas of the j th petal can be ignored iff they are all set to the state $|0\rangle$. Let us say that such a petal is *active*. Otherwise, if at least one ancilla of the j th petal is $|1\rangle$, it becomes *inactive* because the corresponding clause is already satisfied. An inactive petal contributes a factor of $2^{k-1} - 2$ to the rank. For a fixed subset A of active petals, the generic rank is

$$R(A) = (a + 2)(2^{k-1} - 2)^{d-a}, \quad a \equiv |A|. \quad (4)$$

Therefore, the generic rank for the (d, k) -sunflower is

$$\begin{aligned} R &= \sum_{A \subseteq \{1, \dots, d\}} R(A) \\ &= \sum_{a=0}^d \binom{d}{a} (a + 2)(2^{k-1} - 2)^{d-a} \\ &= 2(2^{k-1} - 1)^d \left(\frac{d}{2^k - 2} + 1 \right). \end{aligned}$$

This completes the proof.

The second useful gadget is a *nosegay*, as shown in Fig. 2. It consists of a single edge, where each of its vertices has some number of additional edges attached to it. We will restrict our attention to the case $k = 3$. Then an (a, b, c) -nosegay has a , b , and c additional edges, for a total of $a + b + c + 1$ edges and $3 + 2(a + b + c)$ vertices.

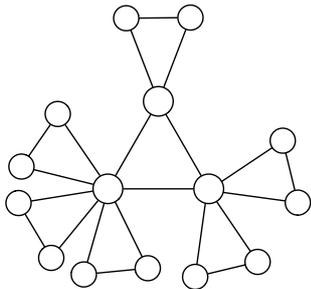


Figure 2: The $(1, 2, 3)$ -nosegay.

Computing the generic rank of a nosegay involves two technical ingredients. Firstly, we use the fact that a generic three-qubit pure state can be transformed by a product of local invertible operators to the GHZ state [12]. Using this fact one can transform an instance of quantum 3-SAT to a “direct sum” of instances of *stoquastic 3-SAT* studied in [13]. An instance of stoquastic 3-SAT is defined by a family of projectors $\phi = (\Pi_1, \dots, \Pi_m)$ which have real non-negative matrix elements in the computational basis, such that every projector acts on at most 3 qubits. The satisfying rank of a stoquastic 3-SAT can be computed using the techniques developed in [13]. It leads to the following result.

Lemma 3. *Let $R_{(a,b,c)}$ denote the generic rank of the (a, b, c) -nosegay. Then*

$$R_{(a,b,c)} = 3^{a+b+c-3} [(a+6)(b+6)(c+6) - (a+3)(b+3)(c+3)].$$

Proof. Let us label the qubits of the central triangle by 1, 2 and 3. These qubits have a , b and c hanging triangles attached to them respectively. Each triangle represents a generic forbidden 3-qubit state.

Let us also define the $[a, b, c]$ -nosegay: it coincides with the (a, b, c) -nosegay except that each hanging triangle is replaced by a hanging edge. Accordingly, the $[a, b, c]$ -nosegay has $n = 3 + a + b + c$ qubits. Each hanging edge represents a generic forbidden 2-qubit state while the central triangle represents a generic forbidden 3-qubit state.

Let $R_{(a,b,c)}$ and $R_{[a,b,c]}$ be the generic ranks of the (a, b, c) -nosegay and the $[a, b, c]$ -nosegay. Repeating

the arguments used to compute the generic rank of sunflowers in Lemma 2, we get

$$R_{(a,b,c)} = \sum_{p=0}^a \sum_{q=0}^b \sum_{r=0}^c 2^{a+b+c-p-q-r} \times \binom{a}{p} \binom{b}{q} \binom{c}{r} R_{[p,q,r]}. \quad (5)$$

In the rest of the section we prove that

$$R_{[a,b,c]} = (a+2)(b+2)(c+2) - (a+1)(b+1)(c+1) \quad (6)$$

which after simple algebra yields

$$R_{(a,b,c)} = 3^{a+b+c-3} [(a+6)(b+6)(c+6) - (a+3)(b+3)(c+3)]. \quad (7)$$

We shall start from using the symmetry of the $[a, b, c]$ -nosegay to bring the forbidden states into a canonical form such that the forbidden state associated with the central triangle is

$$|v_{1,2,3}\rangle = \frac{1}{\sqrt{2}} (|0, 0, 0\rangle - |1, 1, 1\rangle) \quad (8)$$

while the forbidden state associated with any hanging edge (i, j) is the singlet,

$$|v_{ij}\rangle = \frac{1}{\sqrt{2}} (|0, 1\rangle - |1, 0\rangle). \quad (9)$$

We claim that any set of generic forbidden states can be mapped to the ones defined in (8), (9) by applying invertible local operators (ILO) to every qubit. Indeed, it was shown by Dür, Vidal, and Cirac [12] that a generic 3-qubit state is ILO-equivalent to the GHZ state which is in turn ILO-equivalent to the state $|v_{1,2,3}\rangle$ defined in (8). Note that applying an ILO at this step maps the forbidden states on the hanging edges to some entangled 2-qubit states. We can convert them to singlets by applying proper ILO to the free end of every hanging edge. One can easily see that the dimension of the satisfying subspace is invariant under ILO, so it suffices to compute $R_{[a,b,c]}$ with forbidden states defined by (8) and (9).

Our arguments will rely on the fact that the canonical forbidden states define an instance of *stoquastic 3-SAT* studied in [13]. An instance of stoquastic 3-SAT is defined by a family of projectors $\phi = (\Pi_1, \dots, \Pi_m)$ which have real non-negative matrix elements in the computational basis, such that every projector acts on at most 3 qubits. A state $|\psi\rangle$ is a satisfying assignment for ϕ iff $\Pi_a |\psi\rangle = |\psi\rangle$ for every $a = 1, \dots, m$.

In our case the instance ϕ is defined by a family of projectors

$$\begin{aligned} \Pi_{1,2,3} &= \mathbb{1} - |v_{1,2,3}\rangle\langle v_{1,2,3}|, \\ \Pi_{i,j} &= \mathbb{1} - |v_{ij}\rangle\langle v_{ij}|, \end{aligned}$$

where $\Pi_{1,2,3}$ acts on the central triangle and $\Pi_{i,j}$ acts on every hanging edge (i, j) . A direct inspection shows that matrix elements of the above projectors have the following properties:

- Any off-diagonal matrix element belongs to the set $\{0, 1/2\}$
- Any diagonal matrix element belongs to the set $\{1, 1/2\}$

A family of projectors with such properties defines an instance ϕ of *simplified stoquastic 3-SAT* [13, Section 6.3]. In particular, it was shown in [13] that the number of satisfying assignments of ϕ is equal to the number of connected components of a graph $G = (V, E)$, where $V = \{0, 1\}^n$ and $(x, y) \in E$ iff there exists a projector $\Pi_a \in \phi$ such that $\langle x | \Pi_a | y \rangle = 1/2$. (A satisfying assignment associated with a connected component $V_\alpha \subseteq V$ is the uniform superposition of all vertices in V_α , see [13] for details.) It remains to count connected components in the graph G associated with the $[a, b, c]$ -nosegay.

Let us partition the qubits of the nosegay into 3 disjoint subsets A, B, C such that qubit 1 and the free ends of all hanging edges attached to it form A , qubit 2 and the free ends of all hanging edges attached to it form B , and the remaining qubits form C . By definition,

$$\begin{aligned} |A| &= n_a = a + 1, \\ |B| &= n_b = b + 1, \\ |C| &= n_c = c + 1. \end{aligned}$$

Using the projectors that live on the hanging edges we conclude that $(x, y) \in E$ whenever the restrictions of x and y onto any of the subsets A, B, C have the same Hamming weight. Leaving out the projector $\Pi_{1,2,3}$ temporarily we can thus label the connected components of G by triples of integers

$$\begin{aligned} (\alpha, \beta, \gamma), \quad & 0 \leq \alpha \leq n_a, \\ & 0 \leq \beta \leq n_b, \\ & 0 \leq \gamma \leq n_c, \end{aligned}$$

where for any vertex $x \in V$ we define α, β and γ as the Hamming weight of $x|_A, x|_B$, and $x|_C$ respectively. Adding the projector $\Pi_{1,2,3}$ adds extra edges to the graph G , which glue together some components according to the following rules:

$$\begin{aligned} (\alpha, \beta, \gamma) &\sim (\alpha + 1, \beta + 1, \gamma + 1), \\ (\alpha, \beta, \gamma) &\sim (\alpha - 1, \beta - 1, \gamma - 1). \end{aligned}$$

Thus the number of connected components of G is the same as the number of diagonals parallel to the axis $(1, 1, 1)$ in a cube of size $[0, n_a] \times [0, n_b] \times [0, n_c]$. This yields (6) and completes the proof.

5 Upper Bounds on the Critical Density

Let H be a k -uniform hypergraph with n vertices and m edges, corresponding to a formula with m clauses. If we can partition H into a set of sunflowers, where there are n_d sunflowers of each degree d , then by Lemmas 1 and 2, the generic rank of its satisfying subspace is bounded by

$$R_{\text{sat}}^{\text{gen}} \leq 2^n \prod_{d=1}^{\infty} \left(\frac{S(d, k)}{2^{1+d(k-1)}} \right)^{n_d}.$$

Similarly, if we can partition H into a set of noseays, where there are $n_{a,b,c}$ noseays of type (a, b, c) , then the generic rank of its satisfying subspace is bounded by

$$R_{\text{sat}}^{\text{gen}} \leq 2^n \prod_{a,b,c=0}^{\infty} \left(\frac{R_{(a,b,c)}}{2^{3+2(a+b+c)}} \right)^{n_{a,b,c}}.$$

In order to partition the original hypergraph of clauses H into sunflowers we use the following simple algorithm: at each step, choose a uniformly random vertex of H , declare it and its edges to be a sunflower, and remove them from H . Similarly, in order to partition H into noseays, we choose a uniformly random edge of H , declare it and the edges it shares a vertex with to be a noseay, and remove them and its vertices from the hypergraph. If we can show that with high probability (over the execution of the partitioning algorithm and over the distribution of H) the resulting upper bound on $R_{\text{sat}}^{\text{gen}}$ is exponentially small in n for some value of α , it would imply that α is an upper bound on the quantum satisfiability threshold. These ideas are formalized in Theorems 4 and 5 below, which prove the upper bound Eq. (1).

Theorem 4. *Let H be a random 3-uniform hypergraph with n vertices and $m = \alpha n$ edges. If $\alpha > 3.894$, then with high probability the corresponding quantum 3-SAT problem is unsatisfiable.*

Proof. Let H be a k -uniform hypergraph with n vertices and m edges, corresponding to a formula with m clauses. If we can partition H into a set of sunflowers, where there are n_d sunflowers of each degree d , then by Lemmas 1 and 2, the generic rank of its satisfying

subspace is bounded by

$$\begin{aligned}
 R_{\text{sat}}^{\text{gen}} &\leq 2^n \prod_{d=1}^{\infty} \left(\frac{S(d, k)}{2^{1+d(k-1)}} \right)^{n_d} \\
 &= 2^n \prod_{d=1}^{\infty} \left(\left(1 - \frac{1}{2^{k-1}} \right)^d \left(\frac{d}{2^k - 2} + 1 \right) \right)^{n_d} \\
 &= 2^n \prod_{d=1}^{\infty} \left(\left(\frac{3}{4} \right)^d \left(\frac{d}{6} + 1 \right) \right)^{n_d} \tag{10}
 \end{aligned}$$

where in the last line we set $k = 3$.

Clearly (10) is minimized if we have a small number of sunflowers of high degree. Ideally, we would like to characterize the best possible such partition. For now, we content ourselves with the partition resulting from the following simple algorithm: at each step, choose a random vertex, declare it and its edges to be a sunflower, and remove them from the graph.

We can carry out this algorithm in continuous time, by assigning each vertex a uniformly random index $t \in [0, 1]$ and removing vertices in the order of decreasing t . In that case, by the time we remove a vertex v with degree t , its sunflower includes those edges whose other vertices all have index less than t . Since this is true of each of its clauses independently with probability t^{k-1} , and since the original degree distribution of v is Poisson with mean $k\alpha$, the degree of v 's sunflower at the moment when it is removed is Poisson-distributed with mean $k\alpha t^{k-1}$. Integrating over t , the expected number of vertices whose sunflowers have degree d is na_d , where

$$\begin{aligned}
 a_d &= \int_0^1 \frac{e^{-k\alpha t^{k-1}} (k\alpha t^{k-1})^d}{d!} dt \\
 &= \int_0^1 \frac{e^{-3\alpha t^2} (3\alpha t^2)^d}{d!} dt \\
 &= \frac{\Gamma(d + 1/2) - \Gamma(d + 1/2, 3\alpha)}{2\sqrt{3\alpha} d!}, \tag{11}
 \end{aligned}$$

where $\Gamma(a, z) = \int_z^\infty x^{a-1} e^{-x} dx$ is the incomplete Gamma function.

We then upper bound $R_{\text{sat}}^{\text{gen}}$ by cutting off the product above $d_{\text{max}} = 100$, ignoring the effect of the tiny fraction of sunflowers of greater degree. Standard Azuma-type inequalities tell us that, with high probability, the number of sunflowers of degree d is $a_d n + o(n)$ for all $d \leq d_{\text{max}}$. Also, with high probability there are less than $\log n$ pairs of edges which share more than one vertex. Their neighborhoods consist of sunflowers with two petals stuck together. We claim that such a sunflower has lower rank than a normal one, but in any case pretending that these petals are

not stuck together only changes the rank by a constant factor, and the effect of $\log n$ such steps changes the rank by $\text{poly}(n)$. So, with high probability,

$$R_{\text{sat}}^{\text{gen}} \leq \text{poly}(n) \left[2 \prod_{d=0}^{d_{\text{max}}} \left(\left(\frac{3}{4} \right)^d \left(\frac{d}{6} + 1 \right) \right)^{a_d} \right]^n,$$

and therefore

$$\begin{aligned}
 \lim_{n \rightarrow \infty} \frac{1}{n} \ln R_{\text{sat}}^{\text{gen}} &\leq \ln 2 \\
 &\quad + \sum_{d=0}^{d_{\text{max}}} a_d \left(d \ln \frac{3}{4} + \ln \left(\frac{d}{6} + 1 \right) \right).
 \end{aligned}$$

If we set $\alpha = 3.894$, we find that this limit is -1.372×10^{-4} , so $R_{\text{sat}}^{\text{gen}}$ is exponentially small. Thus with high probability in H , $R_{\text{sat}}^{\text{gen}} = 0$ with probability 1 in the clause vectors, and the formula is unsatisfiable.

Next, we improve this result by partitioning the graph into nosebags instead of sunflowers. Although the analysis is slightly harder, the algorithm is equally simple.

Theorem 5. *Let H be a random 3-uniform hypergraph with n vertices and $m = \alpha n$ edges. If $\alpha > 3.594$, then with high probability the corresponding quantum 3-SAT problem is unsatisfiable.*

Proof. At each step we choose a uniformly random edge, declare it and the edges it shares a vertex with to be a nosebag, and remove them and its vertices from the hypergraph. The remaining hypergraph has 3 fewer vertices. Moreover, if we condition on how many edges it has, it is uniformly random in the model where edges are chosen with replacement. This allows us to model this process with differential equations [14, 15].

If t is the number of steps we have taken so far, then there are $n - 3t$ remaining vertices. Let m denote the number of remaining edges. Its expected change on each step is

$$\mathbb{E}[\Delta m] = -1 - \frac{9m}{n - 3t}.$$

Now we write $m = \mu n$ and $t = \tau n$, and rescale this to give a differential equation:

$$\frac{d\mu}{d\tau} = -1 - \frac{9\mu}{1 - 3\tau}. \tag{12}$$

Changing variables to the fraction $\nu = 1 - 3\tau$ of vertices remaining, this is

$$\frac{d\mu}{d\nu} = \frac{1}{3} + \frac{3\mu}{\nu}. \tag{13}$$

With the initial condition $\mu(1) = \alpha$, the solution to this is

$$\mu(\nu) = \frac{\nu}{6} ((6\alpha + 1)\nu^2 - 1). \quad (14)$$

This becomes zero when ν_0 vertices are left, where

$$\nu_0 = \frac{1}{\sqrt{6\alpha + 1}}.$$

At that point, there are no edges left, and the algorithm stops.

With high probability, for any $\nu > \nu_0$, the number of edges remaining when there are νn vertices left is $m(\nu) = \mu(\nu)n + o(n)$. Similar to the proof of Theorem 4, summing over the $(1 - \nu_0)n/3 + o(n)$ steps of the algorithm then gives

$$\lim_{n \rightarrow \infty} \frac{1}{n} \ln R_{\text{sat}}^{\text{gen}} \leq \ln 2 + \frac{1}{3} \int_{\nu_0}^1 \mathbb{E}_{a,b,c} \left[\ln \frac{R_{(a,b,c)}}{2^{3+2(a+b+c)}} \right] d\nu$$

where $R_{(a,b,c)}$ is given by Lemma 3, and where a , b , and c are chosen according to independent Poisson distributions with mean $3\mu/\nu$.

Strictly speaking, the theorem from [15] on which we rely to approximate this discrete random process with this differential equation only applies as long as $m = \Omega(n)$. So, we run the differential equations until $\mu = \epsilon$ for some small $\epsilon > 0$, and ignore the remaining $\epsilon n + o(n)$ edges.

If we set $\alpha = 3.594$ and upper bound the expectation over a , b , and c by ignoring terms where any of them is greater than 50, then evaluating the resulting integral numerically we find that limit is -1.601×10^{-4} . Again $R_{\text{sat}}^{\text{gen}}$ is exponentially small, so these formulas are unsatisfiable with high probability.

There are a number of potential ways to improve this result. First, we can achieve a better partition of the graph into gadgets by prioritizing high-degree vertices. Analogous to [16], we can analyze the resulting partition using a system of coupled differential equations, using the configuration model to keep track of the random graph conditioned on its degree distribution. For the sunflower, this gives a bound of 3.689—a significant improvement over Theorem 4, but not as good as Theorem 5. We have not attempted a partition into nose-gays that prioritizes high-degree clauses.

Second, we have no obligation to consider partitions into sunflowers or nose-gays that can be found in polynomial time. We could also use non-algorithmic proofs that a desirable partition exists. These algorithms simply happen to be both efficient and easy to analyze.

Thirdly, we could use notions of local maximality which have been successful in the classical case, but it is not obvious how to apply these in the quantum setting. When is an entangled satisfying state locally maximal?

In the rest of the paper we use the partition into sunflowers to prove an upper bound on the quantum k -SAT threshold for general k , see Eq. (2). We have made no attempt to optimize this bound beyond the simplest possible argument, but it establishes that the quantum threshold is strictly less than the classical one for all $k \geq 6$.

Theorem 6. *Let $b \approx 0.573$ be the unique positive root of the equation $\ln 2 - 2b + \ln(b + 1) = 0$. Then for all $k \geq 3$, if $\alpha \geq 2^k b$ then with high probability the corresponding quantum k -SAT problem is unsatisfiable.*

Proof. First we rewrite (10) as follows:

$$R_{\text{sat}}^{\text{gen}} \leq 2^n \left(1 - \frac{1}{2^{k-1}}\right)^m \prod_{d=1}^{\infty} \left(\frac{d}{2^k - 2} + 1\right)^{n_d}, \quad (15)$$

since $\sum_d n_d d = m$. Treating the product as a harmonic mean over the vertices, bounding it as an arithmetic mean, and using the fact that the mean degree of a sunflower is $\mathbb{E}[d] = \alpha$ then gives

$$\begin{aligned} \frac{1}{n} \ln R_{\text{sat}}^{\text{gen}} &\leq \ln 2 + \alpha \ln \left(1 - \frac{1}{2^{k-1}}\right) \\ &\quad + \mathbb{E} \left[\ln \left(\frac{d}{2^k - 2} + 1\right) \right] \\ &\leq \ln 2 + \alpha \ln \left(1 - \frac{1}{2^{k-1}}\right) + \ln \left(\frac{\alpha}{2^k - 2} + 1\right). \end{aligned} \quad (16)$$

Rearranging, applying the Taylor series of $\ln(1 - x)$, and setting $\alpha = 2^k b$, we have

$$\begin{aligned} \frac{1}{n} \ln R_{\text{sat}}^{\text{gen}} &\leq \ln 2 + (\alpha - 1) \ln(1 - 2^{1-k}) \\ &\quad + \ln(2^{-k} \alpha + 1 - 2^{1-k}) \\ &< \ln 2 - (\alpha - 1)(2^{1-k} + 2^{1-2k}) \\ &\quad + \ln(2^{-k} \alpha + 1) - \frac{2^{1-k}}{2^{-k} \alpha + 1} \\ &= \ln 2 - 2b + \ln(b + 1) + 2^{1-k} \left(1 - b - \frac{1}{b + 1}\right) \\ &= 2^{1-k} \left(1 - b - \frac{1}{b + 1} + 2^{-k}\right) < 0 \end{aligned}$$

for all $k \geq 3$ since $1 - b - 1/(b + 1) < -1/8$.

We note that Laumann et al. [8] improved the constant b to $1/2$ by evaluating the integral (11) for general k .

In contrast, the classical k -SAT threshold is known to be $(1 - o(1)) 2^k \ln 2$ [9]. Since $b < \ln 2$, it follows

that the quantum threshold is less than the classical one for sufficiently large k . In fact, explicit lower bounds on the classical threshold for finite k from [9] are greater than the upper bounds on the quantum threshold obtained by setting (16) to zero for $k \geq 6$. Using the approach of Theorem 4 improves this to $k \geq 5$.

After this work was completed, we became aware of a remarkable result by Ambainis, Kempe, and Sattath [17] who, using a quantum version of the Lovász Local Lemma, improved the lower bound on the quantum satisfiability threshold to $\alpha_c^q = \Omega(2^k/k^2)$. This raises the following open question: is α_c^q asymptotically proportional to 2^k , as it is in the classical case? If so, what is the constant of proportionality?

Acknowledgments

S.B. received support from the DARPA QUEST program under contract no. HR0011-09-C-0047 and is grateful to CWI for hospitality while this work was being done. C.M. and A.R. are supported by the NSF under grant CCF-0829931, and by the DTO under contract W911NF-04-R-0009.

References

- [1] S. Bravyi, Efficient algorithm for a quantum analogue of 2-SAT. Preprint, [quant-ph/0602108](#).
- [2] L.G. Valiant, The complexity of enumeration and reliability problems. *SIAM J. Comput.* **8**(3) 410–421 (1979).
- [3] Alexis Kaporis, Lefteris Kirousis, Efthimios Lalas, Selecting complementary pairs of literals. Proc. LICS Workshop on Typical Case Complexity and Phase Transitions, 2003.
- [4] Mohammad Hajiaghayi and Gregory Sorkin, The satisfiability threshold for random 3-SAT is at least 3.52. Preprint, (2003).
- [5] J. Díaz, L. Kirousis, D. Mitsche, and X. Pérez-Giménez, A new upper bound for 3-SAT. Proc. FSTTCS 2008, 163–174.
- [6] M. Mezard, G. Parisi, and R. Zecchina, *Science* **297**, 812 (2002).
- [7] C.R. Laumann, R. Moessner, A. Scardicchio, and S.L. Sondhi. Phase transitions and random quantum satisfiability. Preprint, [arXiv:0903.1904](#).
- [8] C.R. Laumann, A.M. Läuchli, R. Moessner, A. Scardicchio, and S.L. Sondhi, On product, generic and random generic quantum satisfiability. Preprint, [arXiv:0910.2058](#).
- [9] D. Achlioptas and Y. Peres, The threshold for random k -SAT is $2^k \ln 2 - O(k)$. Proc. STOC 2003, 223–231.
- [10] M. Molloy, Cores in random hypergraphs and Boolean formulas. *Random Struct. Algorithms* **27**(1) 124–135 (2005).
- [11] V. Chvátal and B. Reed. Mick gets some (the odds are on his side), Proc. 33rd Symposium on the Foundations of Computer Science, 620–627 (1992).
- [12] W. Dür, G. Vidal, and J. I. Cirac, Three qubits can be entangled in two inequivalent ways. *Phys. Rev. A* **62**, 062314 (2000).
- [13] S. Bravyi and B. Terhal, Complexity of stoquastic frustration-free Hamiltonians. Preprint, [arXiv:0806.1746](#).
- [14] N. Wormald, Differential equations for random processes and random graphs. *Annals of Applied Probability* **5**, 1217–1235 (1995).
- [15] N. Wormald, The differential equation method for random graph processes and greedy algorithms. Lectures on *Approximation and Randomized Algorithms*, 73–155, Warsaw (1999).
- [16] D. Achlioptas and C. Moore, “Almost all graphs of degree 4 are 3-colorable.” *Journal of Computer and System Sciences*, **67** (2003) 441–471. Invited paper in special issue for STOC 2002.
- [17] A. Ambainis, J. Kempe, and O. Sattath, A Quantum Lovasz Local Lemma, Preprint, [arXiv:0911.1696](#).