

Clustering in Hilbert space of a quantum optimization problem

S. C. Morampudi,^{1,2} B. Hsu,³ S. L. Sondhi,³ R. Moessner,² and C. R. Laumann¹

¹*Department of Physics, Boston University, Boston, Massachusetts 02215, USA*

²*Max Planck Institut für Physik komplexer Systeme, 01187 Dresden, Germany*

³*Department of Physics, Princeton University, Princeton, New Jersey 08544, USA*

(Received 1 April 2017; published 5 October 2017)

The solution space of many classical optimization problems breaks up into clusters which are extensively distant from one another in the Hamming metric. Here, we show that an analogous quantum clustering phenomenon takes place in the ground-state subspace of a certain quantum optimization problem. This involves extending the notion of clustering to Hilbert space, where the classical Hamming distance is not immediately useful. Quantum clusters correspond to macroscopically distinct subspaces of the full quantum ground-state space which grow with the system size. We explicitly demonstrate that such clusters arise in the solution space of random quantum satisfiability (3-QSAT) at its satisfiability transition. We estimate both the number of these clusters and their internal entropy. The former are given by the number of hard-core dimer coverings of the core of the interaction graph, while the latter is related to the underconstrained degrees of freedom not touched by the dimers. We additionally provide numerical evidence suggesting that the 3-QSAT satisfiability transition may coincide with the product satisfiability transition, which would imply the absence of an intermediate entangled satisfiable phase.

DOI: [10.1103/PhysRevA.96.042303](https://doi.org/10.1103/PhysRevA.96.042303)

I. INTRODUCTION

The concept of a “rugged energy landscape” plays a central role in many parts of statistical physics and optimization theory. The idea is that in glassy systems, this energy landscape in configuration space becomes labyrinthine and difficult to explore with local dynamics. Optimal configurations disappear into deep, narrow valleys which are hard to find amidst the hills and numerous shallow valleys. Classically, this geometric intuition has been at least partially quantified by studying the statistical “clustering” of low-energy configurations into distinct locally connected regions separated by a macroscopic Hamming distance (which measures distance between two solution strings by looking at the number of characters which differ). Such clustering has been used to understand a large range of classical phenomena including the slow down of combinatorial optimization algorithms in certain regimes [1–4], dynamical freezing in spin glasses [5], stable associative memories in neural networks [6], and even improved error correction algorithms in classical codes [7].

Although clustering has been argued to play a role in quantum-mechanical systems (see below), it is less clear how to construct an intrinsic geometric understanding of clustering in Hilbert space. There are several essential difficulties. A low-energy quantum-mechanical state space is a continuous linear vector space where the simplest notion of distance is the Fubini-Study metric which measures the overlap between two states. However, this ascribes the same distance to states which differ by a single spin flip as to those which are macroscopically rearranged. This clearly fails to capture the much larger distance between the states according to local dynamics and local physical expectation values. On the other hand, a purely discrete distance such as the Hamming weight defined by the minimal support of operators mapping between two states is far too sensitive to small continuous deformations in Hilbert space. Moreover, quantum mechanics allows superpositions of macroscopically distinct states. How should one measure the distances between such “cat” states?

Finally, one would like to extend these notions to allow “local degeneracy” within a cluster.

To overcome these difficulties, we propose to focus on short-range correlated states [8] whose connected correlations decay at large physical separation [5]. That is,

$$\langle \psi | O_i O_j | \psi \rangle \xrightarrow{|i-j| \rightarrow \infty} \langle \psi | O_i | \psi \rangle \langle \psi | O_j | \psi \rangle \quad (1)$$

for local operators O_i and O_j centered on sites i and j , where $|i-j|$ measures the interaction graph distance between i and j . Such states are as far as possible from cat states (macroscopic superpositions). In order for a low-energy subspace to be clustered, we take as prerequisite that it must contain a large number of such short-range correlated states which are macroscopically distinct, i.e., for which all local operators have vanishing off-diagonal matrix elements when the number of qubits $N \rightarrow \infty$ (which we refer to as the thermodynamic limit in the following). This definition is analogous to that used to identify distinct symmetry breaking vacua (such as in ferromagnets), but here we do not expect that the short-range correlated states be related by symmetries.

This set of reference states may not provide a complete or unique basis for the low-energy subspace, especially if there are local degeneracies. Rather, it is natural to consider the state space clustered if it can be decomposed into a sum of macroscopically distinct subspaces (which become orthogonal in the thermodynamic limit as we describe below), each of which can be constructed from local operations acting on a reference short-range correlated state.

A prominent setting where classical clustering arises is in random ensembles of the canonical NP -complete problem of satisfiability (k -SAT). Satisfiability consists of determining whether a given k -body classical Hamiltonian composed of a sum of classical projectors has zero-energy (“satisfying”) states [1]. Quantum satisfiability (k -QSAT) generalizes classical k -SAT by allowing the terms to be general noncommuting projectors [9]. Random k -QSAT thus provides a natural place to look for quantum clustering. Random k -QSAT is known




FIG. 1. Illustration of clusters of solutions (grey discs) in Hilbert space at the product satisfiability transition α_p in 3-QSAT. Random quantum satisfiability exhibits a satisfiable phase with nonentangled satisfying states (PRODSATs) at small clause density α and an unsatisfiable (UNSAT) phase at large clause density. For k sufficiently large, these are separated by a satisfiable phase with only entangled satisfying states, but for $k = 3$, the numerics in Sec. IV suggest that the transition from PRODSAT to UNSAT is direct at $\alpha_p \approx 0.91$.

to exhibit both satisfiable (SAT) and unsatisfiable (UNSAT) phases as in classical satisfiability. However, the phases show finer structure due to the possibility of quantum entanglement. More precisely, the SAT phase in general contains two distinct phases—one where the ground states include product states (PRODSAT) [10,11] and another where the only ground states present are entangled (ENTSAT) [12]. For a more detailed review of this phase diagram, see [13–15].

In this paper, we show that quantum clustering arises in the zero-energy space of 3-QSAT at the critical point out of the PRODSAT phase. Our argument is constructive, as we associate the cluster subspaces with discrete product state reference states on an appropriately defined core of the critical instances. These reference states are further in correspondence with dimer coverings of the core, which we show are both exponentially numerous and macroscopically distinct—leading to the same properties for the reference states. Finally, each reference state extends into its own exponentially large satisfying subspace when additional degeneracy associated with the noncore qubits (the “hair”) is taken into account.

In the bulk of the paper, we make these arguments more precise. In Sec. II, we review QSAT and its various phases, and explain the mapping of product ground states to dimer coverings. We then show how this leads to quantum clustering in the 3-QSAT ground space at the PRODSAT critical point. In Sec. IV, we provide numerical evidence suggesting that there is a direct transition to an UNSAT phase from the PRODSAT phase in 3-QSAT—that is, that the ENTSAT phase does not exist at $k = 3$ (see Fig. 1). This helps to motivate our focus on the PRODSAT transition, and is also of independent interest in pinning down the phase diagram of random k -QSAT. In Sec. V, we compute the entropy of clusters S_{core} (i.e., the logarithm of the number of cluster subspaces) using cavity techniques on the associated dimer covering problem. In Sec. VI, we provide an estimate of the internal entropy of a cluster S_{hair}

by considering the contribution of the noncore degeneracy. Finally, we conclude briefly in Sec. VII.

There have been several streams of previous work on clustering in random quantum-mechanical systems. The most direct connections are provided by introducing weak quantum terms in classical models. The clustering properties of the classical energy landscape then provide the backdrop for the tunneling induced by the quantum terms, with consequences for quantum spin-glass dynamics [16–19] and the adiabatic algorithm [20–24]. These treatments take the classical clustering properties as a starting point and do not try to extend them to an intrinsic notion of clustering with finite transverse field. On the other hand, the replica treatment of certain mean-field spin-glass models exhibits replica symmetry breaking in imaginary time even at finite fields [16,25]. In the classical case, replica symmetry breaking is often identified with clustering in configuration space, so this provides a possible approach to identifying clustering in quantum models. However, this microscopic interpretation of replica symmetry breaking is not self-evident, especially in the quantum case. Finally, the solvable AKLT spin-glass model lifts the classical clustering exhibited by a finite temperature vector spin glass into a quantum-mechanically degenerate set of macroscopically distinct ground states [26]. This work bears the closest resemblance to the structure that we report in this paper.

Several recent works have explored the computational difficulty of identifying whether two given low-energy states of a local Hamiltonian can be connected by a sequence of k -local operations which remain at low energy [27,28]. The essential result is that determining whether such low-energy paths exists is quite difficult in the worst case—QMA complete, to be precise. While this notion of path traversal is closely related to clustering, it requires more detailed control of the excitation energetics than we have in random QSAT. It would be very interesting to develop a statistical picture of such paths in this model.

II. TECHNICAL BACKGROUND

The rank 1, k -QSAT Hamiltonian on N qubits is given by

$$H = \sum_{m=1}^M \Pi_m = \sum_{m=1}^M |\phi_m\rangle\langle\phi_m|, \quad (2)$$

where each clause, Π_m , is a rank 1 projector acting on a group of k qubits labeled by $m = (m_1, m_2, \dots, m_k)$. Pictorially, we represent H graphically through an interaction graph G where the clauses are represented by squares and the qubits by circles, as in Fig. 2. Each clause Π_m projects onto one state $|\phi_m\rangle$ in the local 2^k -dimensional Hilbert space, imposing an energy cost on any state which overlaps that direction. Such a Hamiltonian has a zero energy state $|\Psi\rangle$ if and only if $|\Psi\rangle$ is simultaneously annihilated by all of the projectors Π_m . In this case, we say that H is *satisfied* by the state $|\Psi\rangle$. Deciding whether a given k -QSAT Hamiltonian is satisfiable is QMA₁ complete for $k \geq 3$ [9,29], modulo certain technical restrictions.

We focus on the prototypical random k -QSAT ensemble introduced in [11]. There are two sources of randomness:




FIG. 2. Interaction graph representation of an instance of 3-QSAT. Projectors (blue squares) act on the three qubits (brown circles) to which they are connected. In a product satisfiable instance, dimers (salmon ovals) cover every projector but may leave some qubits free (monomers, circled). On a critical core (inside green dashed loop), however, a dimer covering leaves no qubit free. The only free qubits (circled) lie on the hair of the graph (outside green dashed loop).

(1) the discrete choice of interaction graph G and (2) the continuous choice of projectors Π^m associated to each interaction. The latter choice is particularly powerful: *generic* choices [30] of projectors reduce quantum satisfiability to a graph, rather than Hamiltonian, property [11]. More precisely, for fixed G , the dimension $R_G = |\ker H|$ of the satisfying subspace is *almost always* minimal with respect to the continuous choice of $\Pi^m \in \mathbb{C}\mathbb{P}^{2^k-1}$. We refer to this as the “geometrization” property. The satisfying dimension R_G may be lower bounded using the quantum Shearer formula [13] and in certain limits it is conjectured that this formula gives a tight, and thus generic, result.

For generic QSAT instances, zero energy *product* states exist if and only if the graph G admits “dimer coverings,” which match qubits (circles) to interactions (squares) covering all of the interactions (see Fig. 2) [10]. When the graph is fully packed, every interaction is covered by exactly one dimer. Away from full packing, there are uncovered qubits or interactions which are “monomers” on the interaction graph. The correspondence between product ground states and dimer coverings is clearest in the case of product projectors $\Pi^m = \Pi_1^m \otimes \dots \otimes \Pi_k^m$. Such projectors are satisfied if and only if one of the single qubit factors Π_i^m is satisfied, which in turn specifies that the qubit on which Π_i^m acts is in a local state orthogonal to Π_i^m . A dimer covering provides a matching between projectors m and the choice of qubit that satisfies m . More detailed arguments are required to show that these satisfying product states may be extended to generic (nonproduct) projectors [10]. The upshot is that there is a mapping from dimer coverings DCs of the interaction graph to product states |DC>.

We follow previous work and choose the random interaction graph G according to an Erdős-Renyi ensemble with clause density $\alpha = M/N$. That is, each of the possible $\binom{N}{k}$ interac-

tions is placed with independent probability $p = M/\binom{N}{k}$. In this ensemble, each clause has exactly k neighboring qudits while qudit i has d_i neighboring squares where d_i is a random degree which is Poisson distributed with mean $\bar{d}_i = k\alpha$. The clause density α provides a tuning parameter to explore k -QSAT instances ranging from extremely underconstrained at small α to severely overconstrained at large α .

One of the most important geometric features of the Erdős-Renyi interactions graphs is the existence of a core above a critical α_{hc} . The core $G' \subset G$ is the maximal subgraph of G such that every qubit has degree at least two [31]. The core may be constructed efficiently from a given graph G by recursively removing the “dangling” degree 1 qubits along with the neighboring interactions until no degree 1 qubits remain [32]. We refer to the part of the graph removed by this process (i.e., $G - G'$) as the *hair* of G (see Fig. 2).

III. QUANTUM CLUSTERING AT THE PRODSAT TRANSITION

In this section, we show the ground space of random QSAT is clustered at the critical point α_p out of the PRODSAT phase.

To recap, for quantum clustering, we seek a set of short-range correlated reference states each of which generate macroscopically distinct subspaces through local operations at zero energy. More specifically, we need the following: (i) a set of reference states whose connected correlators of local operators vanish and from which subspaces can be generated through local operations. In QSAT, these reference states correspond to product states as described below; (ii) local operators should have vanishing off-diagonal matrix elements between states in different subspaces.

We construct the reference states in two steps. First, we choose a product state |DC> on the core of the interaction graph. These are in one-to-one correspondence with dimer coverings (DCs) of the core. Second, we extend this dimer covering to the full interaction graph and correspondingly extend the reference state to a product state on the full graph. The extension is not unique, even among product states, but this nonuniqueness corresponds to local degeneracy. It is precisely this degeneracy which generates the subspace $K_{\text{hair}}^{\text{DC}}$ associated to each reference state. This leads to a clustering decomposition of the full ground space K as

$$K = \sum_{\text{DC}} |\text{DC}\rangle_{\text{core}} \otimes K_{\text{hair}}^{\text{DC}}. \quad (3)$$

Since the reference states are product states, their connected correlators vanish, i.e., $\langle \hat{O}_i \hat{O}_j \rangle \rightarrow \langle \hat{O}_i \rangle \langle \hat{O}_j \rangle$ whenever \hat{O}_i and \hat{O}_j have nonintersecting support, and certainly when $|i - j| \rightarrow \infty$. These then satisfy the first condition for clustering.

If we consider the situation when the dimer covering is fully packed on the core ($M_c = N_c$, i.e., $\beta = \beta_c = 1$), the only way to go from one dimer configuration on the core to another is to rearrange along a loop on the graph due to the absence of monomers (free qubits). It is important to note that the presence of monomers on the hair (Fig. 2) does not allow us to rearrange the fully packed dimer configuration on the core (analogous to a “backbone”). Then the typical overlap between two states

$|\text{DC}_1\rangle$ and $|\text{DC}_2\rangle$ is

$$\overline{\log\langle\text{DC}_1|\text{DC}_2\rangle} \approx -\gamma\mathcal{L}[\text{DC}_1,\text{DC}_2], \quad (4)$$

where $\mathcal{L}[\text{DC}_1,\text{DC}_2]$ is the length of the loops formed by superposing the two dimer coverings onto each other, and γ is related to the $O(1)$ overlap of two random vectors on the Bloch sphere. Now if the length of the loops increase with system size, then the overlap vanishes in the thermodynamic limit and the states are orthogonal. More importantly, the expectation values of local operators also similarly decay and have vanishing off-diagonal components between the states in the thermodynamic limit.

So the question is whether there is a regime in 3-QSAT where we only have large loops. In fact, this occurs at the transition point out of the PRODSAT phase. Although there is a finite extent of clause density where we have only large loops, it is only at the transition point that we have fully packed dimer coverings of the core. These cannot be rearranged to each other through local changes such as hopping unpaired spins (monomers) to neighboring positions. We show in Sec. V that there is an extensive number of dimer coverings on the core $(|\text{DC}\rangle_{\text{core}})$ with a corresponding entropy (S_{core}) at the transition point. When we add the hair back to the core, it generates a continuous manifold of states around each dimer covering, thus enlarging each point to a cluster $(K_{\text{hair}}^{\text{DC}})$. We estimate the internal entropy of this cluster (S_{hair}) in Sec. VI.

IV. NEW NUMERICAL BOUNDS ON UNSAT TRANSITION FOR $k = 3$

The question of the properties of the end point of the PRODSAT phase goes hand in hand with that about the nature of the adjacent phase. The two generic possibilities are that there is either an (entangled) satisfiable phase, ENTSAT, as happens for k -QSAT for sufficiently large k [12,13], or in its absence, a phase which is unsatisfiable, UNSAT, as happens for $k = 2$ [11]. For numerical studies, small k are more tractable than large ones, so that the question of the existence of an ENTSAT phase for $k = 3$ is of considerable practical relevance.

For $k = 3$ -QSAT, satisfying product states are known to persist up to the threshold $\alpha_{\text{ps}} \approx 0.91$, where the core of the interaction graph has density $\beta_c = M_c/N_c = 1$ [10]. Previous numerical diagonalization studies have reported that the SAT phase might persist slightly further to a transition at $\alpha \approx 1 \pm 0.06$ (quoted error bars) [10].

Here, we report a more focused study which reinforces the possibility that the UNSAT transition actually coincides with the disappearance of product states. More precisely, our numerical results below suggest an upper bound on the UNSAT transition $\beta \leq \beta_c(1 + 1/17)$, which corresponds to $\alpha \leq 0.97$, while we were not able to resolve a stable ENTSAT regime intermediate between the PRODSAT and UNSAT ones.

The exponential dependence of Hilbert space dimension on size N remains a significant barrier to finite-size exact diagonalization studies—even with Moore’s law, this only leads to a linear increase in tractable system size with (human rather than CPU) time.

In order to make progress, we exploit monotonicity, which states that a given QSAT instance is UNSAT if any subgraph is

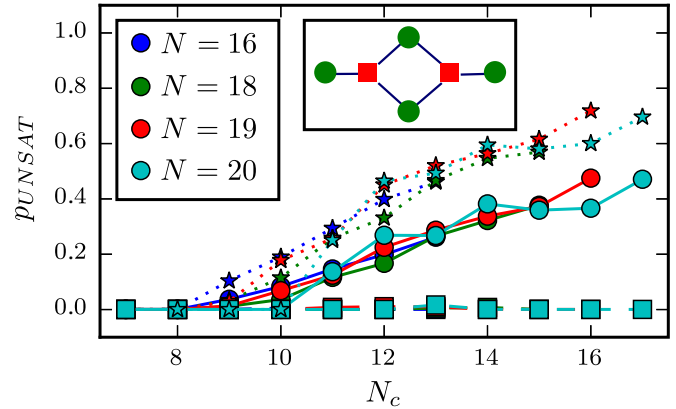


FIG. 3. Probability that a random $k = 3$ core with $M_c = N_c + 1$ is UNSAT. Circles indicate probability across all decidable generated instances; instances with minifans (squares) are essentially always satisfiable at these sizes so that those without (stars) contain all of the UNSAT instances. Inset: A minifan motif: two projectors (red squares) share two qubits (green circles).

UNSAT. This holds because adding a projector increases the number of constraints to be satisfied, and so cannot heal the violation of any existing projectors.

Since we are interested in separating α_c from α_{ps} , we focus on random instances whose cores are just barely overconstrained as far as product states are concerned, i.e., $M_c = N_c + 1$. These may be generated by repeatedly sampling (N, M) graphs and stripping them to their cores until such instances are found. By monotonicity, reattaching the stripped projectors can only make SAT instances UNSAT, and not vice versa.

We use a sparse diagonalization routine on instances of generic projector Hamiltonians in order to determine whether or not these cores are satisfiable—by geometrization, the outcome does not depend on the realization of the generic projectors with high probability, so that our averaging focuses on generating different random cores.

The resulting data for $N = 10\text{--}25$, $N_c = 8\text{--}17$, with $N_{\text{samp}} = 50\text{--}500$ samples per (N, N_c) are shown in Fig. 3. Note that, as N_c increases, we find that the probability of the core being UNSAT increases monotonically. There is no additional dependence on N within statistical fluctuations. This is particularly striking as the effective density of the core $\beta = 1 + 1/N_c$ is monotonically *decreasing* with N_c , so that upon increasing N , one moves closer to the critical value $\beta_c = 1$.

Although for the system sizes available, $N_c \leq 17$, the fraction of UNSAT instances never reaches above $p_{\text{UNSAT}} \approx 0.5$, there is a considerable finite-size dependence of p_{UNSAT} which shows no clear sign of asymptotically reaching a plateau for these system sizes.

Further evidence that the core goes UNSAT at $\beta = 1$ is provided by the observation that a large fraction of the SAT instances owe their existence to the presence of a geometrical motif in the random graph which we call a “minifan,” a subgraph of G in which a pair of clauses shares two qubits (see Fig. 3, inset).

The density of minifans in the random graph vanishes with increasing N as $\propto 1/N$ and thus disappears in the thermodynamic limit. Nonetheless, at accessible sizes, minifans are still quite common and even at $N_c = 17$ roughly 30% of the generated cores have at least one minifan as a subgraph. Crucially, at these sizes, when a core contains a minifan, it is satisfiable (squares, Fig. 3) and thus filtering them out to better approximate the behavior of the large N ensemble raises the probability of UNSAT significantly (stars) and by itself reduces the fraction of SAT instances by about two-fifths for the instances with the largest N . It seems reasonable to expect that identifying and culling more complex finite-size motifs will further lead to an enhancement of p_{UNSAT} .

Taken together, this does leave a tiny sliver in clause density α for the possible existence of an ENT SAT phase at $k = 3$. However, these results provide little evidence of an ENT SAT phase for the systems sizes studied, while being fully consistent with an UNSAT phase setting in at the critical value quoted above.

V. ENTROPY OF CLUSTERS

In this section, we estimate the entropy of dimer coverings S_{core} on the core of a random interaction graph G , as this provides an estimate of the entropy of reference product states on the core. The core is itself a random interaction graph with degree distribution given by a truncated Poisson distribution with mean $\lambda(\alpha)$ (see Appendix for more details) [33].

A. Pauling estimate

Let us construct a simple Pauling-type estimate of S_{core} on a core with N_c nodes of fluctuating degrees $d_i, i = 1 \dots N_c$, M_c k clauses, and an appropriate core clause density, $\beta = M_c/N_c$. There are a total of kM_c links which can be either occupied or unoccupied. Since each clause must be covered by one dimer, we have k allowed configurations out of the 2^k configurations at each clause and $d_i + 1$ allowed configurations out of 2^{d_i} configurations at each node. Thus, the total number of allowed configurations can be estimated

$$\# \approx 2^{kM_c} \left(\frac{k}{2^k}\right)^{M_c} \prod_i \left(\frac{1+d_i}{2^{d_i}}\right) \quad (5)$$

$$\approx k^{M_c} \prod_i \left(\frac{1+d_i}{2^{d_i}}\right). \quad (6)$$

Taking a logarithm and using $kM_c = k\beta N_c = \sum_i d_i$,

$$\frac{S_{\text{core}}}{N_c} = \beta \log k + \frac{1}{N_c} \sum_i \log \left(\frac{1+d_i}{2^{d_i}}\right). \quad (7)$$

When $\beta > 1$, we no longer have any product state solutions since there are more nodes than clauses and hence no more dimer coverings. Plugging in numbers, the entropy of product state solutions at the transition point for $k = 3$ (see the Appendix) and $\beta = 1$ is $S_{\text{core}}/N_c \approx 0.37$.

B. Cavity calculation

As the core is itself a random interaction graph, its local structure is treelike and we use the cavity method to count the




FIG. 4. The cavity method introduces messages between elements of the graph which correspond to probabilities. We introduce additional degrees of freedom on the bonds (triangles) which indicate whether or not the bond is covered by a dimer.

number of dimer coverings. The following treatment is similar to the counting of matchings in random graphs undertaken in [34], suitably generalized to the interaction graph structure of k -QSAT. Recall that in a valid dimer covering, we require that every constraint a is covered but not every spin i needs to be covered. We define variables S_α on the bonds α which denote whether a particular bond is covered by a dimer ($S_\alpha = 1$) or not ($S_\alpha = 0$), and also a fugacity λ to the dimer [Fig. 4(b)]. Hence we associate a partition function to the system through

$$\mathcal{Z} = \sum_{\{S_\alpha\}} \left[\prod_{a,i} \mathbb{I} \left(\sum_{\beta \in \partial a} S_\beta \leq 1 \right) \mathbb{I} \left(\sum_{\beta \in \partial i} S_\beta \leq 1 \right) \right] \lambda^{\sum_\alpha S_\alpha}, \quad (8)$$

where the term inside the $[]$ enforces the hard-core constraint on dimers. In the limit of large λ , the partition function is dominated by terms corresponding to valid dimer coverings provided they exist.

We evaluate the free energy corresponding to this partition function through the cavity method since the interaction graph is locally treelike near the transition. We define the appropriate cavity probabilities

$$P'_{i \rightarrow \alpha}[S_\alpha] = \sum_{\substack{\{S_\beta\}, \\ \beta \in \partial i \setminus \alpha}} \mathbb{I} \left(\sum_{\gamma \in \partial i} S_\gamma \leq 1 \right) \lambda^{S_\alpha} \prod_{\beta} P_{\beta \rightarrow i}[S_\beta], \quad (9)$$

$$P'_{\alpha \rightarrow i}[S_\alpha] = P_{a \rightarrow \alpha}[S_\alpha], \quad (10)$$

$$P'_{a \rightarrow \alpha}[S_\alpha] = \sum_{\substack{\{S_\beta\} \\ \beta \in \partial a \setminus \alpha}} \mathbb{I} \left(\sum_{\gamma \in \partial a} S_\gamma \leq 1 \right) \lambda^{S_\alpha} \prod_{\beta} P_{\beta \rightarrow i}[S_\beta], \quad (11)$$

$$P'_{\alpha \rightarrow a}[S_\alpha] = P_{i \rightarrow \alpha}[S_\alpha]. \quad (12)$$

After normalizing, we define the messages [Fig. 4(b)]

$$q_{i \rightarrow \alpha} := P_{i \rightarrow \alpha}[1] = \frac{\lambda}{1 + \lambda + \sum_{\beta \in \partial i \setminus \alpha} \frac{l_{\beta \rightarrow i}}{1 - l_{\beta \rightarrow i}}}, \quad (13)$$

$$l_{\alpha \rightarrow i} := P_{a \rightarrow \alpha}[1] = q_{a \rightarrow \alpha}, \quad (14)$$

$$q_{a \rightarrow \alpha} := P_{a \rightarrow \alpha}[1] = \frac{\lambda}{1 + \lambda + \sum_{\beta \in \partial a \setminus \alpha} \frac{l_{\beta \rightarrow a}}{1 - l_{\beta \rightarrow a}}}, \quad (15)$$

$$l_{\alpha \rightarrow a} := P_{\alpha \rightarrow a}[1] = q_{i \rightarrow \alpha}. \quad (16)$$

Here, $q_{i \rightarrow \alpha}$ and $l_{\alpha \rightarrow a}$ denote the probability for bond α to be occupied ($= 1$) in the absence of a ; and the other two messages $l_{\alpha \rightarrow i}$ and $q_{a \rightarrow \alpha}$ denote the probability for $\alpha = 1$ in the absence of i . Although two of the messages are redundant, we have retained them as the resulting equations are more symmetric.

The free energy of the system can be expressed in terms of these messages and consists of five contributions coming from the three kinds of vertices ($a = \square$, $i = \circ$, $\alpha = \triangle$) and the two types of bonds ($i\alpha = \circ - \triangle$, $a\alpha = \square - \triangle$),

$$F_a = \log \left[\prod_{\beta \in \partial a} (1 - l_{\beta \rightarrow a}) + \sum_{\beta \in \partial a} \left(l_{\beta \rightarrow a} \prod_{\gamma \in \partial a \setminus \beta} (1 - l_{\gamma \rightarrow a}) \right) \right], \quad (17)$$

$$F_i = \log \left[\prod_{\beta \in \partial i} (1 - l_{\beta \rightarrow i}) + \sum_{\beta \in \partial i} \left(l_{\beta \rightarrow i} \prod_{\gamma \in \partial i \setminus \beta} (1 - l_{\gamma \rightarrow i}) \right) \right], \quad (18)$$

$$F_\alpha = \log \left[(1 - q_{a \rightarrow \alpha})(1 - q_{i \rightarrow \alpha}) + \frac{q_{a \rightarrow \alpha} q_{i \rightarrow \alpha}}{\lambda} \right], \quad (19)$$

$$F_{i\alpha} = \log \left[(1 - l_{\alpha \rightarrow i})(1 - q_{i \rightarrow \alpha}) + \frac{l_{\alpha \rightarrow i} q_{i \rightarrow \alpha}}{\lambda} \right], \quad (20)$$

$$F_{a\alpha} = \log \left[(1 - l_{\alpha \rightarrow a})(1 - q_{a \rightarrow \alpha}) + \frac{l_{\alpha \rightarrow a} q_{a \rightarrow \alpha}}{\lambda} \right]. \quad (21)$$

The total free energy is then given by

$$F = \sum_a F_a + \sum_i F_i + \sum_\alpha F_\alpha - \sum_{i\alpha} F_{i\alpha} - \sum_{a\alpha} F_{a\alpha}. \quad (22)$$

We can also express the average occupation of a constraint a by a dimer ($\langle n_a \rangle := N_{\text{dimer}}^c / M^c$) through

$$\langle n_a \rangle = \frac{\sum_{\beta \in \partial a} l_{\beta \rightarrow a} \prod_{\gamma \in \partial a \setminus \beta} (1 - l_{\gamma \rightarrow a})}{\sum_{\beta \in \partial a} l_{\beta \rightarrow a} \prod_{\gamma \in \partial a \setminus \beta} (1 - l_{\gamma \rightarrow a}) + \prod_{\beta \in \partial a} (1 - l_{\beta \rightarrow a})} \quad (23)$$




FIG. 5. Cavity estimate of core entropy per core spin variable (top) and average occupation of constraints (bottom) as a function of core density β at various fugacities λ [increasing (decreasing) from top to bottom in top (bottom) figure].

and hence we can obtain the entropy density of dimer coverings at a given clause density on the core β as

$$\frac{S_{\text{core}}}{N^c} = \frac{F}{N^c} - \frac{N_{\text{dimer}}^c}{N^c} \log \lambda = \frac{F}{N^c} - \beta \langle n_a \rangle \log \lambda. \quad (24)$$

For regular random graphs with $k = d = 3$ (corresponding to $\beta = 1$), we can solve the equation exactly to obtain

$$\langle n_a \rangle \approx 1 - 0.47\lambda^{-0.5} + 0.06\lambda^{-1} + O(\lambda^{-1.5}), \quad (25)$$

$$\frac{S_{\text{core}}}{N^c} \approx 0.29 + 0.94\lambda^{-0.5} - 0.06\lambda^{-1} + O(\lambda^{-1.5}). \quad (26)$$

For Erdos-Rényi random graphs, we use population dynamics to find a fixed-point solution to the cavity equations. We initialize a population of 10^4 random messages and iterate the cavity equations for 4000 steps after which convergence is achieved and we obtain a fixed-point distribution. The free energy is then sampled from this fixed-point distribution to obtain the entropy density S_{core}/N^c and average occupation $\langle n_a \rangle$. The entropy density corresponding to exact dimer coverings is obtained in the limit of large fugacity λ (see Fig. 5).

In principle, one can bypass the fugacity and instead directly impose a hard constraint in the partition function that every interaction is covered by a dimer. This treatment leads to instabilities in the cavity equations and the associated population dynamics. This is unsurprising as the partition

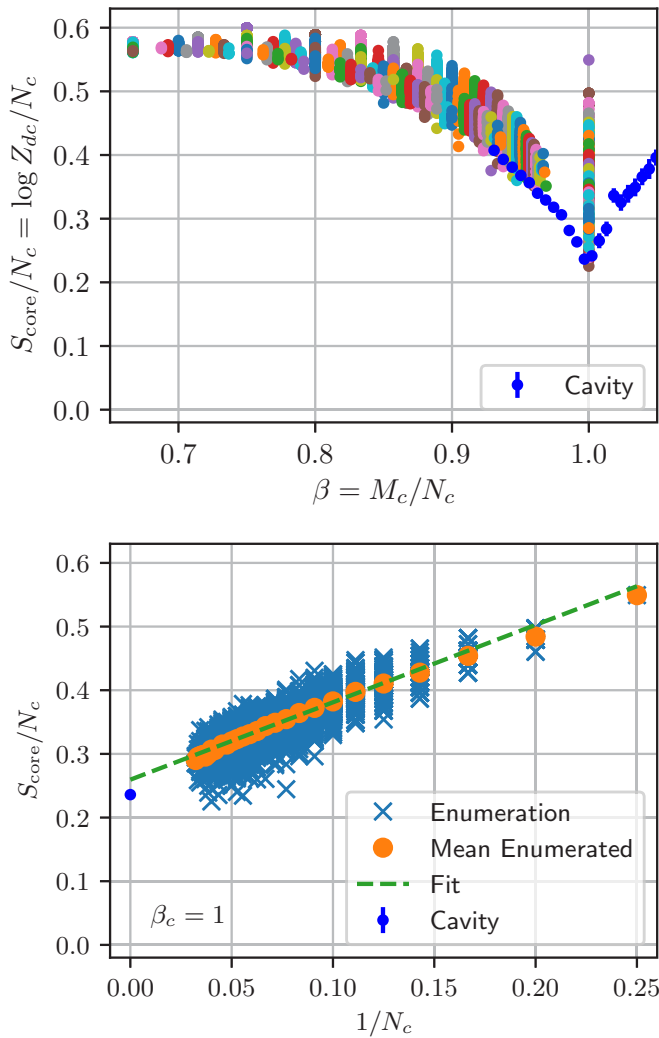


FIG. 6. Entropy of dimer coverings of the core as a function of core density β (top) and at the critical point ($\beta = \beta_c = 1$) as a function of core size N_c (bottom). Top: The different color markers correspond to instances from cores of sizes up to $N_c = 35$. At a given density, larger sizes flow to smaller S_{core} , although this is hard to see in the figure. The cavity estimate is the value extrapolated in the limit of infinite fugacity λ . Bottom: The finite size scaling of S_{core}/N_c at the critical point is consistent with a fit linear in $1/N_c$.

function itself is exponentially large at the critical point $\beta = 1$ and formally zero beyond it since there are no coverings. The introduction of a finite fugacity allows the system to find the maximal covering that fits even for $\beta > 1$ and accordingly undergo a simple first-order transition at the critical point. Analogous instabilities arise in the treatment of the matching problem on random graphs [34].

Finally, for small system sizes, we also use an exhaustive enumeration procedure to find all possible dimer coverings. We find agreement between the two algorithms and compute S_{core} for various values of α_c (see Fig. 6). For $\beta < \beta_c = 1$, the covering problem on the core is under constrained. There are a number of dimer coverings that leave spins uncovered and hence many *small* dimer flips can connect differing possible dimer coverings through monomers representing uncovered qubits. This situation can be contrasted with the one at $\beta = 1$

where we have equal numbers of constraints and qubits so that every node in the core is covered by a dimer (Fig. 2). In this fully packed situation, there is still a macroscopic number of macroscopically distinct dimer coverings with an entropy density $S/N_c \approx 0.23$. Any two dimer coverings differ along *large* loops in the graph. Finally, for $\beta > 1$, there are more clauses than variables and the entropy corresponds to dimer coverings which leave some clauses uncovered and result in a density of monomers. In this situation there are no valid dimer coverings of the graph. (We note in passing that while there are no dimer coverings beyond this point, it has been found that product states can still give good approximations to the ground states of QSAT [14].)

VI. INTERNAL ENTROPY OF A CLUSTER

Now that we have estimated the entropy due to the macroscopically distinct product states on the core (analogous to backbones), we can construct the full state by reattaching the hair and satisfying the projectors on it by constructing dimer coverings on the hair [Eq. (3)]. Unlike dimer coverings on the core, there are free spins on the hair which allow local rearrangements (Fig. 2). However, these free spins do not allow rearrangement on the core since the core is fully packed. The total zero energy entropy at the critical point is thus given by sum of the entropy of clusters and the internal entropy of each cluster due to degeneracy associated to the hair,

$$S_{\text{total}} = S_{\text{core}} + S_{\text{hair}}. \quad (27)$$

For a given dimer covering on the core, the hair entropy has two components—a discrete geometric component due to rearrangement of the dimer coverings and a “zero mode” contribution due to the continuous manifold of satisfying product states associated even to a fixed dimer covering due to the free spins. A naive estimate of the geometric entropy is given by

$$S_{\text{hair}}^{\text{geometric}} = S_2[1 - (M - M_c)/(N - N_c)] \approx 0.5N_h, \quad (28)$$

where N_h is the number of spins on the hair.

Below, we estimate the zero mode entropy corresponding to a fixed dimer covering on the full graph.

Zero mode entropy

Restricting the allowed states to be product states in QSAT gives us PSAT [14]. Zero modes in the PSAT problem correspond to finite dimensional, zero energy submanifolds of product state space $(\mathbb{C}\mathbb{P}^1)^N$. In the quantum world, arbitrary superpositions of the zero energy product states satisfy the related QSAT problem. Thus, the linear span of the (nonlinear) product state manifold is contained in the satisfying space for QSAT and its dimension provides a lower bound on the QSAT degeneracy of a cluster.

In this section, we provide an entropic estimate for the dimension of the quantum zero energy space due to the presence of PSAT zero modes. To warm up, we consider a connected instance of QSAT whose clauses project onto product states and select a zero energy state associated with a particular dimer covering. In this case, the $N - M$ uncovered spins are completely free while the M covered spins are

completely fixed by the dimer covering. The span of this manifold of product states is thus precisely the Hilbert space of $d = N - M$ free qubits, with dimension 2^d .

Let us now try to count this 2^d dimensional space by local arguments. We take a particular zero energy state $|\Omega\rangle$ and choose local bases such that it corresponds to the all \uparrow state, $|\Omega\rangle = |\uparrow\uparrow\dots\uparrow\rangle$. Let $\{w_i^\alpha\}_{\alpha=1\dots d}$ be an orthonormal basis for the $d = N - M$ dimensional kernel of the linearized zero energy condition. The satisfying product states in a small neighborhood of $|\Omega\rangle$ are then parametrized by d complex coordinates δc_α :

$$\begin{aligned} |\Psi(\delta c_\alpha)\rangle &= \prod_{j=1}^N (|\uparrow\rangle + \delta c_\alpha w_j^\alpha |\downarrow\rangle) \\ &= |\Omega\rangle + \delta c_\alpha \sum_j w_j^\alpha b_j^\dagger |\Omega\rangle \\ &\quad + \frac{1}{2} \delta c_\alpha \delta c_\beta P \sum_{j,k} w_j^\alpha w_k^\beta b_j^\dagger b_k^\dagger |\Omega\rangle + \dots, \end{aligned} \quad (29)$$

where we have introduced a collection of bosonic spin flip operators $b_j^\dagger = |\downarrow\rangle_j \langle \uparrow|$ and the projector P implements a hard-core constraint on the bosons. In this language, each order n in the expansion in δc lives in the n boson sector of the full Hilbert space.

Since each order is algebraically independent in δc , the span of the product states in the neighborhood of $|\Omega\rangle$ is the direct sum of the span of the n -boson states at each order. A generic state at order n can be written

$$\frac{1}{n!} \left(\sum_{\delta c} \psi_{\delta c} \delta c_{\alpha_1} \dots \delta c_{\alpha_n} \right) P w_{j_1}^{\alpha_1} \dots w_{j_n}^{\alpha_n} b_{j_1}^\dagger \dots b_{j_n}^\dagger |\Omega\rangle, \quad (30)$$

where the sum runs over an arbitrary finite collection of δc vectors. This expression has two parts: the left-hand bracketed expression is a symmetric rank n tensor over \mathbb{C}^d (i.e., an element of $\text{Sym}^n \mathbb{C}^d$). This tensor of coefficients is a completely general symmetric tensor since the product tensors of the form $\delta c \otimes \delta c \otimes \dots \otimes \delta c$ span this space [35]. The second part is a state constructed from n bosons placed into any of d modes with a hard-core constraint imposed by P .

In the special case of product projectors with d free spins, the mode functions w^α are completely localized on lattice sites. By relabeling the sites, we can write this $w_j^\alpha = \delta_j^\alpha$ and then simplify the product state expansion

$$\sum_{\delta c} \psi_{\delta c} |\Psi(\delta c)\rangle \propto |\Omega\rangle + \left(\sum_{\delta c} \psi_{\delta c} \delta c_\alpha \right) b_\alpha^\dagger |\Omega\rangle \quad (31)$$

$$+ \frac{1}{2} \left(\sum_{\delta c} \psi_{\delta c} \delta c_\alpha \delta c_\beta \right) P b_\alpha^\dagger b_\beta^\dagger |\Omega\rangle + \dots \quad (32)$$

With this expansion, we see that the state space at order n is precisely that of n hard-core bosons on d lattice sites. This has dimension

$$D_n = \binom{d}{n} \quad (33)$$

and we find a total dimension

$$D = \sum_n D_n = 2^d, \quad (34)$$

recovering precisely the counting found for d free spins above.

More generally, for generic nonproduct projectors, we expect the mode matrix w_j^α to have nonzero entries for all α and j on the hair. The mode matrix might also show some additional structure due to the presence of disconnected clusters of hair, but we do not consider this aspect here. In this case, the dimension of the space spanned at order n is no larger than the space $\text{Sym}^n \mathbb{C}^d$ spanned by the coefficient tensor in δc^α and the space of n hard-core bosons in N_h sites. The relevant number of sites here is the number of sites on the hair (N_h) since the core is fully packed and rearrangements on the core and hair do not affect each other,

$$D_n \leq \min \left\{ \binom{n+d-1}{n}, \binom{N_h}{n} \right\}. \quad (35)$$

Since the binomial coefficients monotonically increase with the top argument, the minimizer is given by $\binom{n+d-1}{n}$ for $n+d-1 \leq N_h$. Thus, the total dimension is

$$D = \sum_n D_n \leq \sum_{n=0}^{N_h-d} \binom{n+d-1}{n} + \sum_{n=N_h-d+1}^{N_h} \binom{N_h}{n}. \quad (36)$$

If we now assume that $d = \gamma N_h$, we may estimate this dimension to leading exponential order in N_h by steepest descent. The first term is

$$\begin{aligned} \sum_{n=0}^{N_h-d} \binom{n+d-1}{n} &\sim N_h \int_0^{1-\gamma} dx \binom{(x+\gamma)N_h}{xN_h} \\ &\sim N_h \int_0^{1-\gamma} dx e^{N_h f(x)}, \end{aligned} \quad (37)$$

where the exponent function follows from Stirling's formula as

$$f(x) = - \left[x \log \frac{x}{x+\gamma} + \gamma \log \frac{\gamma}{x+\gamma} \right]. \quad (38)$$

As $f(x)$ is monotonically increasing on the range of the integral, the integral is dominated by the right end point:

$$e^{N_h f(1-\gamma)} = e^{N_h S_2(\gamma)}, \quad (39)$$

where $S_2(\gamma)$ is the entropy of a single coin flip with probability γ .

Similarly, the second term is

$$\begin{aligned} \sum_{n=N_h-d+1}^N \binom{N_h}{n} &\sim N_h \int_{1-\gamma}^1 dx \binom{N_h}{xN_h} \\ &\sim N_h \int_{1-\gamma}^1 dx e^{N_h S_2(x)}. \end{aligned} \quad (40)$$

The entropy function S_2 has a maximum at $x = \frac{1}{2}$ corresponding to an entropy $\log(2)$. If $\gamma > \frac{1}{2}$, the integral includes this saddle point and the result is $\sim 2^{N_h}$ to leading exponential order. If $\gamma < \frac{1}{2}$, the integral is dominated by its left end point and we have

$$N_h \int_{1-\gamma}^1 dx e^{N_h S_2(x)} \sim e^{N_h S_2(1-\gamma)}. \quad (41)$$

Putting this together, we find the QSAT entropy per spin due to the zero modes to be

$$\frac{S_{\text{zero}}}{N_h} = \frac{1}{N_h} \log D \leq \begin{cases} \log(2) & \gamma > \frac{1}{2} \\ S_2(\gamma) & \gamma < \frac{1}{2} \end{cases}. \quad (42)$$

The result is nonrigorous as it relies on analyzing only a leading-order expansion near a given product state, but it provides a nontrivial conjecture for the dimension of the space of states spanned by the PSAT zero modes. The estimated entropy per spin is significantly greater than the $\log(2)\gamma$ associated with having $d = \gamma N_h$ strictly free spins.

VII. CONCLUSIONS AND OUTLOOK

We have generalized the notion of clustering to quantum systems. Unlike the classical definition involving solutions far separated by a Hamming distance, here we define a space as clustered if it can be decomposed into a set of orthogonal subspaces which are macroscopically distinct. That is, matrix elements of local operators vanish between states in the different subspaces. On the other hand, each of these subspaces can be constructed from local operations on a reference short-range correlated state whose connected correlations decay at large distances. We thus identify each subspace as a cluster in the quantum ground space.

We have provided evidence that there is a direct transition from the PRODSAT to UNSAT phase in 3-QSAT. Moreover, quantum clustering occurs in the ground-state space at the critical point separating these two phases. Our argument is essentially constructive as we can identify reference product states from which the cluster subspaces may be generated. These reference product states are in one-to-one correspondence with dimer coverings on the core of the interaction graph.

The entropy of clusters is

$$S_{\text{core}} \approx 0.14N. \quad (43)$$

Each cluster carries an additional internal entropy S_{hair} , which arises from two contributions. First, for a fixed reference dimer covering, there is entropy due to the zero modes,

$$S_{\text{zero}} = N_h S_2(0.2) \approx 0.2N \quad (44)$$

(using $N_h/N \approx 0.4$). This is significantly more than would be expected from the naive estimate that each free spin on the hair contributes $\log(2)$ which would give an entropy $\approx 0.06N$.

The second contribution to the internal entropy arises due to monomer rearrangements on the hair. While it is possible to estimate this hair entropy from calculations analogous to those presented in Sec. VB on the full interaction graph geometry, it is clear that the zero mode entropy and this additional geometric entropy interact. It would be very interesting to develop a theory of this interaction. Here, we content ourselves with a simple upper bound for the total entropy due to the hair by considering all of the hair spins to be completely unconstrained,

$$S_{\text{hair}} < N_h \log(2) \approx 0.28N. \quad (45)$$

The quantum clustering in this model is particularly striking as it provides a precise quantum parallel between the structure of solution space for 1RSB classical models, such as k -XORSAT [33], and the structure of the zero energy Hilbert space for k -QSAT. More broadly speaking, we have used the discrete nature of the solution space of k -QSAT to identify a classical quantity—dimer coverings of the core of the interaction graph—based on which clustering appears naturally in the space of satisfying solutions embedded in the full Hilbert space. It will be interesting to see in what further settings clustering appears.

In particular, it is intriguing to ask if there exists a more “intrinsically” quantum-mechanical definition of clustering. One natural approach would be to aim for a prescription of equivalence classes of ground states based on their connectedness via a set of natural local operations. Given the wide range of possible “landscapes” in—and the continuous nature of—Hilbert space, this may very well turn up much richer structures than even the hierarchical clustering discussed for the solution spaces of classical optimization problems [5].

ACKNOWLEDGMENTS

C.R.L. acknowledges the hospitality of the International Center for Theoretical Physics, Trieste, the Max-Planck Institute for the Physics of Complex Systems, Dresden, the Aspen Center for Physics, the Kavli Institute for Theoretical Physics, and the support of the Sloan Foundation through a Sloan Research Fellowship and the NSF through Grant No. PHY-1656234. S.L.S. acknowledges the support of the John Templeton Foundation. The work in Dresden was supported by the German Science Foundation (DFG) via Grant No. SFB 1143.

APPENDIX: CORE DATA

We summarize a few results derived in [33] regarding the core of random interaction graphs G . The core is the maximal subgraph of G on which all qubit nodes have degree at least 2. The quoted results follow from analyzing the leaf removal algorithm applied to the random interaction graph. In the notation of [33], the clause density α is γ and the number of spins per clause k is p .

The degree distribution (per node of G) on the core is Poissonian for degrees $d \geq 2$:

$$P_c(d) = \begin{cases} 0 & \text{for } d = 0, 1 \\ e^{-\lambda^*(\alpha)} \frac{\lambda^*(\alpha)^d}{d!} & \text{for } d \geq 2 \end{cases}, \quad (\text{A1})$$

where the parameter λ^* corresponds to the largest solution of the equation:

$$e^{-\lambda^*} - 1 + \left(\frac{\lambda^*}{k\alpha}\right)^{1/(k-1)} = 0. \quad (\text{A2})$$

The total number of nodes in the core is

$$N_c(\alpha) = N \sum_{d \geq 2} P_c(d) = N[1 - (1 + \lambda^*)e^{-\lambda^*}] \quad (\text{A3})$$

and the total number of clauses

$$M_c(\alpha) = N \frac{\lambda^*}{k} (1 - e^{-\lambda^*}). \quad (\text{A4})$$

For large α , the core takes over most of the graph up to exponentially small corrections: $\lambda^* \approx k\alpha$ and $N_c = N[1 - (1 + k\alpha)e^{-k\alpha}]$, $M_c = N\alpha(1 - e^{-k\alpha})$.

For reference, at the critical point for $k = 3$,

$$\lambda^*(\alpha_c \approx 0.917) \approx 2.149. \quad (\text{A5})$$

-
- [1] M. Mézard, T. Mora, and R. Zecchina, *Phys. Rev. Lett.* **94**, 197205 (2005).
- [2] R. Mulet, A. Pagnani, M. Weigt, and R. Zecchina, *Phys. Rev. Lett.* **89**, 268701 (2002).
- [3] S. Cocco, O. Dubois, J. Mandler, and R. Monasson, *Phys. Rev. Lett.* **90**, 047205 (2003).
- [4] O. C. Martin, M. Mézard, and O. Rivoire, *Phys. Rev. Lett.* **93**, 217205 (2004).
- [5] M. Mezard, *Spin Glass Theory And Beyond: An Introduction To The Replica Method And Its Applications*, World Scientific Lecture Notes in Physics (World Scientific, Singapore, 1986).
- [6] J. J. Hopfield, *Proc. Natl. Acad. Sci. USA* **79**, 2554 (1982).
- [7] A. Montanari, *Eur. Phys. J. B* **23**, 121 (2001).
- [8] Such states have been called “pure” or “clustering” states in the literature, but we will avoid this terminology to prevent confusion with the quantum-mechanical definition of pure states and the notion of clustering defined in solution space.
- [9] S. Bravyi, [arXiv:quant-ph/0602108](https://arxiv.org/abs/quant-ph/0602108).
- [10] C. R. Laumann, A. M. Läuchli, R. Moessner, A. Scardicchio, and S. L. Sondhi, *Phys. Rev. A* **81**, 062345 (2010).
- [11] C. R. Laumann, R. Moessner, A. Scardicchio, and S. L. Sondhi, *Quantum Inf. Comput.* **10**, 0001 (2010).
- [12] A. Ambainis, J. Kempe, and O. Sattath, *J. ACM* **59**, 24 (2012).
- [13] O. Sattath, S. C. Morampudi, C. R. Laumann, and R. Moessner, *Proc. Natl. Acad. Sci. USA* **113**, 6433 (2016).
- [14] B. Hsu, C. R. Laumann, A. M. Läuchli, R. Moessner, and S. L. Sondhi, *Phys. Rev. A* **87**, 062334 (2013).
- [15] S. Bravyi, C. Moore, and A. Russell, [arXiv:0907.1297](https://arxiv.org/abs/0907.1297).
- [16] Y. Y. Goldschmidt and P.-Y. Lai, *Phys. Rev. Lett.* **64**, 2467 (1990).
- [17] C. R. Laumann, A. Pal, and A. Scardicchio, *Phys. Rev. Lett.* **113**, 200405 (2014).
- [18] C. Baldwin, C. Laumann, A. Pal, and A. Scardicchio, *Phys. Rev. Lett.* **118**, 127201 (2017).
- [19] C. Laumann, A. Scardicchio, and S. L. Sondhi, *Phys. Rev. B* **78**, 134424 (2008).
- [20] S. Knysh and V. N. Smelyanskiy, [arXiv:cond-mat/0602257](https://arxiv.org/abs/cond-mat/0602257).
- [21] S. Knysh and V. N. Smelyanskiy, *Phys. Rev. E* **78**, 061128 (2008).
- [22] B. Altshuler, H. Krovi, and J. Roland, *Proc. Natl. Acad. Sci. USA* **107**, 12446 (2010).
- [23] L. Foini, G. Semerjian, and F. Zamponi, *Phys. Rev. Lett.* **105**, 167204 (2010).
- [24] V. Bapst, L. Foini, F. Krzakala, G. Semerjian, and F. Zamponi, *Phys. Rep.* **523**, 127 (2013).
- [25] J. Ye, S. Sachdev, and N. Read, *Phys. Rev. Lett.* **70**, 4011 (1993).
- [26] C. R. Laumann, S. A. Parameswaran, S. L. Sondhi, and F. Zamponi, *Phys. Rev. B* **81**, 174204 (2010).
- [27] S. Gharibian and J. Sikora, in *Automata, Languages, and Programming: 42nd International Colloquium, ICALP 2015, Kyoto, Japan, July 6-10, 2015, Proceedings, Part I*, edited by M. M. Halldórsson, K. Iwama, N. Kobayashi, and B. Speckmann (Springer, Berlin, 2015), pp. 617–628.
- [28] D. Gosset, J. C. Mehta, and T. Vidick, *Quantum* **1**, 16 (2017).
- [29] D. Gosset and D. Nagaj, Quantum 3-SAT Is QMA1-Complete, in *Proceedings of the 2013 IEEE 54th Annual Symposium on Foundations of Computer Science (FOCS '13)* (IEEE Computer Society, Washington, DC, 2013), pp. 756–765.
- [30] We use the term *generic* to refer to the continuous choice of projectors and *random* to refer to the choice of graph G .
- [31] In much of the relevant literature, this core is referred to as a hypercore as, strictly speaking, G' is the maximal subhypergraph of the hypergraph G where the interactions are viewed as hyperedges. In this paper, we eschew this unnecessarily hyperactive terminology.
- [32] We also remove any isolated degree 0 qubits.
- [33] M. Mézard, F. Ricci-Tersenghi, and R. Zecchina, *J. Stat. Phys.* **111**, 505 (2003).
- [34] L. Zdeborová and M. Mézard, *J. Stat. Mech.: Theory Exp.* (2006) P05003.
- [35] We wish to show that states of the form $\delta c \otimes \delta c \dots \otimes \delta c$ span the space $\text{Sym}^n \mathbb{C}^d$. This space is isomorphic to the space of degree n polynomials in the d variables e_1, \dots, e_d representing a basis of the space \mathbb{C}^d . In particular, this space is spanned by monomials of the form $e_1^{n_1} \dots e_d^{n_d}$ where $\sum n_i = n$ for non-negative integers n_i . If we show that these monomials are in the span of states of the form $f(a_i) = (\sum_i a_i e_i)^n$, then we are done. Observe that derivatives of f with respect to the coefficients a^i are in the span of the image of f . But $\frac{\partial^n f}{\partial^{n_1} a_1 \partial^{n_2} a_2 \dots \partial^{n_d} a_d} = n! e_1^{n_1} \dots e_d^{n_d}$, so we are done.