

This is the accepted manuscript made available via CHORUS. The article has been published as:

Approximating random quantum optimization problems

B. Hsu, C. R. Laumann, A. M. Läuchli, R. Moessner, and S. L. Sondhi

Phys. Rev. A **87**, 062334 — Published 25 June 2013

DOI: [10.1103/PhysRevA.87.062334](https://doi.org/10.1103/PhysRevA.87.062334)

Approximating random quantum optimization problems

B. Hsu,^{1,2} C. R. Laumann,³ A. M. Läuchli,⁴ R. Moessner,² and S. L. Sondhi¹

¹*Department of Physics, Princeton University, Princeton, NJ 08544*

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

³*Department of Physics, Harvard University, Cambridge, MA 02138, USA*

⁴*Institute for Theoretical Physics, University of Innsbruck, 6020 Innsbruck, Austria*

(Dated: June 6, 2013)

We report a cluster of results regarding the difficulty of finding approximate ground states to typical instances of the quantum satisfiability problem k -QSAT on large random graphs. As an approximation strategy, we optimize the solution space over ‘classical’ product states, which in turn introduces a novel autonomous classical optimization problem, PSAT, over a space of continuous degrees of freedom rather than discrete bits. Our central results are: (i) The derivation of a set of bounds and approximations in various limits of the problem, several of which we believe may be amenable to a rigorous treatment. (ii) A demonstration that an approximation based on a greedy algorithm borrowed from the study of frustrated magnetism performs well over a wide range in parameter space, and its performance reflects structure of the solution space of random k -QSAT. Simulated annealing exhibits metastability in similar ‘hard’ regions of parameter space. (iii) A generalization of belief propagation algorithms introduced for classical problems to the case of continuous spins. This yields both approximate solutions, as well as insights into the free energy ‘landscape’ of the approximation problem, including a so-called dynamical transition near the satisfiability threshold. Taken together, these results allow us to elucidate the phase diagram of random k -QSAT in a two-dimensional energy-density–clause-density space.

CONTENTS

I. INTRODUCTION

I. Introduction	1
II. Model	3
III. Low and high temperature expansions	4
A. Low temperature	4
B. High temperature	5
IV. Finite-size numerics	6
A. Greedy local quench	6
B. Simulated annealing	8
C. Exact diagonalization	8
V. Cavity approach	8
A. Cavity method	8
B. Cavity thermodynamic results	10
C. Dynamic instability	10
VI. Concluding remarks	11
Acknowledgments	12
A. Exact singletons	12
1. Singleton clause	12
2. Single product clause	13
B. Uniform measure on \mathbb{CP}^{n-1}	13
C. Hypercore data	14
References	14

Ever since the realization that quantum mechanics could be used to solve computationally difficult problems, there has been an intense effort to understand quantum computers and quantum algorithms. Naturally, much of the algorithmic research front has been focused on those that give a substantial speed up over their classical counterparts. Complexity theory suggests, however, that there are many natural optimization problems that are hard even for a quantum computer to solve. These include both classical NP-complete problems such as satisfiability (SAT) and intrinsically quantum QMA-complete problems such as quantum satisfiability (QSAT) which require a quantum computer even to represent the relevant state space. These hardness results suggest the need for approximate strategies where one settles for less-than-globally-optimal solutions. Unfortunately, this too may be hard as evidenced by the classical PCP theorem [1, 2].

The PCP theorem, in this setting, implies that there are classical energy functions for which it is NP-complete not only to find the ground state but even to find highly excited states with an energy *density* below some threshold. This result has astonishing dynamical consequences. Under the assumption that nature cannot solve NP-complete problems, there are *no* dynamics which can relax these systems to low temperature. This constitutes an abstract argument for the existence of glassiness which, in an amusing parallel with the usual statistical analysis of spin glass theory, makes no reference to the specific dynamics of the system.

As we expect that quantum computers find NP-complete problems just as troubling as classical computers, the classical problems that are hard to approximate

remain so even exploiting quantum dynamics. However, this leaves open the question of how hard intrinsically *quantum* (QMA-complete) problems [3–5] might be to approximate, i.e. of the precise statement of a quantum analog of the PCP theorem. One possibility derives from the usual intuition of quantum statistical mechanics that a finite energy density corresponds to a finite temperature which renders quantum mechanics essentially classical. If this intuition governs then we would expect to be able to efficiently represent approximate finite energy density states for QMA-complete problems on a classical computer and thus the approximation problem would “only” be in NP and thus NP-complete. The other, more interesting, possibility is that the above is not true and that there exist quantum Hamiltonians whose low energy *density* states are both hard to find and impossible to represent with an efficient classical description. Here even the approximation problem would be in QMA. Of course these possibilities are not exclusive—possibly there is a “phase transition” as a function of the energy density where the problem goes from NP to QMA as the energy density is lowered.

In this paper we take an approach to this question which fits naturally within the quantum statistical mechanics of random systems wherein, unlike in the Computer Science setting, we do not attempt to make statements about all possible instances of a problem but instead pick a natural measure on the space of problems and try to investigate the properties of instances which are typical with respect to it. Specifically, we consider the approximation problem for the low energy states of a canonical QMA-complete problem: k -body quantum satisfiability (k -QSAT). As in previous work [6–9], we consider a uniform random ensemble of k -QSAT instances with M clauses, N qubits and clause density $\alpha = M/N$ held fixed in the thermodynamic limit.

The current state of knowledge and our general program with regard to this ensemble is summarized in Fig. 1 where we consider a “phase diagram” in a plane labeled by α and the energy per clause. It is known that typical instances are satisfiable (i.e. have strictly zero energy ground states) by product states at low density, $\alpha < \alpha_{ps}$, satisfiable by entangled states for intermediate densities $\alpha_{ps} < \alpha < \alpha_s$, and strictly unsatisfiable for high densities $\alpha > \alpha_s$. For $k = 3$, of maximum interest in this paper, $\alpha_{ps} = 0.91 \dots$ and $\alpha_s = 1.0 \pm 0.06$ [7, 9]. [10]. The questions we would like to answer are denoted by a set of conjectured “phases” in the Fig. 1. Specifically, these are first to establish the boundaries of the regions where

- it is possible to use product states to achieve the required energy densities
- it is possible to use non-product classical states to achieve the required energy densities
- it is essential to use quantum states to achieve the required energy densities




FIG. 1. (Color online) Heuristic ‘hardness’ phase diagram from approximating random QSAT. PS regions have product state witnesses, C have classical non-product and Q have only quantum witnesses. “Easy” and “Hard” refer to conjectural algorithm-dependent hardness transitions for finding such witnesses.

- it is not possible to achieve the specified energy densities

By “classical states” we mean states for which energies can be evaluated in a time $\text{poly}(N)$ by a classical computer. Additionally, there is the question of how long it takes to find such states. In principle for each category of states there can be algorithm dependent easy and hard regions. It is also a challenge to determine these boundaries. Needless to say, it is not known *a priori* whether all of these regions actually exist (apart from the product state and inaccessible regions). We do know though that the random ensemble exhibits a finite surface/volume ratio for subgraphs of its core; this ‘expansion’ property makes it resistant to simple divide-and-conquer approximation strategies that apply to finite dimensional graphs and thus leaves open the possibility of hard-to-approximate phases.

In this paper we take a step towards answering the above questions by attempting to find the boundaries for the product state regions. Working over the variational class of product states for QSAT defines a classical, continuous spin, energy functional on the same interaction graph. We dub this the PSAT model (see Eq. (2)). The analysis of this model is the central technical goal of this paper. By introducing a classical Gibbs ensemble with temperature $T = 1/\beta$ onto PSAT, we can use statistical mechanical techniques adapted from cavity analysis [11, 12] to probe the structure of the finite energy density landscape of this classical model. We note that cavity analysis of the associated classical model is similar in spirit to the cavity analysis of the coherent state representation of the AKLT spin glass [13, 14], but the quantum-classical mapping involved here is approximate. We will extract complementary information regarding the low en-

ergy structure by direct numerical optimization of finite-size instances using both a local greedy quench algorithm and simulated annealing.

Using these approaches, our aims are two-fold. First, we extract the ground state energy density of PSAT as a function of density α , as this defines the energy density below which QSAT has no product state approximation. Second, we accumulate evidence for the existence of a classically ‘hard’ phase at energy densities approaching the ground state boundary from above. This phase exhibits metastability in the quench dynamics at finite size and a ‘dynamical replica symmetry breaking’ instability in our cavity analysis [11]. In passing, we will provide a characterization of the low and high temperature thermodynamics of PSAT.

Finally we note that our work here is complementary to more rigorous partial results on the quantum PCP problem. The development of a quantum gap amplification procedure [15] provided a quantum analogue of the central step in Dinur’s proof of the classical result [16]. However, the quantum no-cloning theorem has stymied attempts to complete a quantum PCP proof along those lines. Rather, the field has turned to constructing classical approximation strategies and classical approximate witnesses to restricted classes of Hamiltonians in an attempt to constrain possible quantum PCPs. For example, a now classic result is that *commuting* 2-local Hamiltonians have only short-range entangled ground states [17] which possess efficient classical representations and no topological order. Several workers have extended these considerations to finite energy density witnesses to some classes of $k > 2$ -local commuting Hamiltonians [18–21] under various restrictions. Of particular relevance to us is the result that commuting Hamiltonians wherein a small fraction of the terms can be dropped to render the interaction complex 1-localizable possess classically representable states of low energy density [20, 21]. We note that the random ensemble we study in this work is indeed 1-localizable but not commuting. Ground states constructed for commuting Hamiltonians may be perturbatively extended to finite energy density states for almost-commuting Hamiltonians [22]. Finally, for general (non-commuting) quantum Hamiltonians, the rigorous results have relied on constructing guarantees that product states [23] or coarse grained product states (in the 2-local case) [24] exist which provide good approximations to the ground state energy under certain geometric constraints on the interaction graph. Of these results, we note that low degree expanders as we study in this paper are precisely in the class that evades existing efficient approximation guarantees.

In the balance of the paper we do the following. We describe in the next section the random QSAT model, its relationship to the classical PSAT model and various pieces of useful terminology. Section III discusses various limiting expressions for the PSAT thermodynamic quantities. Some details of these calculation can be found in the appendices. These analytic results serve as a useful

test of the numerics. Next, we present results from our greedy local quench, simulated annealing and exact diagonalization algorithms in Section IV. In Section V we analyze the product state phase with the cavity method where we find some evidence for a dynamical instability. Finally, we give concluding remarks in Section VI. As an aid to the reader, we collect several results on the Haar measure on the space \mathbb{CP}^{n-1} in Appendix B and hypercores in Appendix C.

II. MODEL

The QSAT Hamiltonian on N qubits is given by

$$\tilde{H} = \sum_m \Pi^m \quad (1)$$

where Π^m is a k -local rank 1 projector associated with the (hyper-)edge $m = (m_1, m_2, \dots, m_k)$ of an interaction graph G . Π^m projects the state of k qubits onto one ‘direction’ in their 2^k -dimensional joint Hilbert space, exacting on energy cost given by the size of the projection. 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 \tilde{H} is *satisfiable* and that each of the *clauses* Π^m is satisfied by the state $|\Psi\rangle$. We note that deciding whether a given QSAT Hamiltonian \tilde{H} is satisfiable is QMA₁-complete for $k \geq 3$ [5, 25], subject to the technical restriction that the ground state energy of \tilde{H} , if non-zero, is larger than a polynomially small promise gap $\delta = 1/\text{poly}(N)$. As we expect that the energetics of QSAT in the ensembles we consider are extensive, the promise gap should be satisfied with high probability and we will not worry about it further.

The classical PSAT Hamiltonian on N vector spins is given by the variational energy of the QSAT Hamiltonian \tilde{H} over the class of qubit product states:

$$H(\{\hat{n}_i\}) = \sum_m \langle \{\hat{n}_i\}_{i \in \partial m} | \Pi^m | \{\hat{n}_i\}_{i \in \partial m} \rangle \quad (2)$$

$$= \sum_m | \langle \phi^m | \{\hat{n}_i\}_{i \in \partial m} \rangle |^2 \quad (3)$$

Here, $|\hat{n}_i\rangle$ is the spin-1/2 coherent state pointing in the direction of the unit vector \hat{n}_i at site i and $|\phi^m\rangle$ is the k -local state onto which the projector Π^m projects. In complex coordinates, the Hamiltonian takes on the k -quadratic form

$$H = \sum_m | \phi_{a_1 \dots a_k}^m z_{m_1}^{a_1} \dots z_{m_k}^{a_k} |^2 \quad (4)$$

where the a indices run over the two components of the spinor z . In this representation, the Hamiltonian is a simple polynomial, but the coordinates z are constrained to be normalized $z^a (z^a)^* = 1$ and the local phase rotation $z^a \rightarrow e^{i\theta} z^a$ is redundant. As usual, a convenient choice

for the spinor z associated to the direction \hat{n} is

$$z = \begin{pmatrix} \cos(\theta/2) \\ \sin(\theta/2)e^{i\phi} \end{pmatrix} \quad (5)$$

where θ, ϕ are polar coordinates for the direction \hat{n} .

The ground state energy of the PSAT Hamiltonian clearly gives a variational upper bound to the associated QSAT ground state energy. More generally, the *free energy* of PSAT at temperature $T = 1/\beta$ gives a variational upper bound to the *free energy* of QSAT at the same temperature. This can be seen most readily by considering the QSAT partition function in the coherent state basis

$$\begin{aligned} Z_Q &= \text{Tr } e^{-\beta \tilde{H}} = \int \prod_i D\hat{n}_i \langle \{\hat{n}_i\} | e^{-\beta \tilde{H}} | \{\hat{n}_i\} \rangle \\ &= \int \prod_i D\hat{n}_i \sum_{\epsilon} |\langle \{\hat{n}_i\} | \epsilon \rangle|^2 e^{-\beta \epsilon} \end{aligned} \quad (6)$$

and comparing this to the classical partition function

$$\begin{aligned} Z_P &= \int \prod_i D\hat{n}_i e^{-\beta \langle \{\hat{n}_i\} | \tilde{H} | \{\hat{n}_i\} \rangle} \\ &= \int \prod_i D\hat{n}_i e^{-\beta \sum_{\epsilon} |\langle \{\hat{n}_i\} | \epsilon \rangle|^2 \epsilon}. \end{aligned} \quad (7)$$

The variational bound

$$F_Q(T) \leq F_P(T) \quad (8)$$

then follows from the convexity of the exponential function. We note that the normalization of the spinor space has been absorbed into the measure $D\hat{n}_i$. This variational relationship implies that the finite temperature results on PSAT calculated in the remainder of this paper may always be interpreted as direct variational results for the finite temperature QSAT model.

Finally, let us define the random ensemble for k -P/QSAT in more detail. There are two sources of randomness in the ensemble: 1. the discrete choice of interaction graph G and 2. the continuous choice of projectors Π^m associated to each edge. The latter choice is particularly powerful: *generic* choices of projectors reduce quantum satisfiability to a graph, rather than Hamiltonian, property [6]. More precisely, for fixed G , the dimension $R_G = |\ker H|$ of the satisfying subspace of H is *almost always* minimal with respect to the choice of projectors Π^m . Thus, if G can be frustrated by *some* choice of projectors, it is frustrated for almost all such choices. We refer to this as the ‘‘geometrization’’ property.

Also of particular interest to us is the result that zero energy product states for generic QSAT instances exist if and only if the factor graph G has ‘dimer coverings’ between its variables and clauses which cover all of the clauses [7]. We review this construction in some detail below (see Sec. III A). Throughout this paper, we use the term *generic* to refer to the continuous choice of projectors and *random* to refer to the choice of the graph.

As for the discrete choice of random interaction graph G , we follow previous work and use the Erdős-Renyi ensemble with clause density α , in which each of the potential $\binom{N}{k}$ edges appears with probability $p = \alpha N / \binom{N}{k}$. One of the most important geometric features of such interaction graphs is the existence of a (hyper-)core above a critical α_{hc} . The hypercore $G' \subset G$ is the maximal subgraph of G such that every node has degree at least two. One can show that all non-core clauses in $G - G'$ may be satisfied by local product states, whether or not G' is satisfiable, thus all of the nonzero energetics are associated with the core.

III. LOW AND HIGH TEMPERATURE EXPANSIONS

We begin our study of the PSAT model by considering the leading order low and high temperature expansions for the classical Gibbs ensemble,

$$Z = \int \prod_i D\hat{n}_i e^{-\beta H(\{\hat{n}_i\})}. \quad (9)$$

These limits provide non-trivial consistency checks for the more sophisticated numerical and cavity analyses in later sections. As the model is classical, the low temperature expansion follows from equipartition, even in the presence of disorder averaging. However, one must be careful to take account of ‘zero modes’. A classical zero mode corresponds to a degenerate continuous manifold of ground states which therefore do not contribute to the specific heat. We work this out in Sec. III A below.

The disorder-averaged high temperature expansion is a straightforward exercise in combinatorics. We summarize the leading order details in Sec. III B.

A. Low temperature

Generic PSAT has zero energy satisfying states when the interaction graph G possesses ‘dimer coverings’: matchings of spins to clauses such that every clause is covered by exactly one dimer and no spin is used more than once [7]. Let us review the construction of these states. The counting is easiest if we assume that the projectors are all of product form

$$\phi_{a_1 \dots a_k}^m = \phi_{a_1}^{m_1} \dots \phi_{a_k}^{m_k} \quad (10)$$

where each 2-spinor ϕ^m is individually normalized. A zero energy product state may be associated to each dimer covering of the factor graph: if site i is paired with clause m by the dimer covering, we use site i to satisfy clause m by making it orthogonal to the relevant factor (ϵ^{ab} is the antisymmetric tensor)

$$z_i^a = \epsilon^{ab} \phi_b^i. \quad (11)$$

This leaves $N - M$ spinors completely unspecified – these are zero modes of the system. Meanwhile, any deviation of the matched spinor at site i leads to a quadratic energy cost for generic choices of the unspecified spinors, so that there are M complex harmonic modes, or $2M$ real harmonic modes. Recalling that the ground state energy is 0, equipartition produces

$$U \approx TM = T\alpha N. \quad (12)$$

This counting persists for generic projectors (not of product form). To recap [7], let $\{z_i\}$ be a zero energy ground state of H . Since $H(\{z_i\}) = 0$, the z_i must satisfy

$$\phi_{a_1 \dots a_k}^m z_{m_1}^{a_1} \dots z_{m_k}^{a_k} = 0 \quad \forall m = 1 \dots M \quad (13)$$

By rotating the local bases, we can assume that the ground state $z_i = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$ for all sites i , in which case these equations simply reduce to $\phi_{0 \dots 0}^m = 0$. Perturbing $z_i \rightarrow \begin{pmatrix} 1 \\ \delta w_i \end{pmatrix}$, we find (to leading order)

$$\sum_{j=1}^k \phi_{0 \dots 1 \dots 0}^m \delta w_{m_j} = 0 \quad (14)$$

The dimension of the kernel of this set of M linear equations in N unknowns δw_i gives the number of (complex) zero modes. For factor graphs G with dimer coverings, this linear map is generically surjective so the dimensional of the kernel is $N - M$ and we recover the counting from the product projector case.

For graphs that are not product satisfiable, the ground state energy is non-zero on the hypercore of the graph G and we expect no zero modes associated with the core spins. Nonetheless, there may still be zero modes on the non-core spins. An equipartition based estimate of the low temperature energy density in the non-product satisfiable phase looks like:

$$U \approx M\epsilon_0 + T(N_c + (M - M_c)) \quad (15)$$

where ϵ_0 is the ground state energy per clause ($\epsilon_0 > 0$), N_c is the number of spins on the core, M_c is the number of clauses on the core. For random interaction graphs G , the geometric quantities $N_c(\alpha)$ and $M_c(\alpha)$ are known [26] and we have quoted representations for them in Appendix C.

Of course, this estimate only applies at sufficiently low temperatures that all of the non-zero modes indeed look harmonic.

B. High temperature

The high temperature expansion for the disorder averaged free energy is straightforward:

$$\begin{aligned} \bar{F} &= -\frac{1}{\beta} \overline{\log Z(0)} + \overline{\langle H \rangle} - \frac{\beta}{2!} \overline{\langle H^2 \rangle} - \overline{\langle H \rangle^2} \\ &+ \frac{\beta^2}{3!} \overline{\langle H^3 \rangle} - 3 \overline{\langle H \rangle \langle H^2 \rangle} + 2 \overline{\langle H \rangle^3} \\ &- \frac{\beta^3}{4!} \overline{\langle H^4 \rangle} - 4 \overline{\langle H^3 \rangle \langle H \rangle} - 3 \overline{\langle H^2 \rangle^2} + 12 \overline{\langle H^2 \rangle \langle H \rangle^2} - 6 \overline{\langle H \rangle^4} \\ &+ \dots \\ &= -\frac{1}{\beta} C^{(0)} + C^{(1)} - \frac{\beta}{2!} C^{(2)} + \frac{\beta^2}{3!} C^{(3)} - \frac{\beta^3}{4!} C^{(4)} \dots \quad (16) \end{aligned}$$

where $\langle \cdot \rangle$ denotes the infinite temperature average and $\bar{\cdot}$ denotes the disorder average. Normalizing the volume of the spinor space to 4π , we have

$$\log Z(0) = N \log 4\pi \quad (17)$$

independent of disorder realization.

Both the thermal and disorder correlation functions which arise in evaluating the high temperature expansion are with respect to normalized vectors on complex spherical spaces. These satisfy a variant of Wick's theorem which we derive in Appendix B. For example, to evaluate the disorder averaged infinite temperature energy:

$$\begin{aligned} \overline{\langle H \rangle_0} &= \sum_m \overline{\langle E^m \rangle} \\ &= M \overline{\langle |\phi_{a_1 \dots a_k} z_1^{a_1} \dots z_k^{a_k}|^2 \rangle} \\ &= M \overline{\phi_A \phi_B^* \langle z_1^{a_1} (z_1^{b_1})^* \rangle \dots \langle z_1^{a_k} (z_1^{b_k})^* \rangle} \\ &= M \frac{\delta_{AB}}{2^k} \frac{\delta^{a_1 b_1}}{2} \dots \frac{\delta^{a_k b_k}}{2} \\ &= \frac{M}{2^k} \end{aligned}$$

In evaluating the higher order cumulants, two properties are especially important. First, as usual in the high temperature expansion, disconnected clusters vanish. Thus, the second order cumulant is

$$\begin{aligned} C^{(2)} &= \sum_{m,n} \overline{\langle E^m E^n \rangle} - \overline{\langle E^m \rangle \langle E^n \rangle} \\ &= M(c_0^{(2)} + \alpha k^2 c_1^{(2)}) \quad (18) \end{aligned}$$

where the constants

$$c_0^{(2)} = \overline{\langle (E^0)^2 \rangle} - \overline{\langle E^0 \rangle^2} \quad (19)$$

$$c_1^{(2)} = \overline{\langle E^0 E^1 \rangle} - \overline{\langle E^0 \rangle \langle E^1 \rangle} \quad (20)$$

correspond to the single clause energy fluctuations and the overlapping clause energy fluctuations and αk^2 is the

Order	Symbol	Cumulant	Value	Geometry
1	$c_0^{(1)}$	$\langle E^0 \rangle_c$	1×2^{-3}	M
2	$c_0^{(2)}$	$\langle (E^0)^2 \rangle_c$	$7/9 \times 2^{-6}$	M
3	$c_0^{(3)}$	$\langle (E^0)^3 \rangle_c$	$14/15 \times 2^{-9}$	M
4	$c_0^{(4)}$	$\langle (E^0)^4 \rangle_c$	$5078/4455 \times 2^{-12}$	M
	$c_1^{(4)}$	$\langle (E^0)^2 (E^1)^2 \rangle_c$	$-2/243 \times 2^{-12}$	$3M\alpha k^2$

TABLE I. The non-vanishing terms in the high temperature expansion up to order 4 explicitly for $k = 3$. The $\langle \cdot \rangle_c$ indicates the symmetrized joint cumulant of the clause energy variables. The geometric factor indicates the expected number of such terms that contribute to the free energy in the random graph. We break out the factors of 2^{-kn} (n being the order) in the value as we expect the terms to scale as such.

expected number of clauses intersecting a given clause m in one spin.

Second, the disorder average of any power of a single clause energy function produces a number with no thermal fluctuations.

$$\overline{(E^0(z))^n} = \frac{n!}{2^k(2^k + 1) \cdots (2^k + n - 1)} \quad (21)$$

Thus, in any term in the cumulant expansion where all copies of a clause lie in the same thermal average, we may disorder average first and extract the constant. This implies that any cumulant in which a particular clause E^m arises only once vanishes, as ‘constant’ random variables always lead to vanishing joint cumulants. For example, this tells us immediately that $c_1^{(2)} = 0$.

Using these rules, we find that the only terms that contribute to the high temperature expansion to 4th order are those given in Table I. We have included the explicit evaluation of those terms for $k = 3$.

To summarize, the thermodynamic densities to fourth order are

$$F/N = -\frac{1}{\beta} \log 4\pi + \alpha/2^k - \frac{\beta\alpha}{2} c_0^{(2)} + \frac{\beta^2\alpha}{3!} c_0^{(3)} - \frac{\beta^3\alpha}{4!} (c_0^{(4)} + 3\alpha k^2 c_1^{(4)}) \quad (22)$$

$$S/N = \log 4\pi - \beta^2 \frac{\alpha}{2} c_0^{(2)} + \beta^3 \frac{\alpha}{3} c_0^{(3)} - \beta^4 \frac{10\alpha}{4!} (c_0^{(4)} + 3\alpha k^2 c_1^{(4)}) \quad (23)$$

$$U/N = \alpha \left[\frac{1}{2^k} - \beta c_0^{(2)} + \frac{\beta^2}{2} c_0^{(3)} - \frac{\beta^3}{3!} (c_0^{(4)} + 3\alpha k^2 c_1^{(4)}) \right]. \quad (24)$$




FIG. 2. (Color online) Finite energy density phase diagram for 3-QSAT extracted from various numerical methods. Dark blue is QSAT inaccessible region, as estimated by exact diagonalization for size $N = 20$. Medium blue region is PSAT inaccessible region, as estimated by local quench for size $N = 100$. The light blue region indicates where the cavity analysis of PSAT exhibits a dynamical instability. The cavity analysis and simulated annealing data for $N = 80$ roughly agree. More details regarding the numerical techniques and averaging are in the main text.

IV. FINITE-SIZE NUMERICS

The general considerations regarding zero energy states and zero modes provide a surprisingly detailed picture of the satisfiable phase of PSAT, including its low energy density of states in both the satisfiable and non-satisfiable phases. However, these arguments fail to provide any estimate of the ground state energy density when it becomes finite, nor do they help to understand the landscape of the search problem for an algorithm attempting to optimize a particular instance. In this section, we turn to several complementary finite-size numerical studies in order to pin down the PSAT ground state energy density and compare it to the QSAT ground state energy density.

A. Greedy local quench

The most successful numerical line of attack on the PSAT ground state energy density turns out also to be the simplest: the greedy local quench. The algorithm proceeds as follows: Given a PSAT instance H , generate a random initial spin configuration $z_i(0)$ and then iteratively sweep over the sites i in some fixed order. At each step, update the local spin to point in the direction which optimizes the energy of the neighboring clauses:

$$z_i(t+1) = \arg \min_{z_i} \sum_{m \in \partial i} E^m(z_i, \{z_j(t)\}_{j \in \partial m - i}). \quad (25)$$

Here, the integer t keeps track of the number of sweeps in the algorithm. The calculation of the local minimizer is

straightforward. One simply calculates an effective field \vec{h}_m on site i due to each clause m and then orients the spinor z_i in the direction of the net field $\sum_m \vec{h}_m$. As the energy function decreases with every step of the quench, the algorithm is guaranteed to converge toward a local minimum of the energy. We terminate the quench when the relative change in the energy $\delta E/E$ after a full sweep across the N sites is less than a fixed threshold (10^{-4} in all data shown). In applying the quench to the random ensemble, we have generated several hundred instances of size $N = 100$ at varying clause density α , stripped them to their hypercores (as all clauses not in the hypercore can be satisfied independent of the state of the core) and run the quench on each core with ten different initial conditions, taken the lowest resulting energy.

The greedy quench does an excellent job of finding low energy product states in the random 3-PSAT ensemble, as shown in Figure 2. The quench readily finds states with energies equal to zero to machine precision below $\alpha_{ps} = 0.9 \pm 0.06$ (see width of scatter region in Fig. 3 inset). Above this threshold, the found energies begin to grow from zero, in agreement with the rigorous expectation that the ground state energy is no longer precisely zero. We note that the observed error in α_{ps} is consistent with the fluctuations in the core density due to finite-size $N = 100$, which we can roughly estimate $N_c/M_c = 1 \pm \frac{0.7}{\sqrt{N}}$ by noting that $N_c/N = M_c/N = 0.6$ in the thermodynamic limit and assuming that the choice of clauses and sites remaining in the core are selected by an independent random process from the initial random graph with size N .

The observed energy density for $\alpha > \alpha_{ps}$ is both quantitatively very small – even at $\alpha = 8$, the found energies of 0.05 per clause are significantly below the naive high density estimate of $1/8 = 0.125$ per clause – and the critical turn on for α near α_{ps} appears very soft (see inset, Fig. 2). The simplest rigorous upper bound for the large α energy density is to satisfy the first $\alpha_{ps}N$ of the clauses and then upper bound the energy of the remaining clauses by $1/8$. This leads to $E/M \leq (1 - \alpha_{ps}/\alpha)/8$.

A better estimate at large clause density can be made by noting that every spin i is subject to a mean exchange field due to the d_i clauses attached to it, where the degree d_i is a Poisson distributed random variable with mean $k\alpha$. With respect to an infinite temperature spin state, this exchange field is a sum of d_i independent random variables so that it has $O(\sqrt{d_i}) \sim O(\sqrt{k\alpha})$ fluctuations. Thus, by aligning each local spin with its local exchange field, we expect to gain of order $O(\sqrt{k\alpha})$ energy per spin relative to the infinite temperature state. This produces a high density estimate of $E/M = 1/2^k - O(1/\sqrt{\alpha})$. We believe this essential scaling at large α is robust although the coefficient hiding in the O -notation is not so easy to produce – the variance of the exchange field in an infinite temperature state is straightforward to calculate, but the process of tipping spins toward their local exchange field modifies the field on neighbors in a nontrivial




FIG. 3. (Color online) Product state ground state energy found by greedy local quench near critical region $\alpha_{ps} \approx 0.92$ at finite size $N = 100$. Error bars indicate instance-to-instance fluctuations over $N_{samp} \approx 3 - 400$ instances per density α . (inset) Semilogarithmic scatter plot of same data.

manner. In any event, we note that a fit (not shown) to this quadratic form at high clause density $\alpha > 2$ produces good agreement with an extracted coefficient of roughly $0.18 - 0.21/\sqrt{\alpha}$ (at $k = 3$).

The dynamics of the local quench show mild evidence for the presence of metastable minima in the classical PSAT energy landscape above the transition region, as shown in the typical time traces shown in Fig. 4. For densities below $\alpha_{ps} = 0.91$, the typical traces decay continuously to very low values while at higher density, they exhibit a larger variance of found energies for different initial conditions and step-like metastable behavior in the individual trajectories.




FIG. 4. (Color online) Typical energy traces as a function of time for 10 independent sweeps of the greedy quench algorithm in a typical instance at each of $\alpha = 0.88, 0.90, 1.0$ ($N = 100$). The curve in the left panel continue to fall exponentially (linearly on the log-scale) to machine precision (10^{-16}).




FIG. 5. (Color online) Simulated annealing quality and the step size dependence of the best found PSAT energy for $k = 3$ and $N = 16$. The step size controls the maximum change of a single component of the wave function per update trial. In the PRODSAT phase reducing the step size leads to a monotonically decreasing energy. In the UNPRODSAT phase the behavior is irregular.

B. Simulated annealing

As a second approach to finding low energy product states, we implemented a standard simulated annealing protocol to cool the PSAT Hamiltonian to low temperature on randomly generated instances. At each step of this procedure, a randomly sampled spin is perturbed by a small random step and the move is accepted following a Metropolis rule at temperature T . The temperature decays exponentially with a decay constant of 1.003 per 10^4 Monte Carlo sweeps, beginning with temperature 1 and finishing at $T = 10^{-8}$.

The resulting disorder averaged energies $N = 10 - 80$ qubits show extremely limited finite size scaling (not shown) so that we have included only the largest size data ($N = 80$) in Fig. 2. We note that simulated annealing fails to find energies as low as the local quench protocol but agrees well with the low temperature extrapolation of the cavity results (see Sec. V). While it is possible that the annealing results would improve by further optimizing the cooling schedule, we suspect that the timescale for this cooling may become very long in the UNSAT regime. This is indicated by the transition in the behavior of the step size dependence of the simulated annealing results as the clause density is tuned from the SAT to UNSAT phase, see Fig. 5.

C. Exact diagonalization

Finally, we include exact diagonalization results for systems of up to $N = 20$ qubits in Fig. 2. These provide a numerical baseline for the QSAT ground state energy density in the UNSAT regime (striped symbols) and

demonstrate its separation from the best found product states by local quenching. We note that the finite size scaling of these results looks quite converged for $k = 3$ (not shown), as seen in previous studies as well [7], but that for larger k there are still strong finite size effects for the achievable system sizes, as the graphs are quite unlike the thermodynamic random graph at these sizes. This is indeed the central reason that we limit the numerical investigations of this paper to the $k = 3$ case.

V. CAVITY APPROACH

In this section, we look at the PSAT model using so-called cavity techniques [11]. Our motivations for doing so are twofold. Firstly, while the previous finite size numerical work obtained low ground state energy estimates, these were limited to relatively small system sizes of at most $N = 100$ variables. With the cavity approach we can explore the thermodynamic limit where $N \rightarrow \infty$ directly. Secondly, we also observed regions of the parameter space where the finite size algorithms began to slow down. In the local field quench, this slow-down coincided with what appeared to be metastable states, while in the simulated annealing results, the algorithm reached a barrier in finding lower energy ground states (see Figures 5 and 4). However, these observations in a finite size system do not indicate whether these metastable states reflect true transitions in the thermodynamic structure of the configuration space.

The cavity approach allows us to detect the appearance of such states as $N \rightarrow \infty$. This is indicated by the appearance of a dynamical instability where the simplest (replica-symmetric) cavity ansatz breaks down [12]. From an algorithmic perspective, it is unclear whether such a boundary is truly relevant, as indicated by the robust results of the local field quench. However, in the spin glass lore, it is believed that such barriers ultimately limit local search procedures for discrete classical systems and we believe the discovery of similar instabilities in the continuous degree of freedom PSAT system represents an interesting avenue for future detailed exploration.

In the next section, we introduce the cavity approach appropriate to the PSAT problem and provide the thermodynamic quantities that can be obtained from it. After this, we discuss how the replica symmetry broken phase can be detected in the cavity approach.

A. Cavity method

The cavity approach to the classical PSAT model consists of studying the cavity distributions $P_{i \rightarrow m}(\hat{n}_i)$. This is the marginal classical probability distribution for the unit vector \hat{n}_i at site i , in the absence of the clause m . The so-called belief propagation equations follow from assuming that the cavity distributions for the neighbors of the clause m are independent when m is removed. Un-

der this assumption, the cavity distributions as shown in

Fig 6 obey the equations

$$Q_{m \rightarrow i}(\hat{n}_i) = \frac{1}{Z_{m \rightarrow i}} \int \prod_{j \in \partial m - i} (D\hat{n}_j P_{j \rightarrow m}(\hat{n}_j)) e^{-\beta E^m(\{\hat{n}_j\}_{j \in \partial m})} \quad (26)$$

$$P_{i \rightarrow m}(\hat{n}_i) = \frac{1}{Z_{i \rightarrow m}} \prod_{n \in \partial i - m} Q_{n \rightarrow i}(\hat{n}_i) \quad (27)$$

Here, $Q(\hat{n})$ is an intermediate cavity function that is useful primarily as a bookkeeping device. For a k -local interaction graph with M clauses, this provides $2Mk$ functional equations for $2Mk$ unknown cavity distributions $P(\hat{n})$ and $Q(\hat{n})$.

In the thermodynamic limit, we expect that the cavity distributions themselves are *i.i.d* random variables with a self-averaging distribution $\mathbb{P}[P(\cdot)]$ and $\mathbb{P}[Q(\cdot)]$. Under this assumption, we obtain self-consistent equations for the functional distributions $\mathbb{P}[P(\cdot)]$ and $\mathbb{P}[Q(\cdot)]$:

$$\mathbb{P}[P(\cdot)] = \mathbb{E}_{d, Q_n} \delta \left[P(\cdot) - \frac{1}{Z} \prod_{n=1}^d Q_n(\cdot) \right] \quad (28)$$

$$\mathbb{P}[Q(\cdot)] = \mathbb{E}_{E^m, P_j} \delta \left[Q(\cdot) - \frac{1}{Z} \int \prod_{j=1}^{k-1} D\hat{n}_j P_j(\hat{n}_j) e^{-\beta E^m(\{\hat{n}_j\}, \cdot)} \right] \quad (29)$$

These are the so-called replica-symmetric cavity equations for the PSAT model. Here, \mathbb{E} stands for an expectation value with respect to the subscripted quantities, such as d , the onward degree of a spin, and E^m , the choice of projector on the clause. The normalization factors Z implicitly depend on these random variables.




FIG. 6. The definition of the cavity distributions on a locally tree-like interaction graph. Squares indicate clauses, labeled by letters, and spheres the spinor variables, labeled by integers.

In general, the cavity equations (28) cannot be solved exactly for the functional distributions $\mathbb{P}[P(\hat{n})]$, $\mathbb{P}[Q(\hat{n})]$. While the relevant cavity equations are often analytically intractable even in models with discrete local degrees of freedom, such as in classical k -SAT, in PSAT, the cavity messages $P(\hat{n})$ and $Q(\hat{n})$ are general normalized probability distributions on the sphere, and the cavity distributions $\mathbb{P}[\cdot]$ are distributions on that functional space. In the presence of sufficient symmetry, as, for example in rotationally invariant models [14], one can make progress in the symmetric phase analytically, but PSAT has no such rotational symmetry and we must resort to numerical methods throughout the phase diagram.

Thus, we follow established practice[11, 27–29] in studying cavity equations by representing the distributions, $\mathbb{P}[P(\cdot)]$ and $\mathbb{P}[Q(\cdot)]$, by a large *population* of sampled messages $P(\cdot)$ and $Q(\cdot)$. These probability distributions on the sphere in turn are represented by discretized probability vectors P_i , with a finite number $i = 1 \cdots N_{\text{disc}}$ of patches roughly equally sampled over the surface of the sphere. By varying the discretization, we can check for convergence of thermodynamic observables to the continuous degree of freedom $N_{\text{disc}} \rightarrow \infty$ limit.

To solve the cavity equations (28), we employ *popula-*

tion dynamics. We initialize the population of messages $P(\cdot)$ and $Q(\cdot)$ randomly and then proceed iteratively to use the belief propagation equations (26) to generate the new messages from the existing population. As the population size N_{pop} and total number of time steps T approach infinity, we expect the resulting population to converge to a fixed point of the cavity equations [11].

For example, the internal energy per clause follows directly from the expected value of the energy of a randomly added clause in the presence of k sampled cavity messages.

$$U/M = \mathbb{E}_{E^a, P} \left[\frac{1}{Z} \int D\hat{n}_1 D\hat{n}_2 D\hat{n}_3 E^a(\hat{n}_1, \hat{n}_2, \hat{n}_3) e^{-\beta E^a(\hat{n}_1, \hat{n}_2, \hat{n}_3)} P_{1 \rightarrow a}(\hat{n}_1) P_{2 \rightarrow a}(\hat{n}_2) P_{3 \rightarrow a}(\hat{n}_3) \right] \quad (30)$$

The expression for the Bethe free energy per variable $f = F/N$ is somewhat more complicated [11], as it follows from balancing the change in free energy of various graph surgery operations,

$$f(\alpha, \beta) = f_v(\alpha, \beta) + \alpha f_a(\alpha, \beta) - \alpha k f_e(\alpha, \beta) \quad (31)$$

where,

$$\begin{aligned} -\beta f_v &= \mathbb{E}_{\ell, Q} \log \left[\int D\hat{n} Q_{1 \rightarrow i}(\hat{n}) \dots Q_{\ell \rightarrow i}(\hat{n}) \right] \\ -\beta f_a &= \mathbb{E}_{E^a, P} \log \left[\int D\hat{n}_1 D\hat{n}_2 D\hat{n}_3 e^{-\beta E^a(\hat{n}_1, \hat{n}_2, \hat{n}_3)} P_{1 \rightarrow a}(\hat{n}_1) P_{2 \rightarrow a}(\hat{n}_2) P_{3 \rightarrow a}(\hat{n}_3) \right] \\ -\beta f_e &= \mathbb{E}_{P, Q} \log \left[\int D\hat{n} P(\hat{n}) Q(\hat{n}) \right]. \end{aligned} \quad (32)$$

These represent the change in the free energy from the addition of a single clause which has contributions coming from the variables added (f_v), the clause itself (f_a) and each edge connecting variable to edge (f_e). Of course, with the internal energy U and free energy F in hand, it is straightforward to extract the entropy S from the thermodynamic relation $F = U - TS$.

B. Cavity thermodynamic results

We plot the energy density in Fig. 7 and free energy density in Fig. 8 at various temperatures, clause densities and discretizations. The simplest results are at high temperature where the discretization error becomes small (not shown) and we expect detailed agreement with the high temperature expansion. Indeed, the expansion to 4th order, given by equations (22)-(24), works extremely well for $\beta \lesssim 5 - 10$, as can be seen from the (no fit) agreement between the data and the solid lines in both the energy and free energy plots. We view this agreement as confirmation of both the high temperature expansion and the validity of the numerical cavity analysis.

At low temperature, we can fit the cavity results to the equipartition analysis of Sec. III A. In particular, we expect the energy to exhibit T -linear behavior with a zero-temperature intercept giving the ground state energy of

Of course, in practice, we take a fixed population size of $N_{pop} = 10^4$ and iterate until we observe statistical convergence of various moments.

Finally, with the converged messages in hand, one can estimate the various disorder averaged thermodynamic quantities such as the internal energy U and the Bethe free energy, F .

the model and a slope corresponding to the specific heat. We have converged cavity populations down to temperatures of $T = 1/500$ as shown in the inset to Fig. 7. The linear extrapolations provide an estimated ground state energy of roughly 10^{-4} for $\alpha = 0.25, 0.75, 1.00$ while for $\alpha = 1.25, 2, 3, 4$, the ground state energy is of order 10^{-2} . These intercepts can be seen in Figure 2. We note that the statistical error bars in these estimates are quite large relative to the data found in the local quench numerics and, although they in principle exhibit no finite size effects, they do contain discretization error, which we expect to systematically raise the estimated energies and suppress the low temperature specific heat, as can be seen in the inset to Fig. 7. Also, as we will discuss in more detail below, the low temperature-high density cavity calculation exhibits a dynamical instability below which we should hesitate to believe the quantitative results. Nonetheless, the ground state energies predicted by the cavity calculation at $N_{disc} = 24$ are consistent with those found by simulated annealing and, at low density, with the zero expected for $\alpha < \alpha_{ps} = 0.91$.

C. Dynamic instability

In the context of discrete classical systems, the cavity equations (28) exhibit a well studied collection of




FIG. 7. (Color online) Cavity estimate of the internal energy of PSAT as function of temperature with $N_{disc} = 24$ states per Bloch sphere. Error bars reflect sampling error in calculating E from the converged population. Solid lines indicate the high temperature expansion of the energy to 4th order. Solid (hollow) markers indicate data in the replica symmetric (broken) phase. (inset) Zoomed in at low temperature of same data, including results for both $N_{disc} = 24$ states (■) and $N_{disc} = 12$ states (●). Solid (dashed) lines indicate T -linear fits to the $N_{disc} = 24$ ($N_{disc} = 12$) data.




FIG. 8. (Color online) The free energy, F/N as a function of α, T for qubits that have been discretized into 24 states. Point labeling is the same as in Figure 7 and likewise, hollow points indicate data collected ignoring the dynamical replica instability. (inset) As a function of β , we show the fit to the high temperature (small β) approximation for the free energy (cf. equations (22) - (24)).

instabilities that are believed to indicate the proliferation of solutions to the underlying belief propagation (BP) equations (26) on large instances [11, 30]. In turn, each of the many, extensively distinct solutions to the BP equations has been argued to correspond to a distinct thermodynamic pure state in the Gibbs ensemble. In the replica treatment of spin glass models, these instabilities are associated with replica symmetry break-

ing, and this language has been largely imported into the cavity formalism. From an algorithmic point of view, replica symmetry breaking indicates the appearance of metastable states that are ‘thermodynamically separated’. Such metastable configurations are believed to trap local search and annealing algorithms, frustrating the approach to lower temperatures – a behavior that we observed in our finite-size numerics in the high clause density regime.

Fortunately, while the full treatment of replica symmetry breaking requires the solution of a yet more complicated hierarchy of equations than (28), it is possible to detect the presence of replica symmetry breaking within the replica symmetric treatment [31, 32]. To perform such a study, one initializes the population dynamics algorithm with two populations that differ by a small random displacement, $\delta\mathbb{P}$, and then evolves each population with an identical sequence of random moves. In the replica symmetric phase, these two initial conditions result in identical distributions for $P_{i \rightarrow m}$ and $Q_{m \rightarrow i}$ at late times. In the presence of replica symmetry breaking, the two distributions diverge. The Lyapunov exponent associated with this process (inset, Fig. 9) provides a robust indication of the underlying instability: where the exponent goes from negative to positive, we identify the dynamic instability phase boundary in the (α, T) plane (Fig. 9). As can be seen best in Fig. 2, the high energy, low clause density regime of PSAT exhibits full replica symmetry while the dynamical instability sets in along a boundary which encloses the low energy density, high clause density regime, frustrating the approach to the ground state in the non-PRODSAT phase.

As our numerical approach relies on discretization of the continuous spins, we have attempted to check the convergence of the instability boundary to a finite limit as N_{disc} becomes large. Indeed, for any discretization we find that there is an $\alpha_c(T)$ where there is a dynamical glass transition. As the discretization is increased, the instability line moves to higher clause density but, for the three discretizations probed, we believe this finite-discretization shift converges to a finite temperature boundary for $\alpha > 0.9$ (see Fig. 9). This is especially striking for, as $T \rightarrow 0$, we find that this transition apparently coincides with the zero temperature PRODSAT transition.[14] We note that this is different from the situation in XORSAT where the dynamical glass transition at zero temperature coincides with the appearance of the hypercore at $\alpha_{hc} \simeq 0.81$ rather than the SAT-UNSAT transition at $\alpha = 0.91$ [33].

VI. CONCLUDING REMARKS

As noted at the outset this paper is fundamentally concerned with the approximation problem for randomly chosen QSAT instances. Specifically, by combining several methods—cavity analysis, simulated annealing and a greedy algorithm—we give an estimate of the limits to

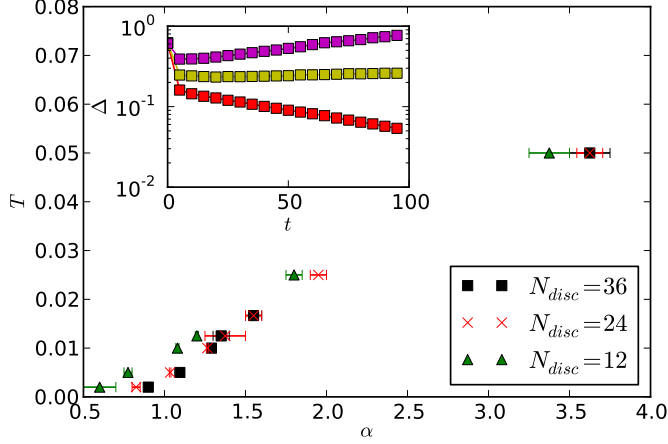


FIG. 9. (Color online) (inset) The sum-of-squares difference between two populations as a function of time in the population dynamics for $N_{disc} = 36$ states and $\beta = 100$. From bottom to top in the inset, we show data for $\alpha = 1.26, 1.28$ and 1.29 . We observed a transition in α between when the populations converge or diverge over time. This boundary is plotted in the main figure. (main) The location of the dynamical instability phase boundary extracted from the Lyapunov exponents of the population dynamics. As the discretization of the Bloch sphere increases, the dynamical instability appears to converge to a finite temperature phase boundary. As $T \rightarrow 0$, the instability line appears to terminate in the PRODSAT transition at $\alpha \simeq 0.91$.

which product states can be used to approximate solutions to 3-QSAT. Among these methods, the cavity analysis shows a dynamical glass transition before the limits of product state approximation are reached but this does not appear to have a marked influence on the performance of our best algorithm, the greedy search – at least for sizes up to $N = 100$. As such it is not clear whether the search for product states exhibits an easy-hard transition before the limits of the product state region are reached. Along the way, we have provided various additional statistical mechanical results on the classical PSAT Hamiltonian which may be of of interest to aficionados of disordered systems.

A natural next step is to study classical wavefunctions that build in some entanglement. It may be possible to analyze bounded-depth local quantum circuits matched to the random graph structure, similar to those studied in [20]. It is unclear how the non-commuting nature of generic QSAT Hamiltonians affects the quality of this approach and whether one needs to move to more general entanglement structures. We expect such considerations and the set of conjectures embodied in our opening figure (Fig. 1) to provide stimulus for future work.

ACKNOWLEDGMENTS

We thank Antonello Scardicchio and Francesco Zamponi for useful discussions. This work was supported in part by NSF grant PHY-1005429 (BH, SLS). CRL acknowledges support by a Lawrence Gollub Fellowship and the NSF through a grant for ITAMP at Harvard University. BH and CRL thank the Max Planck Institut für Physik Komplexer Systeme for their hospitality while this work was being completed.

Appendix A: Exact singletons

1. Singleton clause

The partition function for a single $k = 1$ clause attached to a spin is given by:

$$Z = \int d(\cos \theta) d\phi e^{-\beta |\phi_a z^a|^2} \quad (A1)$$

Choosing the coordinate system to align with the state $|\phi\rangle$, this reduces to

$$Z = \int d(\cos \theta) d\phi e^{-\beta \cos^2(\theta/2)} \quad (A2)$$

$$= 2\pi \int d(\cos \theta) e^{-\beta(1-\cos \theta)/2} \quad (A3)$$

$$= 2\pi e^{-\beta/2} \left. \frac{e^{\beta u/2}}{\beta/2} \right|_{-1}^1 \quad (A4)$$

$$= 4\pi(1 - e^{-\beta})/\beta \quad (A5)$$

The free energy is

$$F = -\frac{1}{\beta} \left[\log 4\pi + \log \frac{1 - e^{-\beta}}{\beta} \right] \quad (A6)$$

The entropy is

$$S = -\frac{\partial F}{\partial T} = \beta^2 \frac{\partial F}{\partial \beta} \quad (A7)$$

$$= \log 4\pi + \log \frac{1 - e^{-\beta}}{\beta} + 1 - \frac{\beta}{e^\beta - 1} \quad (A8)$$

The energy is

$$U = F + S/\beta = \frac{1}{\beta} - \frac{1}{e^\beta - 1}. \quad (A9)$$

These expressions recover the correct limits (high temperature):

$$U \sim \frac{1}{\beta} \left(1 - \frac{1}{1 + \beta/2 + \beta^2/3!} \right) = \frac{1}{2} - \beta/6 + O(\beta^2) \quad (A10)$$

$$S \sim \log 4\pi + O(\beta^2) \quad (A11)$$

and low temperature:

$$U \sim 0 + \frac{1}{\beta} + O(e^{-\beta}) \quad (\text{A12})$$

$$S \sim \log 4\pi/\beta + 1 + O(e^{-\beta}) \quad (\text{A13})$$

To check the high temperature expansion, we have

$$\langle H \rangle_0 = \langle |\phi_a z^a|^2 \rangle_0 = \left\langle \frac{1-u}{2} \right\rangle_0 = \frac{1}{2} \quad (\text{A14})$$

$$\langle H^2 \rangle_0 = \left\langle \frac{1-2u+u^2}{4} \right\rangle_0 = \frac{1}{4} \left(1 + \frac{1}{3}\right) = \frac{1}{3} \quad (\text{A15})$$

$$U \sim \langle H \rangle_0 - \beta(\langle H^2 \rangle_0^c) = \frac{1}{2} - \beta \frac{1}{12} \quad (\text{A16})$$

Here, we note that the infinite temperature measure on u is simply uniform on $[-1, 1]$.

2. Single product clause

For a product projector, the integrals for the partition function may be simplified by rotating the local bases to match those of the projector. Thus,

$$Z = \int d(\cos \theta_1) d\phi_1 \cdots d(\cos \theta_k) d\phi_k e^{-\beta |\phi_a^1 z_1^a|^2 \cdots |\phi_b^k z_k^b|^2} \quad (\text{A17})$$

$$= (2\pi)^k \int du_1 \cdots du_k e^{-\beta \frac{1-u_1}{2} \cdots \frac{1-u_k}{2}} \quad (\text{A18})$$

For $k = 2$, this can be evaluated to the useless form,

$$Z = \frac{(4\pi)^2}{\beta} [\gamma + \Gamma[0, \beta] + \log \beta] \quad (\text{A19})$$

Appendix B: Uniform measure on \mathbb{CP}^{n-1}

The measure on the projectors is the uniform Haar measure on the space \mathbb{CP}^{n-1} . We can represent this measure on an n -component complex vector ϕ_i

$$\begin{aligned} Z &= \int \left(\frac{d\phi_i^* d\phi_i}{2i} \right) \delta(|\phi|^2 - 1) \\ &= \int_{-\infty}^{\infty} \frac{d\lambda}{2\pi} \int \left(\frac{d\phi_i^* d\phi_i}{2i} \right) e^{i\lambda(|\phi|^2 - 1)} \\ &= \int_{-\infty}^{\infty} \frac{d\lambda}{2\pi} e^{-i\lambda} \left(\frac{\pi}{-i\lambda} \right)^n \\ &= \left(\frac{\pi}{-i} \right)^n \int_{-\infty}^{\infty} \frac{d\lambda}{2\pi} e^{-i\lambda k} \frac{1}{\lambda^n} \Big|_{k=1} \\ &= \left(\frac{\pi}{-i} \right)^n \int_{-\infty}^{\infty} \frac{d\lambda}{2\pi} e^{-i\lambda k} \frac{1}{(-1)^{n-1}(n-1)!} \partial_\lambda^{n-1} \frac{1}{\lambda} \Big|_{k=1} \\ &= \left(\frac{\pi}{-i} \right)^n \frac{1}{(-1)^{n-1}(n-1)!} (ik)^{n-1} \text{sign}(k) \Big|_{k=1} \\ &= \pi^n \frac{i}{(n-1)!} \end{aligned}$$

With the partition function in hand we can evaluate the disorder averaged correlators. The two-point correlator is

$$\begin{aligned} \overline{\phi_i^* \phi_j} &= \frac{1}{Z} \int_{-\infty}^{\infty} \frac{d\lambda}{2\pi} \int \left(\frac{d\phi_i^* d\phi_i}{2i} \right) \phi_i^* \phi_j e^{i\lambda(|\phi|^2 - 1)} \\ &= \frac{1}{Z} \int_{-\infty}^{\infty} \frac{d\lambda}{2\pi} e^{-i\lambda} \left(\frac{\pi}{-i\lambda} \right)^n \frac{\delta_{ij}}{-i\lambda} \\ &= \frac{\delta_{ij}}{n} \end{aligned} \quad (\text{B1})$$

The higher order correlators satisfy a slightly modified Wick's theorem:

$$\overline{\phi_i^* \phi_j^* \phi_k \phi_l} = \frac{\delta_{ik} \delta_{jl} + \delta_{il} \delta_{jk}}{n(n+1)} \quad (\text{B2})$$

$$\overline{\phi_1^* \phi_2^* \cdots \phi_k^* \phi_{k+1} \phi_{k+2} \cdots \phi_{2k}} = \frac{\sum_{\text{pairings}} \delta \delta \cdots \delta}{n(n+1) \cdots (n+k-1)} \quad (\text{B3})$$

These averages are equally useful for evaluating high temperature correlators of the two-component spinors z_a .

Appendix C: Hypercore data

We summarize a few results derived in [26] regarding the hypercore of random factor graphs G . The hypercore is the maximal subgraph of G on which all nodes have degree at least 2. The quoted results follow from analyzing the leaf removal algorithm applied to the random graph. In the notation of [26], the clause density α is γ and the number of spins per clause k is p .

The degree distribution on the hypercore is Poissonian for degrees $d \geq 2$:

$$P_c(d) = \begin{cases} 0 & \text{for } d = 0, 1 \\ e^{-\lambda^*(\alpha)} \frac{\lambda^*(\alpha)^d}{d!} & \text{for } d \geq 2 \end{cases} \quad (\text{C1})$$

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

$$e^{-\lambda^*} - 1 + \left(\frac{\lambda^*}{k\alpha} \right)^{\frac{1}{k-1}} = 0. \quad (\text{C2})$$

The total number of nodes in the hypercore is

$$N_c(\alpha) = N \sum_{d \geq 2} P_c(d) = N \left[1 - (1 + \lambda^*)e^{-\lambda^*} \right] \quad (\text{C3})$$

and the total number of edges (clauses)

$$M_c(\alpha) = N \frac{\lambda^*}{k} (1 - e^{-\lambda^*}). \quad (\text{C4})$$

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

-
- [1] S. Arora, C. Lund, R. Motwani, and M. Szegedy, Colloquium on Computational Complexity, **5** (1998).
 - [2] S. Arora and B. Barak, *Complexity Theory: A Modern Approach* (Cambridge University Press, 2009).
 - [3] A. Y. Kitaev, in *AQIP'99* (DePaul University, 1999).
 - [4] D. Aharonov and T. Naveh, (2002), arXiv:quant-ph/0210077v1 [quant-ph].
 - [5] S. Bravyi, (2006), arXiv:quant-ph/0602108v1 [quant-ph].
 - [6] C. R. Laumann, R. Moessner, A. Scardicchio, and S. L. Sondhi, Quant. Inf. and Comp., **10**, 0001 (2010).
 - [7] C. R. Laumann, A. M. Läuchli, R. Moessner, A. Scardicchio, and S. L. Sondhi, Phys. Rev. A, **81**, 062345 (2010).
 - [8] S. Bravyi, C. Moore, and A. Russell, (2009), arXiv:0907.1297v2 [quant-ph].
 - [9] A. Ambainis, J. Kempe, and O. Sattath, in *42nd Annual ACM Symposium on Theory of Computing* (2009) 0911.1696v1.
 - [10] Current rigorous arguments provide $\alpha_{ps} = 0.91\dots$ and that $\alpha_s < 3.59$ for $k = 3$ body interactions. The estimate of $\alpha_s = 1.0 \pm 0.06$ comes from small-size numerics. The separation of α_{ps} and α_s and the accompanying entangled-SAT phase can be shown rigorously for $k \geq 12$.
 - [11] M. Mézard and A. Montanari, *Information, physics, and computation* (Oxford University Press Inc., New York, NY, USA, 2009).
 - [12] M. Mézard and G. Parisi, Euro. Phys. J. B, **20**, 217 (2001).
 - [13] C. R. Laumann, A. Scardicchio, and S. L. Sondhi, Phys. Rev. B, **78**, 134424 (2008).
 - [14] C. R. Laumann, S. A. Parameswaran, S. L. Sondhi, and F. Zamponi, Phys. Rev. B, **81**, 174204 (2010).
 - [15] D. Aharonov, I. Arad, Z. Landau, and U. Vazirani, (2008), 30 pages, 2 figures, arXiv:0811.3412v1 [quant-ph].
 - [16] I. Dinur, J. ACM, **54**, 12 (2007).
 - [17] S. Bravyi and M. Vyalii, Quant. Inf. and Comp., **5**, 187 (2005).
 - [18] D. Aharonov and L. Eldar, in *2011 IEEE 52nd Annual Symposium on Foundations of Computer Science—FOCS 2011* (IEEE Computer Soc., Los Alamitos, CA, 2011) pp. 334–343.
 - [19] D. Aharonov and L. Eldar, **quant-ph** (2013), arXiv:1301.3407.
 - [20] M. B. Hastings, (2012), arXiv:1201.3387.
 - [21] M. H. Freedman and M. B. Hastings, (2013), arXiv:1301.1363.
 - [22] I. Arad, Quant. Inf. and Comp., **11**, 1019 (2011).
 - [23] S. Gharibian and J. Kempe, **quant-ph** (2011), arXiv:1101.3884v1.
 - [24] F. G. S. L. Brandao and A. W. Harrow, Preprint. (2013).
 - [25] D. Gosset and D. Nagaj, (2013), arXiv:1302.0290.
 - [26] M. Mézard, F. Ricci-Tersenghi, and R. Zecchina, J. Stat. Phys., **111**, 505 (2003).
 - [27] L. Zdeborová and F. Krzakala, Phys. Rev. E, **76**, 031131 (2007).
 - [28] L. Zdeborová, in *Workshop on Quantum Statistical Mechanics, Computation and Information* (ICTP, Trieste, Italy, 2010).
 - [29] L. Zdeborová and M. Mezard, Phys. Rev. Lett., **101**, 078702 (2008).
 - [30] F. Krzakala, A. Montanari, F. Ricci-Tersenghi, G. Semerjian, and L. Zdeborova, P. Natl. Acad. Sci. USA, **104**, 10318 (2007).
 - [31] A. Montanari, G. Parisi, and F. Ricci-Tersenghi, J Phys A: Math. Gen., **37**, 2073 (2004).
 - [32] S. Mertens, M. Mezard, and R. Zecchina, arXiv:cs/030902 (2003).
 - [33] F. Altarelli, R. Monasson, G. Semerjian, and F. Zamponi (2009).