Classical-quantum mixing in the random 2-satisfiability problem

Ionut-Dragos Potirniche,[1, 2] C. R. Laumann,[3, 4] and S. L. Sondhi[2, 4]

[1]*Department of Physics, University of California, Berkeley, California, 94720-7300, USA*
[2]*Department of Physics, Princeton University, Princeton, New Jersey 08544, USA*
[3]*Department of Physics, University of Washington, Seattle, Washington 98195, USA*
[4]*Max-Planck-Institut für Physik Komplexer Systeme, 01187 Dresden, Germany*
(Dated: August 22, 2018)

Classical satisfiability (SAT) and quantum satisfiability (QSAT) are complete problems for the complexity classes NP and QMA, respectively, and are believed to be intractable for both classical and quantum computers. Statistical ensembles of instances of these problems have been studied previously in an attempt to elucidate their typical, as opposed to worst-case, behavior. In this paper, we introduce a new statistical ensemble that interpolates between classical and quantum. For the simplest 2-SAT–2-QSAT ensemble, we find the exact boundary that separates SAT and UNSAT instances. We do so by establishing coincident lower and upper bounds, in the limit of large instances, on the extent of the UNSAT and SAT regions, respectively.

The potential power of quantum computers drives the immense effort to build and understand them. There are two primary theoretical approaches to characterizing the precise extent of this potential. Complexity theory [1] proceeds by identifying so-called 'complete' problems which are the hardest problems in a given class, such as NP. Classically, under the widely believed conjecture that P ≠ NP, *all* algorithmic approaches to NP-complete problems such as satisfiability (SAT) must fail on at least some subset of worst-case instances. Over the last decade, these venerable considerations have been extended to the quantum case, where QMA-complete problems are now believed to be intractable and to capture the intrinsic differences between quantum and classical computing [2]. The natural quantum generalization of SAT, so-called quantum satisfiability (QSAT) [3, 4], is conveniently such a QMA$_1$-complete problem [5].

In the second approach, the introduction of an appropriate measure on the instances of a problem generates a question in statistical Physics. Instead of worrying about worst-case complexity, we attempt to understand the structure of problems that are typical with respect to the measure. For example, phase transitions, which arise as functions of parameters controlling the measure, often signal the regimes where the most complicated problems may be found [6]. This approach builds on the seminal insight of Fu and Anderson [7] that the intractability of NP-complete problems is a form, indeed an extreme one, of (spin) glassiness. Ensembles of both classical SAT [8, 9] and, more recently, QSAT [10–12] have been studied in this fashion.

In this paper we build on the second approach and introduce new ensembles that interpolate between the SAT and QSAT ones. The mixture provides a convenient framework for characterizing the crossover from classical to quantum search complexity. For example, the classical PCP theorem [13] shows that it is computationally hard

to approximately determine the ground state of the SAT problem, while it is still an open question whether an analogous hardness result applies in the quantum case. The interpolation allows the study of the crossover in entanglement properties of low-energy states, which may bear on this question. Similarly, in statistical physics, the mixed ensemble can shed light on entanglement phase transitions in spin-glass models with quenched disorder: the classical problem has no entangled solutions, whereas the mixed one exhibits entangled states in the UNSAT regime.

Specifically for the 2-SAT–2-QSAT interpolation, we show that there is a sharp phase boundary (Fig. 1) that we determine rigorously by deriving coincident lower and upper bounds on the extent of the UNSAT and SAT regions, respectively. The interest in this interpolation flows considerably from a "geometrization" theorem that applies to the QSAT limit [10]. As this result is not widely known, we begin with a quick review of the relevant background, which will also enable a proper definition of the problem studied herein. We then present our central technical results on the phase boundary, and we close with some remarks on the lessons learned and future directions.

SAT, QSAT, and the mixed ensemble: An instance of k-QSAT is defined by a positive-semidefinite Hamiltonian on N qubits given by the sum of M k-body interactions, $H = \sum_{m=1}^{M} \Pi_m$. Here, each interaction $\Pi_m = |\phi\rangle \langle\phi|_m$ projects onto a particular state in the local Hilbert space of the k qubits associated with interaction m. The computational problem is to decide whether H has a ground state of energy zero. If so, the problem is SAT, and if not it is UNSAT [14]. If the states $|\phi\rangle_m$ are computational basis states, we recover classical k-SAT as a special case. We can now summarize the statistical ensembles used in previous work. These involve the uniform Erdős-Rényi measure [15–17], parametrized by the
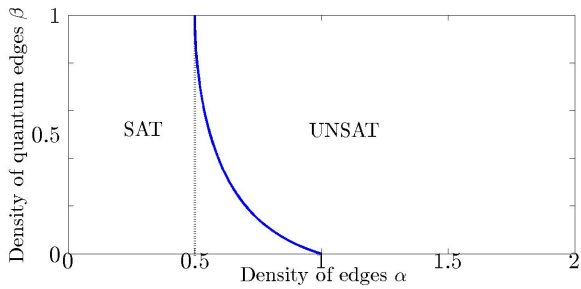
FIG. 1. The phase boundary in the $\alpha - \beta$ plane that separates the SAT and the UNSAT regimes for the mixed classical-quantum problem. The dashed line indicates the emergence of a giant component in the ER graph (the percolation phase transition). $\beta = 0$ corresponds to the 2-SAT problem and $\beta = 1$ corresponds to the 2-QSAT problem.

'clause density' $\alpha = M/N$, over the set of k-hypergraphs representing the interactions in H. The interaction associated with each hyperedge is likewise uniformly chosen from the 2^k projectors in classical SAT or from the Haar measure over the rays in the 2^k-dimensional Hilbert space for QSAT. We observe that the latter makes classical SAT instances highly non-generic within the QSAT ensemble, and so these two can be expected to behave very differently, and indeed they do [8, 9].

The classical ensemble has been studied intensively. The broad picture that has emerged is that for all $k \geq 2$ there is a sharp SAT-UNSAT transition in the $N \to \infty$ limit as a function of α. For $k \geq 3$, there are additional transitions in the SAT regime wherein the structure of the solution space changes. For the quantum ensemble, it is known that there is a SAT-UNSAT transition for $k \geq 2$ [10, 11, 18, 19] and that there is at least one sharp transition involving the growth of entanglement in the satisfying states for $k \geq 12$ [11, 20]. A remarkable result to come out of the quantum generalization is a "geometrization theorem" [10] wherein uniformly chosen quantum projectors on *any* graph exhibit the same dimension of the satisfying manifold with probability 1. This reduces the generic QSAT decision problem to a purely graph-theoretic question! The identification of this implicit graph-theoretic property for $k \geq 3$ and understanding its computational difficulty is an outstanding problem.

As advertised above, we initiate a new approach by introducing ensembles that interpolate between the fully classical and quantum regimes. We do this by constraining each realization to include a fraction β of uniformly chosen quantum projectors and $1 - \beta$ uniformly chosen classical ones. As β varies between 0 and 1, we pass from the classical ensemble to the quantum one. As the classical ensemble does not exhibit geometrization and typically has larger satisfying manifolds, the interpolation has the potential to shed light on the emergence of

geometrization in the quantum limit and thence on precisely where in the ensemble one might look for genuinely difficult quantum cases.

As a first step in this program, we study the case of 2-SAT/QSAT. We generate the underlying 2-graphs by drawing edges between any two sites with probability $\alpha N / \binom{N}{2}$. In the thermodynamic limit $N \to \infty$, this generates an Erdős-Rényi (ER) random graph with $M = \alpha N$ expected edges. For each edge $m = 1...M$ we label it 'quantum' with probability β and 'classical' with probability $1 - \beta$: we write $e_m \in \{Q, C\}$ if the edge between sites $(m, m+1)$ is quantum or classical. The purely classical and quantum limits are very well understood. At $\beta = 0$ [21–23] there is a sharp SAT-UNSAT transition at $\alpha_c = 1$, while for $\beta = 1$ there is a SAT-UNSAT transition at $\alpha_q = 1/2$ [10]. The quantum transition coincides with the emergence of a giant component in the underlying random graph [16].

The snip-core: The primary tool in our analysis is snipping qubits out of an interaction graph G. Classically, a node i such that $G = G' \cup \{i\}$ is *snippable* if all of the bonds connected to it agree about the bit arrangement they locally disfavor. All clauses attached to such a snippable node i can then be trivially satisfied by assigning the appropriate value to qubit i without reference to the state on G'. Thus, these bonds can be snipped from the graph G to produce a smaller graph G', which is SAT if and only if the original G is.

This definition extends naturally to the mixed classical-quantum problem ($\beta \neq 0$): a degree-1 site i with a quantum projector attaching it to site j of G' is snippable. From Bravyi's construction [3], we know that the satisfying state for G' can be written as a product state $|\psi_j\rangle \otimes |\Psi_{G' \setminus \{j\}}\rangle$. If the quantum edge attaching site i disfavors the state $|\phi_{j,i}\rangle$, then we can use the transfer matrix $T_\phi = \epsilon \phi_{j,i}^\dagger$ (ϵ is the 2×2 Levi-Civita symbol) to find $|\xi_i\rangle = T_\phi |\psi_j\rangle$ such that $\langle \phi_{j,i}| |\psi_j\rangle \otimes |\xi_i\rangle = 0$. Therefore, $|\xi_i\rangle \otimes |\psi_j\rangle \otimes |\Psi_{G' \setminus \{j\}}\rangle$ is a satisfying state for $G = G' \cup \{i\}$, which shows that a degree-1 site with a quantum projector attached is snippable.

However, if a generic quantum edge attaches to a site i of degree greater than 1, it cannot be locally satisfied by the state on i without reference to the rest of the graph (with probability 1). Thus, any site of degree at least 2 with a quantum edge attached is unsnippable.

For a random instance G of the mixed problem, we can iteratively remove snippable sites and the incident edges in a similar fashion to the "leaf removal" algorithm [24]. When there are no snippable qubits left, this algorithm stops and we end up with a unique maximal *snip-core*. Clearly, G is SAT iff its snip-core is SAT. Moreover, if the snip-core is empty, then G is SAT. Behind our considerations lies the result that when G is SAT, we can always find a satisfying product state [3]. We will now analyze the structure and probability of non-empty snip-cores in order to determine the SAT-UNSAT boundary.

Upper bound on SAT region: The basic idea in this part is to identify an UNSAT motif that must be present on all snip-cores for $\alpha > \alpha_c(\beta)$, thus establishing an upper bound on the extent of the SAT region. Toward that end, we first note that the simplest motif that is unsnippable is the *unsnippable loop*. A loop with classical and quantum edges is said to be *unsnippable* if all of its sites are unsnippable [see Fig. 2(a)]. We can find a simple example in the fully classical problem: a loop that dislikes 01 on each edge. While this loop is unsnippable, it is SAT: it has exactly two satisfying states comprised of all sites 0 or all sites 1. The same is true of the other 2^L such classical unsnippable loops of length L that are equivalent under "gauge" transformations. It can be shown that the fully quantum loop also has two SAT states [3, 10]. Moreover, based on the Geometrization Theorem, we can conclude that the dimension of the kernel of the fully quantum problem is a lower bound for the mixed problem, whereas the fully classical one is an upper bound: we start from the classical unsnippable loop and slowly turn classical projectors into generic quantum ones; this can only decrease the degeneracy of satisfying states. But since both kernels have dimension 2, we conclude that the mixed classical-quantum unsnippable loop always has exactly two linearly independent SAT states. Finally, an unsnippable but UNSAT motif can be constructed by decorating the unsnippable loop with two unsnippable cross-links—strings joining two different points on the loop whose interior sites are unsnippable [Fig. 2(b)] that penalize the two SAT states.

We now turn to the probability of finding UNSAT unsnippable loops with cross-links. Quite generally, the expected number of subgraphs A in the ER ensemble on N nodes is

$$\mathbb{E}(\# \text{ of } A) = \frac{N!}{(N - |A|)!\text{Aut}(A)} p^{e(A)}, \qquad (1)$$

where $|A|$ and $e(A)$ represent the number of vertices and edges of A, respectively; $\text{Aut}(A)$ is the number of automorphisms of A and $p = \alpha N / \binom{N}{2}$. For a loop of length L we have $|L| = e(L) = L$ and $\text{Aut}(L) = 2L$. Introducing the probability that a given loop is unsnippable, we find the expected number of unsnippable loops $\#_{\text{uns}}(L)$ is

$$\#_{\text{uns}}(L) = \binom{N}{L} \frac{L!}{2L} \left(\frac{2\alpha}{N-1} \right)^L p(\text{L is unsnippable}). \qquad (2)$$

For any fixed length L and as $N \to \infty$, $\#_{\text{uns}}(L)$ scales as $\mathcal{O}(N^0)$. More generally, the number of subgraphs with $m = e(A) - |A|$ cross-links vanishes in the thermodynamic limit as $\mathcal{O}(N^{-m})$. This result holds irrespective of the form of $p(A \text{ is unsnippable})$ since this probability has no explicit dependence on N. It follows that in order to get a non-vanishing number of UNSAT motifs we must consider *giant loops* whose size scales with N as $L = lN$ ($0 < l \le 1$).

To calculate the number of unsnippable giant loops we need the last factor in Eq. (2). For a collection $\{e_i\}$ of M edges, the probability that site i is unsnippable is $1 - \delta_{e_i,C}\delta_{e_{i+1},C}/2$. In words, if it is connected to a quantum edge then it is unsnippable with probability 1; otherwise it is unsnippable with probability 1/2. Also, from the definition of the random ensemble we have $p(e_i = Q) = \beta$ and $p(e_i = C) = 1 - \beta$. Hence, for a loop of length L

$$p(\text{L is unsnippable}) = \sum_{\{e_i\}} \prod_{i=1}^{L} p(e_i) \left(1 - \frac{\delta_{e_i,C}\delta_{e_{i+1},C}}{2} \right). \qquad (3)$$

Using a standard transfer matrix technique and focusing solely on the dominant eigenvalue λ_+ which controls the result for large loops, we find that

$$p(\text{L is unsnippable}) \approx \left(\frac{1 + \beta + \sqrt{-7\beta^2 + 10\beta + 1}}{4} \right)^L. \qquad (4)$$

With this in hand, we return to Eq. (2) and use Stirling's approximation to find the extensive part of the entropy,

$$S_{\text{uns}}(l) = N \left[l \left(\log(2\alpha\lambda_+) - 1 \right) - (1 - l) \log(1 - l) \right]. \qquad (5)$$

Giant unsnippable loops proliferate exponentially in N if the entropy function is positive for some $1 \ge l > 0$. Since $S_{\text{uns}}(0) = 0$ and $S'_{\text{uns}}(l) = N \left[\log(2\alpha\lambda_+) + \log(1 - l) \right] \le N \log(2\alpha\lambda_+)$, we see that $S_{\text{uns}}(l)$ is a negative and decreasing function on $0 < l \le 1$ for $2\alpha\lambda_+ < 1$. However, for $2\alpha\lambda_+ > 1$, the entropy goes positive for small l and large numbers of giant unsnippable loops emerge. As we find $e^{\mathcal{O}(N)}$ loops on just N sites, the loops must intersect and overlap repeatedly. It follows that any given giant loop is covered by a finite density of cross-links and therefore the probability of finding an UNSAT unsnippable loop with cross-links approaches 1 as $N \to \infty$ for α greater than

$$\alpha_c(\beta) = \frac{2}{1 + \beta + \sqrt{-7\beta^2 + 10\beta + 1}}. \qquad (6)$$

Hence, we conclude that that $\alpha_c(\beta)$ is an upper bound on the extent of the SAT region at any fixed β.

Lower bound on UNSAT region: We will now show that for $\alpha < \alpha_c(\beta)$ the snip-core is always SAT so that $\alpha_c(\beta)$ is also a lower bound on the extent of the UNSAT region. We do this by showing that the snip-core must contain one of a finite list of motifs [25] and that all of these *except* SAT unsnippable loops are not present as $N \to \infty$ for $\alpha < \alpha_c(\beta)$.

To show that there is a finite list of necessary motifs, we start from an arbitrary site in the snip-core which we label 1 and walk along any edge which we label e_1 that connects it to a site 2 [see Fig. 2(c)]. At step 2, we walk along an edge e_2 to site 3 such that 2 is unsnippable

with respect to edges e_1 and e_2. Thus, e_2 can either be quantum or, if it is classical, it has to disagree with e_1 on site 2. From this point onward, we take further steps as follows:

- At step k: we move along an edge e_k that connects it to $k+1$ such that site k is unsnippable.

- Iterate until we pass through a site i twice.

Since the size of the snip-core is finite at any given N, then each such path must be self-intersecting. Therefore, the algorithm stops in either of the following scenarios:

1. The path returns to the starting point ($i = 1$) and we end up with a loop. There are two subcases:

    (a) If all sites on the loop have degree 2 then the loop must be unsnippable [Fig. 2(a)].

    (b) If there exists a site on the loop that has a degree of at least 3 then we walk one step away from the loop starting at that site as before and get a "lasso motif" [Fig. 2(c)].

2. The path crosses itself at a site $i \neq 1$ and we encounter the same lasso motif.

At this point, we continue from the open end of the lasso and generate unsnippable sites as before. Again, our path necessarily returns and touches the lasso. When this happens, we end up with one of three motifs: an unsnippable loop with one cross-link [Fig. 2(d)], a "figure eight" [Fig. 2(e)] or a "dumbbell" [Fig. 2(f)]. Together with the unsnippable loop [Fig. 2(a)], these constitute the set of structures of which at least one *must* be present on each non-empty disconnected component of a snip-core. Strictly speaking, we should classify somewhat more finely by specifying the unsnippability of each loop passing through the degree 3 and 4 sites in these motifs. But that only changes our estimates below by constant factors as $N \to \infty$, so we refrain from exhibiting these details here.

Now for the frequency of occurrence of these motifs. In the limit $N \to \infty$ for any fixed number of sites L, the expected number of loops is $O(N^0)$, while the expected numbers of the other three are $O(1/N)$ and hence vanish. As before, we are led to examine giant versions of these graphs. The most optimistic case assumes that all the individual legs of the motifs are large and diverge with N. In this case, the expected numbers take the form

$$\mathbb{E}(\# \text{ motifs}) = c \binom{N}{L-1} \frac{(L-1)!}{a} \left(\frac{2\alpha}{N-1}\right)^L \lambda_+^L, \quad (7)$$

where a is the number of automorphisms ($a = 4$ for the figure eight and the dumbbell and $a = 2$ for the loop with a single cross-link) and c is an $\mathcal{O}(1)$ number (dependent on β) associated with the precise unsnippability of vertices that have degree at least 3 alluded to above.

For $\alpha < \alpha_c(\beta)$ these expected numbers vanish also for large motifs of size $L = lN^\gamma$ as $N \to \infty$ for any l and $0 < \gamma \leq 1$. Using Stirling's approximation, the $\mathcal{O}(N^\gamma)$ part of the entropy $S(l) = \log(\mathbb{E}(\# \text{ motifs}))$ is approximately

$$S(l) = lN^\gamma \log(2\alpha\lambda_+) - lN^\gamma - N(1 - lN^{\gamma-1}) \log(1 - lN^{\gamma-1}). \quad (8)$$

Once again, $S(0) = 0$ and the derivative $S'(l) = N^\gamma \left[\log(2\alpha\lambda_+) + \log(1 - lN^{\gamma-1})\right] \leq N^\gamma \log(2\alpha\lambda_+)$ for any l. We see that for $\alpha < \alpha_c(\beta)$, $S(l)$ is a negative and decreasing function so the expected number of motifs vanishes in the thermodynamic limit as $e^{-\mathcal{O}(N^\gamma)}$.

Hence, the loop with cross-link, figure eight, and dumbbell are entirely absent for $\alpha < \alpha_c(\beta)$ and the snip-core is either empty or composed entirely of unsnippable loops. Such snip-cores are SAT so we can conclude that our starting graphs are SAT for $\alpha < \alpha_c(\beta)$ with probability 1 and thus $\alpha_c(\beta)$ is a lower bound on the extent of the UNSAT region at fixed β. Putting together the upper and the lower bounds, we conclude that $\alpha_c(\beta)$ from Eq. (6) represents the exact location of the phase boundary between SAT and UNSAT in the $\alpha - \beta$ plane, as shown in Fig. 1.

Concluding remarks: In the preceding, we introduced a new family of mixed classical-quantum ensembles for the k-SAT/QSAT problems, and we established the exact phase diagram for the simplest member of this family with $k = 2$. We note that the shape of the phase boundary is consistent with what is known about the limits. The quantum limit is insensitive to the choice of projectors, and we find that the restriction of a dilute concentration of projectors to classical values barely shifts the phase boundary. The classical limit is sensitive to the choice of projectors, and the phase boundary near it is maximally sensitive to the inclusion of a dilute set of quantum projectors. A problem that we leave open is the nature of scaling near the phase boundary. One basic question concerns the scaling of the probability P_{SAT} that a random instance is SAT. Work on the classical problem by Bollobás *et al.* [23] showed that this has the scaling form $P_{\text{SAT}}(N, \alpha) = f((\alpha - \alpha_c)/N^{1/3})$. Work on the ER ensemble by Bollobás [26] and Łuczak [27] implies that the same scaling holds for the purely quantum problem. Qualitatively, our work makes sense of this apparent coincidence and implies that this scaling will hold everywhere along the phase boundary—that everywhere there exists a single transition involving the proliferation of unsnippable loops with the mix of projectors along such loops changing continuously with β. However, more work is required to show this rigorously and to try and extract more detailed information on the scaling functions near the transition. Finally, we intend to examine similar ensembles for $k = 3$. The current work suggests that perturbing about the classical limit by introducing a dilute set of quantum projectors into the Hamiltonian
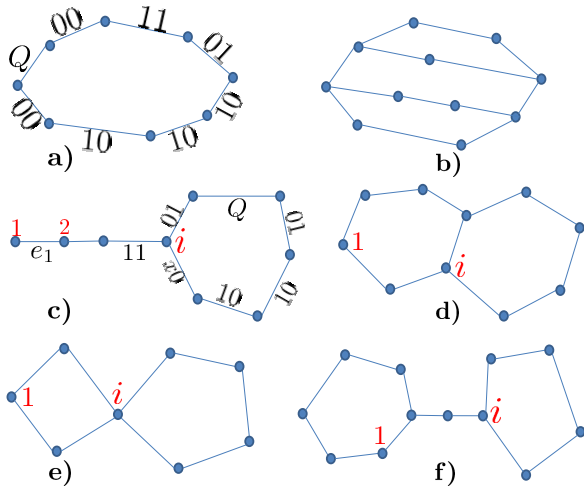
FIG. 2.
(a) An unsnippable loop that contains both classical and quantum edges.
(b) A loop with two cross-links. If all the strings are unsnippable and the cross-links penalize the loop's two linearly independent satisfying states, this motif becomes UNSAT.
(c) The "lasso" motif. We start from a random site 1 and move along an unsnippable path until we pass through the same site i twice. The loop can either be unsnippable ($x = 1$) or snippable at site i ($x = 0$).
(d) The loop with a single cross-link. We encounter this motif if we traverse the dangling branch of the lasso and end up at a site located on the initial loop.
(e) The "figure eight." We encounter this motif if we traverse the dangling branch of the lasso and end up at the same starting point i.
(f) The "dumbbell." We encounter this motif if we traverse the dangling branch of the lasso and pass through the same point twice (located outside of the initial loop).

could be informative in quantum mechanical perturbation theory.

[1] S. Arora and B. Barak, *Complexity Theory: A Modern Approach* (Cambridge University Press, Cambridge, 2009).
[2] A. Yu, A. Kitaev, H. Shen, and M. Vyalyi, *Classical and quantum computation*, Graduate Studies in Mathematics, Vol. 47 (American Mathematical Society, Providence, RI, 2002).
[3] S. Bravyi, arXiv:quant-ph/0602108v1 .
[4] D. Gosset and D. Nagaj, in *Foundations of Computer Science (FOCS), 2013 IEEE 54th Annual Symposium on* (2013) pp. 756–765.
[5] While both NP and QMA complete problems are believed to be hard for both classical and quantum computers to solve, the latter are such that a classical computer cannot even verify the solution efficiently.
[6] A. Hartmann and M. Weigt, *Phase Transitions in Combinatorial Optimization Problems* (Wiley-VCH Verlag GmbH and Co. KGaA, Weinheim, 2005).
[7] Y.-T. Fu and P. Anderson, *Lectures in the Sciences of Complexity* (Addison-Wesley, Reading, MA, 1989).
[8] M. Mézard and A. Montanari, *Information, Physics, and Computation* (Oxford Graduate Texts, New York, 2009).
[9] S. Kirkpatrick and B. Selman, Science **264**, 1297 (1994).
[10] C. R. Laumann, R. Moessner, A. Scardicchio, and S. L. Sondhi, Quant. Inf. and Comp. **10**, 0001 (2009).
[11] C. Laumann, A. Läuchli, R. Moessner, A. Scardicchio, and S. Sondhi, Physical Review A **81** (2010), 10.1103/PhysRevA.81.062345.
[12] B. Hsu, C. R. Laumann, A. M. Läuchli, R. Moessner, and S. L. Sondhi, Physical Review A **87** (2013), 10.1103/PhysRevA.87.062334.
[13] S. Arora, C. Lund, R. Motwani, M. Sudan, and M. Szegedy, J. ACM **45**, 501 (1998).
[14] The ground state energy is at least a promised gap Δ above zero if it is not precisely zero.
[15] J. Schmidt-Pruzan and E. Shamir, Combinatorica **5**, 81 (1985).
[16] P. Erdős and A. Rényi, Magyar Tud. Akad. Mat. Kutató Int. Közl **5**, 17 (1960).
[17] B. Bollobás, *Modern Graph Theory* (Springer, New York, 1998).
[18] S. Bravyi, C. Moore, and A. Russell, arXiv:0907.1297 .
[19] N. de Beaudrap, arXiv:1403.1588 .
[20] A. Ambainis, J. Kempe, and O. Sattath, J. ACM **59**, 24:1 (2012).
[21] A. Goerdt, Journal of Computer and System Sciences **53**, 469 (1996).
[22] E. Friedgut and J. Bourgain, Journal of the American Mathematical Society **12**, 1017 (1999).
[23] B. Bollobás, C. Borgs, J. Chayes, J. Kim, and D. Wilson, Random Structures and Algorithms **18**, 201 (2001).
[24] M. Mézard, F. Ricci-Tersenghi, and R. Zecchina, Journal of Statistical Physics **111**, 505 (2003).
[25] We use motif to mean a fixed subgraph up to isomorphisms, irrespective of its frequency of occurrence. This is midway between the everyday meaning of the term and the more technical graph theoretic usage which makes reference to its frequency.
[26] B. Bollobás, Transactions of the American Mathematical Society **286**, 257 (1984).
[27] T. Łuczak, Random Structures and Algorithms **1**, 287 (1990).