

A Theory of Frustrated Degeneracy

*MARTIN LOEBL

Department of Applied Mathematics

and

Institute for Theoretical Computer Science (ITI)

Charles University

Malostranské nám. 25, 11800 Praha, Czech Republic

loebl@kam.ms.mff.cuni.cz

JAN VONDRÁK

Department of Applied Mathematics

Charles University

Malostranské nám. 25, 11800 Praha, Czech Republic

vondrak@kam.ms.mff.cuni.cz

Abstract

We describe some observations on how degeneracy of toroidal and 3-dimensional spin glasses depends on frustration.

1 Introduction

In this paper we investigate the maximum *ground state degeneracy* of 2D and 3D Ising models. The general Ising model is defined

*Partially supported by the Project LN00A056 of the Czech Ministry of Education and by GAUK 158 grant. currently visiting School of Mathematics, Georgia Institute of Technology, Atlanta, GA 30332-0160, U.S.A.

by an undirected graph $G = (V, E)$ and a coupling constant J_{ij} assigned to each edge $\{i, j\}$. A physical state of the system is given by a *spin assignment* $\sigma : V \rightarrow \{\pm 1\}$ which has the corresponding energy

$$E(\sigma) = - \sum_{\{i,j\} \in E} J_{ij} \sigma_i \sigma_j.$$

A state with the minimum possible energy is called a *ground state*. The question we study is how many different ground states there are for a given model. The number of ground states determines the *ground state degeneracy*. Note that the ground states exactly correspond to minimum edge-cuts of G and so the ground state degeneracy determines the number of minimum edge-cuts of G .

In particular, we study square toroidal lattices (which we will call simply square lattices) and cubic lattices, both with the restriction that the coupling constants are ± 1 . We investigate how the maximum degeneracy depends on frustration of basic building blocks of these lattices, *plaquettes* (unit squares) and (unit) *cubes*. We will call these basic building blocks *stones*. Our approach is similar for square and cubic lattices.

Let the edges of a (square, cubic) lattice be covered by edge-disjoint stones P_1, \dots, P_n . Then the collection P_1, \dots, P_n is called a *cover* of the lattice. A cover is *satisfiable* if there exists a spin assignment S_0 such that in S_0 , each P_i is in its ground state (the local energy calculated for the Ising model restricted to P_i is minimum possible).

Then S_0 is clearly a ground state of the whole lattice, and if S is another ground state then restriction of S to each P_i must be a ground state of P_i . A lattice will be called *extremal* if it has a satisfiable cover.

Covers exist for both square toroidal and cubic lattices of even size: for square lattices, the plaquettes in the cover form a checkerboard pattern (and there are two distinct covers), while for cubic lattices, there are four covers each containing one fourth of the cubic cells.

The origin of our research was a question of Bruno Nachtergaele

[1], who asked whether the maximum degeneracy (the number of distinct ground states) for the square lattices is attained when all the plaquettes are frustrated (a plaquette is *frustrated* if it has an odd number of negative bonds, and *happy* otherwise), and what the maximum degeneracy is. Note that the lattice where all the plaquettes are frustrated is extremal.

We will show that the answer to the first question is negative: higher degeneracy is obtained when one cover of the square lattice consists of frustrated plaquettes and the complementary cover consists of happy plaquettes. We will also give upper and lower bounds for the maximum degeneracy of extremal lattices, both square and cubic. We conjecture that the maximum degeneracy for general square and cubic lattices is the same as for extremal ones. Our computer experiments so far support this.

Pure States: At the end of introduction let us illustrate the usefulness of the notion of covers by an observation concerning infinite square grids with coupling constants ± 1 distributed uniformly at random. Let us call these grids *RIS grids*. A state of an RIS grid is called *pure* if its energy cannot be locally decreased by changing one spin. Observe that RIS grids have an unbounded number of pure states.

Indeed, let \mathcal{L} be a RIS grid and let \mathcal{C} be its cover. Any state of \mathcal{L} which induces a pure state in each plaquette of \mathcal{C} is a pure state of \mathcal{L} . Cover \mathcal{C} may be built up starting from a row of disjoint plaquettes by attaching the remaining plaquettes one by one to the already constructed part by two vertices connected by an edge. We will construct pure states along with this construction of \mathcal{C} . Assume an initial part W of \mathcal{C} is given together with a set \mathcal{S} of states of W which induce a pure state in each plaquette of W . Let plaquette P be added to W . Observe that any spin assignment to a pair of vertices of P connected by an edge may be extended into a pure state of P . Hence any state of \mathcal{S} may be extended into a state of $W \cup P$ so that each plaquette is in its pure state. Moreover if P is frustrated and the attaching edge is satisfied by a state s of \mathcal{S} then s may be extended in 3 ways so that P is in a pure state. This

situation happens with non-zero probability, and so \mathcal{L} must have an unbounded number of pure states.

2 Square lattices

Any happy plaquette has 2 ground states. For example, the ground states on a plaquette with four positive bonds are the two spin assignments where all the spins are either up or down. In general, for any ground state there is a complementary ground state which is obtained by reversing all the spins.

Any frustrated plaquette has 8 ground states. Due to the spin-flipping symmetry, we can choose one spin arbitrarily and the ground state can be extended to the other 3 spins in 4 ways. However, if we choose 2 spins independently, the possible extensions depend on our choice of the spins.

Assume we choose spin values for vertices i, j where $\{i, j\}$ is an edge of the frustrated plaquette. If $\{i, j\}$ is a positive edge and the spins of i, j are equal, then there are 3 ground states, and if the spins are different then there is 1 ground state. If $\{i, j\}$ is a negative edge and the spins of i, j are the same, there is 1 ground state, and if the spins are different then there are 3 ground states. In other words, a *satisfied* edge can be extended in 3 ways, while a non-satisfied edge can be extended in only 1 way.

However, if we choose spin values for vertices i, j where $\{i, j\}$ is not an edge ("diagonal" vertices), the ground state can be always extended to the other two vertices in 2 different ways, regardless of the first two spin values.

We observe that among the extremal lattices, maximum degeneracy is attained when all the plaquettes in the cover (which is satisfiable) are frustrated. Indeed, let P_i be a happy plaquette in the cover. Then we can change it into a frustrated plaquette by altering one of its edges and it holds that any ground state on the original lattice remains a ground state on the new one. Therefore, the number of ground states on an extremal lattice can only in-

crease if we turn all the happy plaquettes in the satisfiable cover into frustrated ones.

We consider a cover of frustrated plaquettes and we try to estimate the number of ground states. In order to do so, we introduce an algebraic formalism which transforms the number of ground states into the trace of a certain linear operator.

2.1 A frustrated strip

In order to explain our formalism let's first consider a simpler situation where we have only a single column of plaquettes P_1, \dots, P_n . Let's construct the spin assignments that satisfy all these plaquettes (which we will call *optimal assignments*). Note that a groundstate of a single column of plaquettes doesnot need to be an optimal assignment. We are counting the number of optimal assignments rather than the number of groundstates in order to introduce our method.

We start from the topmost pair of spins (a_1, b_1) and choose arbitrary values for them. There are several ways in which we can choose the values for the next pair of spins (a_2, b_2) , so that we get a ground state for plaquette $P_1 = (a_1, b_1, b_2, a_2)$. As mentioned above, we have 3 options if the edge $\{a_1, b_1\}$ has been satisfied and only 1 option otherwise. However, it will be important how the optimal assignments to (a_2, b_2) depend on (a_1, b_1) . We encode this dependence in a 4×4 matrix A which is indexed by the spin assignments to $\alpha = (a_1, b_1)$ and $\beta = (a_2, b_2)$. There are four possibilities for each pair of spins and we will consider them as indices, in the order $(++, -+, +-, --)$. We define

- $a_{\beta\alpha} = 1$ if $[\alpha, \beta]$ satisfies the plaquette
- $a_{\beta\alpha} = 0$ otherwise

For example, consider a plaquette P where the bond $\{b_1, b_2\}$ is negative while the other 3 bonds are positive. The corresponding

matrix would be

$$A = \begin{pmatrix} 1 & 0 & 1 & 1 \\ 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \\ 1 & 1 & 0 & 1 \end{pmatrix}$$

It is important to realize how this matrix works as a linear operator. If we take the unit vector \mathbf{e}_α then the vector $\mathbf{x} = A\mathbf{e}_\alpha$ is the α -column of the matrix which describes the optimal assignments to spins (a_2, b_2) , on the condition that the values of spins (a_1, b_1) are given by α . We can also say that the β -component of the vector $\mathbf{x} = A\mathbf{e}_\alpha$ equals the number of optimal assignments of plaquette $P_1 = P$ where $\alpha = (a_1, b_1)$ and $\beta = (a_2, b_2)$.

Let $A = A_1$ and let A_2 be the matrix describing plaquette P_2 . Let (a_3, b_3) be the edge of P_2 disjoint with P_1 . Then we analogously get that the β -component of the vector $\mathbf{y} = A_2 A_1 \mathbf{e}_\alpha$ equals the number of optimal assignments of the column of two plaquettes P_1, P_2 where $\alpha = (a_1, b_1)$ and $\beta = (a_3, b_3)$.

Now suppose we have constructed all the optimal assignments up to spins (a_k, b_k) , the number of these assignments is encoded in a *optimal vector* \mathbf{y} (whose α -coordinate is the number of optimal assignments with the spin values of (a_k, b_k) given by α), and A_k is the matrix corresponding to plaquette P_k . Then multiplying the vector \mathbf{y} by A_k realizes the *extension* of our optimal assignments to P_k - the vector $A_k \mathbf{y}$ encodes the number of optimal assignments with respect to the spin values of (a_{k+1}, b_{k+1}) .

Inductively, for a strip of n plaquettes we get that

$$g(\alpha, \beta) = \mathbf{e}_\beta^T A_n A_{n-1} \dots A_1 \mathbf{e}_\alpha$$

is the number of optimal assignments on plaquettes P_1, \dots, P_n , with value α on the topmost pair (a_1, b_1) and value β on the bottom pair (a_{n+1}, b_{n+1}) . In case the strip of plaquettes is periodic, we just identify the top with the bottom and we get that the total number of optimal assignments is

$$g = \sum_{\alpha} \mathbf{e}_\alpha^T A_n A_{n-1} \dots A_1 \mathbf{e}_\alpha = \text{Tr}(A_n \dots A_1).$$

For example, let's consider a periodic strip of n equal frustrated plaquettes characterized by the matrix A above. Then the number of optimal assignments is

$$g = \text{Tr}(A^n) = \sum_{i=1}^4 \lambda_i^n$$

where $\lambda_1, \lambda_2, \lambda_3, \lambda_4$ are the eigenvalues of A . Since the largest eigenvalue of A is $\lambda_1 = 1 + \sqrt{2}$ (and the other eigenvalues are ± 1 and $1 - \sqrt{2}$), the number of optimal assignments grows approximately as $(1 + \sqrt{2})^n$.

2.2 Square lattices

Now let's extend this formalism to square lattices. Let \mathcal{L} be $L \times L$ square lattice and let \mathcal{C} be its cover which is satisfiable. Since \mathcal{C} is satisfiable, any groundstate of \mathcal{L} is an optimal assignment for \mathcal{C} and vice versa. The approach will be similar to the previous case of a single column of plaquettes but here we will need much richer vector space because the vector components will count ground states with a given configuration on the entire boundary of an area partially covering the lattice. Specifically, we will need to encode the configurations of L independent spins.

How do we construct the ground states? We start from a horizontal row of L vertices and proceed downward by adding edge-disjoint plaquettes of \mathcal{C} which always share 2 vertices with the area covered so far. We are keeping track of the configuration on the "bottom" boundary \mathcal{B} of the presently covered area. The initial boundary is the "top" row and it grows downward. Let $S = \{-1, 1\}^L$. Each element of S corresponds to a spin assignment of a row $\mathcal{B} = (v_1, \dots, v_L)$ of L vertices. Any non-negative integer vector indexed by the elements of S will be called a *ground state vector*. A ground state vector has 2^L components each of which corresponds to a fixed spin configuration on \mathcal{B} . The α -component describes how many times the configuration α appears on \mathcal{B} as the boundary of a ground state on the area bounded by \mathcal{B} .

Now how does adding a new plaquette affect the boundary \mathcal{B} and a given ground state vector? Two vertices on \mathcal{B} are replaced by two new vertices and the ground states can be extended in different ways depending on the two relevant spin values.

Specifically, let $P = (a_1, b_1, b_2, a_2)$ be a plaquette of \mathcal{C} located horizontally at position i (that is, its left-hand vertex a_1 is i and its right-hand vertex b_1 is $j = i + 1$, or $j = 1$ if $i = L$). Let $Con(P)$ denote the graph consisting of \mathcal{B} and plaquette P glued by a_1, b_1 to v_i, v_j . Adding the plaquette P modifies \mathcal{B} so that v_i is replaced by a_2 and v_j is replaced by b_2 . Let's denote the modified boundary by \mathcal{B}' .

Next, we want to associate with P a linear operator \mathcal{A}^P such that for any ground state vector \mathbf{z} describing the possible assignments to \mathcal{B} (with multiplicities), $(\mathcal{A}^P \mathbf{z})_\alpha$ equals the number of ground states on $Con(P)$ (extensions of the prescribed configurations on \mathcal{B}) where the spin assignment on \mathcal{B}' is α . As before, we have a 4×4 matrix A associated with P , but here we must modify it into a $2^L \times 2^L$ matrix that can operate on the vector \mathbf{z} . Therefore we associate with P a linear operator \mathcal{A}^P such that

- $\mathcal{A}_{\alpha\beta}^P = a_{\alpha(i,j)\beta(i,j)}$, if α and β are equal in all their components beside those corresponding to spins i and j
- $\mathcal{A}_{\alpha\beta}^P = 0$ otherwise

In other words, we take the original plaquette operator, place it in the appropriate dimension, and extend it (tensor-wise) by identity in the other dimensions. This is exactly what we need, because if \mathbf{z} is the ground state vector for an area with a given bottom boundary, $\mathcal{A}^P \mathbf{z}$ is the ground state vector for the area with the new plaquette P added at the bottom.

Finally, we can repeat the argument from the previous section. We start from any configuration of spins and we keep adding plaquettes until we arrive at the bottom row of vertices, which we identify with the top row, so we require that the top and bottom configurations are the same. If the plaquette operators are $\mathcal{A}_1, \dots, \mathcal{A}_p$

(in the order as they have been added), the total number of ground states is

$$g = \text{Tr}(\mathcal{A}_p \mathcal{A}_{p-1} \dots \mathcal{A}_1).$$

2.3 Upper Bound

Now we apply the expression of ground state degeneracy by means of linear operators to bound the maximum possible degeneracy. We normalize the degeneracy with respect to lattice size and we study the *ground state entropy*

$$\epsilon = \frac{\log_2 g}{n}$$

where n is the number of lattice vertices and g is the number of ground states. Note that the number of ground states can range between 1 and 2^n so we always have $\epsilon \in [0; 1]$ and the parameter gives an impression of what portion of all the states are ground states.

We use the following characteristic of a matrix.

Definition 2.1 *The first singular value of a matrix is defined as*

$$\sigma_1(A) = \max\{\|A\mathbf{x}\|; \|\mathbf{x}\| = 1\}.$$

The following well-known lemma describes how $\sigma_1(A)$ can be evaluated efficiently.

Lemma 2.2

$$\sigma_1(A) = \sqrt{\lambda_1(A^T A)}$$

where $\lambda_1(A^T A)$ is the largest eigenvalue of $A^T A$ (whose eigenvalues are real and non-negative).

Lemma 2.3 *If $\mathcal{A}_1, \dots, \mathcal{A}_p$ are linear operators of dimension d then*

$$\text{Tr}(\mathcal{A}_p \mathcal{A}_{p-1} \dots \mathcal{A}_1) \leq d \sigma_1(\mathcal{A}_p \mathcal{A}_{p-1} \dots \mathcal{A}_1).$$

Proof: The trace of a linear operator is the sum of its d eigenvalues. If the eigenvalues of $\mathcal{A}_p \mathcal{A}_{p-1} \dots \mathcal{A}_1$ are $\lambda_1, \dots, \lambda_d$, we have

$$|\mathrm{Tr}(\mathcal{A}_p \mathcal{A}_{p-1} \dots \mathcal{A}_1)| \leq \sum_{i=1}^d |\lambda_i| \leq d \sigma_1(\mathcal{A}_p \mathcal{A}_{p-1} \dots \mathcal{A}_1).$$

The elements of $\mathcal{A}_1, \dots, \mathcal{A}_p$ are natural numbers, though, so the trace is real and non-negative and we can omit the absolute value. \square

Lemma 2.4 *For any two matrices A, B :*

$$\sigma_1(AB) \leq \sigma_1(A)\sigma_1(B).$$

Proof: Let $\|\mathbf{x}\| = 1$ and $\mathbf{y} = \frac{B\mathbf{x}}{\|B\mathbf{x}\|}$:

$$\|AB\mathbf{x}\| = \|A\mathbf{y}\| \|B\mathbf{x}\| \leq \sigma_1(A)\sigma_1(B).$$

\square

Corollary 2.5

$$\mathrm{Tr}(\mathcal{A}_p \mathcal{A}_{p-1} \dots \mathcal{A}_1) \leq d \sigma_1(\mathcal{A}_p) \dots \sigma_1(\mathcal{A}_1).$$

Now we just need to calculate $\sigma_1(\mathcal{A})$ for the operators corresponding to frustrated plaquettes. As the operator has been constructed by tensor multiplication of a 4×4 matrix A with identity, its singular value is the same as that of the matrix A . By elementary calculation, the reader can verify that for any frustrated plaquette,

$$\sigma_1(A) = 1 + \sqrt{2}.$$

Thus we obtain (for an $L \times L$ lattice, we have $n = L^2$ vertices and $p = \frac{n}{2}$ plaquettes in the cover):

$$\mathrm{Tr}(\mathcal{A}_p \dots \mathcal{A}_1) \leq 2^L \sigma_1(\mathcal{A}_p) \dots \sigma_1(\mathcal{A}_1) = 2^L (1 + \sqrt{2})^p.$$

For the entropy we get

$$\epsilon = \frac{\log_2 \text{Tr}(\mathcal{A}_p \dots \mathcal{A}_1)}{n} \leq \frac{1}{L} + \frac{1}{2} \log_2(1 + \sqrt{2}) < 0.636$$

(for L large enough).

This is already a non-trivial upper bound but we can do better than this. So far, we have not taken advantage of the dependencies between neighboring plaquettes, and our approach is based only on the fact that the plaquettes in the cover are frustrated. Now we will improve on this.

The idea is to group the operators in the product $\mathcal{A}_p \dots \mathcal{A}_1$ into uniform "gadgets" of fixed size whose eigenvalues can be obtained numerically. The most convenient gadget proves to be the "pyramid" operator which is organized as follows.

Definition 2.6 *The pyramid operator $\mathcal{P}^{(k)}$ is a product of $\frac{k(k+1)}{2}$ plaquette operators of the form*

$$\mathcal{P}^{(k)} = \mathcal{A}^{(1,2)} \mathcal{A}^{(3,4)} \dots \mathcal{A}^{(2k-1,2k)} \mathcal{A}^{(2,3)} \mathcal{A}^{(4,5)} \dots \mathcal{A}^{(2k-2,2k-1)} \dots \mathcal{A}^{(k,k+1)}$$

where $\mathcal{A}^{(i,i+1)}$ is a plaquette operator located at the spin position as indicated. This corresponds to a pyramid of plaquettes pointed upward. Similarly, the transpose operator to $\mathcal{P}^{(k)}$ corresponds to a pyramid pointed downward.

If the horizontal size of the lattice is divisible by $2k$ and the vertical size is divisible by $k+1$, the lattice cover can be decomposed exactly into a union of $q = \frac{n}{k(k+1)}$ pyramids $\mathcal{P}^{(k)}$. Then the number of ground states can be written as

$$g = \text{Tr}(\mathcal{P}_q^{(k)} \mathcal{P}_{q-1}^{(k)} \dots \mathcal{P}_1^{(k)})$$

where half of the operators are "transposed pyramids" but this does not affect the eigenvalues. Nonetheless, the pyramid operators are not all the same, depending on the specific choice of frustrated plaquettes, and in order to find the maximum singular values we

have to test all the combinations. This is too tedious to do by hand and we have used a computer to calculate the largest possible singular value for the pyramid operator $\mathcal{P}^{(k)}$, which we denote by $s(k)$. Here are the values of $s(k)$ as far as we have been able to compute them:

- $s(1) = 2.414$
- $s(2) = 11.597$
- $s(3) = 114.14$
- $s(4) = 2368.2$
- $s(6) = 9501461.5$
- $s(8) = 745550617441.5$

For the entropy, we get

$$\epsilon \leq \frac{\log_2 s(k)}{k(k+1)}.$$

In case the lattice cannot be covered entirely by the pyramids, we cover as large a portion as possible and leave the remaining plaquettes as individual operators. As the number of remaining individual plaquettes is $O(L) = o(n)$, this will not affect the asymptotic behavior of the entropy.

The pyramids of increasing size yield decreasing upper bounds. We obtain the best result for $\mathcal{P}^{(8)}$ which yields

$$\epsilon \leq 0.548.$$

Let us repeat here that this is the upper bound for all square lattices with a satisfiable cover. It is not known whether it can be exceeded by a lattice which does not have a satisfiable cover but it does not seem very likely.

It is also worth mentioning that the maximum value is not achieved for a lattice where *all* the plaquettes are frustrated, as

one might have guessed. We get the optimum when the plaquettes in the satisfiable cover are frustrated while the plaquettes in the complementary cover are all happy! Thus we conjecture that the square lattice with maximum degeneracy is the one with one frustrated cover and one happy cover. We call this lattice *semifrustrated*.

2.4 Lower bound

Now we are going to present a lower bound as well which is based on the same method as the upper bound. We express the number of ground states again as the trace of a linear operator which is a product of operators corresponding to plaquettes in the cover. This time, we group the plaquette operators in a different way, so that we are able to estimate the trace from below.

Definition 2.7 *The "double-row" operator \mathcal{R} is the product of operators corresponding to two consecutive rows of plaquettes in the cover.*

The "brick" operator $\mathcal{Q}^{(k)}$ is the product of k plaquettes in one row and $k - 1$ plaquettes in the next row:

$$\mathcal{Q}^{(k)} = \mathcal{A}^{(1,2)} \mathcal{A}^{(3,4)} \dots \mathcal{A}^{(2k-1,2k)} \mathcal{A}^{(2,3)} \mathcal{A}^{(4,5)} \dots \mathcal{A}^{(2k-2,2k-1)}.$$

Now let's consider a lattice of width $2kl$ where one cover is frustrated and the complement is happy. Instead of the double-row operator \mathcal{R} , we take the tensor product of l "bricks" $\mathcal{Q}^{(k)}$. (This corresponds to laying the bricks next to each other horizontally.) Note that there is one plaquette missing at the end of each brick. Equivalently, we can imagine *virtual plaquettes* in these places; a virtual plaquette is in its ground state if its top and bottom pairs of spins are the same (i.e., the plaquette only copies the two spin values to the next row). Note that if we replace the virtual plaquettes by suitable frustrated plaquettes (whose set of ground states comprises those of a virtual plaquette), the existing ground states will be preserved. Thus for the sake of a lower bound, we can assume that

the double-row operator \mathcal{R} for our lattice is the tensor product of l bricks $\mathcal{Q}^{(k)}$ and we can ignore the missing plaquettes.

Finally, we estimate the trace of an operator by applying this lemma.

Lemma 2.8 *Let μ denote the largest eigenvalue of a matrix A and assume that the other eigenvalues are strictly smaller in the absolute value. Then for any $m \geq 1$:*

$$\text{Tr}(A^m) = \mu^m(1 + o(1)).$$

Proof: Denote the eigenvalues of A by $\mu, \lambda_2, \dots, \lambda_n$. Then

$$\text{Tr}(A^m) = \mu^m + \sum_{i=2}^n \lambda_i^m = \mu^m + (n-1)o(\mu^m) = \mu^m(1 + o(1))$$

(n is constant here while m tends to infinity). □

The purpose of this lemma is that if the largest eigenvalue is unique, we can approximate the trace of A^m simply by μ^m . Here A stands for the double-row operator, μ for its largest eigenvalue and $m = \frac{L}{2}$ is the number of double rows we have to multiply. Since we consider the double-row operator as a tensor product of l "bricks" $\mathcal{Q}^{(k)}$, its eigenvalues are products of the respective eigenvalues of $\mathcal{Q}^{(k)}$ and the largest eigenvalue is simply $\mu = \lambda^l$ where λ is the largest eigenvalue of $\mathcal{Q}^{(k)}$. It only remains to check that the other eigenvalues of $\mathcal{Q}^{(k)}$ are strictly smaller which implies the same for A as well. If that holds, we can write

$$g \geq \text{Tr}(A^m) = \lambda^{lm}(1 + o(1)).$$

Once we know the value of λ for a specific brick operator $\mathcal{Q}^{(k)}$, we can estimate the ground state entropy for a lattice composed of such bricks. Recall that we have a lattice of size $L \times L$ where $L = 2kl$ and $L = 2m$. So we get

$$\epsilon = \frac{\log_2 g}{L^2} \geq \frac{\log_2 \lambda(\mathcal{Q}^{(k)})}{4k} + o\left(\frac{1}{L^2}\right).$$

The error term can be made arbitrarily small by choosing a large enough lattice.

By numerical calculation, we have obtained

- $\lambda(Q^{(2)}) = 10.225$
- $\lambda(Q^{(4)}) = 200.02$
- $\lambda(Q^{(8)}) = 76862.9$
- $\lambda(Q^{(10)}) = 1503611.1$

We have also verified that the remaining eigenvalues are strictly smaller so we can use these findings for a lower bound on the ground state entropy. Again, the best bound has been obtained for the largest brick operator $Q^{(10)}$:

$$\epsilon \geq 0.513$$

We summarize our results concerning the maximum degeneracy in the following theorem.

Theorem 2.9 *The maximum ground state entropy of an extemal square lattice satisfies*

$$0.513 \leq \epsilon \leq 0.548$$

where both bounds hold for semifrustrated lattices as well.

2.5 Degeneracy of a Fully Frustrated Lattice

The degeneracy of a lattice where *all* plaquettes are frustrated can be estimated more precisely by means of perfect matchings. We explain the method briefly here so we can compare the result with our bounds.

Let $\mathcal{L} = (V, E)$ be a square $(L \times L)$ lattice where all the plaquettes are frustrated. Hence both complementary covers of \mathcal{L} are satisfiable simultaneously and a spin assignment is a ground state

if and only if it satisfies exactly 3 edges of each plaquette of \mathcal{L} . Let $\mathcal{L}^* = (V^*, E^*)$ be the geometric dual of \mathcal{L} , i.e. the vertex set V^* of \mathcal{L}^* corresponds to the faces of the toroidal embedding of \mathcal{L} , and if $e \in E$ is an edge on the boundary of faces F_1, F_2 then its *dual edge* e^* connects vertices F_1, F_2 of \mathcal{L}^* .

The square grid has a well-known property that it is self-dual, i.e. \mathcal{L}^* is isomorphic to \mathcal{L} . Let s be a ground state of \mathcal{L} and let $D(s)$ be the set of edges which are not satisfied by s . Then $D^* = \{e^*; e \in D\}$ forms a perfect matching of \mathcal{L}^* . Hence we get that the number of ground states of \mathcal{L} is bounded from above by the number of perfect matchings of \mathcal{L} .

In fact, it may be shown that the number of perfect matchings of \mathcal{L} equals the sum of the numbers of ground states of four lattices $\mathcal{L}_0, \mathcal{L}_1, \mathcal{L}_2, \mathcal{L}_3$, such that $\mathcal{L}_0 = \mathcal{L}$ and $\mathcal{L}_1, \mathcal{L}_2, \mathcal{L}_3$ are obtained from \mathcal{L} by flipping the signs of the coupling constants on a horizontal boundary H and/or a vertical boundary V . Clearly, such a flipping does not alter the frustration of any plaquette so the four lattices are still fully frustrated. The number of ground states of one of these lattices is at least one fourth of the number of perfect matchings which implies that the two quantities are equal up to a factor of 4 (at least for one of the four lattices).

The number of perfect matchings of \mathcal{L} was calculated analytically by Kasteleyn [2] and that gives us the following:

Theorem 2.10 *The ground state entropy of a square ($L \times L$) lattice where each plaquette is frustrated is approximately 0.4206336 (as L tends to infinity).*

Note that the degeneracy of a fully frustrated lattice is substantially lower than that of a semifrustrated lattice. However, there are still two important questions left unsolved:

- What is the precise degeneracy of the semifrustrated lattice?
- Can it be exceeded by the degeneracy of non-extremal lattices?

3 Cubic lattices

Now we move to cubic (3-dimensional) lattices. Here many of the methods for planar or toroidal lattices fail as they rely on planar embeddings and the concept of duality. Cubic lattices cannot be embedded in any surface of bounded genus and there is no obvious way to transform the ground states into a more convenient combinatorial structure such as perfect matchings.

But surprisingly, most of the tools developed in this paper work without much modification for cubic lattices as well. Let's see how we extend our concepts to the 3D lattice. First, let's define a *periodic cubic* lattice.

Definition 3.1 *A periodic cubic lattice of size $L \times L \times L$ is a graph (V, E) where*

- $V = \{[x, y, z]; 0 \leq x, y, z < L\}$
- $E = \{\{[x, y, z], [x+1, y, z]\}, \{[x, y, z], [x, y+1, z]\}, \{[x, y, z], [x, y, z+1]\}; 0 \leq x, y, z < L\}$ (addition modulo L)

We cover this lattice by elementary cells which are edge-disjoint. Instead of a plaquette, the basic element of our cover is an elementary *cube*. This is a well-known graph containing 8 vertices and 12 edges. The cubes can be indexed by 3 coordinates just like the vertices; we adopt the convention that cube $C_{x,y,z}$ contains the vertices $[x, y, z], [x+1, y, z], \dots, [x+1, y+1, z+1]$. Observe that if we have a periodic cubic lattice of size $2L \times 2L \times 2L$, its edges can be covered by $2L^2$ cubes, for example $\{C_{2x,2y,2z}, C_{2x+1,2y+1,2z+1}; 0 \leq x, y, z < L\}$. There are four such covers, every cube belongs to exactly one cover.

Now we turn to the possible ways of frustration of a single cube. The cubes can be categorized by the frustration of their 6 faces (which is defined in the same way as in the 2D case):

- All 6 faces are happy - then the cube has only 2 ground states (with all the edges satisfied); if one spin is fixed, the ground state is uniquely determined.

- 4 faces are happy, 2 adjacent faces are frustrated - then there are 2 ground states (with 1 unsatisfied edge); if one spin is fixed, the ground state is determined.
- 4 faces are happy, 2 opposite faces are frustrated - then there are 8 ground states (with 2 unsatisfied edges); any spin assignment to one of the frustrated faces that satisfies 3 of its edges determines a ground state.
- 2 faces are happy, 4 faces are frustrated - then there are 4 ground states (with 2 unsatisfied edges); any assignment to one of the happy faces that satisfies all its edges can be extended in 2 ways to obtain a ground state.
- All 6 faces are frustrated - then there are 16 ground states (with 3 unsatisfied edges); any assignment to a face that satisfies 3 of its edges can be extended in 2 ways to obtain a ground state.

3.1 Upper Bound

Again, we employ the formalism of linear algebra. However, here we need a ground state vector that captures the configurations of the whole layer of L^2 vertices. The vector space we use has dimension 2^{L^2} and the basic operators \mathcal{A}^C correspond to elementary cubes of the lattice. We can envision the vertices relevant to the ground state vector as a horizontal layer of L^2 vertices. Multiplying the ground state vector by \mathcal{A}^C means extending the ground states to a cube attached by 4 vertices to the layer and exchanging the 4 vertices in the layer for 4 new vertices. The layer moves gradually downward until it wraps around and reaches its initial position. The result is the same as in the 2D case:

$$g = \text{Tr}(\mathcal{A}_p^C \mathcal{A}_{p-1}^C \dots \mathcal{A}_1^C).$$

When we study the first singular value of each of the different operators \mathcal{A}^C , we find that it is maximal for the fully frustrated

cube F where

$$\sigma_1(\mathcal{A}^F) = 2.$$

Therefore we get

$$g \leq 2^p$$

and as the number of vertices is $n = 4p$,

$$\epsilon = \frac{\log_2 g}{4p} \leq 0.25.$$

We improve this estimate when we study "pyramid operators" (analogous to 2D pyramids) consisting of 5 cubes: a base of 4 disjoint cubes and a top cube in the center, connected to each of the base cubes in one vertex. The optimal pyramid operator is a combination of 5 fully frustrated cubes with the respective linear operator P and its first singular value

$$\sigma_1(P) = 20.078.$$

When a cubic lattice is covered by such pyramids (and their reversed copies), we get q pyramids for $n = 20q$ vertices. The number of ground states is then

$$g \leq 20.078^q$$

In general, the lattice cannot be covered by pyramids entirely, but the portion of the lattice that has to be covered by individual cubes is asymptotically irrelevant. This yields the following upper bound:

$$\epsilon = \frac{\log_2 g}{20q} \leq 0.222.$$

3.2 Lower bound

Unfortunately, here we cannot apply the method of "brick laying" because it always leaves some holes in the cover. In the 2D case, we called them "virtual plaquettes" and we argued that substituting them with real plaquettes does not destroy any ground states.




Figure 1: Lattice cell with a free spin

However, in 3D we would need an elementary cube that satisfies the conditions of a "virtual cube". That is, any state such that the assignments to the top and bottom faces are equal should be a ground state. There is no such cube: It would have to have 16 different ground states which is true for the fully frustrated cube, but the structure of its ground states is different.

The best lower bound we could find is the following:

Consider a cubic lattice of even size in each dimension, built up from a $2 \times 2 \times 2$ cell which is repeated periodically. Every vertex with all 3 coordinates even is a "free spin" which is incident with 3 negative edges and 3 positive edges (see Fig. 1; the thick edges are negative). All other edges are positive. It can be seen that a state with all spins equal is a ground state, because any $2 \times 2 \times 2$ cell must contain at least 3 unsatisfied edges. (This is demonstrated by the 3 disjoint frustrated cycles in each cell, marked by dotted lines in the picture.) Therefore a state with all spins positive is a ground state. Furthermore, when we flip any "free spin", the energy of the state remains the same, because these spins are incident with 3 satisfied and 3 unsatisfied edges. Therefore we can choose the value of every free spin arbitrarily and still obtain a ground state.

If we build up a lattice by repeating this gadget k times in each dimension, we get a $2k \times 2k \times 2k$ lattice. This lattice has $g \geq 2^{k^3}$ ground states and for the ground state entropy, we get

$$\epsilon = \frac{\log_2 g}{(2k)^3} \geq 0.125.$$

To summarize, we have proved

Theorem 3.2 *The maximum ground state entropy of an extremal cubic lattice satisfies*

$$0.125 \leq \epsilon \leq 0.222.$$

References

- [1] B. Nachtergaele, personal communication 1999.
- [2] P. W. Kasteleyn, The statistics of dimers on a lattice, *Physica* 27, 1209-1225.