

Height Representation, Critical Exponents, and Ergodicity in the Four-State Triangular Potts Antiferromagnet

Cristopher Moore¹ and M. E. J. Newman¹

Received June 2, 1999; final December 20, 1999

We study the four-state antiferromagnetic Potts model on the triangular lattice. We show that the model has six types of defects which diffuse and annihilate according to certain conservation laws consistent with their having a vector-valued topological charge. Using the properties of these defects, we deduce a $(2+2)$ -dimensional height representation for the model and hence show that the model is equivalent to the three-state Potts antiferromagnet on the Kagomé lattice and to bond-coloring models on the triangular and honeycomb lattices. We also calculate critical exponents for the ground-state ensemble of the model. We find that the exponents governing the spin-spin correlation function and spin fluctuations violate the Fisher scaling law because of constraints on path length which increase the effective wavelength of the spin operator on the height lattice. We confirm our predictions by extensive Monte Carlo simulations of the model using the Wang-Swendsen-Kotecký cluster algorithm. Although this algorithm is not ergodic on lattices with toroidal boundary conditions, we prove that it is ergodic on lattices whose topology has no noncontractible loops of infinite order, such as the projective plane. To guard against biases introduced by lack of ergodicity, we perform our simulations on both the torus and the projective plane.

KEY WORDS: Antiferromagnets; Potts models; topological defects; Monte Carlo methods; ergodicity; critical exponents; height representations.

I. INTRODUCTION

There has been considerable interest over the last few years in the properties of classical spin systems possessing highly degenerate ground states.

¹ Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, New Mexico 87501.

Many such models, including ice models,^(1, 2) the triangular Ising antiferromagnet,^(3, 4) and dimer models^(5, 6) have been found to have ground state ensembles which display critical properties such as algebraically decaying spin-spin correlations and divergent fluctuations in the order parameter. It is now known that these properties are associated with the existence of interface or “solid-on-solid” representations for the models, in which sites can be assigned heights h which vary smoothly over the lattice and which can be mapped onto the states of the spins or other microscopic variables in a simple way. If we assume that this height field behaves as a Gaussian surface, which is justified if the model is in its rough phase, then the critical behavior follows in a straightforward fashion.⁽⁷⁻⁹⁾ The values of the critical exponents are related to the stiffness K of the surface and the wavelength of the appropriate operator as a function of height.

A few models have been studied for which h is vector rather than scalar and the ideas above generalize to this case also.⁽¹⁰⁻¹³⁾ In this paper we look at one particular model of this type, the four-state antiferromagnetic Potts model on the triangular lattice. This turns out to be an especially lucid example of a model with a vector height, being defined on a simple Bravais lattice and, as we will show, possessing a very straightforward height representation.

Our approach is to study the properties of the ground state ensemble by considering first the behavior of defects in the model at zero temperature. We show that there are six distinct types of non-trivial defects, and from the conservation laws that govern their collisions we deduce that they have vector charges that point to the corners of a regular hexagon. We use this observation to derive a Burgers vector for the model and hence show that when no defects are present the system has a two-dimensional height representation. The defects then correspond to screw dislocations on a $2+2$ dimensional lattice and we predict that pairs of them will be attracted or repelled with an entropic Coulomb force proportional to the dot product of their charges.

We use our height representation to deduce a number of facts about the four-state model. First we show that it is equivalent to the three-state Potts antiferromagnet on the bonds of a honeycomb lattice. This equivalence has also been derived using a different approach by Baxter,⁽¹⁴⁾ but the derivation given here is nonetheless instructive because it respects the symmetries of the system under permutation of states in a way that Baxter's does not. By a simple geometrical construction we show further that the model is equivalent to the $q=3$ antiferromagnet on the Kagomé lattice, a model which has previously been studied by Huse and Rutenberg⁽¹⁰⁾ who found a height representation equivalent to ours. Employing results due to Kondev and Henley,⁽¹⁵⁾ our model is then also equivalent to the fully-packed loop model on the honeycomb lattice with a loop fugacity of 2. We

generalize these equivalences to several other cases, including one related to loop models on the square lattice.

The existence of a height representation also implies, as mentioned above, that the ground state ensemble is critical and we have verified this by Monte Carlo simulation. Simulation of this model is not trivial, since no single-spin-flip algorithm is ergodic and the best-known cluster algorithm, that of Wang, Swendsen and Kotecký,^(16, 17) is not ergodic under toroidal boundary conditions⁽¹⁸⁾ and has not been proved to be ergodic for any other case. (For models defined on bipartite lattices, however, it is known to be ergodic.^(13, 19, 20)) Here we make use of our height representation to prove for the first time that the algorithm is in fact ergodic for the $q=4$ Potts model on any lattice with triangular plaquets and even coordination number which satisfies certain topological conditions. This includes lattices with free boundary conditions or with the topology of a sphere or a projective plane. To take advantage of this result we have performed simulations on the projective plane as well as on the torus. The idea of changing the lattice's large-scale topology to make a sampling algorithm ergodic appears to be new.

The outline of the paper is as follows. In Section II we introduce the model and study the types of defects which occur in it and their interactions. In Section III we derive the height representation of the model and thereby demonstrate the model's equivalence to various others. In Section IV we show how the scaling exponents characterizing the large- r behavior of correlation functions can be calculated from the height representation and in Section V we demonstrate the existence of entropic Coulomb forces between defects. In Section VI we describe the Monte Carlo algorithm we use to simulate the model and in Section VII give the results of our simulations. Finally, in Section VIII we give our conclusions.

II. DEFECTS IN THE FOUR-STATE TRIANGULAR POTTS ANTIFERROMAGNET

The q -state Potts model is a generalization of the Ising model in which a lattice is populated with spins s_i , one on each vertex i , which can take integer values $s_i = 1 \cdots q$. The spin states are also sometimes referred to as "colors," and we will sometimes make use of this metaphor. The energy of a configuration is defined to be proportional to the number of pairs of adjacent sites with the same state

$$H = -J \sum_{\langle ij \rangle} \delta_{s_i s_j}$$

In this paper we consider the Potts model with $q=4$ on the triangular lattice in two dimensions. This model can also be thought of as a discretization of the classical Heisenberg model in which each site has a three-dimensional unit vector spin \mathbf{S}_i :

$$H = -J \sum_{\langle ij \rangle} \mathbf{S}_i \cdot \mathbf{S}_j \quad (2)$$

This is equivalent to (1) up to a rescaling and an additive constant if the \mathbf{S}_i are restricted to the corners of a tetrahedron, since then the dot product of two spins depends only on whether they are the same or different.

In the antiferromagnetic case where $J < 0$, similarly-colored pairs of adjacent sites are energetically unfavorable. We refer to such pairs as “defects.” At zero temperature, the ground state has no defects and consists of the ensemble of four-colorings of the lattice. Since the lattice is three-colorable, we can introduce a finite density of the fourth color and show that the ground state has a non-zero entropy per site. This entropy can be calculated analytically⁽²¹⁾ or closely bounded with series approximations;⁽²²⁾ its exact value is $3\Gamma(\frac{1}{3})^3/4\pi^2 \simeq 1.460998$.

Consider the behavior of the model under a single-spin-flip dynamics at zero temperature, in which we choose a random site and change its color if this will not increase the energy. Such a dynamics allows existing defects to diffuse, interact, and annihilate in a variety of ways, but creates no new defects. We now show that the defects in the model fall into a number of different classes with well-defined properties. The simplest case is that of a

b

defect of the form $a \ a$ in which the sites to either side of the defect pair

b

both have the same state (here denoted b). As Fig. 1 shows, these defects retain this same form when they move. Moreover, it is possible for a defect of this type to disappear entirely if it encounters a neighborhood of the right configuration. This is shown in the rightmost portion of the figure, where the center site in the hexagon can become a 4, resulting in a defect-free configuration. Since these defects can be annealed away with only local moves, and do not require interaction with any other defects, we call them “false”; their density can be expected to fall off exponentially fast in a

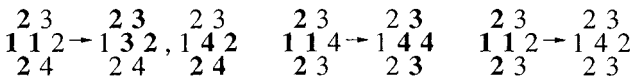


Fig. 1. Defects with the same state on both sides of the frustrated bond we call “false,” because they can be annealed away without interacting with any other defects. The defect, and the sites on each side of it, are shown in bold.

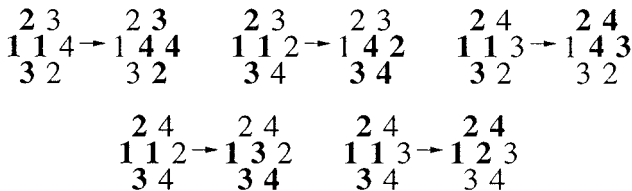


Fig. 2. As a defect diffuses, it stays within one of six equivalence classes, as defined in the text. Here we show the interaction of a defect of type X with all possible neighborhoods (excluding symmetry equivalents). The defect, and the sites on each side of it, are shown in bold.

quench to $T=0$ and so can be ignored where the long-time relaxation of the system is concerned.

b

This leaves us with defects a whose neighboring sites are in different

c

states $b \neq c$. As Fig. 2 shows, these defects are persistent and an isolated defect cannot be annealed away by local moves alone. The rest of our analysis in this section will concentrate on these “true” defects.

The true defects possess two properties which are conserved both during diffusion and in interactions with other defects. The first depends on the states a , b and c which make up the defect and its immediate neighborhood. Since a , b and c are all different for an isolated true defect, they divide the four spin states in the model into two pairs, where a belongs to one pair, and the sites on either side belong to the other. In the

2

case of the defect 1 1, for example, the defect sites belong to the pair $\{1, 4\}$

3

and the adjoining sites to the pair $\{2, 3\}$. There are three distinct ways of dividing up the states in this fashion.

The other conserved property of a true defect is a handedness defined as follows. The states a , b and c , in that order, describe a path which is either clockwise or counter-clockwise on the outside of one of the faces of a tetrahedron whose vertices are labelled with the four spin states as shown in Fig. 3. In order that this property be correctly conserved we must in addition stipulate that it remains the same under 120° rotations of a defect,

2

but changes sign under 60° or 180° ones. Thus for example the defect 1 1

3

considered above has a counter-clockwise (or positive) handedness, while an inversion or 60° rotation gives us a clockwise (or negative) handedness.

Using these two conserved properties, we can divide the 72 possible defects into 6 equivalence classes which we label X , Y , Z and \bar{X} , \bar{Y} , \bar{Z} . Here

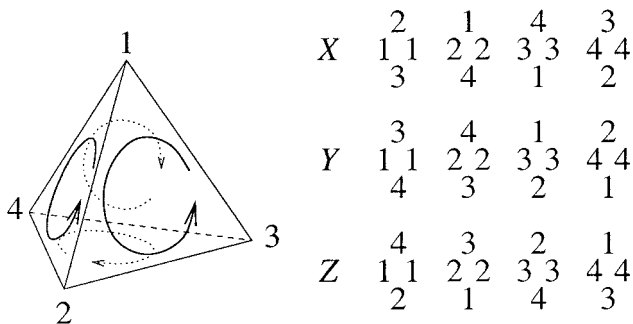


Fig. 3. The defect classes X , Y and Z all correspond to counter-clockwise paths on the outside of one of the faces of the pictured tetrahedron. The defects \bar{X} , \bar{Y} and \bar{Z} are their mirror images and correspond to clockwise paths.

the letters denote the pairing of states and the bars (or absence of them) denote the handedness. Representative members of the classes are shown in Fig. 3.

Considering again a single-spin-flip dynamics, we show in Fig. 4 a selection of possible collisions between various types of defects. Although more complex collisions than these can occur, we always find (and it is proved below) that $X + \bar{X} \rightarrow 0$, $X + Y \rightarrow \bar{Z}$, and similarly for cyclic permutations. If we wish to assign a charge χ to each particle such that $\chi(X) + \chi(Y) + \chi(Z) = 0$, $\chi(X) = -\chi(\bar{X})$, and so on, we can do this with six vectors pointing to the corners of a regular hexagon as in Fig. 5. We adopt the convention that $|\chi| = 1$, $\chi(X) = (1, 0)$, and $\chi(Y) = (-1/2, \sqrt{3}/2)$. The proof that charge is locally conserved goes as follows.

It turns out that the total charge within an area can be expressed quite simply as a sum over the plaquettes making up that area. We start by writing

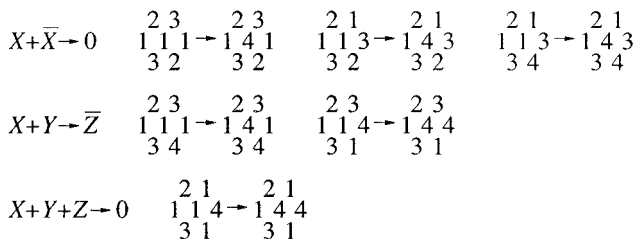


Fig. 4. Examples of collisions between defects. Note that some collisions produce a false defect, which has zero charge.

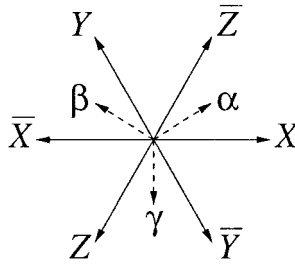


Fig. 5. The vector charges assigned to the defects. The vectors α , β and γ are defined in the text.

the charge inside a diamond $\begin{matrix} a \\ b \\ c \end{matrix}$ as a sum of functions of the upward- and downward-pointing triangular plaquets, $f(a, b, c) + g(d, c, b)$. Since a 180° rotation reverses the charges, we have $g = -f$, and since a 120° rotation keeps them the same, f is symmetric under cyclic permutations, $f(a, b, c) = f(b, c, a) = f(c, a, b)$. Furthermore, since a defect-free plaquet has zero charge, $f(a, b, c) = 0$ if a, b, c are all different, and since rotating the tetrahedron of Fig. 3 around the corner corresponding to spin state a rotates the charge plane of Fig. 5 around the origin, $f(a, a, a) = 0$ for any a .

This just leaves the case where exactly two of a, b and c are equal. Solving the equations $f(2, 1, 1) - f(3, 1, 1) = \chi(X)$, $f(3, 1, 1) - f(4, 1, 1) = \chi(Y)$, and so on, gives

$$f(a, a, b) = \begin{cases} \alpha & \text{if } \{a, b\} = \{1, 2\} \text{ or } \{3, 4\} \\ \beta & \text{if } \{a, b\} = \{1, 3\} \text{ or } \{2, 4\} \\ \gamma & \text{if } \{a, b\} = \{1, 4\} \text{ or } \{2, 3\} \end{cases} \quad (3)$$

where

$$\alpha = \left(\frac{1}{2}, \frac{1}{2\sqrt{3}} \right) \quad (4)$$

$$\beta = \left(-\frac{1}{2}, \frac{1}{2\sqrt{3}} \right) \quad (5)$$

$$\gamma = \left(0, -\frac{1}{\sqrt{3}} \right) \quad (6)$$

as shown in Fig. 5. The reader can also confirm that false defects have zero charge.

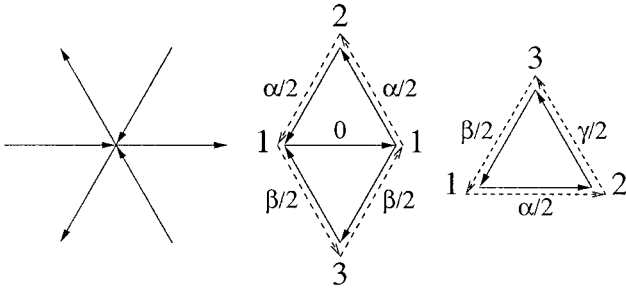


Fig. 6. Left: the direction of the vector \mathbf{t} along the edges of a triangular lattice. Right: the sums $\oint ds \cdot B = \alpha - \beta = 1$ around a diamond enclosing an X defect and $\alpha/2 + \beta/2 + \gamma/2 = 0$ around a defect-free triangle.

If the charge is conserved in all collisions as we have claimed, we should be able to write it as an integral of some quantity around the perimeter of the area we are interested in. It turns out that this is indeed possible if we define a tensor B on each edge of the lattice equal to the outer product $B = \mathbf{t} \otimes \mathbf{E}$ of two vectors, \mathbf{t} and \mathbf{E} . The first of these points along the edge and gives it a direction as in Fig. 6. This ensures that B has the necessary change of sign under 60° rotations. The second is a vector in the charge plane which depends symmetrically on the states a and b at the two ends of the edge thus:

$$\mathbf{E} = \begin{cases} \alpha/2 & \text{if } \{a, b\} = \{1, 2\} \text{ or } \{3, 4\} \\ \beta/2 & \text{if } \{a, b\} = \{1, 3\} \text{ or } \{2, 4\} \\ \gamma/2 & \text{if } \{a, b\} = \{1, 4\} \text{ or } \{2, 3\} \\ 0 & \text{if } a = b \end{cases} \quad (7)$$

Then the charge inside a finite region of the lattice is an integral around a counter-clockwise perimeter

$$\oint ds \cdot B = \sum (\Delta \mathbf{s} \cdot \mathbf{t}) \mathbf{E} \quad (8)$$

In Fig. 6, we show this for a diamond around a single defect and for a defect-free triangle. Since all larger regions can be formed by gluing together diamonds and triangles like these, and since the integral cancels along the shared edges within the region because of the sign change imparted on B by \mathbf{t} , it follows that the charge within any region is correctly given by Eq. (8). This proves our contention that the charge is conserved in all

defect collisions, since the value of the integral around a line completely enclosing any such collision does not change when the collision takes place.

III. HEIGHT REPRESENTATION AND EQUIVALENCE TO OTHER MODELS

The integral in Eq. (8) defines a Burgers vector for a defect in our model. Around any defect-free region, the Burgers vector is zero, and hence on a lattice possessing no defects the integral of B between any two points is path independent. This allows us to define a height representation for the model in which the height h at a site \mathbf{r} is defined as the integral

$$h(\mathbf{r}) = h(0) + \int_0^{\mathbf{r}} d\mathbf{s} \cdot B \tag{9}$$

along any path from 0 to \mathbf{r} , where 0 is a reference site whose height is known. The heights are thus, like the Burgers vector itself, two-dimensional vectors living on a triangular lattice with lattice vectors α , β and γ , which we multiply by $\pm\sqrt{3}$ in order to make the lattice constant 1. In fact, it is straightforward to show that there is a unique mapping of heights onto spin states, which is a four-coloring of the height lattice as shown on the left-hand side of Fig. 7. The particular permutation of the colors in this figure depends on the definition of α , β and γ given in Eqs. (4) through (6) and on the choice of reference site.

Once we have the height representation for the model, there are a number of results which follow. In this section, we use it to demonstrate the equivalence of the ground state ensemble to a number of other models, some of which have been studied previously.

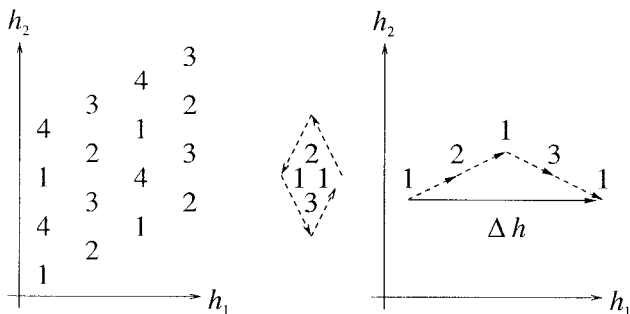


Fig. 7. Left: the mapping from heights to states. Right: integrating around a defect, in this case an X , gives a Burgers vector of $\Delta h = \sqrt{3}(2\alpha - 2\beta) = 2\sqrt{3}\chi$.

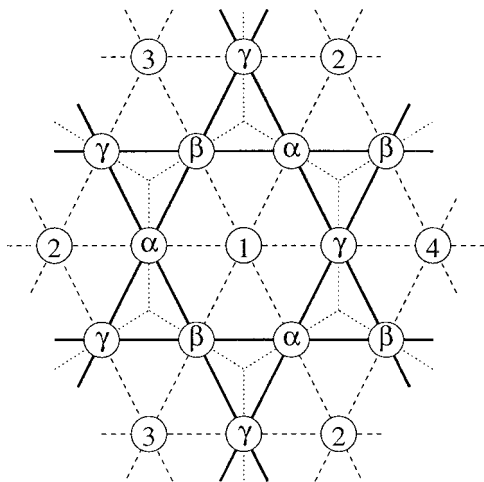


Fig. 8. The relationship between the four-state antiferromagnet on the triangular lattice (dashed) and the three-state antiferromagnet on the vertices of the Kagomé lattice (bold) or the edges of the honeycomb lattice (dotted).

First, imagine coloring the edges of the triangular lattice in a defect-free configuration of the model with three colors α , β and γ according to the height difference along them, as in Fig. 8. If we define two edges as neighboring when they bound the same triangle, then neighboring edges must have different colors since otherwise two of the vertices of the triangle would have the same spin state. Thus the model is equivalent to a three-coloring of the bonds of the triangular lattice in which no two adjacent bonds are the same color.

The reverse mapping is also possible. Since $\alpha + \beta + \gamma = 0$, the change in heights sums to zero around any triangle, and therefore around any closed curve, so the height and therefore the state of every site is well-defined once we have chosen the spin state of one reference site. As there are four choices for this reference state, every configuration of this bond-coloring model corresponds to four ground states of the four-state triangular Potts model.

Since the edges of a lattice are in one-to-one correspondence with the edges of its dual lattice, we can also think of the model as a three-coloring of the edges of the honeycomb lattice, where two edges are neighbors if they share a vertex (see Fig. 8). A simple extension of these mappings is to put a vertex at the midpoint of each edge on the triangular lattice (or the honeycomb lattice) and connect those vertices which fall on neighboring edges. The result, as shown in Fig. 8, is a Kagomé lattice. (In general, this construction is called the “medial graph.”⁽²³⁾) Thus the four-state triangular

Potts antiferromagnet and the three-state one on the Kagomé lattice also have equivalent ground-state ensembles.

A number of these results have appeared previously in one form or another. Huse and Rutenberg⁽¹⁰⁾ found a two-dimensional height representation for the $q=3$ antiferromagnet on the Kagomé lattice, which is equivalent to ours once the equivalence between models demonstrated above is taken into account. Baxter⁽¹⁴⁾ (see also ref. 24) demonstrated the equivalence of the four-state antiferromagnet and the bond-coloring model on the honeycomb lattice using an approach somewhat different from ours. He defined a cyclic ordering $4 \rightarrow 3 \rightarrow 2 \rightarrow 1 \rightarrow 4$, drawing arrows from higher states to lower ones (modulo 4) and leaving edges between states 1 and 3 or 2 and 4 blank. However, since the four-state antiferromagnet is invariant under all permutations of the four states, not just cyclic ones, we feel that the mapping given here better respects the symmetries of the system under such permutations.

Kondev and Henley⁽¹⁵⁾ showed that the bond-coloring model on the honeycomb lattice is also equivalent to a fully packed loop (FPL) model on the honeycomb lattice, where loops are defined as sets of edges alternating between two of the three colors. The loops are then contours of the component of the height perpendicular to the direction corresponding to the third color.⁽²⁵⁾ Each loop can have its colors exchanged without affecting the surrounding configuration. Such loops are said to have a fugacity $n=2$. (When $n=1$, the FPL model is equivalent to the triangular Ising antiferromagnet.⁽²⁶⁾) An exact solution for the ground state entropy of the FPL model on the honeycomb lattice for general n has been given by Batchelor, Suzuki and Yung⁽²⁷⁾ using a Bethe ansatz, and differs from Baxter's solution for the entropy of the $q=4$ triangular Potts antiferromagnet⁽²¹⁾ by exactly $\log 4$, as we would expect given the equivalence demonstrated above.

The equivalence between four-state antiferromagnets, three-state bond-coloring models, and fully-packed loop models applies on other lattices as well. If a lattice has triangular plaquets and vertices with even coordination number, we can define an orientation on the bonds with vectors \mathbf{t} which go either only clockwise or only counter-clockwise around each plaquet. Then we can define $B = \mathbf{t} \otimes \mathbf{E}$ as before, and any path around a defect-free plaquet will have $\Delta h = \pm(\alpha + \beta + \gamma) = 0$, so that a consistent definition of h exists. Furthermore, if we color the bonds with colors α , β and γ then each vertex of the dual lattice has one bond of each color. Bonds of any two colors comprise a set of fully packet loops with fugacity 2, and these are contours of the component of h perpendicular to the remaining color.

As an example, the lattice shown on the left-hand side of Fig. 9 has vertices of coordination number 4 and 8, and the orientation of the bonds

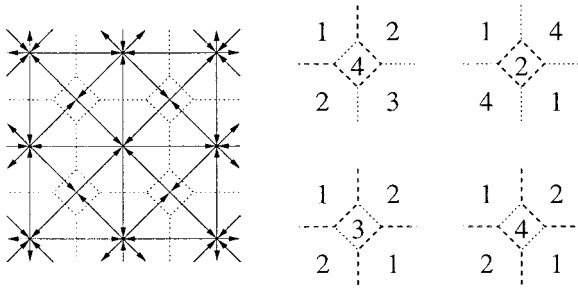


Fig. 9. Left: another lattice on which the four-state antiferromagnet has a height representation (lines with arrows), and the corresponding dual lattice (dotted). Right: the fully-packed loops of α and β bonds (dashed lines) corresponding to four different local spin configurations around one of the square plaquettes. This model is also equivalent to a loop model on the square lattice where loops can collide but not cross.

can be defined as shown. The $q = 4$ Potts antiferromagnet on this lattice is equivalent to a bond-coloring or fully-packed loop model on its dual lattice, the truncated square lattice, whose plaquettes are squares and octagons. This model was studied by Nienhuis⁽²⁸⁾ and Jacobsen⁽²⁹⁾ and is also equivalent to a model on the square lattice where loops can collide but not cross. If we take the trace over the sites with four neighbours, we are left with a model on the square lattice where plaquettes of the form $\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}$ are prohibited, while those of the form $\begin{smallmatrix} a & b \\ b & c \end{smallmatrix}$ and $\begin{smallmatrix} a & b \\ b & a \end{smallmatrix}$ have fugacities 1 and 2 respectively. A similar model where the latter two types of plaquette have equal fugacity was studied by Burton and Henley,⁽¹³⁾ who found a five-dimensional height representation for it.

We close this section by defining another model equivalent to the $q = 4$ antiferromagnet. Let us define a chirality on each plaquette of the triangular lattice according to whether the states at its three corners are oriented like an interior or exterior face of the tetrahedron in Fig. 3, or equivalently, whether the colors on its edges in the three-coloring model above cycle clockwise or anticlockwise. These plaquettes corresponds to the vertices of the honeycomb lattice, and it is not hard to see that 0, 3, or 6 of the vertices of each hexagon must have positive chirality. There are 12 configurations of the four-state triangular antiferromagnet for each state of this model, one for each choice of the states of two adjacent reference sites.

IV. FREE ENERGY AND CALCULATION OF SCALING EXPONENTS

Consider the restricted entropy of the four-state triangular model for a particular value of some (unspecified) coarse-graining of the height field h .

It is not hard to convince oneself that this entropy is lowest when $|\nabla h|$ is large, and highest for configurations that are macroscopically flat. For instance, a three-coloring of the lattice is flat since any two sites with the same spin state have the same height, and the set of configurations in the vicinity of such a three-coloring contributes a large entropy to the ground-state ensemble since every site has a choice of two colors. On the other hand, a four-coloring whose height increases linearly across the lattice corresponds to only one microstate, since no site has any choices at all. Building on considerations such as these we can derive expressions for the scaling exponents of the model. Our presentation follows that of Burton and Henley.⁽¹³⁾

If the model is in its rough phase, the arguments of the previous paragraph suggest that at large length scales it has an effective free energy of the form

$$G = \frac{1}{2} K^{\kappa\lambda} \int \nabla h_\kappa \nabla h_\lambda \, dx \, dy \quad (10)$$

where $K^{\kappa\lambda}$ is a stiffness tensor. However, since the model is invariant under permutations of the four spin states, and since these permutations are equivalent to rotations of the height lattice, K must be a scalar and we have

$$G = \frac{1}{2} K \int (|\nabla h_1|^2 + |\nabla h_2|^2) \, dx \, dy \quad (11)$$

where h_1 and h_2 are components of the height field along any two perpendicular directions in the height space. (In our calculations we have taken h_1 and h_2 along the directions indicated in Fig. 7.) In frequency space this free energy decouples into a sum over independent Gaussians, and the fluctuations in the height are given by

$$\langle |\tilde{h}(\mathbf{q})|^2 \rangle = \frac{1}{K |\mathbf{q}|^2} \quad (12)$$

In real space, this gives spatial correlations between the heights at points a distance $r = |\mathbf{r}_2 - \mathbf{r}_1|$ apart of

$$\langle |h(\mathbf{r}_2) - h(\mathbf{r}_1)|^2 \rangle \simeq \frac{1}{\pi K} \log r + C \quad (13)$$

for large r , where C is a constant.

Quantities such as spin and local magnetization are periodic functions $f(h)$ of the height and hence can be Fourier expanded in the height. We can calculate the spatial correlations in any one such Fourier component $f_g \propto e^{ig \cdot h}$, having frequency g in height space, using the fact that h 's Fourier components are Gaussianly distributed. This gives us

$$\begin{aligned} \langle f_g(\mathbf{r}_1) f_g(\mathbf{r}_2) \rangle &\propto \langle e^{ig \cdot [h(\mathbf{r}_2) - h(\mathbf{r}_1)]} \rangle \\ &= e^{-(1/2) g^2 \langle |h(\mathbf{r}_2) - h(\mathbf{r}_1)|^2 \rangle} \sim r^{-(d-2+\eta)} \end{aligned} \quad (14)$$

for large r where η is the anomalous dimension of the correlation function and d is the dimensionality of the lattice. Given that $d=2$ in the present case and making use of Eq. (13) we then find that

$$\eta = \frac{g^2}{2\pi K} \quad (15)$$

If a quantity has several non-zero Fourier components, then the one with the smallest η —i.e., the longest wavelength in height space—will dominate for large r .

If K is not a scalar, these equations generalize to

$$\langle (h'_\kappa - h_\kappa)(h'_\lambda - h_\lambda) \rangle = \frac{K_{\kappa\lambda}^{-1}}{\pi} \log r + C_{\kappa\lambda} \quad (16)$$

and

$$\eta = \frac{gK^{-1}g^\dagger}{2\pi} \quad (17)$$

For instance, transforming a scalar K to triangular coordinates where basis vectors are an angle $\frac{2}{3}\pi$ apart rather than orthogonal gives a matrix K'

$$K' = K \begin{pmatrix} 1 & -1/2 \\ -1/2 & 1 \end{pmatrix} \quad \text{and} \quad (K')^{-1} = \frac{1}{K} \begin{pmatrix} 4/3 & 2/3 \\ 2/3 & 4/3 \end{pmatrix} \quad (18)$$

Using the equivalence of the bond-coloring model to the fully-packed loop model on the honeycomb lattice, and relating the vortex-antivortex correlation function to the probability that two sites lie on the same loop, Kondev and Henley⁽³⁰⁾ have shown that these models are exactly at their roughening transition and have a stiffness of $K = \frac{2}{3}\pi$. Since our four-state model is equivalent to these models, it has the same value of K . We will use these results below to calculate the scaling exponents of various quantities for comparison with Monte Carlo experiments.

V. FORCES BETWEEN DEFECTS

Since the energy of a pair of defects is 2 regardless of how far apart they are, there is no energy gradient to drive a force between them. However, there is an entropic force, driven by the fact that the presence of a free defect reduces the entropy within an area of radius r by an amount proportional to $\log r$.

In a model with a one-dimensional height representation, a defect with Burgers vector b has an average field around it

$$|\nabla h| = \frac{b}{2\pi r} \quad (19)$$

For large r , Eq. 11 then gives a force between two defects with Burgers vectors b and b' of

$$F = \frac{K bb'}{\pi r} \quad (20)$$

When coupled with a mobility Γ , this gives an average velocity to the defects of

$$v = \Gamma F = \frac{\Gamma K bb'}{\pi r} \quad (21)$$

Such forces have been measured numerically for the three-state Potts model on the square lattice.⁽³¹⁾ The generalization to a higher-dimensional height representation is straightforward. Since the free energy in Eq. (11) is a sum of independent terms in h_1 and h_2 , the force between two defects with Burgers vectors \mathbf{b} and \mathbf{b}' will be proportional to $b_1 b'_1 + b_2 b'_2 = \mathbf{b} \cdot \mathbf{b}'$. Since the Burgers vector for a defect is equal to $2\sqrt{3}$ times the charge χ on that defect (see Section III), this gives us a force of

$$F = \frac{12K \chi \cdot \chi'}{\pi r} \quad (22)$$

In other words, an X and a Y will be attracted to each other, but only half as strongly as an X and an \bar{X} , and an X and a \bar{Y} will be repelled half as strongly as an X and another X .

VI. MONTE CARLO SIMULATION

In order to simulate correctly the properties of the ground-state ensemble of a system, it suffices to find a set of update moves which take

us from one ground state to another without introducing any defects, such that every ground state can be reached from every other in a finite number of moves on a finite lattice. An algorithm based on such a set of moves is said to be ergodic, and it can be shown that any ergodic algorithm will sample all ground states with equal frequency over the course of a long simulation (see ref. 32 for example). Unfortunately, it turns out to be quite difficult to find a suitable ergodic set of moves for the four-state Potts model considered here. To begin with, it is clear that no single-spin-flip dynamics can be ergodic because finite defect-free regions can be pinned under such a dynamics, and remain pinned no matter what happens out-

side them. For instance, in the hexagon $\begin{matrix} & 1 & 2 \\ 3 & 4 & 3 \\ & 2 & 1 \end{matrix}$ every spin has at least one

neighbor of each of the other states, so no spin can change state. Since there is a positive density of such clusters in a random ground state on a large lattice, single-spin-flip dynamics will only explore an exponentially small fraction of the possible configurations.

So we are forced to turn to a cluster update algorithm to simulate this model. The algorithm we use in this paper is the zero-temperature limit of the Wang–Swendsen–Kotecký (WSK) cluster algorithm^(16,17) for Potts antiferromagnets, which is defined as follows. At each step in the simulation we choose two of the four colors on the lattice, identify all connected clusters containing only these two, and, in each cluster independently, either switch the two colors or leave them untouched with probability $\frac{1}{2}$. Clearly this preserves the property that neighbors have differing states, and, as we will show, it is computationally quite efficient.

At finite temperature the WSK algorithm is trivially ergodic, as demonstrated by Wang *et al.* in their original paper.⁽¹⁶⁾ With a little more effort it is also possible to prove that it is ergodic at $T=0$ for all $q > 2$ on the square lattice⁽¹⁹⁾ or more generally on bipartite lattices.^(13,20) The triangular lattice however is not bipartite, and in fact the algorithm is known not to be ergodic at $T=0$ on triangular lattices with toroidal boundary conditions, at least for certain lattice dimensions.⁽¹⁸⁾ In this section we make use of our height representation to prove that the algorithm is ergodic for certain other types of boundary conditions.

Previous proofs of the ergodicity of the WSK algorithm at zero temperature have relied on defining a specific “target configuration” of the spins on the lattice and demonstrating that any given configuration can be transformed into this one in a finite number of reversible moves. For bipartite lattices the target configurations used have been checkerboard colorings. Here we use the same approach to prove the algorithm ergodic on the triangular lattice but, since the lattice is not bipartite and therefore




Fig. 10. A three-coloring of the triangular lattice. There are six such colorings for each choice of three colors.

does not permit checkerboard colorings, we use instead a three-coloring of the lattice as our target configuration.

There are six possible three-colorings of the triangular lattice for each choice of 3 of the 4 colors. We illustrate one of them in Fig. 10. We define a *domain* to be a connected set of sites whose colors coincide with one of these three-colorings. The three colors in such a domain fall on three triangular sublattices with lattice parameter $\sqrt{3}$ times that of the fundamental lattice. Our goal is to add new sites to the domain, singly or in groups, until the domain fills the entire lattice. If we can do this with reversible cluster moves from any given starting state, then our algorithm must be ergodic, since we can get from any state to the target and then from the target to any other state. (It makes no difference which three-coloring we take as our target configuration since we can get from any one to any other in at most three steps of the WSK algorithm—one to get the three colors right and at most two to put the colors on the correct sublattices.)

Our approach is as follows. We choose a site which lies outside our domain, but is adjacent to it. The color of this site, call it a , differs from the color b of the sites inside the domain on the same sublattice. If we can switch the colors a and b everywhere inside the domain, while leaving the color of the new site unchanged, it will now match the other sites on that sublattice and we will have added it to the domain. (Note that the domain now has a new three-coloring, resulting from our switching a and b .) This can be done trivially in a single Monte Carlo move if we make the right choice of two colors for the move, but there is a catch. The problem is that there may be some cluster connecting the new site to a (possibly quite distant) site in the domain via sites of colors a and b . If such a cluster exists, then the new site will get flipped whenever we change colors within the domain, so that its color will always differ from that of the sites on the same sublattice in the domain. An example of such a situation is shown in Fig. 11.

Happily, we can invoke the height representation to prove that any state in which such a path exists must contain at least one defect, and hence that such paths cannot exist in a ground state of the system. To see

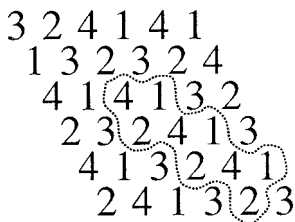


Fig. 12. A configuration on the torus with non-zero Burgers vector around its non-contractible loops. The domain denoted by the line cannot grow monotonically under the WSK algorithm.

it one plaquet at a time. For a loop of this type, the situation depicted in Fig. 11 applies, and the arguments given above are correct. If however the loop is non-contractible, then it is possible for there to be a non-localized defect and we cannot prove the ergodicity of the algorithm. To give an example, we show in Fig. 12 a configuration of the model on a lattice with toroidal boundary conditions. This lattice has two fundamental non-contractible loops, one wrapping around the boundary conditions horizontally, and one vertically. For the configuration shown, these two loops have Burgers vectors of $2\beta + 4\gamma$ and $4\beta + 2\gamma$ respectively, even though there are no localized defects. In fact, the domain inside the dotted line shows how the algorithm can fail. Any attempt to add a new site to the domain by exchanging two colors inside the domain will change the color of the new site as well, because the new site is connected to the domain by paths which wrap around the boundary conditions. This does not actually prove that the WSK algorithm is *not* ergodic on this lattice, only that we cannot prove it to be so using arguments of the type given here in which domains gain sites but do not lose them.

However, Salas and Sokal⁽¹⁸⁾ report that they found a configuration on a 6×6 torus which can only be transformed into a small number of others, thus showing that the algorithm is not in general ergodic on the torus. Huse and Rutenberg, in their study of the three-state model on the Kagomé lattice, noted a similar lack of ergodicity in a loop-flipping algorithm on the torus;⁽¹⁰⁾ it is easy to show that cluster flips in our model are equivalent to concentric sets of loop flips on the Kagomé lattice and vice versa, so these two types of moves are equivalent in terms of what configurations they can reach. While these algorithms may in fact be ergodic on an equivalence class of configurations large enough to sample the system's statistical properties, we have no guarantee of this. It seems imprudent, therefore, to conduct simulations solely on a toroidal lattice.

On lattices with no non-contractible loops, however, the algorithm is, by the arguments given above, ergodic. Examples of such lattices include

the infinite lattice and any finite lattice with free boundary conditions, where opposite edges of the lattice have no connection to each other. Unfortunately, the first of these is impractical for computer simulations, and the second suffers rather dramatically from finite-size effects. A better solution is to perform our simulation on a lattice with periodic boundary conditions, but with a topology chosen so that there are no non-contractible loops. The simplest example is the sphere. Because its Euler characteristic is 2, there is no way to cover a sphere with a regular triangular lattice; some sites have to have less than six nearest neighbors. Fortunately, our proof works not just for the triangular lattice, but for any lattice where both the height representation and the target three-coloring can be defined. As we showed in the previous section, the height representation can be defined for any planar lattice with triangular plaquets and even coordination number, so that we can place vectors \mathbf{t} on the bonds oriented so that they go only clockwise or only counter-clockwise around each plaquet. We can also define a three-coloring on such a lattice by having colors cycle upward or downward along \mathbf{t} . Thus, we can for instance perform a simulation on a lattice tiling the surface of an octahedron where 6 sites have four nearest neighbors, and the WSK algorithm will be ergodic. In Fig. 13 we depict one such lattice and illustrate how the three-coloring behaves near its corners.

Another possible topology is the projective plane—a hemisphere with diametrically opposing points on the equator identified. To create a lattice with this topology, we simply take the upper half of the octahedron depicted in Fig. 13 and identify sites along its square equator as shown in Fig. 14. Since the projective plane has only half as much curvature as the sphere (its Euler characteristic is 1) there are now just three sites with four neighbors rather than six. Thus this lattice is proportionately flatter than

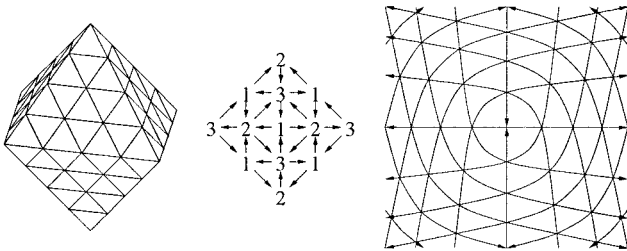


Fig. 13. Left: an octahedral lattice with $L = 4$. The WSK algorithm is ergodic on this lattice. Center: a three-coloring of this lattice around one of the vertices of the octahedron and the definition of the orientation vector \mathbf{t} along the bonds. Right: the structure is that of a triangular lattice everywhere except at the vertices of the octahedron.

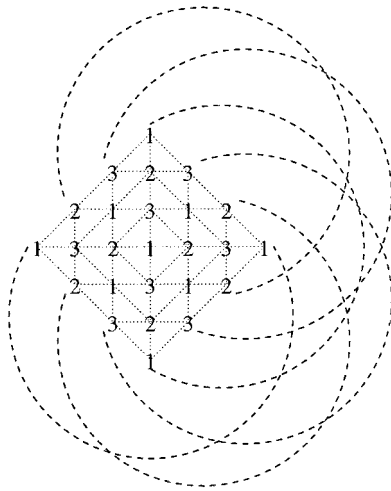


Fig. 14. An $L = 3$ lattice whose topology is that of a projective plane. Pairs of sites connected by dotted lines are identified. Note the consistency of the three-coloring.

the sphere with the same number of sites. Note that the three-coloring is well-defined for any lattice size.

Although the projective plane does possess non-contractible loops, every such loop has the property that going around it twice makes it contractible. Formally, the fundamental group Π_1 for the space is \mathbb{Z}_2 , the integers mod 2. This means that the Burgers vector of any such loop must satisfy $\Delta h + \Delta h = 0$, so $\Delta h = 0$. (Kawamura and Miyashita⁽³³⁾ show that the Heisenberg antiferromagnet on the triangular lattice has defects with a \mathbb{Z}_2 charge, so there are spin systems in which Δh can be non-zero even though $2\Delta h = 0$. This however is not the case with the present system.) Thus non-contractible loops are tolerable as long as they have finite order, i.e., as long as their n th multiple is contractible for some finite n . The sphere and projective plane are the only finite two-dimensional manifolds satisfying this condition, since both tori and Klein bottles have non-contractible loops of infinite order.

While Monte Carlo calculations performed on spheres and projective planes may seem outlandish, they are important in the present case in order to rule out the possibility that lack of ergodicity is introducing a bias into our results. The drawback is that we are forced to introduce a small number of atypical sites into the lattice (those with only four neighbors) which, for example, makes calculation of spatial correlation functions more difficult. In this paper we strike a compromise by performing some (probably non-ergodic) simulations on toroidal lattices, which give excellent statistics

for correlation functions and other quantities, and some on the projective plane, which give poorer statistics, but are more trustworthy. In fact, we find that there are no physical measurements for which the two topologies disagree, so it is possible that the simulations on the torus are “sufficiently ergodic” for our purposes, although we cannot guarantee this.

Just as with the height representation, our proof of ergodicity will work for other lattices as well, whenever we can give each bond a direction such that the bonds around each triangular plaquet go either only clockwise or only counter-clockwise. It can be used for some lattices which have a mixture of ferromagnetic and antiferromagnetic bonds, but which have no frustration. At $T=0$ vertices with a ferromagnetic bond between them can simply be identified, since they must have the same state. Then the WSK algorithm is ergodic on such lattices if the resulting graph fits the conditions above.

VII. RESULTS OF MONTE CARLO SIMULATIONS

We have performed extensive simulations of the four-state triangular model using the WSK algorithm on lattices having the topology both of the projective plane (for which the algorithm is ergodic) and of the torus (for which it probably is not). On the projective plane we simulated systems with linear dimension L equal to a power of two from $L=4$ up to $L=1024$. On the torus the three-coloring of the lattice is only well defined for L a multiple of three, so we simulated systems with $L=6, 12, 24 \dots 768$. In each case simulations ran up to one million Monte Carlo steps for the largest lattices. In order to allow accurate error estimation, and also to examine the efficiency of the algorithm, we first measured the (exponential) correlation time τ in Monte Carlo steps as a function of system size. Figure 15 show the results for both topologies. In each case the dynamic exponent z , defined by $\tau \sim L^z$, takes the value 0.74 ± 0.02 , which is comparable with values for cluster algorithms for ferromagnetic Potts models⁽³⁴⁾. Correlation times for the projective plane are about a factor of 1.3 larger than those for the torus with the same value of L . This is presumably because, as we have defined it, a projective plane of size L has $2L^2 + 1$ sites while a torus has only L^2 . Thus a projective plane has the area of a torus a factor of $\sqrt{2}$ larger, and $2^{z/2} \simeq 1.3$.

In practical terms, since a single step of the WSK algorithm updates a number of spins which scales with the area L^d of the lattice, the CPU time taken to generate a given number of independent lattice configurations scales as $L^{d+z} \simeq L^{2.74}$.

We are interested in measuring the critical exponents of the model. One way of doing this would be first to calculate the stiffness K by measuring

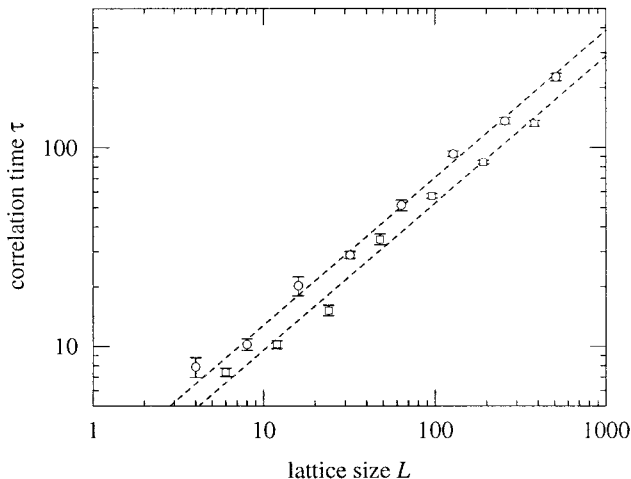


Fig. 15. Correlation times for the WSK algorithm on the projective plane (circles) and the torus (squares) measured in Monte Carlo steps. Both lattices give a dynamic exponent of $z = 0.74 \pm 0.02$.

the fluctuations in the height function and relating them to K using Eq. (12). This approach is taken in refs. 11–13, 35. Then we feed the value of K into an equation such as (17) to calculate a critical exponent. As we will see, however, calculating the exponents from the stiffness involves some surprising subtleties, which make it more desirable and instructive to measure the exponents directly.

In order to measure the critical exponents for the model, we chose two different definitions for the order parameter and measured two-point correlations and fluctuations for each of them. From these results we extracted values for the correlation exponent η by direct measurement, and for the susceptibility exponent γ/ν by finite size scaling. These exponents are not independent; we expect them to be related by the Fisher scaling law $\gamma/\nu = 2 - \eta$.^(36, 37) Our two order parameters were defined as follows:

1. The simplest magnetization measure is just $m_k = N_k - \frac{1}{4}N$ where N_k is the number of spins on the lattice in spin state k . We refer to this as the “unstaggered magnetization,” to distinguish it from our second magnetization measure defined below. The two-point connected correlation for this magnetization, averaged over k , is then $G_c(i, j) = \delta_{s_i, s_j} - \frac{1}{4}$ and the susceptibility is $\chi = \sum_k m_k^2$.

2. We also examined the “staggered magnetization” defined by Huse and Rutenberg⁽¹⁰⁾ for the three-coloring model on the honeycomb lattice, to allow direct comparison with their results. In the context of the present

model, this magnetization may be thought of as a complex number $\sum_l c_l r_l$ where the sum is over bonds on the triangular lattice, c_l represents the color α , β or γ of bond l as defined in Eq. (7), and r_l is a static reference coloring corresponding to the color of bond l in any of the 24 possible three-colorings of the lattice. The staggered magnetization m_i on a site i can be defined as the sum over the bonds l connected to that site and the two-point correlation is then given by $G_c(i, j) = m_i m_j^*$ and the susceptibility by $\chi = |\sum_i m_i|^2$.

The staggered magnetization is in fact slightly the easier of these two to analyze, so we examine this case first. Huse and Rutenberg⁽¹⁰⁾ pointed out that the staggered magnetization has a wavelength $\sqrt{3}$ on the height lattice. Using Eq. (18) to transform from Cartesian coordinates to triangular ones, Eq. (17) then gives $\eta = \frac{4}{3}$ and $\gamma/\nu = \frac{2}{3}$. Figure 16 shows our simulation results for the susceptibility on the projective plane. A least-squares fit gives $\gamma/\nu = 0.71 \pm 0.02$, which is a little greater than the expected value, but, as Fig. 17 shows, there is a clear downward trend in the value with increasing system size. Huse and Rutenberg saw similar behavior in their bond-coloring model (in fact their value for γ/ν is almost exactly the same as ours) and they attributed it to logarithmic corrections to the scaling forms which arise because the height representation is at its roughening transition.

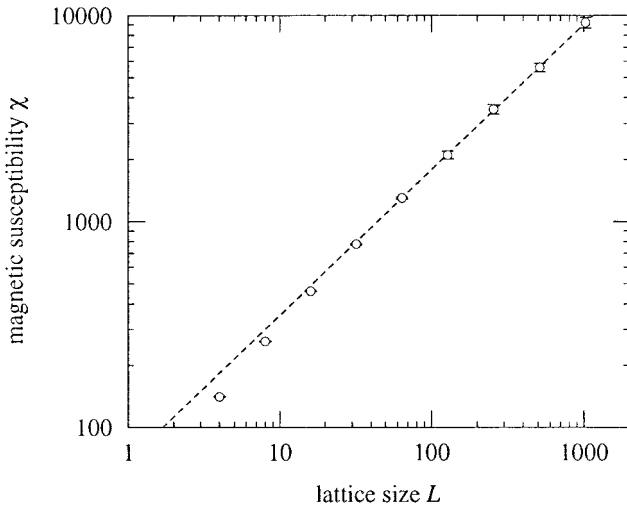


Fig. 16. The susceptibility for the staggered magnetization on the projective plane as a function of lattice size L . The dotted line is a fit to the last five points and gives $\gamma/\nu = 0.71 \pm 0.02$. Theory predicts $\gamma/\nu = \frac{2}{3}$.

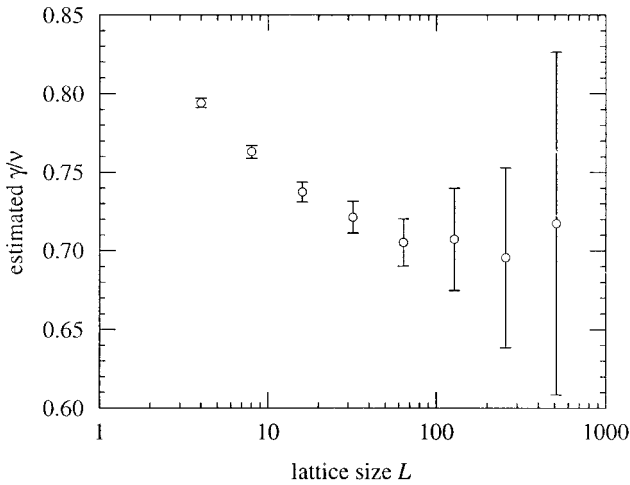


Fig. 17. Estimates of γ/v obtained from fits to simulation results for the staggered magnetic susceptibility on lattices of size ranging from L up to 1024, as a function of L . These tend downward with increasing L , presumably as a result of logarithmic corrections to the scaling form.

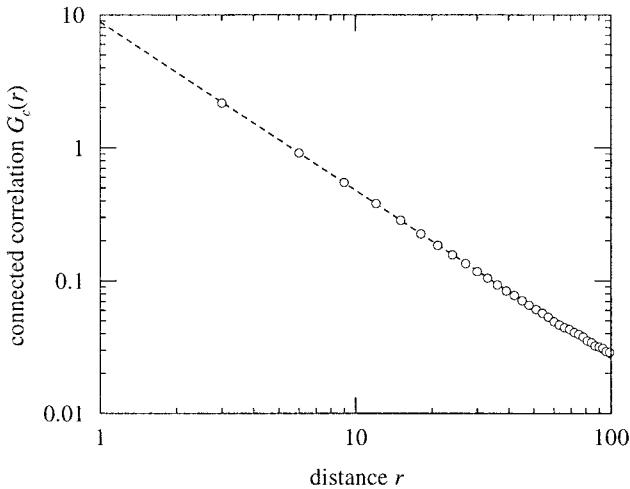


Fig. 18. The connected correlation $G_c(r)$ for the staggered magnetization on the torus. The best power-law fit is indicated by the dotted line and gives $\eta = 1.27 \pm 0.01$. Theory predicts $\eta = \frac{4}{3}$.

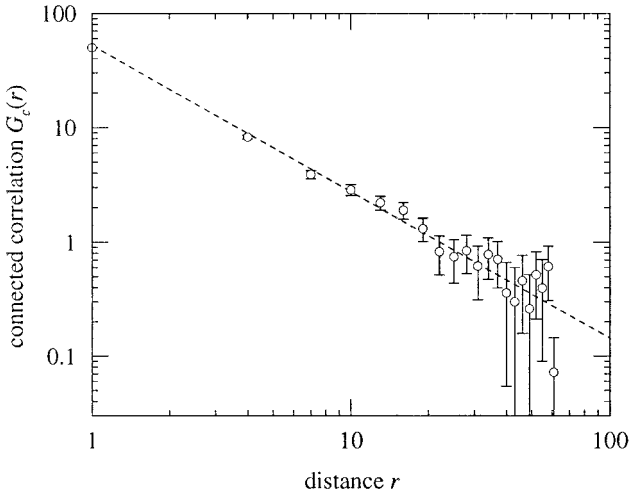


Fig. 19. The connected correlation $G_c(r)$ for the staggered magnetization on the projective plane. While the data are not as good as those for the torus in Fig. 18, we get a compatible value of $\eta = 1.28 \pm 0.09$ for the correlation exponent.

We have also measured spatial correlations in the staggered magnetization on both the torus and the projective plane (Figs. 18 and 19, respectively). The results on the projective plane have larger statistical errors than those on the torus, because of the need to stay well away from sites with local curvature in performing the calculations. Fits to the data yield values of $\eta = 1.27 \pm 0.01$ on the torus and $\eta = 1.28 \pm 0.09$ on the projective plane, in reasonable agreement with theoretical predictions.

The situation is a little more complicated for our other definition of magnetization. Looking at Fig. 7, we see that the mapping from heights to spin states has wavelength 2 on the height lattice, which implies that both the correlation exponent for the spins and the corresponding susceptibility exponent should be equal to 1. Figure 20 shows our data for the susceptibility on the projective plane. A least-squares fit gives $\gamma/\nu = 0.84 \pm 0.01$, which is some way from the theoretical prediction, but, as Fig. 21 shows, the value is increasing steadily with system size, so the discrepancy is again probably due to logarithmic corrections.

However, when we look at the spin-spin correlation function, Fig. 22, we find that $\eta = 0.35 \pm 0.01$, which is nowhere near 1. Our values for the exponents thus appear to violate the Fisher law. The explanation of this result is as follows.

The correlation function shown in Fig. 22 is measured, naturally enough, along one of the three principal directions on the toroidal lattice.

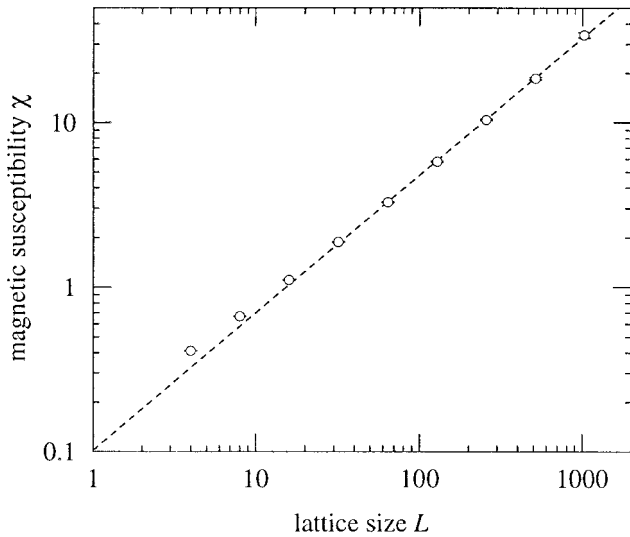


Fig. 20. The susceptibility for the unstaggered magnetization on the projective plane for systems of linear dimension L up to 1024. The dotted line is a fit to the last five points and gives $\gamma/\nu = 0.84 \pm 0.01$, while theory predicts $\gamma/\nu = 1$.

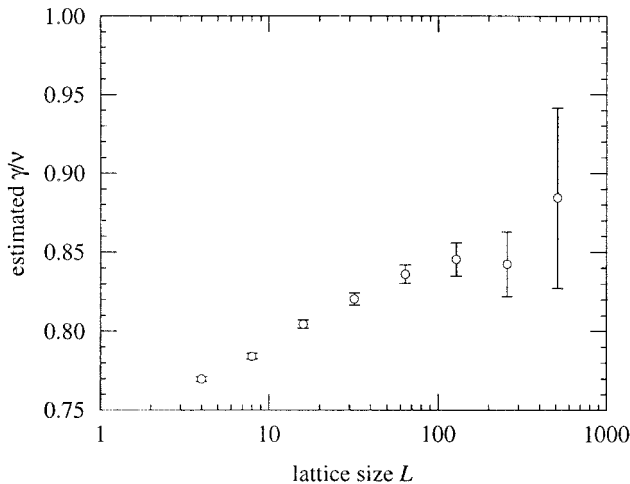


Fig. 21. Estimates of γ/ν obtained from fits to simulation results for the unstaggered magnetic susceptibility on lattices of size ranging from L up to 1024, as a function of L . These tend upward with increasing L .

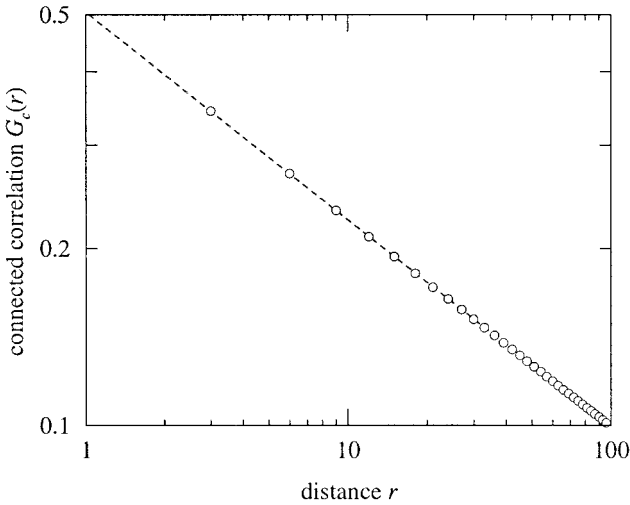


Fig. 22. The connected correlation $G_c(r)$ for the unstaggered magnetization on the torus. The best power-law fit is indicated by the dotted line and gives $\eta = 0.35 \pm 0.01$. Theory predicts $\eta = \frac{1}{3}$.

If two sites lie along such a direction and their spins are in the same state, then their heights are constrained to particular sublattices of the height lattice. For instance, suppose that the distance between them is a multiple of three, and that one site has state 1 and height $(0, 0)$. Then the height of the other site is a multiple of three terms, each equal to α , β or γ , all with the same sign since \mathbf{t} is constant along any of the lattice directions. Since $\alpha + \beta + \gamma = 0$, we can cancel terms in threes until only two kinds of terms are left, say αs and βs . Since the number of these is still a multiple of three, and since α and β (or more precisely, the unit vectors in those directions on the height lattice) have a component of $\frac{1}{2}$ along the h_1 axis, the only heights we can end up with are ones with $h_1 = 3k/2$ for some integer k . This defines a sublattice of the height lattice as shown in Fig. 23.

More generally, once we choose the distance between two sites and the spin state on one of them, the only heights the other site can have and still be in the same spin state are those on one of the three sublattices corresponding to that state. Since the lattice constant of each such sublattice is $2\sqrt{3}$, the effective wavelength of the spin operator is $2\sqrt{3}$ rather than 2. Equation (17) then gives $\eta = \frac{1}{3}$, in good agreement with our simulation results.

Thus the Fisher scaling relation can be violated when an operator is constrained in such a way that its effective wavelength on the height lattice is longer than that for the corresponding susceptibility. It is well-known

We have calculated exact values for a variety of critical exponents. The scaling exponent η for the spin–spin correlation function is of particular interest because the wavelength of the spin operator on the height lattice turns out to be longer than the fundamental periodicity of the height-to-spin mapping due to a constraint on paths connecting sites along a principal lattice direction. This gives a value of $\eta = \frac{1}{3}$ even though the corresponding susceptibility exponent $\gamma/\nu = 1$ —an apparent violation of the Fisher scaling relation.

We have also used the height representation to prove for the first time that the Wang–Swendsen–Kotecký cluster Monte Carlo algorithm is ergodic for the four-state model on the triangular lattice. Our proof however requires that the lattice have a topology which possesses no non-contractible loops of infinite order, and this means that simulations on the torus (which has such loops) are probably not ergodic, calling previous simulations of this and related models into question. Simulations on lattices with free boundary conditions or lattices with the topology of the sphere or the projective plane are, on the other hand, provably ergodic, and we have performed extensive simulations on the projective plane. We find reasonable agreement between the values of the critical exponents measured in these simulations and both theoretical values and values from previous numerical studies. There are, however, significant logarithmic corrections to scaling associated with the fact that the height model is exactly at its roughening transition, and this means that we have to go to extremely large lattice sizes to see the expected behavior.

Note Added. During the course of these investigations, we learned of an unpublished manuscript by Henley which also shows that the four-state antiferromagnet studied here is equivalent to the three-state one on the Kagomé lattice.

ACKNOWLEDGMENTS

We wish to thank Alan Sokal and Jesus Salas for sharing with us their observation that the WSK algorithm is not ergodic on lattices with toroidal boundary conditions, and Michael Lachmann for useful discussions. C.M. also thanks Molly Rose and Spootie the Cat for their support.

REFERENCES

1. H. van Beijeren, Exactly solvable model for the roughening transition of a crystal surface, *Phys. Rev. Lett.* **38**:993–996 (1977).
2. H. J. F. Knops, Renormalization connection between the eight-vertex model and the Gaussian model, *Ann. Phys.* **128**:448–462 (1981).

3. H. W. J. Blöte and H. J. Hilhorst, Roughening transitions and the zero-temperature triangular Ising antiferromagnet, *J. Phys. A* **15**:L631–L637 (1982).
4. B. Nienhuis, H. J. Hilhorst, and H. W. J. Blöte, Triangular SOS models and cubic-crystal shapes, *J. Phys. A* **17**:3559–3581 (1984).
5. W. Zheng and S. Sachdev, Sine-Gordon theory of the non-Néel phase of two-dimensional quantum antiferromagnets, *Phys. Rev. B* **40**:2704–2707 (1989).
6. L. S. Levitov, Equivalence of the dimer resonating-valence-bond problem to the quantum roughening problem, *Phys. Rev. Lett.* **64**:92–94 (1990).
7. M. P. M. den Nijs, M. P. Nightingale, and M. Schick, Critical fan in the antiferromagnetic Potts model, *Phys. Rev. B* **26**:2490–2500 (1982).
8. D. R. Nelson, Defect-mediated phase transitions, in *Phase Transitions and Critical Phenomena*, Vol. 7, C. Domb and J. L. Lebowitz, eds. (Academic Press, New York, 1983).
9. B. Nienhuis, Coulomb gas representations of phase transitions in two dimensions, in *Phase Transitions and Critical Phenomena*, Vol. 11, C. Domb and J. L. Lebowitz, eds. (Academic Press, New York, 1987).
10. D. A. Huse and A. D. Rutenberg, Classical antiferromagnets on the Kagomé lattice, *Phys. Rev. B* **45**:7536–7539 (1992).
11. J. Kondev and C. L. Henley, Four-coloring model on the square lattice: A critical ground state, *Phys. Rev. B* **52**:6628–6639 (1995).
12. R. Raghavan, C. L. Henley, and S. L. Arouh, New two-color dimer models with critical ground states, *J. Stat. Phys.* **86**:517–550 (1997).
13. J. K. Burton, Jr. and C. L. Henley, A constrained Potts antiferromagnet model with an interface representation, *J. Phys. A* **30**:8385–84113 (1997).
14. R. J. Baxter, Colorings of a hexagonal lattice, *J. Math. Phys.* **11**:784–789 (1970).
15. J. Kondev and C. L. Henley, Conformal charge and exact exponents in the $n=2$ fully packed loop model, *Phys. Rev. Lett.* **73**:2786 (1994).
16. J.-S. Wang, R. H. Swendsen, and R. Kotecký, Antiferromagnetic Potts models, *Phys. Rev. Lett.* **63**:109–112 (1989).
17. J.-S. Wang, R. H. Swendsen, and R. Kotecký, Three-state antiferromagnetic Potts models: A Monte Carlo study, *Phys. Rev. B* **42**:2465–2474 (1990).
18. J. Salas and A. D. Sokal, private communication.
19. G. T. Barkema and M. E. J. Newman, Monte Carlo simulation of ice models, *Phys. Rev. E* **57**:1155–1166 (1998).
20. S. J. Ferreira and A. D. Sokal, Antiferromagnetic Potts models on the square lattice: A high-precision Monte Carlo study, cond-mat/9811345.
21. R. J. Baxter, q colourings of the triangular lattice, *J. Phys. A* **19**:2821–2839 (1986).
22. R. Schrock and S.-H. Tsai, Upper and lower bounds for the ground state entropy of antiferromagnetic Potts models, *Phys. Rev. E* **55**:6791–6794 (1997).
23. O. Ore, *The Four-Color Problem* (Academic Press, New York, 1967).
24. R. J. Baxter, S. B. Kelland, and F. Y. Wu, Equivalence of the Potts model or Whitney polynomial with an ice-type model, *J. Phys. A* **9**:397–406 (1976).
25. J. Kondev and C. L. Henley, Geometrical exponents of contour loops on random Gaussian surfaces, *Phys. Rev. Lett.* **74**:4580–4583 (1995).
26. H. W. J. Blöte and B. Nienhuis, Fully packed loop model on the honeycomb lattice, *Phys. Rev. Lett.* **72**:1372–1375 (1994).
27. M. T. Batchelor, J. Suzuki, and C. M. Yung, Exact results for Hamiltonian walks from the solution of the fully packed loop model on the honeycomb lattice, *Phys. Rev. Lett.* **73**:2646–2649 (1994).
28. B. Nienhuis, Critical spin-1 vertex models and $O(n)$ models, *Int. J. Mod. Phys. B* **4**:929–942 (1990).

29. J. L. Jacobsen, On the universality of compact polymers, *J. Phys. A* **32**:5445.
30. J. Kondev and C. L. Henley, Kac-Moody symmetries of critical ground states, *Nucl. Phys. B* **464**:540–575 (1996).
31. C. Moore, M. G. Nordahl, N. Minar, and C. Shalizi, Vortex dynamics and entropic forces in antiferromagnets and antiferromagnetic Potts models, *Phys. Rev. E* **60**:5344–5351 (1999).
32. M. E. J. Newman and G. T. Barkema, *Monte Carlo Methods in Statistical Physics* (Oxford University Press, Oxford, 1999).
33. H. Kawamura and S. Miyashita, Phase transition of the two-dimensional Heisenberg antiferromagnet on the triangular lattice, *J. Phys. Soc. Japan* **53**:4138–4154 (1984).
34. C. F. Baillie and P. D. Coddington, Comparison of cluster algorithms for two-dimensional Potts models, *Phys. Rev. B* **43**:10617–10621 (1991);
35. C. Zeng and C. L. Henley, Zero-temperature phase transitions of an antiferromagnetic Ising model of general spin on a triangular lattice, *Phys. Rev. B* **55**:935–947 (1997).
36. M. E. Fisher, Rigorous inequalities for critical-point correlation exponents, *Phys. Rev.* **180**:594–600 (1969).
37. J. J. Binney, N. J. Dowrick, A. J. Fisher, and M. E. J. Newman, *The Theory of critical Phenomena* (Oxford University Press, Oxford, 1992).
38. J. Cardy, *Scaling and Renormalization in Statistical Physics* (Cambridge University Press, Cambridge, 1996).