From hard spheres to hard-core spins

Grace M. Sommers [1,*] Benedikt Placke [2] Roderich Moessner,[2] and S. L. Sondhi[1]
[1]*Department of Physics, Princeton University, Princeton, New Jersey 08540, USA*
[2]*Max-Planck-Institut für Physik komplexer Systeme, 01187 Dresden, Germany*

A system of hard spheres exhibits physics that is controlled only by their density. This comes about because the interaction energy is either infinite or zero, so all allowed configurations have exactly the same energy. The low-density phase is liquid, while the high-density phase is crystalline, an example of "order by disorder" as it is driven purely by entropic considerations. Here we study a family of hard spin models, which we call hard-core spin models, where we replace the translational degrees of freedom of hard spheres with the orientational degrees of freedom of lattice spins. Their hard-core interaction serves analogously to divide configurations of the many spin system into allowed and disallowed sectors. We present detailed results on the square lattice in $d = 2$ for a set of models with $\mathbb{Z}_n$ symmetry, which generalize Potts models, and their U(1) limits, for ferromagnetic and antiferromagnetic senses of the interaction, which we refer to as exclusion and inclusion models. As the exclusion and inclusion angles are varied, we find a Kosterlitz-Thouless phase transition between a disordered phase and an ordered phase with quasi-long-ranged order, which is the form order by disorder takes in these systems. These results follow from a set of height representations, an ergodic cluster algorithm, and transfer matrix calculations.

I. INTRODUCTION

Systems with constraints but no other interactions are a fascinating corner of statistical mechanics on three counts. First, their equilibrium physics is purely entropic, so any ordering they exhibit is "order by disorder." The canonical example of this somewhat counterintuitive phenomenon is nematic ordering in Onsager's model of thin, hard rods [1]. While at low densities this system lacks orientational order, with rods arranged isotropically, at a critical density the increase in *translational* entropy afforded by aligning the rods outweighs the decrease in *orientational* entropy that such an alignment evidently entails. Second, such systems probe universality in a nontrivial fashion. Viewed as interacting systems, they involve interactions of infinite strength, so many of the standard, perturbative arguments for long-wavelength universality do not directly apply to them. Finally, their dynamics has some simplifying features that have been used in simple cases to reach conclusions that are otherwise difficult. The most famous case of this is a set of results on ergodicity following the seminal work of Sinai [2].

In this paper we study a family of classical spin models in this class of "constraint-only" systems. Our models, most generally, involve of M component spins $\mathbf{S}$ of fixed length $\mathbf{S}^2 = 1$ on a specified lattice. In this paper we do not specify their dynamics but only the pairwise additive energy function

$$H = \sum_{\langle ij \rangle} V(\mathbf{S_i}, \mathbf{S_j}), \tag{1}$$

where $\langle ij \rangle$ indicates that i, j are nearest-neighbor sites and the potential energy function has the form

$$V(\mathbf{S_i}, \mathbf{S_j}) = \begin{cases} 0, & \text{if } |\mathbf{S_i} - \mathbf{S_j}| \geqslant \alpha \\ \infty, & \text{otherwise.} \end{cases} \tag{2}$$

As the parameter α moves between 0 and 2, the neighbors of a given spin are forced to lie outside a solid angle that increases from 0 to eventually cover the entire unit sphere in M dimensions on which the spins live. As the term hard spins is already reserved for spins of fixed length, we refer to these as hard-core spins. With the inequality as above we have an exclusion model of an antiferromagnetic persuasion, while with the sign of the inequality reversed we obtain an inclusion model. Explicitly, we define hard-core inclusion models by the pairwise potential

$$V(\mathbf{S_i}, \mathbf{S_j}) = \begin{cases} 0, & \text{if } |\mathbf{S_i} - \mathbf{S_j}| < \alpha \\ \infty, & \text{otherwise} \end{cases} \tag{3}$$

which is then of a ferromagnetic type. Finally, we will also consider discrete models in which only a finite number of points on the unit sphere are permitted.

One can view the exclusion models as representing a departure from Onsager's original problem in which the translational degrees of freedom are frozen while the orientational degrees of freedom, no longer directors but now spins, face a tunable set of local constraints. By contrast, the well-studied system of hard spheres and disks can be viewed as a departure in which the orientational freedom is removed and translational freedom retained. That system in three dimensions exhibits a first-order phase transition at a critical density above which crystallization occurs as a means to access greater free volume per particle and increase entropy [3]. The hard sphere system has been the subject of much work surrounding its
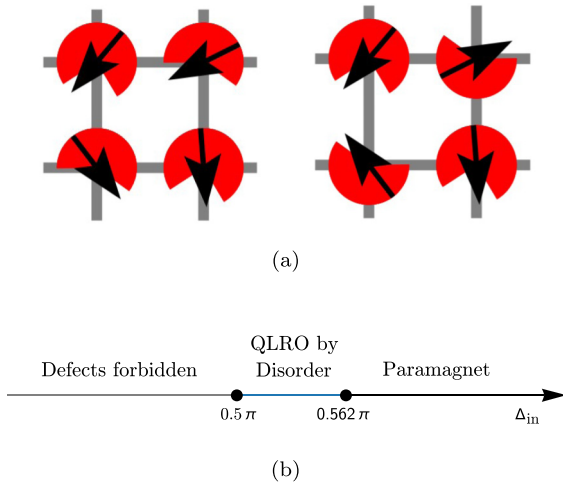
*gsommers@princeton.edu

(a)

(b)

FIG. 1. The classes of hard-core spin models studied in this paper, in which the pairwise potential takes the form of Eq. (2) or (3), specialized to the case $M = 2$ on the square lattice. (a) Valid configurations for (left) an inclusion model with inclusion angle $\Delta = 0.48\pi$ and (right) an exclusion model with exclusion angle $\pi - \Delta = 0.52\pi$. Red shading indicates orientations forbidden by the nearest-neighbor interaction. (b) Sketch of the phase diagram in the XY limit of the inclusion model. Increasing the inclusion angle, defined in Eq. (5), tunes the system from a KT phase, consisting of both a vortex-forbidden region and a quasi-long-ranged order by disorder region, to a paramagnetic phase.

dynamical properties and potential connections to the long-standing problem of vitrification, and in a subsequent paper we will report results on the dynamics of the hard-core spins studied here. Indeed, the question of dynamics is what got us thinking about this family of models in the first place.

Returning to the statics, we focus in this paper on the simplest case of $M = 2$ on the square lattice in $d = 2$. When the spins are continuous, we get a hard-core spin model of XY spins [Fig. 1(a)], and when they are discretized to take N values, we obtain generalizations of Potts models with $\mathbb{Z}_N$ symmetry. In these cases we are able to search for order by disorder and for universal long-wavelength physics dictated by the usual criteria of symmetry and dimensionality. For XY spins, true long-range order is, presumably, ruled out and only algebraically long-ranged order of the Kosterlitz-Thouless (KT) variety [4] is possible, although we are not aware of a proof that actually covers the hard-core limit. For the Z_N cases, it would appear that true long-ranged order as well as quasi-long-ranged order are both possible on grounds of symmetry and as discussed in Ref. [5].

Our principal results are (i) that both the Z_N models for sufficiently large N and the XY limiting model exhibit KT phases both when the constraint parameter forbids vortices but also for a range where they are allowed, thereby furnishing a case of order by disorder; (ii) that none of our Z_N models or their XY limits exhibit a phase with true long-ranged order; and (iii) that for $N \geqslant 4$ and extending into the XY limit, there are always models with nonzero interaction yet sufficiently small (large) exclusion (inclusion) angles which exhibit short-ranged correlations. We portray these results in the phase diagrams in Figs. 1(b) and 3, of which the latter requires some

definitions that we introduce in Sec. II below. These results are obtained by a combination of a set of height representations for our Z_N models, transfer matrix calculations, and finite-size scaling using a cluster algorithm, henceforth referred to as the `reflect` algorithm, whose ergodicity is proven for bipartite lattices.

The paper proceeds as follows. In Sec. II, we define the models and review the theoretical context for analyzing them. We then describe, in Sec. III, the two numerical methods used to construct the phase diagram, the transfer matrix method and the cluster algorithm, as well as the height representations for the vortex-free Z_N models. In Sec. IV, we present results for Z_N models and the XY limit, along with a discussion for how the phase diagram fits into and extends the general understanding of critical spin systems. We conclude in Sec. V.

II. MODELS AND BACKGROUND

A. Models and notation

Equation (2) and its inclusion counterpart, Eq. (3), express the hard constraint between neighboring spins in terms of the norm of the difference between two M-component unit vectors. Alternatively, we could formulate this constraint in terms of the distance along the great circle connecting the two vectors, or the inner product between them. The latter formulation was used in a study of $O(M)$ and RP^{M-1} constraint models for inclusion angles below $\pi/4$ [6]. For $M = 2$ these formulations are all equivalent. We find it is most natural to represent spins as complex exponentials $s_j = \exp(i\theta_j)$, where $\theta_j \in [0, 2\pi)$ is the orientation of the jth spin on the unit circle. In the XY limit, models are parametrized by an angle Δ, with the nearest-neighbor interaction for inclusion models given by

$$V(s_i, s_j) = \begin{cases} 0, & \text{if angle}(s_i, s_j) < \Delta \\ \infty, & \text{otherwise.} \end{cases} \quad (4)$$

This can be viewed as the zero-temperature limit of a step model, which has mainly been studied at finite temperature for $\Delta = \pi/2$ [7–10] but was also examined, via the Migdal approximation, for smaller angles [11].

At zero temperature, a precision study of the phase transition for the "constraint action" described by Eq. (4) found that it belongs to the Kosterlitz-Thouless universality class [12,13]. Where our results overlap with these papers, we are in agreement. However, the present research goes beyond previous works in two major respects.

First, we generalize to Z_N (clock) models, where spins can adopt N orientations at angles $2\pi\sigma/N$ on the unit circle, with $\sigma \in \{0, 1, \dots, N-1\}$ [14]. Using complex notation, we say that the spin at site j is in the state σ_j if $s_j = \exp(2\pi i\sigma_j/N)$. A hard-core inclusion model can then be characterized by a set of three parameters (N, p_{in}, p') where neighboring spins must enclose $< p$ sites on a clock with N orientations:

$$V(s_i, s_j) = \begin{cases} 0, & \text{if angle}(s_i, s_j) < 2\pi p/N \\ \infty, & \text{otherwise} \end{cases} \quad (5)$$

and p' is the number of allowed orientations for spin s_i given a fixed orientation of neighboring spin s_j.

A second generalization of the constraint action comes from inverting the piecewise relation in Eq. (4) to create an exclusion model, wherein neighboring spins must enclose an angle $\geqslant \Delta$. Similarly, Z_N exclusion models are parametrized by (N, p_{ex}, p'), where neighboring spins must enclose a minimum of p sites on the clock, and p' is again the number of allowed orientations for spin s_i given a fixed orientation of spin s_j:

$$p' = N - 2p_{ex} + 1 = 2p_{in} - 1. \qquad (6)$$

The distinction between ferromagnetic inclusion and antiferromagnetic exclusion models on frustrated lattices gives rise to essential differences in the phase diagram [15,16]. In this paper, we instead focus on the square lattice, which for even widths and lengths is bipartite. This feature results in an equivalence between inclusion and exclusion models: if S denotes a valid (zero-energy) inclusion configuration with inclusion angle Δ, then flipping all the spins on one sublattice yields a zero-energy exclusion configuration S' with angle $\pi - \Delta$, as depicted in Fig. 1(a). This equivalence is captured for finite N by the parameter p': an inclusion model with a given set of parameters (N, p') is fully equivalent to an exclusion model with the same (N, p'). We can map both classes of models onto a common set of axes, N vs inclusion angle Δ, by the relation

$$\frac{\pi}{N}(p' - 1) < \Delta \leqslant \frac{\pi}{N}(p' + 1). \qquad (7)$$

The inequality here comes from the fact that when the orientations of spins are discrete, there is ambiguity in how the real-valued inclusion angle is defined. In the following analysis, the most sensible definition of the "effective inclusion angle" will prove to be the midpoint of the interval $\Delta = \pi p'/N$. In the XY limit ($N \to \infty$), the inclusion angle is uniquely defined.

The mapping between inclusion and exclusion means that we can primarily focus our analysis on inclusion models. However, the mapping only exists when the global rotation by π on one sublattice is allowed by the discretization of the spins, namely, for even N. On the other hand, inclusion and exclusion models for odd N probe a complementary set of parameters p' (odd p' for inclusion, even p' for exclusion) as the clock lacks the $\mathbb{Z}_2$ symmetry of the lattice. In these cases, it will be necessary to separately examine exclusion models, to probe the entire parameter space.

One striking example of an exclusion model with no equivalent inclusion model is the $(N, p_{ex}, p') = (3, 1, 2)$ exclusion model, which is simply a three-state zero-temperature Potts antiferromagnet. Indeed, zero-temperature q-state Potts models can be viewed as special cases of our classes of clock models, with $p = 1$: in our notation, an $(N, 1, 1)$ inclusion model is a zero-temperature Potts ferromagnet with $q = N$, while an $(N, 1, N - 1)$ exclusion model is a zero-temperature Potts antiferromagnet (AFM).

Zero-temperature Potts AFMs, which possess macroscopic ground-state entropy, have been the subject of extensive research over the past several decades, through evaluation of chromatic polynomials [17–20] along with formal proofs for the existence [21–24] or nonexistence [25] of entropy-driven long-ranged order on various lattices. A given lattice has a

critical value q_c, such that for $q > q_c$, a q-state Potts AFM is disordered at all temperatures, including $T = 0$; for $q < q_c$ there is a finite-temperature phase transition; and for $q = q_c$ there is a zero-temperature critical point [26]. On the square lattice $q_c = 3$, a rigorous result [27–29] which tells us that the (3,1,2) exclusion model is critical.

Our class of hard-core spin models allows us to generalize these observations about Potts models in two ways. First, by considering families of fixed Δ according to the mapping in Eq. (7), we can analyze trends in ferromagnetic and antiferromagnetic models upon approach to the XY limit. Second, by constructing a height representation we show that the three-state Potts antiferromagnet belongs to a family of models with $p' = 2$ which all have the same critical behavior, and in fact each value of p' hosts a critical family.

B. General remarks about phase diagrams in 2D

Having defined our classes of models, we now situate these models within the broader context of phase transitions in spin systems.

In the absence of a temperature T, the phase of the system is determined instead by the tuning parameters (N, p') (or Δ in the XY limit). Moreover, since all allowed configurations have the same energy, a transition at some critical angle Δ_c must be driven by the quest to maximize entropy. By analogy to the phase transition in systems of hard spheres [30], which crystallize at a critical density as a means to access greater free volume per particle and increase entropy, for inclusion models we would expect ferromagnetic order to set in below Δ_c, as a means to access a greater range of allowed orientations per spin.

In classical systems with short-ranged interactions in fewer than three dimensions, however, a continuous symmetry cannot be broken at finite temperature, by the Mermin-Wagner theorem [31]. In two dimensions, systems instead undergo a Kosterlitz-Thouless (KT) transition [4] to a continuous line of critical points. The critical phase is characterized by quasi-long-ranged order (QLRO), as correlations decay algebraically, rather than exponentially, to zero [32]:

$$G(r) \equiv \langle \mathbf{s}(0) \cdot \mathbf{s}(r) \rangle \propto \frac{1}{r^{d-2+\eta}}, \qquad (8)$$

where $G(r)$ is the spin-spin correlation function, $d = 2$ is the dimensionality, and η is a critical exponent, which varies continuously with temperature in the KT phase. This follows from applying the renormalization group to a Gaussian spin-wave action with irrelevant vortex operators, which become relevant at vortex unbinding where $\eta = \frac{1}{4}$.

Turning now to Z_N models, for which the order parameter has a discrete symmetry, there are three potential phases: paramagnetic, critical, and ordered. A pioneering study of the clock model phase diagram considered a ferromagnetic XY model with symmetry-breaking interactions $h_N \sum_{\mathbf{r}} \cos[N\theta(\mathbf{r})]$ [5]. These interactions become relevant at an effective coupling of $K_{eff} = N^2/8\pi$, implying that within the critical phase,

$$\frac{4}{N^2} \leqslant \eta \leqslant \frac{1}{4}. \qquad (9)$$

Subsequent work [14,33–37] has confirmed that for $N \geqslant 5$ this intermediate phase is bounded by two KT transitions, while for $N \geqslant 4$ the critical phase vanishes, leaving a second-order phase transition between order and disorder. Antiferromagnetic clock models are less well studied, but it was argued that odd N-state antiferromagnetic models belong to the same universality class as $2N$-state ferromagnetic models [38].[1] In both ferromagnetic and antiferromagnetic models, as $N \rightarrow \infty$ we recover the XY model, and the symmetry-broken phase disappears at finite temperature.

The extent to which this canonical lore will hold for constrained systems depends on the assumption that the hard constraint becomes an irrelevant operator upon coarse graining. Although simulations have demonstrated that long-wavelength fluctuations do indeed destroy long-ranged order in two-dimensional (2D) systems of hard disks [40], which instead possess an intermediate hexatic phase with quasi-long-ranged orientational order and a solid phase with quasi-long-ranged translational order [41], the Mermin-Wagner theorem is not guaranteed to hold for models with hard interactions [42,43]. One extension to the theorem has come from analysis of random surface models on the two-dimensional torus, which map a set of vertices to real values $\{x\}$ subject to the nearest-neighbor potential $U(x_i - x_j)$ [44]. A lower bound on the variance of fluctuations logarithmic in the length of the torus is proven for a large class of potentials including the hammock potential, defined as $U(x) = 0$ if $|x| \leqslant 1$ and $U(x) = \infty$ otherwise. We have been informed that a followup work currently in preparation extends this method of proof to 2D XY models with rotationally invariant nearest-neighbor interaction, which include our classes of hard-core models of XY spins [45].

One clear departure of the hard-core spin models from the KT picture is in the region where vortices are explicitly forbidden. For inclusion angles below $\pi/2$, Eq. (5) dictates that the winding number around a plaquette, and hence the winding along any closed loop which does not intersect the boundary, is zero. This amounts to setting the strength of the vortex insertion operator, which is a function of the dual field, identically to zero. This is to be contrasted with the role of defects in the typical KT phase, in which they bind into vortex-antivortex pairs up to the transition temperature [46]. Thus, we might instead call the vortex-forbidden region a Patrascioiu-Seiler phase, after the eponymous constraint [47] applied to a 2D system of XY spins with ferromagnetic interaction in Ref. [48] to induce algebraic correlations at all temperatures, including $T = \infty$.

III. METHODS

This section introduces the three main methods of our study: transfer matrix analysis on semi-infinite cylinders, a

---

[1] The purported correspondence between N-state antiferromagnets and $2N$-state ferromagnets again comes from considering symmetry-breaking interactions with couplings h_N, while for pure clock models ($h_N \rightarrow \infty$), subsequent analysis challenged this interpretation [39]. However, we will find below that the height representation for odd-N exclusion models gives rise to the same lower bound of $1/N^2$ on the critical exponent η as that obtained in Ref. [38].
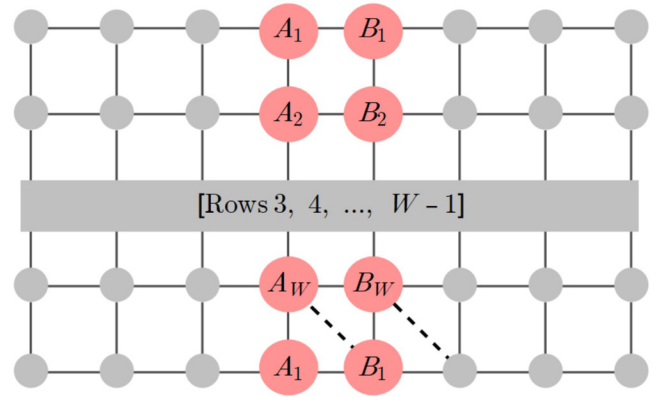

FIG. 2. Nearest-neighbor bonds with periodic boundary conditions on a patch of a torus of width W. The element T_{AB} is the Boltzmann weight due to the energy within and between adjacent columns in spin configurations A and B. The bottom row is a copy of the top row for periodic boundary conditions; for twisted periodic boundary conditions, used for exclusion models when W is odd, the dashed bonds are used instead.

Monte Carlo cluster algorithm on finite lattices of aspect ratio 1, and a height representation within the vortex-forbidden regime. Together, these methods are used to extract η, the critical exponent of the equal-time spin correlation function in the critical phase [Eq. (8)], from which we can assess how these models fit inside our general understanding of 2D spin systems.

A. Transfer matrix method

For tori of length L and width W with the nearest-neighbor interaction defined by Eq. (5), the partition sum can be factored into a product over adjacent columns of W spins. This is encapsulated in the $N^W \times N^W$ transfer matrix T, as shown in Fig. 2. Numbering configurations of W spins by the index $i \in \{1, \ldots, N^W\}$, $T_{ij} = 1$ if there exists a valid state in which i and j are adjacent columns, and $T_{ij} = 0$ otherwise. The matrix T depends on the boundary conditions applied in the transverse direction. We use periodic boundary conditions (PBCs) in all cases except for exclusion models at odd widths. In that case, the periodic boundary conditions are shifted by one lattice constant to make the lattice bipartite, thus circumventing the domain wall of infinite energy cost that would arise in an antiferromagnetic state with regular PBCs. This allows us to extract meaningful data for all numerically feasible widths.

We numerically compute the largest eigenvalues of T. If the eigenvalues are nondegenerate, then in the limit $L \rightarrow \infty$, the free energy per ring is given by [32]

$$f = -\frac{1}{\beta L} \log \mathrm{Tr} T^L = -\frac{1}{\beta L} \log \sum_\Lambda \Lambda^L$$

$$\rightarrow -(1/\beta) \log |\Lambda_1|, \tag{10}$$

where Λ_1 is the eigenvalue of maximum absolute value. For a hard potential, the partition sum is taken over zero-energy configurations only, so from the transfer matrix we can also directly read off the entropy per ring on an infinite cylinder:

$$s = k_B \log |\Lambda_1| = \lim_{L \rightarrow \infty} \log(\Omega/L). \tag{11}$$

The transfer matrix method diagnoses the critical phase via the well-known mapping between the infinite plane and the cylinder in conformal field theories (CFTs) [49]. Recall that discretizing a model on a long, narrow cylinder of finite width and infinite length makes it effectively one dimensional. This destroys (quasi-) long-ranged order, so that, for a system which in two dimensions is critical, Eq. (8) is modulated by an exponentially decaying function whose correlation length is given by the gap between the two largest eigenvalues [50]:

$$G(r) \sim r^{-\eta} e^{r/\xi} \quad \text{where} \quad \xi = (\log|\Lambda_1/\Lambda_2|)^{-1}. \quad (12)$$

The conformal mapping shows that ξ scales linearly with width:

$$\xi(W) = W/A \quad \text{where} \quad A = \pi\eta. \quad (13)$$

Computational constraints limit us to small N and W, for which linear fit to $\xi(W)$ tends to overestimate η. Alternatively, we can calculate a width-dependent exponent $\eta(W)$ by plugging data at different widths into Eq. (13), and then performing a fit to determine η_∞:

$$\eta(W) = \eta_\infty + C/W^2 \quad (14)$$

which tends to underestimate η. Therefore, we mainly use the transfer matrix method as a tool to identify critical models, for which a more accurate estimate of the critical exponents can be determined using the cluster algorithm detailed in the next section. For finite lengths, the correlation function measured by the cluster algorithm can be checked against $G(r)$ as determined by explicit multiplication of the transfer matrix. Most important for our purposes is Eq. (11), by which we identify families of models with the same macroscopic entropy density.

B. Cluster algorithm

To access larger system sizes, and the XY limit, we use a Monte Carlo algorithm that remains ergodic in the critical phase. An algorithm with single-spin updates does not fit these criteria since it cannot unwind defects, and the acceptance rate of single-spin moves becomes prohibitively low as the inclusion angle decreases. To probe configuration space efficiently, a cluster algorithm is required instead. To do so, we need to specify (1) how the cluster is constructed and (2) what operation is applied to this cluster.

A central ingredient in answering question (1) is the concept of pocket Monte Carlo, which constructs a pocket, or bag of elements (in our case, Z_N or XY spins), to which the cluster operation is applied [51]. We remove an item from the pocket by applying the operation to it, which then causes other elements (e.g., the neighbors of the element just transformed) to be added to the pocket. The process terminates when the pocket is empty.

The operation applied to the cluster in Algorithm 1 is a reflection about the axis θ. The move begins by reflecting one spin s_j about this axis and builds the cluster with spins that violate the hard constraint (line 9). This takes inspiration from the pivot cluster algorithm for hard disks [52] and pocket dimer algorithm for nonoverlapping dimers [53], in which the items in the pocket are reflected about a chosen axis or point, and overlapping disks and dimers are added to the pocket.

Algorithm 1. Reflect.

Input : Configuration of spins $\{s_i\}_{i=1}^{N_s}$ at iteration n
Output: Configuration of spins $\{s_i\}_{i=1}^{N_s}$ at iteration $n+1$

1   $\mathcal{P} \leftarrow$ empty bag;
2   $j \leftarrow$ random integer from 0 to $N_s - 1$;
3   add s_i to $\mathcal{P}$;
4   $\theta \leftarrow$ random floating point number from 0 to 2π;
5   **while** $\mathcal{P}$ *is not empty* **do**
6   remove one spin s at random from $\mathcal{P}$;
7      $s \leftarrow$ reflect(s, θ) ;
8   **for** s_k in neighbors of s **do**
9   **if** angle(s_k, s) violates hard constraint **then**
10   add s_k to $\mathcal{P}$;
11   **end**
12   **end**
13 **end**
14 **return** $\{s_i\}_{i=1}^{N_s}$;

However, while these pivot algorithms implement a reflection in real space, the `reflect` algorithm performs the transformation in spin space. In this sense it is more akin to the Wolff algorithm [54], famous for its application to the 2D Ising model but also generalizable to XY spins [55]. Previous works have used this algorithm to simulate XY [12,13] and $O(3)$ [6] inclusion models, but Algorithm 1 can also be applied to exclusion and inclusion clock models, simply by modifying line 4 to sample from a discrete set of reflection axes.

Since the angles between each pair of spins in the cluster are preserved, it is straightforward to prove that the `reflect` algorithm respects detailed balance. Moreover, for the models under consideration the algorithm is ergodic, as proven in Appendix A. A given spin is added to the pocket at most once, so the outer loop is guaranteed to terminate after $\leqslant N_s$ iterations, where N_s is the number of spins. In the worst case, the cluster spans the entire lattice, and the move is a global reflection. This is a problematic aspect of the pivot cluster algorithm when applied to systems of monodisperse disks: the algorithm, which becomes inefficient when clusters are too large, has a percolation threshold below the liquid-solid transition density [51]. Thus, we should be wary of a similar issue occurring for hard-core spins. However, we instead find that the scaling of cluster sizes closely tracks the critical properties of the underlying spin system. Namely, as with the Wolff algorithm applied to Potts models, the average cluster size scales as $\langle s \rangle \propto L^{2-\eta'}$, where $\eta' \approx \eta$, with a relative error $\lesssim 0.01$ for inclusion angles near $\pi/2$. The potential exactness of this relation is explored further in Appendix B.

The critical exponent η is determined from the Fourier spectrum of the spin correlation function, denoted $G(\mathbf{k})$. This spectrum has a peak at $\mathbf{k} = (0, 0)$ for inclusion models, corresponding to the uniform susceptibility χ_u, and $\mathbf{k} = (\pi, \pi)$ for exclusion models, corresponding to the staggered susceptibility χ_s. Given a model whose real-space correlation function decays as a power law [Eq. (8)], the susceptibility on an $L \times L$ lattice scales as $L^{2-\eta}$, times subleading logarithmic corrections [56]. Performing a linear fit to $\log \chi$ vs $\log L$

provides an estimate of η. For small system sizes, the best fit for η is sensitive to the minimum system size L_{min} included in the fit, but we find that the estimated value is roughly stable for $L_{min} \geqslant 24$, and therefore start all fits at $L = 24$. By contrast, in the paramagnetic phase where $G(r)$ decays exponentially to zero, $\chi(L)$ approaches a plateau at a scale set by the correlation length.

We employ a variety of checks to verify that the algorithm converges to the stationary distribution of the system. First, to check that the system decorrelates from its initial state, we run the algorithm from different initial configurations and measure the standard deviation between runs. To generate disordered initial states, we soften the potential into a continuous, differentiable function of the nearest-neighbor angle. The system is initialized in a random configuration, which is likely to have nonzero energy due to violated constraints. The initialization algorithm then alternates between (1) gradient descent to the nearest minimum of the soft potential and (2) a full sweep of zero-temperature single-spin moves, which are accepted only if they reduce the energy. This process is repeated until a zero-energy state is found, which then becomes the initial state for the cluster algorithm or, if no zero-energy state is found after a maximum number of attempts, the ordered state is used instead. For Z_N models, a differentiable soft potential cannot be defined, so we instead initialize in one of the ideal states introduced below. We also use the ideal-state initialization for large systems of XY spins at low inclusion angles ($\Delta < \pi/2$) for which the minimization algorithm generally fails to find a zero-energy state.

A second check is that the spatial correlations determined for Z_N models on tori of small widths should match those determined from the transfer matrix. This comparison was performed for three different models on tori of width 4 and aspect ratios 1, 2, 4, and 8, with a maximum relative error of 2.54×10^{-3} for $G(r) > 0.005$.

To control for the effects of critical slowing down, a single Monte Carlo step was defined as follows. First, we defined one MCS as a single cluster move and calculated the autocorrelation function of the susceptibility. The integrated correlation time was conservatively estimated to scale with system size L no faster than $\sim L^{1/2}$. For comparison, this is not as efficient as the Wolff algorithm for the Ising model, which scales logarithmically or with a very small dynamical exponent near the critical point [57,58], but it is a vast improvement on the $O(L^2)$ critical slowing down observed for local algorithms [59]. One Monte Carlo step (MCS) was then redefined to consist of $CL^{1/2}$ cluster moves, with C chosen such that on the smallest system size considered (8×8), 1 MCS = 8 cluster moves. Through this redefinition, we could then generate at least as many independent samples with increasing system size by running the algorithm for the same number of steps.

Each run consisted of 10^4 MCS of equilibration followed by 10^5 MCS of recording. The correlation time was (much) less than 100 MCS for all models and system size considered, so all error bars were estimated by averaging the data over chunks of 100 MCS and computing the standard error over 1000 chunks. A data bunching analysis verified that these chunks were uncorrelated [59].

C. Height representation for defect-free models

To complement this algorithmic approach, we formulate a height representation for Z_N models with $N > 2(p' - 1)$. This condition forbids defects in the height field, which in the case of inclusion models can be interpreted simply as forbidding vortices and antivortices. In the height representation, also referred to as an interface model, the spin variables are mapped to coarse-grained height variables, an approach that has been used to study a range of critical ground-state ensembles [60,61] including the three-state Potts AFM on the square lattice [26] and the four-state Potts AFM on the triangular [62] and kagome lattices [63].

For our class of inclusion and exclusion models, we map from clock variables $\{\sigma\} = \{0, 1, \dots, N - 1\}$ to heights $\{h\}$ as follows:

(i) At the origin, define

$$h(0) = \sigma(0). \tag{15}$$

(ii) The local change in the height field from site $x = (x_1, x_2)$ to nearest neighbor site y is

$$h(x) - h(y) = \begin{cases} [\sigma(x) - \sigma(y)], & \text{inclusion} \\ [\sigma(x) - \sigma(y) - N/2], & \text{exclusion} \end{cases} \mod N \tag{16}$$

with the modulus chosen such that, in both cases,

$$|h(x) - h(y)| \leqslant (p' - 1)/2. \tag{17}$$

(iii) For $N > 2(p' - 1)$, $\Delta h = 0$ around a plaquette; this is what is meant by the absence of defects. Only if this condition is satisfied can h be uniquely defined:

$$h(x) = \begin{cases} \sigma(x), & \text{inclusion} \\ \sigma(x) - \alpha(x)N/2, & \text{exclusion} \end{cases} (\bmod N) \tag{18}$$

where $\alpha(x) \equiv (x_1 + x_2) \bmod 2$ is 0 on the even sublattice and 1 on the odd sublattice. By the notation $(\bmod N)$ we mean that if the modulo operation is taken on both sides, then the equality holds. For inclusion models, the distinction between h and σ has a simple interpretation: while the spin state σ is only defined modulo N, h has been "lifted" to $\mathbb{Z}$ [64].

We note in passing that the above construction for inclusion models has also appeared in the mathematical literature [65,66], where height fields with $p' = 2m - 1$ are known as m-Lipschitz functions. These works primarily focus on high dimensions, where long-range order can be proven in the vortex-forbidden regime.

Returning to $d = 2$ and a more physically motivated perspective, given a uniquely defined height field h, we posit that it is governed by the effective Hamiltonian

$$F = \int d^2x \left(\frac{K}{2} (\nabla h)^2 + V_{lock}(h) \right), \tag{19}$$

where K is the stiffness and V_{lock} is the so-called locking potential. The latter favors heights on the ideal state lattice $\mathcal{I}$, the set of periodic, macroscopically flat height configurations with maximal entropy (ideal states) [61,67].

In writing the effective Hamiltonian, a typical ground state is assumed to be a patchwork of ideal-state domains. We refer the reader to Ref. [26] for a detailed review of the terminology and briefly recapitulate the main definitions here. Two instances of the same domain X have equal heights modulo

$\tilde{h}$, where $\tilde{h}$ is an element of the repeat lattice $\mathcal{R}$. The repeat lattice is a subgroup of the equivalence lattice $\mathcal{E}$, the set of elements a such that $a + \mathcal{I} = \mathcal{I}$. Thus, V_{lock} has the same periodicity as the equivalence lattice, which in our case has a one-dimensional representation with period $2\pi/g_0$.

If the locking potential is relevant ($\eta_{lock} < 2d = 4$), the model is said to be in the smooth phase, with long-ranged order in the height field. On the other hand, if $\eta_{lock} > 2d$, the locking potential is irrelevant and the interface is rough, with logarithmic correlations in the height field:

$$G_h(x, y) \equiv \langle [h(x) - h(y)]^2 \rangle \sim \frac{1}{\pi K} \log |x - y|. \qquad (20)$$

The critical exponents associated with relevant vertex operators in the original spin model can then be determined by expressing them as periodic functions of the height field and applying Eq. (20). Due to the Gaussian form of the action, an operator O with period $2\pi/g$ scales as

$$\langle O^*(x)O(y) \rangle = \langle \exp\{ig[h(x) - h(y)]\} \rangle$$
$$= \exp[-G_h(x, y)g^2/2] \qquad (21)$$

implying a critical exponent

$$\eta_O = \frac{g^2}{2\pi K}. \qquad (22)$$

The roughening transition therefore occurs when $g_0^2/2\pi K = 4$.

To determine the critical exponent of the spin correlation function, denoted simply as η, we express the magnetization in terms of the height field as

$$e^{\pm 2\pi i h(x)/N} = M_1(x) \pm iM_2(x), \qquad (23)$$

where $\mathbf{M} = (M_1, M_2)$ is the unstaggered magnetization for inclusion models, and the staggered magnetization for exclusion models. Equation (22) then implies

$$\eta = \frac{2\pi}{N^2 K}. \qquad (24)$$

For inclusion, ideal states consist of spins randomly sampled from a set of p' consecutive spin states. Equivalently, for even-N exclusion models, we choose a set S of $(p'+1)/2$ consecutive integers as the spin states on the even sublattice and populate the odd sublattice with the integers $(S + N/2)$ mod N (i.e., rotate the spins on the even sublattice by π), so per Eq. (18), the average height is constant and equal on the two sublattices. In both cases, there are N ideal states, in one-to-one correspondence with the average height mod N. The magnetization has the periodicity of the repeat lattice $\mathcal{R} = N\mathbb{Z}$, whereas the locking potential has the periodicity of the equivalence lattice $\mathcal{E} = \mathbb{Z}$. Thus, for the system to occupy the rough phase, this implies the relation

$$4 < \eta_{lock} = \frac{2\pi}{K} = N^2 \eta \quad \text{(odd } p'\text{)}. \qquad (25)$$

The lower boundary of the critical phase $\eta = 4/N^2$ is consistent with inequality (9) derived for ferromagnetic XY models with clock perturbations.

For exclusion models with even p' (odd N), the absence of $\mathbb{Z}_2$ symmetry leads to an asymmetry in the ideal states: choose
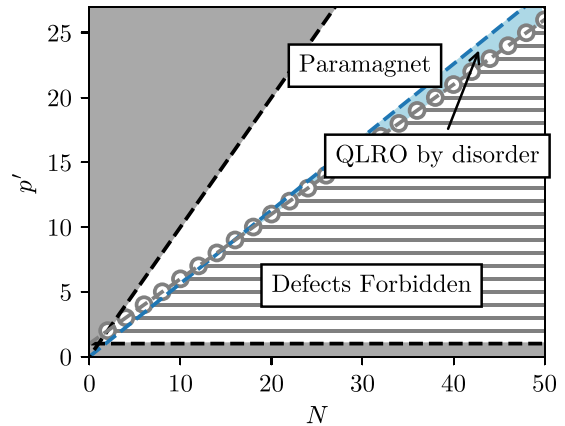

FIG. 3. Phase diagram in the (N, p') plane for Z_N inclusion and exclusion models. p' families are represented by gray horizontal lines, with open circles along the gray dashed line denoting $N = N_c(p')$. Dark gray shaded regions are forbidden, and their black dashed boundaries are the trivial limiting cases: $p' = 1$ is strictly ordered with zero macroscopic entropy density, while $p' = N$ is disordered (unconstrained and noninteracting). Constant inclusion angle corresponds to functions of the form $p' = N\Delta/\pi$; the blue dashed line has $\Delta = \Delta_c \approx 0.56\pi$ and marks the upper boundary ($\eta = \frac{1}{4}$) of the quasi-long-ranged order by disorder phase [Eq. (37)].

a set $S = \{\sigma_i\}$ of $p'/2$ consecutive clock orientations as the allowed states on one sublattice, and a set T of $p'/2 + 1$ integers on the other sublattice: $T = [S + (N-1)/2] \cup \{\sigma_{p'/2} + (N+1)/2\}$. Owing to this asymmetry, there are now $2N$ ideal states, uniquely identified by their average heights mod N. Referring back to Eq. (18), note that $h(x)$ takes integer values on the even sublattice, and half-integer values on the odd sublattice. The ideal-state lattice and equivalence lattices are $\mathbb{Z}/2$, whereas the repeat lattice (which again has the same periodicity as $\mathbf{M}$) is $N\mathbb{Z}$. This implies that within the rough phase

$$4 < \eta_{lock} = \frac{8\pi}{K} = 4N^2\eta \quad \text{(even } p'\text{)}. \qquad (26)$$

Therefore, in contrast with Eq. (25), the lower boundary of the critical phase is characterized by $\eta = 1/N^2$. This agrees with the bound on the critical phase derived from a renormalization group (RG) treatment of odd-N antiferromagnetic clock models with standard action [38]. Interestingly, the RG arguments applied exclusively to $T \neq 0$, whereas the height representation allows us to derive the same bound at $T = 0$, for which these clock models reduce to $p' = 2$ exclusion models.

To summarize, the existence of a height representation indicates that all inclusion and exclusion models with $N > 2(p' - 1)$ are either critical or ordered. The critical exponents of models in the critical phase must satisfy either Eq. (25) or (26). This leaves open two questions: whether any models fall within the smooth phase (long-ranged order), and whether there are critical models which do not admit a height representation.

IV. PHASE DIAGRAM

With the height representation as a guide, we now describe in detail the phase diagram of Z_N models (Fig. 3), using

the `reflect` algorithm to determine the critical exponents. Models described by a height representation are classified into p' families, all in the rough phase except for $p' = 1$. For small p' the transition to the paramagnetic phase is geometric, coinciding with the point at which defects are allowed, but for large p' and in the XY limit, we identify a region of the phase diagram in which vortices are energetically allowed but entropically disfavored, leading to quasi-long-ranged order by disorder. The transition from the critical phase to the paramagnetic phase is a Kosterlitz-Thouless–type transition, driven by the unbinding of vortices at $\eta = \frac{1}{4}$.

A. Families of Z_N models

We begin by assessing trends in the transfer matrix, from which the following definition arises:

Definition 1. (p' family) A class of models with a common p' and a critical value $N_c(p')$ such that all models with $N > N_c(p')$ have the same maximum transfer matrix eigenvalue for each width.

Each family has an assortment of scaling properties. On a torus of finite length, the number of configurations can be factored as

$$\Omega(N, W, L) = N \exp[S(W, L)]. \qquad (27)$$

That is, for a fixed width W and length L, Ω scales proportionally to N. Due to the use of longitudinal periodic boundary conditions implicit in Eq. (10), this equation only strictly holds for small lengths since longer strips can support nontrivial windings.

While Definition 1 classifies p' families in terms of their eigenvalues in the spin representation, members of the same family are related more fundamentally through their height representation. On the square lattice,

$$N_c(p') = 2(p' - 1). \qquad (28)$$

Therefore, the condition $N > N_c(p')$ within a p' family, which for inclusion models amounts to forbidding vortices, coincides with the criterion which allows the height field to be uniquely defined. Defining Δ to saturate the right-hand side of inequality (7), Eq. (28) corresponds to an inclusion angle of $\pi/2$ (gray dashed line in Fig. 3). Below this angle, members of the same p' family admit a height representation obeying the same local height rule [inequality (17)], and thus the same stiffness.

The value of the stiffness determines whether the family occupies the smooth or rough phase. In the latter phase, Eq. (24) implies that ηN^2 is invariant. Since $\xi(W) \propto W/\eta$ at criticality on semi-infinite cylinders of width W, the correlation length factorizes as

$$\xi(W, N, p') \approx N^2 \tilde{\xi}(W, p'). \qquad (29)$$

It must be emphasized that this scaling form is only approximate for finite widths. The universal quantity measured in different members of the same family is the correlation function of the height field $G_h(x - y)$ defined in Eq. (20). This can only be related to a correlation in the spin language if the Gaussian action of the form (19) is assumed. In that case, positing exponential decay on the left-hand side of Eq. (21) in the quasi-1D limit immediately yields Eq. (29). But this is
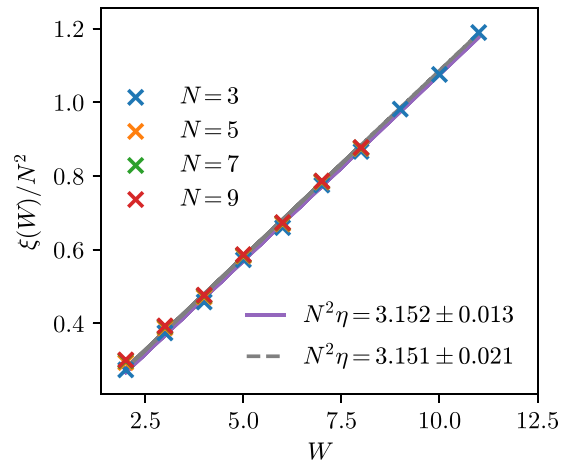

FIG. 4. $\xi(W)/N^2$ for the $p' = 2$, as determined from the transfer matrix for small values of N. Different-colored x's correspond to different members of the family, while the solid (dashed) lines show linear fits for even (odd) widths. Very small widths ($W < 4$) were not included in the fits.

a coarse-grained, long-wavelength description of the system, and should not be expected to hold exactly for small widths. Indeed, as seen in Fig. 4 for the $p' = 2$ family, ξ/N^2 at a given width is a monotonically increasing function of N, although this scaling correction becomes less severe as W increases.

The linear trend with width exhibited in Fig. 4 is indicative of the critical phase. But, just as the data collapse between members of the same family improves with W, so too does the critical exponent determined from a linear fit to the combined data become more accurate as we access larger widths. The largest feasible width is limited by computational constraints, which thus limits the accuracy of our estimate of η. For the $p' = 2$ family, a linear fit of the form $\xi(W)/N^2 = BW + C$ to data points with $W \geqslant 4$ for even (odd) widths yields $\eta N^2 = 3.152 \pm 0.013$ ($\eta N^2 = 3.152 \pm 0.021$). Alternatively, a fit of the form (14) yields $\eta N^2 = 2.994 \pm 0.3$ ($\eta N^2 = 2.996 \pm 0.45$). These fits are broadly consistent with the known exponent of $\eta = \frac{1}{3}$ for the staggered susceptibility of the zero-temperature three-state Potts AFM [61], which belongs to the family.

This not only provides a useful crosscheck of the transfer matrix method, but also points to the power of the height representation: we can immediately bootstrap exact knowledge of the stiffness or approximate determination of η for just one member of the family to a theoretical prediction of η for *all* members of the family. In this sense, the $p' = 2$ family (or, put differently, the ground-state ensembles of odd-N clock models with standard antiferromagnetic cosine action) can be said to descend from the three-state Potts AFM at zero temperature, whose critical exponent places the family safely in the rough phase [Eq. (26), which crucially must be distinguished from Eq. (25) for odd p'].

We will demonstrate below that all $p' > 1$ families occupy the rough phase, and thus the transfer matrix also enables a crude estimate of the central charge c of the associated conformal field theory. For $p' = 2$ the central charge can be deduced from the three-state Potts AFM at zero tempera-
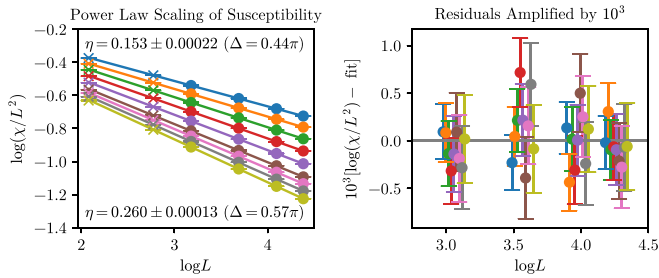
FIG. 5. Power-law scaling of the susceptibility. The left panel shows $\log(\chi/L^2)$ as a function of $\log L$ for $L = 8, 16, 24, 40, 60, 80$. Data from $L = 8$ and 16, marked with x's, were not included in the linear fit. From top to bottom, the inclusion angles are 0.44π, 0.46π, 0.48π, 0.5π, 0.52π, 0.54π, 0.55π, 0.56π, and 0.57π. The right panel shows residuals with respect to the fits, amplified by a factor of 10^3 and with data from different inclusion angles scattered in the horizontal direction to enhance visibility. Error bars are defined as the standard error across 1000 consecutive bunches of 100 MCS each.

ture, for which $c = 1$ is known exactly [17]. In fact, through finite-size scaling of the entropy density, we find $c \approx 1$ for all families examined. This matches our interpretation of p' families as interface models governed by the effective action in Eq. (19), which if the locking potential is irrelevant is a free-boson CFT with $c = 1$ [68]. The same central charge was determined numerically for clock models with $5 \leqslant N \leqslant 8$ obeying the standard action [37].

B. Kosterlitz-Thouless transition

To determine η with greater accuracy, the finite-size scaling of the susceptibility is measured using the `reflect` algorithm. We begin by considering inclusion models in the XY limit, as defined by Eq. (5), and provide evidence of a Kosterlitz-Thouless transition at critical angle Δ_c. Some results in this section overlap with and confirm the findings of Refs. [12,13] but are included for completeness, to contextualize our findings on a broader class of models.

As shown in Fig. 5, the power-law scaling $\chi \sim L^{2-\eta}$ remains a good fit to the susceptibility, with no obvious trend in the residuals, up to $\Delta \approx 0.57\pi$. For larger inclusion angles, the fit becomes increasingly poor; rather than scaling linearly with $\log L$, $\log \chi$ becomes a concave down function of L indicative of the paramagnetic phase. The persistence of quasi-long-ranged order past the angle at which vortices become allowed ($\Delta = 0.5\pi$) lends support to the interpretation that the system becomes a paramagnet when defects become relevant and vortices unbind, the same mechanism which drives the KT transition in the XY model at finite temperature.

To verify this interpretation and pinpoint the critical angle Δ_c at which this transition occurs, we consider three key observables near the apparent transition: the critical exponent η, the second moment correlation length $\xi_{2\mathrm{nd}}$, and the Binder cumulant U. These are discussed in turn.

1. Critical exponents η and η′

Standard KT theory indicates that the most relevant defect operator has $\eta_{vortex} = 1/\eta$, so vortex unbinding occurs at $\eta = \frac{1}{4}$. Therefore, if the transition in our class of inclusion models
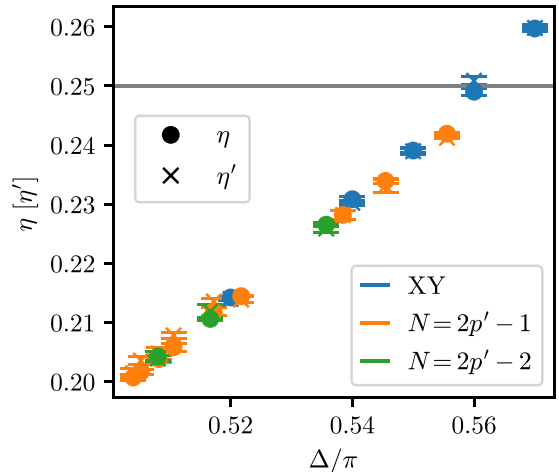

FIG. 6. Critical exponents approaching the KT transition. The gray line in the left panel marks the value of $\eta = \frac{1}{4}$ at vortex unbinding. The critical exponents η, determined from finite-size scaling of the susceptibility, and η', determined from finite-size scaling of the average cluster size, are marked with circles and x's respectively. The plotted error bars show the statistical uncertainty of the log-log fits in Fig. 5. Different colors denote models in the XY limit (blue), the smallest members of p' families, with p' ranging from 5 to 63 (orange), and clock models just outside the p' families, i.e., $N = N_c$, with p' ranging from 15 to 63 (green).

is in the KT universality class, we expect

$$\eta(\Delta = \Delta_c) = 1/4. \tag{30}$$

The threshold value of $\frac{1}{4}$ is indicated by the gray line in Fig. 6, implying a critical angle just above $\Delta = 0.56\pi$. This is roughly consistent with the above finding that the power-law scaling of the susceptibility holds up to $\Delta \approx 0.57\pi$.

Figure 6 contains two other noteworthy features. First, the exponent η' determined from finite-size scaling of the average cluster size, closely tracks η. This is a positive indication of the algorithm's efficiency and is addressed in Appendix B. Second, the critical exponents associated with various clock models, measured by the same method as in Fig. 5, are plotted on the axes η vs Δ/π by defining Δ at the midpoint of the interval in inequality (7), $\Delta = \pi p'/N$. Two sets of models are shown. Marked in orange are models for which $N = 2p' - 1$, which are the smallest members of their respective p' families and are therefore known rigorously to be critical based on the existence of a height representation. Marked in green are those with $N = N_c(p') = 2p' - 2$, which lie just outside the p' families. Both sets of models approach an inclusion angle of 0.5π from above in the limit $N \to \infty$. Strikingly, even for finite N, the two sets of models collapse onto roughly the same trend when plotted vs Δ, and this trend is also consistent with that of the models in the XY limit. The existence of critical models outside the p' families again points to the phenomenon of quasi-long-ranged order (QRLO) by disorder, as the $N = N_c(p')$ models exhibit power-law scaling of the susceptibility (and cluster sizes) despite vortices being allowed. This "QLRO by disorder" regime only exists for sufficiently large p', a statement that will be made quantitative shortly.
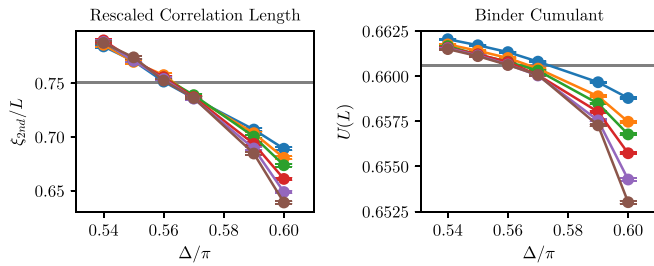
FIG. 7. Data collapse of dimensionless quantities, the rescaled second moment correlation length, and Binder cumulant, near the KT phase transition. System sizes from top (blue) to bottom (brown) are $L = 8, 16, 24, 40, 60,$ and 80. The gray line in the left panel marks the value of $\lim_{L \to \infty} \xi_{2nd}/L = 0.750\,691\,2$ at vortex unbinding [56]; gray line in the right panel marks the value of the Binder cumulant $U = 0.660\,603$ at the KT transition [70].

2. Second-moment correlation length

We now take Eq. (30) as our *definition* of the critical angle, i.e., $\Delta = \Delta_c$ is the angle at which $\eta = \frac{1}{4}$. If the transition is of a Kosterlitz-Thouless type, then this definition of Δ_c should be consistent with that estimated via other metrics. One of these is the second-moment correlation length $\xi_{2nd}(L)$, defined as [56]

$$\xi_{2nd}(L) = \frac{1}{2 \sin(\pi/L)} \sqrt{G(0,0)/G(2\pi/L, 0) - 1}. \quad (31)$$

This quantity takes its name from the fact that it is the second moment with respect to the Fourier-transformed correlation function $G(\mathbf{k})$. In the paramagnet, $\xi_{2nd}(L)$ is independent of system size, whereas in the critical phase, it scales linearly with system size L. Thus, the rescaled correlation length $\xi_{2nd}(L)/L$ is independent of L up to the critical angle, up to subleading corrections at small system sizes. In the thermodynamic limit at the KT transition, this rescaled length takes the value [56]

$$\lim_{L \to \infty} \frac{1}{L} \xi_{2nd}(L, T = T_{KT}) = 0.750\,691\,2. \quad (32)$$

For the XY model, $\xi_{2nd}(L)/L$ was found to quickly approach this limit from below [69].

The hypothesis that our class of models exhibits a KT transition thus yields two predictions: (1) ξ_{2nd}/L should be independent of system size up to Δ_c, and (2) it should take the value given by Eq. (32) at $\Delta = \Delta_c$. (At the system sizes reported in this paper we do not expect this limit to be obtained with great accuracy, but it should fall in the rough neighborhood.) These predictions both approximately hold, as shown in the left panel of Fig. 7. Deviations from $\xi_{2nd}/L \approx$ constant appear above $\Delta \approx 0.57\pi$, and $\xi_{2nd}/L = 0.750\,691\,2$ between 0.56π and 0.57π, consistent with the estimate of Δ_c obtained from Eq. (30).

3. Binder cumulant

The Binder cumulant $U(L)$, defined as

$$U(L) = 1 - \langle |\mathbf{M}|^4 \rangle / 3 \langle |\mathbf{M}|^2 \rangle^2, \quad (33)$$

is, like the rescaled correlation length, asymptotically independent of L within the critical phase. Therefore, for a system

which undergoes a KT transition we should observe data collapse between different system sizes up to $\Delta = \Delta_c$, above which $U(L)$ will decay with L. At the KT transition point, U takes the value [70]

$$\lim_{L \to \infty} U(L, T = T_{KT}) = 0.660\,603. \quad (34)$$

In the XY model, the approach to this limit is sufficiently rapid [$U(L = 32)$ is within 0.02% of the limiting value] that we expect Eq. (34) to be well approximated for the finite system sizes simulated in this study. Our general expectations are confirmed, as shown in the right panel of Fig. 7. Again, $U(L)$ is roughly independent of L within uncertainties up to $\Delta \approx 0.57\pi$, crossing the value given in Eq. (34) at $\Delta \approx 0.56\pi$. Interestingly, as was pointed out in an earlier study of the constraint-only model [13], the $L \to \infty$ limit is approached from above, consistent with the $1/\log L$ scaling corrections of a Gaussian model, but of opposite sign to the scaling corrections measured for the members of the universality class studied in Ref. [70], including the nearest-neighbor XY and Villain models. In that work, a large $1/(\log L)^2$ correction was attributed to the presence of vortices neglected in the spin-wave theory. While it is not clear where our highly nonlinear model falls on the RG flow to the Gaussian theory, we suspect that the strong suppression of defects which persists even above $\pi/2$ could explain the absence of this vortex-driven correction.

In summary, the various measures discussed in this section, i.e., qualitative estimate of the goodness of fit to the susceptibility, the value of η determined from this fit, the system size independence of ξ_{2nd}/L and $U(L)$, and their agreement with Eqs. (32) and (34) in the vicinity of the transition, collectively lend support to the hypothesis that, in the XY limit, our class of inclusion models undergoes a Kosterlitz-Thouless transition between $\Delta = 0.56\pi$ and 0.57π.

C. Scaling of η and QLRO by disorder

Recall that the Kosterlitz-Thouless phase is characterized by a line of critical points, with continuously varying exponents. While in the standard case η is a function of temperature, in the hard-core spin models η is a function of Δ instead. In this section we find that $\eta(\Delta)$ takes a simple functional form, and relate it to the subtler features of the phase diagram in Fig. 3.

In Fig. 8(a) we perform a power-law fit to $\eta(\Delta)$ for inclusion angles ranging from $\Delta = 0.05\pi$ to 0.56π. Strikingly, a single trend

$$\eta(\Delta) \sim \Delta^2 \quad (35)$$

holds for the entire range of Δ. In particular, the functional form of $\eta(\Delta)$ does not appear to change as the system crosses over from the defect-forbidden region ($\Delta < 0.5\pi$) to the "QLRO by disorder" region marked in Fig. 1(b). This is a manifestation of universality in that the critical exponent is not sensitive to whether vortices are strictly forbidden by the hard constraint. This perhaps makes the simplicity of Eq. (35) all the more surprising, as it relates the scaling exponent of a macroscopic observable and the microscopic parameter Δ in such a straightforward way. Assuming that Eq. (35) holds yields a more precise estimate of the critical angle Δ_c, by
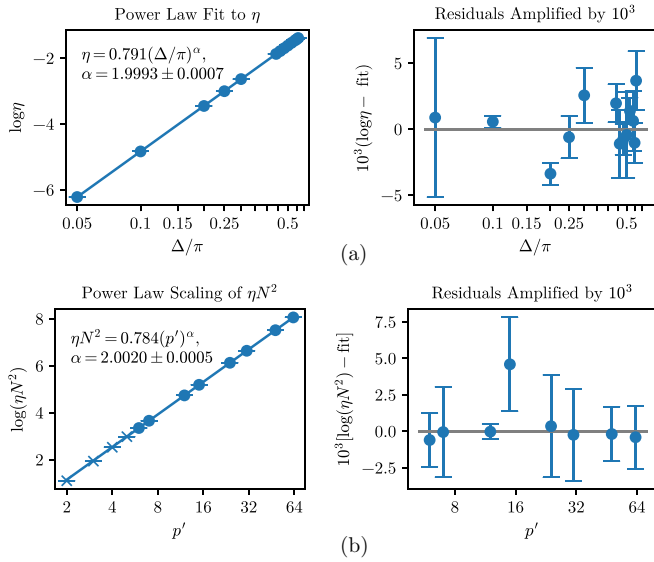
FIG. 8. Scaling of the critical exponent η in the critical phase. (a) Power-law scaling as a function of Δ in the XY limit. Error bars are derived from the statistical uncertainties in the power-law fits to the susceptibility. The right panel shows residuals with respect to the fit $\eta = 0.79\Delta^{1.999}$, amplified by a factor of 10^3. The Δ/π axis uses a logarithmic scale. (b) Scaling of ηN^2 as a function of p' for p' families of Z_N models. Error bars are statistical uncertainties from power-law fits to the susceptibility of one model within the p' family (except $p' = 2$ where the theoretical value of 3 is used). The p' axis uses a logarithmic scale. The right panel shows the residuals with respect to the fit $\eta N^2 = 0.784p'^2$, starting at $p' = 6$.

solving for the angle at which $\eta = \frac{1}{4}$. A single parameter fit fixes the constant of proportionality and implies

$$\Delta_c = (0.561\,92 \pm 0.000\,14)\pi. \tag{36}$$

It must be emphasized that the error bars here only capture the statistical uncertainty of the fit to Eq. (35) and not any of the uncertainties associated with the limited number of independent samples, logarithmic corrections to scaling for small system sizes, and so on. In particular, taking logarithmic corrections into account in the fit to susceptibility for system sizes sizes up to $L = 4096$ yielded a value for Δ_c that is 0.5% greater than our estimate [13]. Nevertheless, it is remarkable that we could come this close to the precision determination of Δ_c from such a simple ansatz for $\eta(\Delta)$.

Returning to clock models, in Fig. 8(b) we fit ηN^2, the invariant quantity within p' families, to a power law as a function of p' and find

$$\eta N^2(p') = 0.784(p')^\alpha, \quad \alpha = 2.0020 \pm 0.0005 \tag{37}$$

for a fit beginning at $p' = 6$. This is equivalent to Eq. (35) once we make the identification $\Delta = \pi p'/N$. This equivalence was partly anticipated by Fig. 6, in which clock models, both inside and outside their respective p' families, followed a similar trend in η as models in the XY limit near an inclusion angle of 0.5π. But given the presence of the symmetry-breaking h_N interaction which forces the spins to be discretized, it is not obvious beforehand that models of Z_N and XY spins will

exhibit the same scaling once we have mapped the parameters $(N, p') \leftrightarrow \Delta$. Thus, a few explanatory remarks are in order.

First, the monotonic increase of ηN^2 as a function of p' implies the absence of a symmetry-broken (smooth) phase in the phase diagram in Fig. 3. (We exclude $p' = 1$, which is trivially ferromagnetic or antiferromagnetic and maps to an interface of constant height.) If there existed a region of (N, p') parameter space that possessed long-ranged order, it would overlap with the defect-forbidden region, and thus we could identify some p' family with $\eta_{lock} < 4$. For odd p', this would imply [Eq. (25)] $\eta N^2 < 4$, while for even p', this would imply [Eq. (26)] $\eta N^2 < 1$. Since small p' families are in the rough phase, as determined both from finite-size scaling of the transfer matrix and from the `reflect` algorithm, the monotonicity of $\eta N^2(p')$ implies that *all* p' families are in the rough phase. This is a peculiar result, as nothing in our class of models forbids a smooth phase *a priori*. But, the fact that the locking potential becomes strongly irrelevant as p' increases is also internally consistent with the observation that, for p' of order 10 and above, there is no real distinction in terms of critical exponents between Z_N models and XY models of commensurate inclusion angle. This is reminiscent of a known feature of the phase diagram of Z_N models with standard action: while the Z_5 model exhibits a lower transition temperature between the QLRO and paramagnetic phase as well as stronger finite-size effects in the partition function zeros and helicity modulus, these quantities are roughly independent of N for $N \geqslant 6$, collapsing onto the XY limit [36,37].

Having ruled out a smooth phase, there remain two other potential phases for clock models: the critical phase and the paramagnetic phase. As indicated by the unshaded region in Fig. 3 for a fixed $N \geqslant 4$ there exists a sufficiently large $p' < N$ such that the corresponding (N, p') model is nontrivially paramagnetic. This phase includes, for example, the q-state Potts antiferromagnets, which are disordered for $q \geqslant 4$ on the square lattice [26]. This leaves open the question of whether all models outside the p' families are disordered, or whether there exists a regime, as in the XY limit, in which defects are allowed but irrelevant.

For small p', there is no such regime, a feature that can be understood via the height representation. A defect in the height field of winding number 1 has $\Delta h = N$, which, when substituted into the effective free energy in Eq. (19), implies that the most relevant vortex operator has critical exponent [64]

$$\eta_{vortex} = \frac{N^2 K}{2\pi} = \frac{1}{\eta}. \tag{38}$$

Within a p' family, if $\eta > \frac{1}{4}$ for some N, this implies that the vortex operator is relevant, and quasi-long-ranged order is preserved only because the strength of the operator is identically zero. In that case, the transition from QLRO to the paramagnet as a function of N for fixed p' is not Kosterlitz-Thouless in nature but geometric, occurring exactly at $N = N_c(p')$ when defects become allowed. This is the case, for example, for $p' = 2$, as the smallest member of the family, the three-state Potts AFM, has $\eta = \frac{1}{3}$. Within the range of p' exhibiting this geometric transition, the corresponding value of ηN^2 also deviates slightly from the trend in Eq. (37). Including the data

points with $p' < 6$ leads to a greater reduced χ^2 value and a noticeable trend in the residuals indicative of deviations from a power law, although it does not significantly change the predicted exponent. These deviations from the trend possibly reflect the greater influence of the locking potential at low p', as well as the fundamental contrast between the geometric nature of the transition for low p' and the KT transition for XY spins.

As p' increases, however, the effective inclusion angle of the smallest member of the p' family decreases, and the associated critical exponent decreases. Evaluating Eq. (37) at $N = N_c(p') + 1$ indicates that for $p' \geqslant 5$, all members of the corresponding p' family have $\eta < \frac{1}{4}$. This allows for critical models outside the family. The minimum value of p' for which this QLRO by disorder phase exists can be estimated by positing that the critical exponent in the vortex-allowed critical region obeys the same scaling form as in the vortex-forbidden region, and evaluating Eq. (37) at $N = N_c(p')$:

$$\eta(N_c, p') = 0.784(p')^2/(2p' - 2)^2 \qquad (39)$$

which implies a crossover at $p' \approx 8.7$. Since p' only takes integer values, this indicates that the $(N, p') = (16, 9)$ model is the smallest p' critical model not described by a height representation.

V. CONCLUSIONS

It is a marvel of statistical mechanics that systems with drastically different microscopic physics can exhibit the same macroscopic behavior near critical points. The class of "constraint-only" spin models examined in this paper serve as an extreme example of this phenomenon. The divergent nature of the potential means that temperature is not well defined, and the ordering and phase transitions are driven solely by entropy. It is not immediately obvious that these phenomena, examined here for models of Z_N and XY spins on the square lattice, will fall into the same universality classes as their finite-temperature cousins and, yet, when the dust settles, they do. That said, there are many interesting details, most notably the existence of a vortex-free and yet nonlinear "spin-wave" regime and that the long-range ordered phases that *can* exist for Z_N models do not arise in our models. We have investigated our models using height representations, transfer matrix calculations, and an ergodic cluster algorithm, arriving at the phase diagram shown in Fig. 3.

Our focus in this paper has been on the equilibrium properties of this class of spin models. Yet, the simplified nature of the interaction, constructed by analogy to that of hard spheres and disks, also imbues the spin models with interesting dynamics. This dynamics will be addressed in a following paper.

Note added. Recently, we came across Ref. [71] which has applied "replica-symmetry-breaking" techniques to related models on random graphs, and it will also be instructive to make contact with this line of work.

APPENDIX A: ERGODICITY OF THE `reflect` ALGORITHM

In this Appendix we prove that Algorithm 1 is ergodic for inclusion models on any lattice. Furthermore, ergodicity of exclusion models on bipartite lattices is proven. While this has already been proven elsewhere [6], the alternate proof provided here will also serve to introduce the notation employed in the following Appendix, for a closer study of the algorithm's properties.

Since each move is reversible, in proving ergodicity it suffices to prove that a chosen target configuration can be attained from any initial state via a finite sequence of cluster moves [61,62].

Theorem 1. Let T be a ferromagnetic target configuration on a lattice of N_s spins, with $\theta = 0$ for all spins. Given any initial configuration A which respects the inclusion constraint, A can be transformed to T by a series of $\leqslant N_s$ iterations of the `reflect` algorithm.

Proof. Let S denote the (possibly empty) domain of spins on which A agrees with T. To transform A into T, we successively add spins to this domain, with each move adding $\geqslant 1$ spin to the domain and leaving the spins in S unchanged.

A cluster move starts by adding a random spin s to the pocket $\mathcal{P}$. Without loss of generality, let $s = \exp(2i\theta)$, where $0 < \theta < \pi$. A reflection about the axis θ will align s with the ferromagnetic domain, that is, it will add s to S. If C denotes the set of all spins that participate in the cluster move, i.e., all spins that are at some point added to the pocket $\mathcal{P}$, then it suffices to prove that $C \cap S = \emptyset$.

For convenience, we redefine our angles with respect to the axis θ, so that $s = \exp(i\theta)$, and all spins in S are aligned with $-\theta$. For reasons that will become clear in the next section, each spin on the lattice can be written in the form $s_i = \exp(i\sigma_i\phi_i)$, where $\sigma_i = \pm 1$ is an Ising variable, and $\phi_i \in [0, \pi]$. Then, the proposed reflection flips the Ising variable while leaving ϕ unchanged.

A spin s_j is added to the pocket if and only if one of its neighbors s_i is reflected and if, after the reflection, the pair violates the inclusion constraint. This leads to the following proposition:

Proposition 1. Consider two neighboring spins $s_i = \exp(i\sigma_i\phi_i)$, $s_j = \exp(i\sigma_j\phi_j)$, where $\phi \in [0, \pi]$, $\sigma = \pm 1$. If, after reflecting s_i about the axis $\theta = 0$, i.e., $s_i \to s_i^*$, the pair (s_i^*, s_j) violates the inclusion constraint, then $\sigma_i = \sigma_j$.

For, since each move starts from a valid configuration, s_i and s_j must initially enclose an angle less than Δ:

$$\cos \Delta < \cos(\sigma_j \phi_j - \sigma_i \phi_i) = \cos \phi_i \cos \phi_j + \sigma_i \sigma_j \sin \phi_i \sin \phi_j \tag{A1}$$

while, after reflecting the spin s_i, the inclusion constraint will only be violated if the inner product satisfies

$$\cos \Delta > \cos(\sigma_j \phi_j + \sigma_i \phi_i) = \cos \phi_i \cos \phi_j - \sigma_i \sigma_j \sin \phi_i \sin \phi_j. \tag{A2}$$

Inequalities (A1) and (A2) can only hold simultaneously if $\sigma_i = \sigma_j$.

Every spin in the set C can be viewed as the end point of a directed path of spins $s_1 \to s_2 \ldots \to s_n$ where s_i gets reflected, which then violates the hard constraint with its neighbor s_{i+1}, which is then added to the pocket, and so on. Since the cluster move starts by transforming s, every such path can be traced back to s, and since all spins on the path have the same Ising variable due to Proposition 1, it follows that every spin in C has $\sigma = 1$. Recalling our convention that all spins in the ferromagnetic domain S have a negative Ising variable, Theorem 1 immediately follows. ∎

The same proof applies for inclusion Z_N models, the only difference being that the set of initial allowed configurations is restricted to discretized spin states.

Theorem 1 is independent of the choice of lattice. On bipartite lattices, the algorithm is also ergodic for exclusion models, including those which lack a direct mapping to inclusion models, namely, odd-N clock models. In this case, the target configuration consists of $\theta = 0$ on the A sublattice and $\theta = \pi(N-1)/N$ on the B sublattice, approaching an antiferromagnet in the XY limit. Using the same definitions as before, with angles defined with respect to the axis of reflection, we arrive at the following corollary to Proposition 1.

Proposition 2. Consider two neighboring spins $s_i = \exp(i\sigma_i \phi_i)$, $s_j = \exp(i\sigma_j \phi_j)$, where $\phi \in [0, \pi]$, $\sigma = \pm 1$. If, after reflecting s_i about the axis $\theta = 0$, i.e., $s_i \to s_i^*$, the pair (s_i^*, s_j) violates the exclusion constraint, then $\sigma_i = -\sigma_j$.

This implies that C contains only $\sigma = 1$ ($\sigma = -1$) on the A (B) sublattice, whereas S contains only $\sigma = -1$ ($\sigma = 1$) on the A (B) sublattice. Once again, $C \cap S = \emptyset$, so the target configuration can be attained in $\leqslant N_s$ moves.

APPENDIX B: MAPPING TO RANDOM CLUSTER MODEL

1. Fortuin-Kasteleyn mapping

An appealing feature of the `reflect` algorithm is that the average cluster size near the transition scales with the same exponent, within error bars, as the susceptibility, as shown in Fig. 9. This property is observed for "good" cluster algorithms, such as the Wolff algorithm for Ising and Potts models [54] and a generalized geometric cluster algorithm used to study lattice gases and Ising models at constant magnetization [72,73]. Conversely, the pivot cluster algorithm is inefficient for systems of hard disks, and even more so for hard rods, near the transition because the percolation threshold and transition density do not match [51]. Thus, agreement between η and η' is a measure of the algorithm's efficiency, indicating the extent to which the clusters produced by the algorithm are "physical" clusters mediating the transition [74].
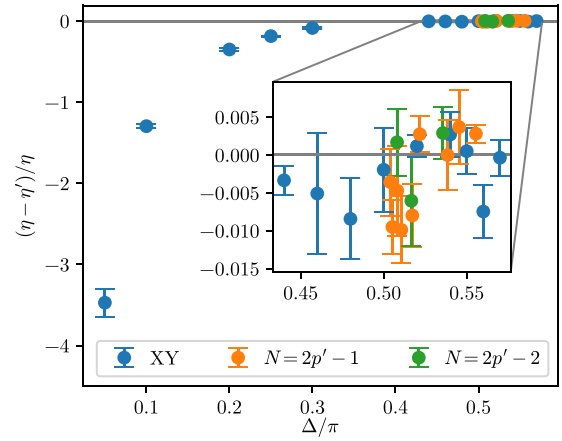

FIG. 9. Discrepancy between η and η' as a function of inclusion angle. Vertical axis is the relative error $(\eta - \eta')/\eta$, with error bars estimated from the statistical uncertainty of the power-law fit used in finite-size scaling. Data from clock models were plotted by defining $\Delta = \pi p'/N$.

The scaling of cluster sizes in the Wolff algorithm follows from an exact mapping between the Ising model and the Fortuin-Kasteleyn random cluster model [75–77]. This equivalence is constructed from a joint probability distribution over Ising spins $\{\sigma_i\}$, which live on the sites of the lattice, and bond variables $\{b_{ij}\}$ defined on the nearest-neighbor links. Each bond is either occupied ($b_{ij} = 1$) or empty ($b_{ij} = 0$). A given bond configuration decomposes the lattice into clusters, where two sites (i, j) are said to be in the same cluster if and only if they are connected by a path of occupied bonds, denoted $i \leftrightarrow j$. The Swendsen-Wang algorithm uses this bond configuration to generate the next spin state by independently flipping each cluster with probability $\frac{1}{2}$ [78]. On the other hand, the Wolff algorithm selects just one of these clusters and flips it with probability 1; this single cluster is constructed by choosing a random site and adding neighbors to the cluster with the conditional probability $p(b_{ij} = 1|\{\sigma\})$. This leads to two questions: (1) Can the scaling of clusters in the `reflect` algorithm likewise be explained through an exact mapping to a bond model? (2) If such a mapping exists, what drives the systematic discrepancy between η and η' at lower inclusion angles seen in Fig. 9?

As a first step toward such a mapping, we define an Ising variable ("pseudospin") on each site and relate the pseudospin correlation function to the correlation function of the original spins, following a similar procedure to that in Refs. [47,48,79]. As in Appendix A, we let $\sigma = \text{sgn}(\varphi)$, where $\varphi \in [-\pi, \pi]$ is the angle with respect to the axis of reflection θ. Then, expressing the angle $\varphi = \sigma \phi$ where the "auxiliary spin" ϕ is in the interval $[0, \pi]$, a reflection about axis θ flips the Ising variable while leaving the auxiliary spin unchanged. The inclusion models considered in this paper belong to a class of ferromagnetic Hamiltonians whose pair potential $V(\varphi_i, \varphi_j)$ obeys the inequality

$$V(\phi_i, \phi_j) = V(-\phi_j, -\phi_j) \leqslant V(\phi_i, -\phi_j) = V(-\phi_i, \phi_j). \tag{B1}$$

That is, given a choice of auxiliary spins ϕ_i, ϕ_j, the pair potential between like Ising pseudospins is less than or equal

to the pair potential between opposite Ising pseudospins. This suggests the following definition of the bond occupation probability, conditioned on both $\{\sigma\}$ and $\{\phi\}$:

$$p(b_{ij} = 1 | \{\sigma\}, \{\phi\}) = \delta_{\sigma_i \sigma_j} p_{ij}(\{\phi\}), \qquad (B2)$$

where

$$p_{ij}(\{\phi\}) = 1 - \exp[V(\phi_i, \phi_j) - V(-\phi_i, \phi_j)]. \qquad (B3)$$

Note that p_{ij} depends only on the auxiliary spins; the conditional bond occupation probability in Eq. (B2) depends on the Ising pseudospin only through the Kronecker delta function. This allows us to define the joint probability distribution of the $\{\sigma\}$, $\{b\}$ variables, conditioned on a given configuration of auxiliary spins, as

$$p(\{\sigma\}, \{b\} | \{\phi\}) = Z^{-1} \prod_{\langle ij \rangle} [(1 - p_{ij})\delta_{b_{ij},0} + p_{ij}\delta_{b_{ij},1}\delta_{\sigma_i \sigma_j}], \qquad (B4)$$

where Z is the partition function used to normalize the probability distribution. Summing out over the bond variables yields the marginal distribution for $\{\sigma\}$ conditioned on $\{\phi\}$:

$$p(\{\sigma\} | \{\phi\}) = Z^{-1} \prod_{\langle ij \rangle} e^{V(\phi_i, \phi_j)} \prod_{\langle ij \rangle} e^{-V(\sigma_i \phi_i, \sigma_j \phi_j)}. \qquad (B5)$$

The first product depends only on the auxiliary variables while the second product is simply the Boltzmann weight associated with a given spin configuration; supplemented with an appropriately normalized definition of the $\{\phi\}$ probability distribution as $p(\{\phi\}) \propto \prod_{\langle ij \rangle} \exp[-V(\phi_i, \phi_j)]$, we recover the partition function of the original Hamiltonian. Equation (B5) further implies the conditional bond occupation probability as defined in Eq. (B2). The marginal distribution for the bond variables also has a simple form [80], but in the present context we are most interested in the resulting form of the pseudospin probability distribution conditioned on the bond variables. As with the Fortuin-Kasteleyn (FK) representation of the Ising model, Ising spins belonging to the same cluster are aligned, while different clusters are uncorrelated. Summing over all possible bond configurations, the expectation value of the Ising variable correlation is

$$\langle \sigma_i \sigma_j \rangle = p(i \leftrightarrow j), \qquad (B6)$$

where $p(i \leftrightarrow j)$ is the probability, over all possible bond configurations, that i and j belong to the same cluster. Equation (B6) is the central identity to the FK mapping.

While this construction works for any ferromagnetic Hamiltonian defined on XY spins (and can straightforwardly be generalized to antiferromagnetic Hamiltonians by considering a staggered correlation function), it is most useful in the context of our hard-core inclusion models. In this case, the bond configuration is fully determined by the spin configuration

$$p_{ij}(\phi) = \begin{cases} 1, & \text{angle}(-\phi_i, \phi_j) \geqslant \Delta \\ 0, & \text{angle}(-\phi_i, \phi_j) < \Delta. \end{cases} \qquad (B7)$$

In words, this equation says that the bond $\langle ij \rangle$ is occupied if and only if reflecting spin s_i about the chosen axis would cause the inclusion constraint with s_j to be violated. This is precisely the condition for adding a spin to the cluster in line 8 of Algorithm 1. Thus, the `reflect` algorithm identifies and
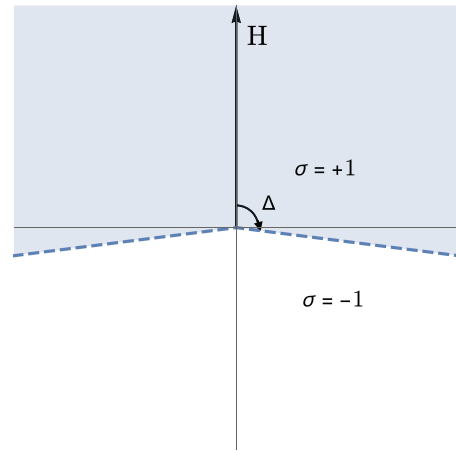

FIG. 10. Sketch of the relationship between the Ising variable σ and the piecewise function $C(\vartheta)$ defined in Eq. (B8). For the Ising pseudospins, $\sigma = 1$ for spins pointed above the x axis, and $\sigma = -1$ for spins pointed below the x axis. For the function $C(\vartheta)$, we define ϑ with respect to an external magnetic field $\mathbf{H}$ oriented along the positive y axis. Then, given an inclusion angle Δ, spins pointing in the blue shaded region have $C(\vartheta) = 1$ and spins pointing in the unshaded region have $C(\vartheta) = -1$. For the inclusion angle shown, $C(\vartheta) = \sigma$ everywhere except in the narrow shaded region below the x axis.

transforms one of the clusters in the bond configuration $\{b\}$, which in turn is determined by the original spin state and the randomly chosen axis of reflection.

To relate Eq. (B6) to the XY spin correlation function $\langle \mathbf{s}_i \cdot \mathbf{s}_j \rangle$, consider instead the correlation function $\langle C(\vartheta_i)C(\vartheta_j) \rangle$, where

$$C(\vartheta) = \begin{cases} 1, & |\vartheta| < \Delta \\ -1, & |\vartheta| \geqslant \Delta \end{cases} \qquad (B8)$$

with ϑ the angle with respect to an external field $\mathbf{H}$. This correlation function was used to study the phase diagram of a finite-temperature step model with $\Delta = \pi/2$ in Refs. [7,8], in which it was speculated that the associated susceptibility scales with the same exponent as that associated with $\langle \mathbf{s}_i \cdot \mathbf{s}_j \rangle$. To justify this speculation, note that $C(\vartheta)$ can be expanded as a cosine series on the interval $[-\pi, \pi]$, yielding the correlation function

$$\langle C(\vartheta_i)C(\vartheta_j) \rangle = \sum_{n=1}^{\infty} a_n^2 \langle \cos(n\vartheta_i) \cos(n\vartheta_j) \rangle, \qquad (B9)$$

where a_n is the Fourier coefficient of the nth term in the series and charge neutrality ensures that cross terms vanish [46]. Since the associated critical exponent of each term in the series scales as n^2, Eq. (B9) is dominated by the first term, which is proportional to $\langle \mathbf{s}_i \cdot \mathbf{s}_j \rangle$. The justification for using the $C(\vartheta)$ correlation function as a proxy for the XY spin correlation function is especially strong at an inclusion angle of $\Delta = \pi/2$, for which the coefficients a_n are identically zero for even n. In this case, the subleading correction is $\langle \cos(3\vartheta_i) \cos(3\vartheta_j) \rangle$ which, although relevant, is strongly suppressed compared to the leading term $\langle \cos \vartheta_i \cos \vartheta_j \rangle$.

At $\Delta = \pi/2$, $C(\vartheta)$ is also in exact correspondence with the Ising variable σ, if we align $\mathbf{H}$ at an angle of $\pi/2$ with respect

to the axis of reflection (Fig. 10). Equation (B6) then implies

$$\langle C(\vartheta_i)C(\vartheta_j)\rangle = p(i \leftrightarrow j) \quad (\Delta = \pi/2). \tag{B10}$$

For inclusion angles near $\pi/2$, such as those in the neighborhood of the Kosterlitz-Thouless transition, Eq. (B10) approximately holds. Then, summing over lattice sites i, j to obtain the susceptibility and recalling that the `reflect` algorithm selects one cluster in the random bond configuration with probability proportional to the size of the cluster, we arrive at the final result

$$L^{2-\eta} \sim \chi \sim \langle s \rangle \sim L^{2-\eta'}. \tag{B11}$$

Equation (B11) justifies the close agreement of η and η' near the critical angle $\Delta_c \approx 0.56\pi$, which is fortuitously close to the angle at which Eq. (B10) holds exactly. On the other hand, the sharp deviation of η' from η for low inclusion angles seen in Fig. 9 is explained by the fact that far away from $\Delta = \pi/2$, the approximate relation between the pseudospin correlation function and the $C(\vartheta)$ correlation function no longer holds.

2. Scaling of the cluster-size distribution

To further assess whether the `reflect` algorithm chooses the appropriate physical clusters, we measure the distribution of cluster sizes at $\Delta = 0.5\pi$. Defining $n^*(s)$ as the probability density function of hitting on a cluster of size s, we obtain a bimodal distribution, with a large peak at small cluster sizes and a secondary peak at large cluster sizes. This indicates that the scaling of the average cluster size, from which the exponent η' is determined, results from a subtle interplay between these two peaks.

The cluster-size distribution is typically studied in the context of the Swendsen-Wang algorithm, for which scaling forms have been derived, through the FK mapping to a critical percolation problem, to describe observables such as (1) the percolation probability $\langle P_\infty \rangle$, which exhibits the same finite-size scaling $\sim L^{-\beta/\nu}$ as the net magnetization, (2) the distribution of the size of the largest (spanning) cluster in each bond configuration, whose mean scales as L^{d_F} where d_F is the fractal dimension, and (3) the number per site of clusters of size s [74,81]. The scaling form of this third observable must be modified when considering a single-cluster algorithm, such as the Wolff and `reflect` algorithms. Under the hypothesis that the inclusion model at $\Delta = 0.5\pi$ maps, via the above construction, to the critical point of some unknown percolation problem, we posit the scaling form [82]

$$n^*(s, L) \sim s^{-\tau+1} f(s/L^{d_F}), \tag{B12}$$

where $\tilde{n}(x)$ is a universal scaling function. To verify this scaling form, in Fig. 11 the cluster-size distribution is multiplied by $s^{\tau-1}$ and plotted as a function of s/L^{d_F} for system sizes ranging from $L = 8$ to 128. The exponents d_F and τ are related to the exponent η' by summing Eq. (B12) over s to obtain the average cluster size:

$$L^{2-\eta'} \sim \langle s \rangle = \sum_s n^*(s)s \sim L^{d_F(3-\tau)}$$

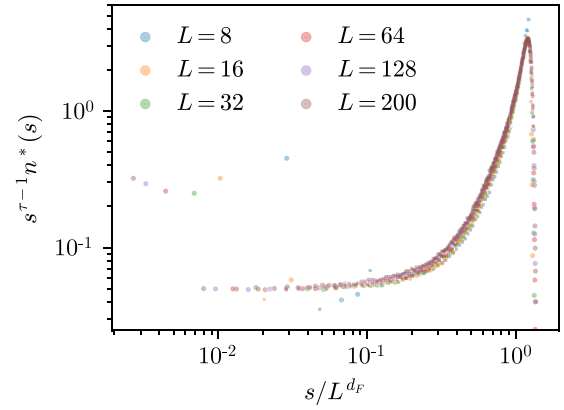$$\Rightarrow d_F = \frac{2-\eta'}{3-\tau}. \tag{B13}$$

FIG. 11. Rescaled probability distribution $\tilde{n}(s/L^{d_F})$ of cluster sizes at $\Delta = 0.5\pi$. The horizontal axis is cluster size s rescaled by L^{d_F}, where $d_F = 1.901$ is the fractal dimension determined from Eq. (B15), using $\eta = 0.198$. The vertical axis is $n^*(s)s^{\tau-1}$. Optimal data collapse was achieved using a cluster exponent of $\tau - 1 = 1.089$, obtained from a power-law fit described in the text. The distribution was measured over $\approx 8 \times 10^5$ cluster moves for $L = 8$, $\approx 1.1 \times 10^6$ cluster moves for $L = 16$, $\approx 1.65 \times 10^6$ cluster moves for $L = 32$, $\approx 2.3 \times 10^6$ cluster moves for $L = 64$, $\approx 3.2 \times 10^6$ cluster moves for $L = 128$, and $\approx 4 \times 10^6$ moves for $L = 200$.

Assuming hyperscaling, $d\nu = \gamma + 2\beta$ [32], the cluster exponent τ can be related to the fractal dimension d_F via

$$\tau = 1 + d/d_F \tag{B14}$$

which, when substituted into Eq. (B13) with $d = 2$, implies

$$d_F = 2 - \frac{\eta'}{2} = 2 - \frac{\eta}{2}. \tag{B15}$$

Using the critical exponent $\eta = 0.198$ at $\Delta = 0.5\pi$, this implies a fractal dimension $d_F = 1.901$. As a test of internal consistency, d_F can be estimated via finite-size scaling of s_{max}, the cluster size at which $s^{\tau-1}n^*(s, L)$ is a maximum. A fit to the form $s_{max} \sim L^{d_F}$ yields $d_F = 1.902 \pm 0.003$.

Substituting Eq. (B15) into (B14) yields

$$\tau = 1 + \frac{4}{4-\eta} \tag{B16}$$

which, for $\eta = 0.198$, implies a cluster exponent of $\tau = 2.052$. Again, we can verify the internal consistency of the scaling form by fitting $n^*(s) \sim s^{-\tau+1}$ in the regime $1 \ll s \ll L^{d_F}$ [81]. At the accessible system sizes, however, such a fit is prone to error, so we instead measure τ through a power-law fit to the value of $n^*(s, L)$ measured at the position of the secondary peak at system size L vs the position of the peak. This yields the estimate $\tau = 2.089$, which is used to achieve the data collapse between different system sizes in Fig. 11. The good data collapse at large cluster sizes, particularly for $L \geqslant 16$, provides evidence in favor of the assumed scaling form as the leading behavior of the distribution.

Taken together, Figs. 9 and 11 illustrate the connection between the scaling exponents of the original spin model and the cluster model derived from the `reflect` algorithm. Combined with the ergodicity proven in Appendix A, this suggests that the algorithm is, in the same sense as the original
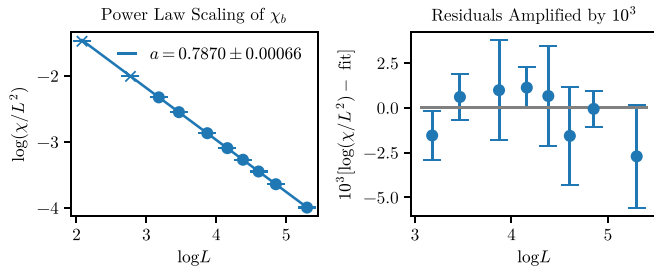
FIG. 12. Finite-size scaling of the bond susceptibility, defined in Eq. (B17), for $L = 8, 16, 24, 32, 48, 64, 80, 100, 128, 200$. Only data with $L \geqslant 24$ were included in the fit. Error bars were determined from the standard error across 100 or 1000 chunks. Right panel shows the residuals with respect to the fit, amplified by a factor of 10^3.

Swendsen-Wang/Wolff algorithm for the $d = 2$ Ising model, optimally adapted to our model at the value $\Delta = 0.5\pi$.

Our value of τ is quite close to the exact value 187/91 for short-ranged percolation in $d = 2$, which leads us to conjecture that our algorithm is constructing clusters precisely at the critical point of that problem. Further, if we use the exact value of τ for the scaling collapse, the results are equally impressive by eye.

3. Long-ranged bond correlations

We note that the connection to short-ranged percolation is somewhat unexpected as the underlying spin correlations are long ranged. Decomposing the entire lattice into clusters in the manner of Swendsen-Wang helps clarify this picture. The first moment of the bond distribution is measured to be $\langle b \rangle = \frac{1}{2}$, in agreement with the percolation threshold p_c of uncorrelated bond percolation on the square lattice on the grounds of duality [83]. However, in contrast to that short-ranged percolation problem, we do find evidence of QLRO in the bond correlations. Figure 12 shows the finite-size scaling of the bond susceptibility, defined analogously to the spin susceptibility as

$$\chi_b(L) = \frac{1}{L^2} \sum_{\substack{\langle ij \rangle \\ \langle kl \rangle}} [\langle b_{ij} b_{kl} \rangle - \langle b \rangle^2]. \qquad (B17)$$

We find that this scales as a power law $\chi_b \sim L^{2-a}$, implying algebraic decay of the bond correlations with exponent a:

$$\langle b_{ij} b_{kl} \rangle - \langle b \rangle^2 \propto 1/r^a, \qquad (B18)$$

where r is the distance between bonds $\langle ij \rangle$ and $\langle kl \rangle$ on the lattice. The estimated value of a, determined from scaling up to $L = 200$, is

$$a = 0.7870 \pm 0.000\,66. \qquad (B19)$$

The Harris criterion applied to long-range correlated percolation suggests that correlations of the form of Eq. (B18) are irrelevant for $a > 2/\nu$, where ν is the percolation correlation length exponent for the short-ranged percolation problem [84]. In this regime, the critical exponents at the percolation threshold are unaffected by the presence of long-range bond correlations. When the correlations become relevant ($a < 2/\nu$), the critical exponents are expected to deviate from the uncorrelated percolation exponents in an a-dependent fashion.

Given that $\nu = \frac{4}{3}$ for short-ranged 2D percolation, Eq. (B19) implies that the quasi-long-ranged bond correlations are in fact relevant. However, simulation of long-range correlated site percolation on the square lattice [85] indicates that while other critical exponents depend on a, the fractal dimension d_F (and in turn, τ) remains consistent with its uncorrelated value as a decreases. Similarly, on the triangular lattice d_F was found to be independent of a down to $a = \frac{2}{3}$, below which d_F increases continuously [86]. Thus, we conclude that at $\Delta = 0.5\pi$, the cluster model produced by the `reflect` algorithm maps to standard percolation in $d = 2$, with relevant quasi-long-ranged correlations that do not affect the scaling of average or largest cluster size. To further investigate this interpretation, in future research it will be worthwhile to measure other critical exponents of the percolation model that are sensitive to these correlations.

[1] L. Onsager, The Effects of Shape on the Interaction of Colloidal Particles, Ann. N.Y. Acad. Sci. **51**, 627 (1949).

[2] Y. G. Sinai, On the foundations of the ergodic hypothesis for a dynamical system of statistical mechanics, in *Doklady Akademii Nauk*, Vol. 153 (Russian Academy of Sciences, Moscow, 1963), pp. 1261–1264.

[3] D. Frenkel, Order through entropy, Nat. Mater. **14**, 9 (2015).

[4] J. M. Kosterlitz and D. J. Thouless, Ordering, metastability and phase transitions in two-dimensional systems, J. Phys. C: Solid State Phys. **6**, 1181 (1973).

[5] J. V. José, L. P. Kadanoff, S. Kirkpatrick, and D. R. Nelson, Renormalization, vortices, and symmetry-breaking perturbations in the two-dimensional planar model, Phys. Rev. B **16**, 1217 (1977).

[6] M. Hasenbusch, $O(N)$ and RP^{N-1} models in two dimensions, Phys. Rev. D **53**, 3445 (1996).

[7] A. J. Guttmann and G. S. Joyce, Critical behaviour of an isotropic spin system. I, J. Phys. C: Solid State Phys. **6**, 2691 (1973).

[8] A. J. Guttman and A. Nymeyer, Critical behaviour of an isotropic spin system. II, J. Phys. A: Math. Gen. **11**, 1131 (1978).

[9] I.-Hsiu Lee and R. E. Shrock, Critical behavior of a new class of statistical-mechanical models with continuous variables and discretely valued interactions, Phys. Rev. B **36**, 3712 (1987).

[10] A. C. Irving and R. Kenna, Critical behavior of the two-dimensional step model, Phys. Rev. B **53**, 11568 (1996).

[11] M. N. Barber, Migdal transformations of O(2)-symmetric spin Hamiltonians, J. Phys. A: Math. Gen. **16**, 4053 (1983).

[12] W. Bietenholz, M. Bögli, F. Niedermayer, M. Pepe, F. G. Rejón-Barrera, and U. J. Wiese, Topological lattice actions for the 2d XY model, J. High Energy Phys. 03 (2013) 141.

[13] W. Bietenholz, U. Gerber, and F. G. Rejón-Barrera, Berezinskii-Kosterlitz-Thouless transition with a constraint lattice action, J. Stat. Mech.: Theory Exp. (2013) P12009.

[14] J. L. Cardy, General discrete planar models in two dimensions: Duality properties and phase diagrams, J. Phys. A: Math. Gen. **13**, 1507 (1980).

[15] T. Obuchi and H. Kawamura, Spin and chiral orderings of the antiferromagnetic XY model on the triangular lattice and their critical properties, J. Phys. Soc. Jpn. **81**, (2012).

[16] J.-P. Lv, T. M. Garoni, and Y. Deng, Phase transitions in XY antiferromagnets on plane triangulations, Phys. Rev. B **87**, 024108 (2013).

[17] R. Shrock and S. H. Tsai, Ground state entropy of potts antiferromagnets and the approach to the 2D thermodynamic limit, Phys. Rev. E **58**, 4332 (1998).

[18] S. C. Chang and R. Shrock, Ground state entropy of the potts antiferromagnet on triangular lattice strips, Ann. Phys. **290**, 124 (2001).

[19] J. L. Jacobsen, J. Salas, and A. D. Sokal, Transfer Matrices and Partition-Function Zeros for Antiferromagnetic Potts Models. III. Triangular-Lattice Chromatic Polynomial, J. Stat. Phys. **112**, 921 (2003).

[20] J. L. Jacobsen, J. Salas, and C. R. Scullard, Phase diagram of the triangular-lattice Potts antiferromagnet, J. Phys. A: Math. Theor. **50**, 345002 (2017).

[21] R. Peled and Y. Spinka, Rigidity of proper colorings of $\mathbb{Z}^d$, arXiv:1808.03597.

[22] R. Kotecký, A. D. Sokal, and J. M. Swart, Entropy-driven phase transition in low-temperature antiferromagnetic potts models, Commun. Math. Phys. **330**, 1339 (2014).

[23] Y. Huang, K. Chen, Y. Deng, J. L. Jacobsen, R. Kotecký, J. Salas, A. D. Sokal, and J. M. Swart, Two-dimensional Potts antiferromagnets with a phase transition at arbitrarily large q, Phys. Rev. E **87**, 012136 (2013).

[24] R. Peled and Y. Spinka, Long-range order in discrete spin systems, arXiv:2010.03177.

[25] L. A. Goldberg, M. Jalsenius, R. Martin, and M. Paterson, Improved Mixing Bounds for the Anti-Ferromagnetic Potts Model on $\mathbb{Z}^2$, LMS J. Comput. Math. **9**, 1 (2006).

[26] J. Salas and A. D. Sokal, The 3-State Square-Lattice Potts Antiferromagnet at Zero Temperature, J. Stat. Phys. **92**, 729 (1998).

[27] N. Chandgotia, R. Peled, S. Sheffield, and M. Tassy, Delocalization of uniform graph homomorphisms from $\mathbb{Z}^2$ to $\mathbb{Z}$, arXiv:1810.10124.

[28] H. Duminil-Copin, M. Harel, B. Laslier, A. Raoufi, and G. Ray, Logarithmic variance for the height function of square-ice, arXiv:1911.00092.

[29] G. Ray and Y. Spinka, Proper 3-colorings of $\mathbb{Z}^2$ are Bernoulli, arXiv:2004.00028.

[30] B. J. Alder and T. E. Wainwright, Phase transition for a hard sphere system, J. Chem. Phys. **27**, 1208 (1957).

[31] N. D. Mermin and H. Wagner, Absence of Ferromagnetism or Antiferromagnetism in One- or Two-Dimensional Isotropic Heisenberg Models, Phys. Rev. Lett. **17**, 1133 (1966).

[32] N. Goldenfeld, *Lectures on Phase Transitions and the Renormalization Group* (Addison-Wesley, Boston, 1992).

[33] S. Elitzur, R. B. Pearson, and J. Shigemitsu, Phase structure of discrete Abelian spin and gauge systems, Phys. Rev. D **19**, 3698 (1979).

[34] J. Fröhlich and T. Spencer, The Kosterlitz-Thouless transition in two-dimensional Abelian spin systems and the Coulomb gas, Commun. Math. Phys. **81**, 527 (1981).

[35] O. Borisenko, G. Cortese, R. Fiore, M. Gravina, and A. Papa, Numerical study of the phase transitions in the two-dimensional z(5) vector model, Phys. Rev. E **83**, 041120 (2011).

[36] D.-H. Kim, Partition function zeros of the p-state clock model in the complex temperature plane, Phys. Rev. E **96**, 052130 (2017).

[37] Z.-Q. Li, L.-P. Yang, Z. Y. Xie, H.-H. Tu, H.-J. Liao, and T. Xiang, Critical properties of the two-dimensional q-state clock model, Phys. Rev. E **101**, 060105(R) (2020).

[38] J. L. Cardy, Antiferromagnetic clock models in two dimensions, Phys. Rev. B **24**, 5128 (1981).

[39] M. den Nijs, String melting of the floating phase in antiferromagnetic clock models, Phys. Rev. B **31**, 266 (1985).

[40] B. Illing, S. Fritschi, H. Kaiser, C. L. Klix, G. Maret, and P. Keim, Mermin-Wagner fluctuations in 2D amorphous solids, Proc. Natl. Acad. Sci. USA **114**, 1856 (2017).

[41] M. Engel, J. A. Anderson, S. C. Glotzer, M. Isobe, E. P. Bernard, and W. Krauth, Hard-disk equation of state: First-order liquid-hexatic transition in two dimensions with three simulation methods, Phys. Rev. E **87**, 042134 (2013).

[42] B. Simon and A. D. Sokal, Rigorous entropy-energy arguments, J. Stat. Phys. **25**, 679 (1981).

[43] B. I. Halperin, On the Hohenberg-Mermin-Wagner Theorem and Its Limitations, J. Stat. Phys. **175**, 521 (2019).

[44] P. Miłoś and R. Peled, Delocalization of two-dimensional random surfaces with hard-core constraints, Commun. Math. Phys. **340**, 1 (2015).

[45] R. Peled (private communication).

[46] E. Fradkin, *Field theories of Condensed Matter Physics* (Cambridge University Press, Cambridge, 2013).

[47] A. Patrascioiu and E. Seiler, Phase structure of two-dimensional spin models and percolation, J. Stat. Phys. **72**, 1407 (1993).

[48] M. Aizenman, On the Slow Decay of O(2) Correlations in the Absence of Topological Excitations: Remark on the Patrascioiu-Seiler Model, J. Stat. Phys. **77**, 351 (1994).

[49] J. L. Cardy, Universal amplitudes in finite-size scaling: Generalisation to arbitrary dimensionality, J. Phys. A: Math. Gen. **18**, L757 (1985).

[50] P. Nightingale and H. Blöte, The relation between amplitudes and critical exponents in finite-size scaling, J. Phys. A: Math. Gen. **16**, L657 (1983).

[51] W. Krauth, Cluster Monte Carlo algorithms, in *New Optimization Algorithms in Physics* (John Wiley & Sons, Ltd, 2004) Chap. 2, pp. 5-22.

[52] C. Dress and W. Krauth, Cluster algorithm for hard spheres and related systems, J. Phys. A: Math. Gen. **28**, L597 (1995).

[53] W. Krauth and R. Moessner, Pocket Monte Carlo algorithm for classical doped dimer models, Phys. Rev. B **67**, 064503 (2003).

[54] U. Wolff, Collective Monte Carlo Updating for Spin Systems, Phys. Rev. Lett. **62**, 361 (1989).

[55] E. Luijten, Introduction to Cluster Monte Carlo Algorithms, Lect. Notes Phys. **703**, 349 (2006).

[56] M. Hasenbusch, The two-dimensional XY model at the transition temperature: A high-precision Monte Carlo study, J. Phys. A: Math. Gen. **38**, 5869 (2005).

[57] C. F. Baillie and P. D. Coddington, Comparison of cluster algorithms for two-dimensional Potts models, Phys. Rev. B **43**, 10617 (1991).

[58] J. Du, B. Zheng, and J. S. Wang, Dynamic critical exponents for Swendsen-Wang and Wolff algorithms obtained by a nonequilibrium relaxation method, J. Stat. Mech.: Theory Exp. (2006) P05004.

[59] W. Krauth, *Statistical Mechanics: Algorithms and Computations*, Vol. 13 (Oxford University Press, Oxford, 2006).

[60] 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**, 14935 (1997).

[61] J. K. Burton and C. L. Henley, A constrained Potts antiferromagnet model with an interface representation, J. Phys. A: Math. Gen. **30**, 8385 (1997).

[62] C. Moore and M. E. Newman, Height Representation, Critical Exponents, and Ergodicity in the Four-State Triangular Potts Antiferromagnet, J. Stat. Phys. **99**, 629 (2000).

[63] D. A. Huse and A. D. Rutenberg, Classical antiferromagnets on the kagomé lattice, Phys. Rev. B **45**, 7536 (1992).

[64] J. Cardy, J. L. Jacobsen, and A. D. Sokal, Unusual Corrections to Scaling in the 3-state Potts Antiferromagnet on a Square Lattice, J. Stat. Phys. **105**, 25 (2001).

[65] R. Peled and Y. Spinka, A condition for long-range order in discrete spin systems with application to the antiferromagnetic Potts model, arXiv:1712.03699.

[66] R. Peled, High-Dimensional Lipschitz Functions Are Typically Flat, Ann. Probab. **45**, 1351 (2017).

[67] J. Kondev and C. L. Henley, Four-coloring model on the square lattice: A critical ground state, Phys. Rev. B **52**, 6628 (1995).

[68] J. D. Qualls, Lectures on conformal field theory, arXiv:1511.04074.

[69] Y. Komura and Y. Okabe, Large-scale Monte Carlo Simulation of Two-Dimensional Classical XY Model Using Multiple GPUs, J. Phys. Soc. Jpn. **81**, 113001 (2012).

[70] M. Hasenbusch, The binder cumulant at the Kosterlitz–Thouless transition, J. Stat. Mech.: Theory Exp. (2008) P08003.

[71] H. Yoshino, Disorder-free spin glass transitions and jamming in exactly solvable mean-field models, SciPost Phys. **4**, 40 (2018).

[72] J. R. Heringa and H. W. J. Blöte, Geometric symmetries and cluster simulations, Phys. A (Amsterdam) **254**, 156 (1998).

[73] J. R. Heringa and H. W. J. Blöte, Geometric cluster Monte Carlo simulation, Phys. Rev. E **57**, 4976 (1998).

[74] M. D. De Meo, D. W. Heermann, and K. Binder, Monte carlo study of the ising model phase transition in terms of the percolation transition of "physical clusters", J. Stat. Phys. **60**, 585 (1990).

[75] C. M. Fortuin and P. W. Kasteleyn, On the random-cluster model: I. Introduction and relation to other models, Physica (Amsterdam) **57**, 536 (1972).

[76] C. M. Fortuin, On the random-cluster model: II. The percolation model, Physica (Amsterdam) **58**, 393 (1972).

[77] C. M. Fortuin, On the random-cluster model: III. The simple random-cluster model, Physica (Amsterdam) **59**, 545 (1972).

[78] R. H. Swendsen and J.-S. Wang, Nonuniversal Critical Dynamics in Monte Carlo Simulations, Phys. Rev. Lett. **58**, 86 (1987).

[79] A. Patrascioiu and E. Seiler, Percolation and the Existence of a Soft Phase in the Classical Heisenberg Model, J. Stat. Phys. **106**, 811 (2002).

[80] R. G. Edwards and A. D. Sokal, Generalization of the Fortuin-Kasteleyn-Swendsen-Wang representation and Monte Carlo algorithm, Phys. Rev. D **38**, 2009 (1988).

[81] P. Hou, S. Fang, J. Wang, H. Hu, and Y. Deng, Geometric properties of the Fortuin-Kasteleyn representation of the Ising model, Phys. Rev. E **99**, 042150 (2019).

[82] P. W. Leung and C. L. Henley, Percolation properties of the Wolff clusters in planar triangular spin models, Phys. Rev. B **43**, 752 (1991).

[83] K. Christensen, *Percolation Theory* (MIT, 2002).

[84] A. Weinrib, Long-range correlated percolation, Phys. Rev. B **29**, 387 (1984).

[85] S. Prakash, S. Havlin, M. Schwartz, and H. E. Stanley, Structural and dynamical properties of long-range correlated percolation, Phys. Rev. A **46**, R1724 (1992).

[86] K. J. Schrenk, N. Posé, J. J. Kranz, L. V. M. van Kessenich, N. A. M. Araújo, and H. J. Herrmann, Phys. Rev. E **88**, 052102 (2013).