Universality class of the two-dimensional polymer collapse transition

Adam Nahum
Department of Physics, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA
(Received 26 November 2015; published 6 May 2016)

The nature of the θ point for a polymer in two dimensions has long been debated, with a variety of candidates put forward for the critical exponents. This includes those derived by Duplantier and Saleur for an exactly solvable model. We use a representation of the problem via the $CP^{N-1}\sigma$ model in the limit $N \to 1$ to determine the stability of this critical point. First we prove that the Duplantier-Saleur (DS) critical exponents are robust, so long as the polymer does not cross itself: They can arise in a generic lattice model and do not require fine-tuning. This resolves a longstanding theoretical question. We also address an apparent paradox: Two different lattice models, apparently both in the DS universality class, show different numbers of relevant perturbations, apparently leading to contradictory conclusions about the stability of the DS exponents. We explain this in terms of subtle differences between the two models, one of which is fine-tuned (and not strictly in the DS universality class). Next we allow the polymer to cross itself, as appropriate, e.g., to the quasi-two-dimensional case. This introduces an additional independent relevant perturbation, so we do not expect the DS exponents to apply. The exponents in the case with crossings will be those of the generic tricritical O(n) model at $n = 0$ and different from the case without crossings. We also discuss interesting features of the operator content of the CP^{N-1} model. Simple geometrical arguments show that two operators in this field theory, with very different symmetry properties, have the same scaling dimension for any value of N (or, equivalently, any value of the loop fugacity). Also we argue that for any value of N the CP^{N-1} model has a marginal odd-parity operator that is related to the winding angle.

DOI: 10.1103/PhysRevE.93.052502

I. INTRODUCTION

One of the most elegant ideas in polymer physics is de Gennes's mapping between long polymer chains and the O(n) field theory in the limit $n \to 0$ [1]. The large-scale geometry of a chain in a good solvent, or a lattice self-avoiding walk, is described by the critical O(n) model. If the solvent quality is reduced, the monomers effectively attract each other and eventually the polymer collapses into a compact object via a phase transition known as the θ point. In de Gennes's correspondence the θ point maps to the tricritical O(n) model [2]. This has upper critical dimension 3, so in three dimensions the θ-point polymer is ideal (up to logarithmic corrections). The nature of the θ point in two dimensions is much more interesting and, surprisingly, not fully understood.

In two dimensions we must distinguish two kinds of models according to whether or not we allow the polymer to cross itself (Fig. 1). Most of the theoretical and numerical work has focused on models without crossings; we discuss these first. A key development was the derivation by Duplantier and Saleur of exact critical exponents for a particular honeycomb lattice model, in which polymer conformations have a relationship with percolation cluster boundaries [3,4]. Let us call the corresponding renormalization group (RG) fixed point the Duplantier-Saleur (DS) fixed point. The fact that the honeycomb lattice model is only solvable at a fine-tuned point (where the correspondence with percolation holds) led to debate about whether the DS exponents captured the *generic* critical behavior at the θ point, even for noncrossing polymers. For example, Blöte and Nienhuis proposed another solvable model for the θ point [5] (which has recently attracted new interest [6]), with different exponents, and argued that it should be more stable in the RG sense than the model solved by Duplantier and Saleur. On the other hand, numerical results seem to indicate that the DS exponents are robust against changes of the model [7–10]; see in particular Ref. [11]. Further complicating the issue, models are known that initially appeared to behave similarly to the DS polymer, but later turned out to show different universal behavior with anomalously large finite-size effects [12–16].

The question of what the generic universal behavior is for the collapse transition has remained unresolved until now. In this paper we address it using a representation of the DS universality class via a σ model with SU(N) symmetry [17] in the limit $N \to 1$. We show that the the DS exponents are robust for noncrossing polymers. The critical exponents of the original honeycomb lattice model [3,4] can arise in a generic noncrossing model, such as the interacting self-avoiding walk (ISAW), without the need for fine-tuning.

At the same time, there is an apparent paradox that we must resolve. At first sight one reaches contradictory conclusions about the stability of the DS point by analyzing different popular models that share the same field theory description and at first sight are in the same universality class. We explain why this naive symmetry analysis gives misleading results. We connect this with the fact that one of the models suffers from fine-tuning related to an Ising-like order parameter defined in Ref. [18].

To obtain the above we classify the allowed perturbations of simple models for the θ point that show the DS exponents, making use of mappings to concrete lattice field theories [12,19–22]. The lattice field theories for these models have SU(N) symmetry. This symmetry is enhanced compared to more generic polymer models: This is a manifestation of fine-tuning of the Boltzmann weights for the polymer. Any generic perturbation to the polymer's interactions breaks the symmetry to a subgroup. However, that does not in itself imply that the DS *fixed point* is fine-tuned. The SU(N) symmetry may be restored in the infrared even when it is broken microscopically. We argue that this symmetry enhancement

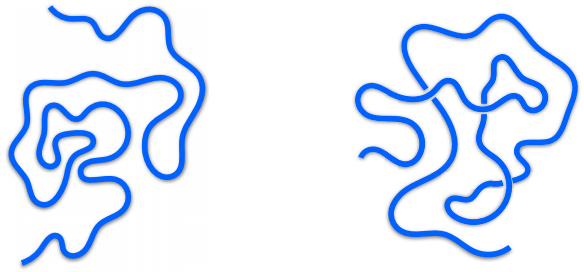

FIG. 1. In two dimensions, a basic topological distinction is between models in which the polymer chain cannot cross itself (left) and those in which it can (right).

under RG is what happens for generic models in the DS universality class. The question of the robustness of the DS exponents is therefore related to the number of relevant symmetry-breaking perturbations. (For polymers, the potential complication is that a given model may be mappable to lattice field theory in multiple ways and an ill-chosen mapping may conceal the full symmetry.)

A generic description of the θ point should have two relevant perturbations. Although for the polymer we need only tune one interaction parameter to reach the θ point, the field theory is automatically tuned to criticality by taking the length of the polymer to be large. We show that polymer models that are truly in the DS universality class indeed have two relevant perturbations (when crossings are not allowed) and the fine-tuned model mentioned above has three. In order to show that no other relevant or marginal perturbations can play a role, we are also led to analyze various features of the relevant σ model, the CP^{N-1} model.

A physical polymer system in two dimensions or quasi-two-dimensions may allow for crossings, where one part of the polymer chain lies on top of another part (Fig. 1). These may have an important effect at large length scales even if energetically disfavored at small scales. Crossings are known to lead to new universal behavior in completely packed two-dimensional (2D) loop models [23,24]. Here too, crossings may be shown to destabilize the DS fixed point. Therefore, in the case with crossings we do not expect the DS exponents to apply.

Further, we argue that the tricritical O(n) behavior expected by de Gennes will only be seen when crossings are allowed. That is, the DS universality class (and the collapse transition of the ISAW) should not be identified with the generic tricritical O(n) model, contrary to what is often assumed. A subtlety here is that at first sight certain of the models we discuss have O(n) symmetry even when the polymer cannot cross itself. However, we point out that a higher symmetry is revealed in these models by mapping them to field theory in a different way.

The best studied model with crossings is the collapse point of the interacting self-avoiding trail [12–16,23,24]. This model is in many ways analogous to the honeycomb lattice model solved by Duplantier and Saleur. It also has enhanced symmetry; in this case SO($N \to 1$), which should be regarded as larger than the O($n \to 0$) of a generic model. Unfortunately, this critical point turns out to be infinitely fine-tuned [12], so it is certainly not the generic θ point in the case with crossings. [Unlike the case with SU(N) symmetry, here there is an infinite number of relevant symmetry-breaking perturbations.] With a more generic choice of interactions for the trail, the tricritical O($n \to 0$) behavior should be seen. We are not aware of any exact results for more generic models with crossings (see Ref. [25] for a numerical study), and this is an interesting subject for future research. It was suggested by de Gennes, on the basis of the smallness of the coefficients in the $3-\epsilon$ expansion, that the tricritical $O(n)$ exponents may be close to mean field values even in two dimensions [2].

The field theory that is central to our analysis is the CP^{N-1} nonlinear σ model with a Θ term at $\Theta = \pi$. In quantum condensed matter this theory is familiar from the Heisenberg spin-1/2 chain and its SU(N) generalizations [26,27]. Its relationship with 2D loop models for loops with fugacity N has been discussed extensively [17,28,29]. Here we are interested in the limit $N \to 1$, which *a priori* describes a soup of many loops rather than a single polymer. However, a well-known trick [3–5,30] is to isolate a single marked loop and integrate out (i.e., ignore) all the others. At $N = 1$ the marked loop turns out to be governed by a local Boltzmann weight, as appropriate to a polymer. To study generic interactions for the polymer, we must change the interactions for this marked loop without modifying the weights for the soup of background loops. This corresponds to introducing various anisotropies in the σ model. This strategy was pursued for the RP^{N-1} σ model describing the interacting self-avoiding trail in Refs. [12,19]. In that case the effect of the perturbations is simpler to analyze (because the governing fixed point is free [23]), but the logic is the same.

The analysis will lead us to examine the operator content of the CP^{N-1} σ model. We find some features that are surprising from the point of view of field theory but transparent from the loop-gas perspective. For example, a simple geometrical argument shows that two operators in the field theory with very different properties under spatial and SU(N) symmetries (and different numbers of spatial derivatives) are forced to have the same scaling dimension for any N. This is related to the results of Refs. [31,32] on the symmetry algebra of these models. (The operator product expansions of these operators are also constrained by geometrical arguments.) Finally, we show that the σ model has an odd-parity operator whose scaling dimension is fixed by a relation with the loops' winding angles.

Outline

The theme of the remainder is the collapse transition in various settings, but much of the material is relevant to the CP^{N-1} model more generally. Here is an overview. Section II reviews the basic models and tools we will need (the first half of this section will be familiar to many readers) with our results presented in subsequent sections. Section III describes the operators in the CP^{N-1} model that are most important for the discussion of collapse. A full demonstration that these are the only important operators is deferred until Sec. VII. Section IV shows that the archetypal honeycomb model, and by extension any model in the DS universality class, is stable to arbitrary perturbations of the interactions. Section V considers

a well-known model on the square lattice (which we refer to as model T), in order to resolve an apparent paradox about the stability of the DS point. Section VI argues that models with crossings (Fig. 1, right) will have non-DS exponents and discusses some other features of models with crossings (the special case of smart walks). Section VII uses simple geometrical arguments to pin down the scaling dimensions of some interesting operators in the CP^{N-1} model (or its supersymmetric cousin the $CP^{N+k-1|k}$ model), specifically odd-parity variants of the two- and four-leg operators. This also confirms that our classification of perturbations in the polymer problem is complete. Technical results necessary for Secs. IV–VII are given in the Appendixes, including the lattice mappings that underlie our analysis of perturbations and an aspect of the supersymmetric formulation.

II. BACKGROUND: MODELS AND FIELD THEORIES

In this section we review the polymer models we consider and their relations with loop gases and field theory.

A. Honeycomb model

Usually models for a single polymer can be thought of as loop gases in the limit where the loop fugacity n tends to zero [1]. (As usual it will be convenient to consider a closed ring polymer rather than an open chain.) The unusual feature of solvable models in the DS universality class is that they allow a different type of mapping that is instead between the polymer model and a loop gas at fugacity $N = 1$. The loops in this gas are essentially cluster boundaries in critical percolation. The correspondence is that the Boltzmann weight of a given polymer conformation is proportional to the probability of a loop with that conformation appearing in the loop gas.

The model of Refs. [3,4] for the collapse transition maps to the much-studied gas of nonintersecting loops on the honeycomb lattice [33,34]:

$$Z_{\text{honeycomb}} = \sum_{\substack{\text{loop}\\ \text{configs}}} x^{\text{length}} N^{\text{No. of loops}} = \sum_{\substack{\text{colored}\\ \text{loop configs}}} x^{\text{length}}. \qquad (1)$$

The superscript "length" is the total length of the loops. For the second equality we have assumed N to an integer, allowing us to obtain the fugacity N by summing over N possible colors for each loop. We will be interested in the 'dense' phase (i.e., x larger than a critical value), in particular at $x = 1$ and $N = 1$.

It is useful to regard the loops as cluster boundaries, as in Fig. 2, left. Given a loop configuration, the coloring of the hexagons is unique up to a global exchange of white and black: We may, for example, sum over both choices, which simply multiplies $Z_{\text{honeycomb}}$ by 2. Viewing the loops as cluster boundaries shows that there is a natural convention for orienting them: We declare that the loops encircle black clusters in a counterclockwise direction, as in Fig. 2. The fact that the loops in Eq. (1) are "secretly" oriented has crucial consequences for the continuum theory [17]. Viewing the loops as cluster boundaries also shows that at $N = 1$ and $x = 1$ the above loop gas is nothing but uncorrelated site percolation on the triangular lattice, which is critical since black and white hexagons are equiprobable.

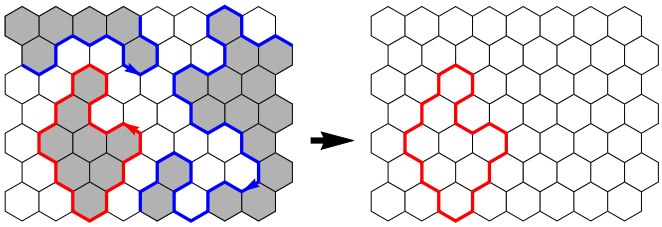

FIG. 2. Correspondence between loop gas on the honeycomb lattice (percolation) and the polymer model. A randomly chosen loop from the former is statistically equivalent to a ring polymer governed by the Boltzmann weight in Eq. (2). (All strands form closed loops; the left panel is part of a larger configuration.)

The above theory at $N = 1$ describes a soup of many loops, rather than a single polymer. However, at $N = 1$ and $x = 1$ there is a well-known mapping to a partition function for the latter. Crudely, the point is that a loop picked at random from the loop gas (Fig. 2) is statistically equivalent to a ring polymer with certain interactions. (These interactions are local, due to the short-range correlations in the percolation problem.) Taking a system on a finite lattice, say, with periodic boundary conditions, the polymer partition function is

$$Z_{\text{polymer}} = \sum_{\substack{\text{polymer}\\ \text{configs}}} \left(\frac{1}{2}\right)^{\text{No. of hexagons visited by polymer}}. \qquad (2)$$

The weight may be seen as the combination of a weight per unit length together with attractive interactions of a certain kind (for a given length, more compact configurations are favored since they visit fewer hexagons). These interactions are such that the polymer is tuned to the collapse point. For example, the mapping to the loop gas implies that the fractal dimension of the polymer is $d_f = 7/4$ [35], which is in between the self-avoiding walk value ($d_f = 4/3$) and the value in the collapsed phase ($d_f = 2$).

The introduction of the loop colors in Eq. (1) gives a useful way of formalizing the connection between the gas of many loops and the polymer model [30]. We write

$$N = 1 + n \qquad (3)$$

and label the N possible colours for each loop by

$$a = 0, \ldots, n. \qquad (4)$$

We distinguish loops of color $a = 0$, which we refer to as background loops, from loops of color $a = 1, \ldots, n$, which we refer to as polymers. Informally, the point is that integrating out the background loops in a configuration with a single polymer gives the desired weight in Eq. (2). Furthermore, since each polymer has a statistical weight n, we can use a replicalike limit $n \to 0$ to isolate configurations with a single polymer. To reiterate, $N \to 1$ is fugacity for loops in the loop gas, while $n \to 0$ is the fugacity for polymers.

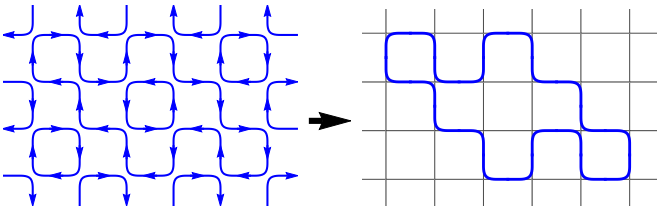

FIG. 3. Correspondence between the completely packed loop model and polymer model. A randomly chosen loop from the former is statistically equivalent to a ring polymer governed by the Boltzmann weight in Eq. (8). [The orientations assigned to the links (left) are fixed, not fluctuating degrees of freedom.]

Explicitly, expanding $Z_{\text{honeycomb}}$ in n gives

$$Z_{\text{honeycomb}} = \left(\sum_{\substack{\text{background} \\ \text{loop configs}}} 1 \right) + n \left(\sum_{\substack{\text{configs with} \\ \text{1 polymer}}} 1 \right) + \cdots . \quad (5)$$

The first term is proportional to the sum over percolation configurations. Absorbing this trivial constant into the definition of $Z_{\text{honeycomb}}$ and performing the sum over the configurations of the background loops in the second term, we obtain

$$Z_{\text{honeycomb}} = 1 + n Z_{\text{polymer}} + \cdots . \quad (6)$$

Z_{polymer} is the first θ-point model we will consider. More generally, we also wish to consider the space of models close to this one. We will show that the DS θ-point behavior of this model is robust: The universality class of the collapse transition remains the same if we slightly change the form of the interactions. If we wish to consider the introduction of crossings this model is not very convenient, so we will be led to consider models on the square lattice.

B. Square lattice model (model *T*)

The second model derives from the well-known completely packed loop model on the square lattice (Fig. 3). Configurations are generated by choosing the pairing of links at each node (Fig. 4). The partition function is

$$Z_{\text{CPL}} = \sum_{\substack{\text{loop} \\ \text{configs}}} N^{\text{No. of loops}} = \sum_{\substack{\text{colored} \\ \text{loop configs}}} 1. \quad (7)$$

Loops in this model always turn at nodes. Therefore, if we assign fixed directions to the links of the square lattice by the arrow convention in Fig. 3, the loops acquire consistent orientations. (This oriented square lattice is known as the L lattice.) The loop gas is again in the dense phase, and the loops have the same universal properties as those in the honeycomb model. When $N = 1$ there is again a correspondence with a polymer model. A loop picked at random from the gas (Fig. 3, right) is governed by the effective polymer partition function

$$Z_{\text{model } T} = \sum_{\substack{\text{polymer} \\ \text{configs}}} \left(\frac{1}{2} \right)^{\text{No. of nodes visited}} . \quad (8)$$

The configurations appearing in the partition function are constrained to Turn at each node (Fig. 3, right), so we refer to this as model T. Later on we will relax this constraint. This model is well known [5,9,18,36].

The large-scale properties of a polymer ring in this model are identical to those in Eq. (2). However, we will not refer to this model as being in the DS universality class. This is because some universal properties differ, despite the fact that the same field theory applies in each case. This will be discussed below (Sec. V D). In particular, correlators for *open* chains are related to field theory correlators in slightly different ways in the two cases and as a result the entropic exponent γ for the partition function of an open chain takes a different value in model T and the honeycomb model [9,36].

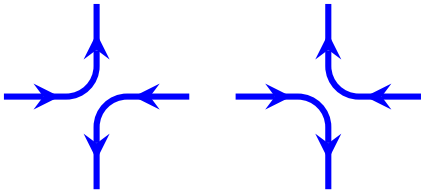

FIG. 4. In the completely packed loop model each node has two possible configurations. (Link orientations are fixed by the convention in Fig. 3, left.)

C. The σ model for loop gases

Loop gases with fugacity N map to nonlinear σ models for N-component fields. The best-known example of this is the relationship between the honeycomb model of Eq. (1) and the high-temperature expansion of a modified O(N) lattice magnet [33,34]. However, the true global symmetry in these models is in fact larger, namely, SU(N), as a result of the fact that the loops do not cross [17,28]. In the next subsection we will see explicitly how this arises on the lattice. Heuristically, the key point is that in the models without crossings there are natural prescriptions, discussed above, for orienting the loops. The appearance of oriented loops signals that we should be working with a *complex* N-component field

$$\vec{Z} = (Z_0, \ldots, Z_{N-1}), \quad |\vec{Z}|^2 = 1. \quad (9)$$

One may think of this as follows. If we treat the 2D space as the Euclidean space-time for a 1+1D quantum problem, the theory with the complex field describes N colors of charged bosons, labeled $a = 0, \ldots, n$. The loops are simply the world lines of these bosons. In addition to the color index labeling the species, they carry an orientation that distinguishes particles from antiparticles.

The appropriate field theory for $\vec{Z}$ turns out to have the U(1) gauge symmetry $\vec{Z}(x) \to e^{i\phi(x)} \vec{Z}(x)$ [17,20,28]. (This is related to the fact that the orientation of a given loop in Fig. 3, left, is not free to fluctuate.) Therefore, it is useful to introduce the gauge-invariant field

$$Q_{ab} = Z_a Z_b^* - N^{-1} \delta_{ab}. \quad (10)$$

The traceless matrix Q parametrizes complex projective space, CP^{N-1} (and satisfies a nonlinear constraint, since $|\vec{Z}|^2 = 1$). The field theory describing the nonintersecting loop gas is the

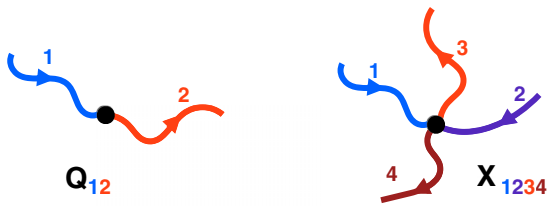

FIG. 5. Schematic representation of some components of the two- and four-leg operators. These operators enforce vertices with strands of specified color emanating from them.

CP^{N-1} nonlinear σ model with a topological Θ term [17]:

$$\mathcal{L}_{\mathrm{CP}^{N-1}} = \frac{K}{2}\,\mathrm{tr}(\vec{\nabla} Q)^2 + \frac{\Theta}{2\pi}\epsilon_{\mu\nu}\,\mathrm{tr}\, Q\vec{\nabla}_\mu Q\vec{\nabla}_\nu Q. \qquad (11)$$

The coefficient Θ is equal to π [37]. This σ model flows, for sufficiently large bare stiffness K and for $N \leqslant 2$, to a nontrivial fixed point that describes the dense phase of the loop gas. The regime of interest to us is $N = 1 + n$, with n infinitesimal, so N must be treated as variable in the spirit of the replica trick. An alternative is to formulate a supersymmetric version of the σ model [17]; for our purposes the two approaches are equivalent.

The σ model captures correlation functions in the loop gases and by extension in the polymer models. It is useful to keep in mind the heuristic picture of the loops as world lines of $\vec{Z}$. So, for example, the operator $Q_{12} = Z_1 Z_2^*$ absorbs an incoming world line of color index $a = 1$ and emits an outgoing one of color $a = 2$, as illustrated in Fig. 5. In the next subsection we will make this more precise on the lattice.

Recall the distinction between background loops ($a = 0$) and polymer loops ($a = 1, \ldots, n$). We make a corresponding splitting of the components of $\vec{Z}$,

$$\vec{Z} = (Z_0, \vec{Z}_\perp), \quad \vec{Z}_\perp = (Z_1, \ldots, Z_n), \qquad (12)$$

with world lines of Z_0 and $\vec{Z}_\perp$ corresponding to background and polymer loops, respectively. $\vec{Z}_\perp$ has a vanishing number of components in the limit of interest, namely, $N \to 1$ or $n \to 0$. "Watermelon" correlation functions for the polymer may be expressed in terms of $\vec{Z}_\perp$.

The field theory $\mathcal{L}_{\mathrm{CP}^{N-1}}$ is appropriate to the polymer model Z_{polymer}, which derives from a loop gas in which the polymer and background loops are on exactly the same footing. However, a general perturbation of the Boltzmann weight for the polymer will, when translated back to the loop gas, break the symmetry between the polymer and the background loops [12]. Correspondingly, the Lagrangian will be perturbed by operators $\mathcal{O}_i$ which reduce the SU(N) symmetry to something smaller:

$$\mathcal{L} = \mathcal{L}_{\mathrm{CP}^{N-1}} + \sum_i \lambda_i \mathcal{O}_i. \qquad (13)$$

D. Lattice field theories

To make the connection between the loop gases and the σ model concrete, we will need lattice field theories that (i) map exactly to the loop gas and (ii) turn into the σ model upon coarse-graining [12,21].

1. Completely packed model

First consider the completely packed model on the square lattice, Eq. (7) (see Refs. [12,22,39] for more detail). We take a model with N-component complex vectors $\vec{Z}$ located on the links of the lattice, with fixed length $|\vec{Z}|^2 = N$. The Boltzmann weight is a product over terms for each node. Denoting the two outgoing links at a given node by o, o' and the two incoming links by i, i', we have

$$Z_{\mathrm{CPL}} = \mathrm{Tr} \prod_{\mathrm{nodes}} [(\vec{Z}_o^\dagger \vec{Z}_i)(\vec{Z}_{o'}^\dagger \vec{Z}_{i'}) + (\vec{Z}_{o'}^\dagger \vec{Z}_i)(\vec{Z}_o^\dagger \vec{Z}_{i'})], \qquad (14)$$

where Tr denotes the integral over the $\vec{Z}$s with the length constraint. Note that the two terms at each node correspond to the two ways of pairing up the links at that node shown in Fig. 4. Therefore, expanding out the product over nodes generates the sum over loop configurations, with each loop decorated with a product of $\vec{Z}^\dagger \vec{Z}$ factors. In a loose notation where the links on a given loop are denoted by $1 \ldots, \ell$ as we go around the loop in the direction of its orientation, we have

$$Z_{\mathrm{CPL}} = \sum_{\mathrm{configs}} \prod_{\mathrm{loops}} \mathrm{Tr}(\vec{Z}_1^\dagger \vec{Z}_\ell) \cdots (\vec{Z}_3^\dagger \vec{Z}_2)(\vec{Z}_2^\dagger \vec{Z}_1). \qquad (15)$$

Using $\mathrm{Tr}\, Z_a Z_b^* = \delta_{ab}$ to integrate out the $\vec{Z}$, we find that each loop has a single color index a that is conserved along its length. Therefore, Eq. (14) is equal to the partition function of the loop model, $Z = \sum_{\text{colored loop configs}} 1$.

The above theory has SU(N) global symmetry and U(1) gauge symmetry under independent phase rotations on each link, $\vec{Z}_\ell \to e^{i\phi_\ell}\vec{Z}_\ell$. One may show that the continuum limit of this lattice field theory is the CP^{N-1} Lagrangian of Eq. (11) with $\Theta = \pi$ [40]. This agrees with the field theory derived for the loop model by first taking an anisotropic limit that maps it to a quantum spin chain [17].

Inserting operators on the links modifies the graphical expansion. For example, if we insert Q_{cd} on a link, the integral over $\vec{Z}$ on that link is modified from $\mathrm{Tr}\, Z_a Z_b^* = \delta_{ab}$ to $\mathrm{Tr}\, Z_a Z_b^* Q_{cd} = (n\delta_{ad}\delta_{bc} - \delta_{ab}\delta_{cd})/(n+1)$. It follows that inserting Q_{12} forces the color of the incoming part of the strand passing through the link to be 1 and the color of the outgoing segment to be 2. The correlation function $\langle Q_{12}(l) Q_{21}(l)\rangle$ then contains only configurations in which the links l and l' lie on the same loop. That is, Q_{ab} is a lattice two-leg operator.

2. Honeycomb model

The lattice field theory for the honeycomb model given in Ref. [20] (see also Ref. [28]) is very similar to the lattice magnet of Nienhuis and co-workers [33,34], but includes a U(1) gauge field. The role of this gauge field is to fix the relative orientations of the loops in accordance with the cluster boundary convention in Fig. 2, which leads to adjacent loops being oppositely oriented.

The spins of the lattice magnet are again complex vectors $\vec{Z} = (Z_0, \ldots, Z_{N-1})$ with length $|\vec{Z}|^2 = N$, but are now located at the sites i of the honeycomb lattice. The gauge field is an angular degree of freedom $U_{ij} = e^{ia_{ij}}$, which is located on the links (with $U_{ij} = U_{ji}^*$). The partition function we

need is

$$Z = \mathrm{Tr} \prod_{\text{hexagons } H} \left(1 + \prod_{\langle ij \rangle \in \mathrm{H}} U_{ij} \right) \prod_{\langle ij \rangle} (1 + x U_{ij} \vec{Z}_i^\dagger \vec{Z}_j + \text{c.c.}). \tag{16}$$

In the product $\prod_{\langle ij \rangle \in H} U_{ij}$, the links are oriented counterclockwise around the hexagon.

This model allows a graphical expansion similar to the previous one, showing its equivalence to $Z_{\text{honeycomb}}$ in Eq. (1). The graphical expansion involves not only loops [which come from the expansion of the product over links in Eq. (16)] but also shaded hexagons (which come from the expansion of the product over hexagons). The shaded hexagons make up the black clusters in Fig. 2 and the loops are glued to the boundaries of these clusters once we integrate over U_{ij}.

Equation (16) has the same gauge and global symmetries as Eq. (14). The microscopic field content is different because of the presence of the fluctuating gauge field, but this does not change the coarse-grained Lagrangian. [In fact, the continuum σ model (11) admits an equivalent formulation with $\vec{Z}$ coupled to a dynamical continuum gauge field A_μ. In this formulation the Θ term is simply proportional to the integral of the flux $\epsilon_{\mu\nu} \vec{\nabla}_\mu A_\nu$. This term arises naturally from coarse graining the first term in Eq. (16) [20].]

III. RELEVANT OPERATORS

Having completed our review of the formalism, we now return from the lattice to the continuum field theory (Sec. II C) in order to introduce the operators that will be pertinent to our discussion of stability. These will be components of the two- and four-leg operators (Fig. 5), both of which are relevant if added to the action.

As noted above, the two-leg operator is essentially the matrix Q_{ab} defined in Eq. (10) [41]. It transforms in the adjoint representation of SU(N), which has dimension $N^2 - 1$, and its RG eigenvalue in the $N = 1$ limit is $y_2 = 7/4$. A lattice version of this operator can, for example, be defined on a link of the completely packed loop model as discussed above [42], and its two point function is proportional to the probability that two links lie on the same loop.

The four-leg operator comes in two types, as we will discuss in Sec. VII A, with different behavior under parity (spatial reflections). Only the even-parity operator is important for the RG flows we consider, since spatial symmetry prevents the odd-parity operator from appearing in the action. This even-parity four-leg operator $X_{aa'bb'}$ is essentially $Z_a Z_{a'} Z_b^* Z_{b'}^*$, with trace terms subtracted to ensure it transforms irreducibly under SU(N) [19]:

$$X_{aa'bb'} = Z_a Z_{a'} Z_b^* Z_{b'}^* - \frac{\delta_{ab} Z_{a'} Z_{b'}^* + (3 \text{ terms})}{N+2} + \frac{\delta_{ab}\delta_{a'b'} + \delta_{ab'}\delta_{a'b}}{(N+1)(N+2)}. \tag{17}$$

The operator $X_{aa'bb'}$ is symmetric under $a \leftrightarrow a'$ and under $b \leftrightarrow b'$. Graphically, it is a four-leg vertex with incoming directed lines of color a,a' and outgoing directed lines of color b,b' (Fig. 5). A lattice version of this operator may be written down in the completely packed loop model, but will not be needed here. X has the RG eigenvalue $y_4 = 3/4$ at $N = 1$ and it forms an irreducible representation of SU(N) whose dimension is

$$d_X = \frac{N^2(N-1)(N+3)}{4}. \tag{18}$$

Q and X above are the only operators in the σ model that are invariant under both spatial rotations and parity and are relevant at $N = 1$. We defer the demonstration of this to Sec. VII. In order to show that no further relevant or marginal operators can appear as perturbations to the action, we identify the full set of operators with dimensions $x = x_4$ and $x = 2$: We confirm that this set is complete using the results of Read and Saleur on the counting of states in the spectrum of the supersymmetric σ model [17]. We find that all other perturbations are excluded by parity symmetry.

When we perturb the Boltzmann weight for the polymer the global symmetry will be reduced to a subgroup of SU(N), and the operators above will split into more than one representation of the reduced symmetry. Four operators will play a role in the discussion of RG flows below:

$$Q_{00} = -\sum_{a=1}^{n} Q_{aa}, \quad \mathcal{A} \equiv -X_{0000} = -\sum_{a,b=1}^{n} X_{abab},$$
$$\mathcal{S} \equiv -\sum_{a=1}^{n} X_{aa00} + \text{c.c.}, \quad \mathcal{C} \equiv -\sum_{a,b=1}^{n} X_{aabb}. \tag{19}$$

The effects of these perturbations are summarised heuristically in Fig. 6 and will be explained in the following sections.

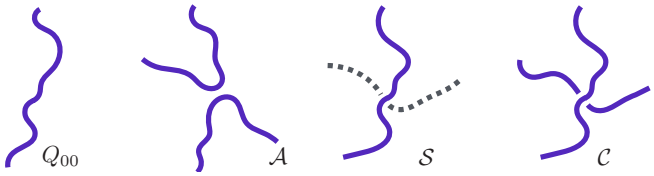

FIG. 6. Schematic interpretation of the four relevant perturbations in Eq. (19). Q_{00} is the leading perturbation when we change the variable conjugate to the length of the polymer. $\mathcal{A}$ can be independently tuned by modifying the strength of attraction between nearby monomers. In the loop gas interpretation, $\mathcal{S}$ is a crossing between the polymer strand and a background strand (dashed line). In generic models in the DS universality class this perturbation does not play any role. However, in model T, which is fine-tuned in a certain sense, $\mathcal{S}$ corresponds to introducing nodes where the polymer does not turn (straight segments). The perturbation $\mathcal{C}$ arises when the polymer is allowed to cross itself.

IV. PERTURBING THE HONEYCOMB MODEL

First we show that the collapse transition in the honeycomb polymer model (2) remains in the DS universality class even when the interactions between monomers are slightly perturbed. (For the present we retain the constraint that the polymer cannot cross itself.)

Recall that the polymer is a world line of the field $\vec{Z}_\perp$. Therefore, we might expect that the effect of changing the interactions between monomers will simply be to add local interactions for $\vec{Z}_\perp$. Neglecting terms with derivatives, the

perturbed Lagrangian will then be of the form

$$\mathcal{L} = \mathcal{L}_{\mathrm{CP}^{N-1}} + V(|\vec{Z}_\perp|^2), \tag{20}$$

where V is an arbitrary potential. This expectation is correct (though see the next section). In Appendix A we confirm this explicitly for an arbitrary perturbation of the interactions, using the lattice field theory in Eq. (16).

The perturbation in Eq. (20) results in the symmetry breaking (recall $N = 1 + n$)

$$\mathrm{SU}(N) \to \mathrm{U}(n). \tag{21}$$

The remaining U(n) rotates the components of $\vec{Z}_\perp$.

As discussed in the previous section, any relevant perturbations allowed by spatial symmetry must be sums of components of Q or X, with RG eigenvalue $y_2 = 7/4$ or $y_4 = 3/4$, respectively. Taking into account U(n) symmetry, there are only two linearly independent possibilities

$$Q_{00} = -|\vec{Z}_\perp|^2 + n/N, \quad \mathcal{A} \equiv -\sum_{a,b=1}^{n} X_{abab}. \tag{22}$$

These appear when, for example, we change the monomer fugacity (weight per unit length) or the strength of self-attraction ($\mathcal{A}$ stands for attraction). Increasing the monomer fugacity favors links of color $a = 1, \ldots, n$ over those of color $a = 0$ and so naturally generates a positive mass for Z_0 or equivalently a negative mass for $\vec{Z}_\perp$. In a coarse-grained picture, increasing the polymer's self-attraction means increasing the weight for a meeting of four polymer legs, explaining the appearance of the four-leg operator X_{abab} with color indices a and b both greater than zero. Note that, by virtue of the tracelessness of Q and X, operators like $\sum_{a>0} Q_{aa}$, or $\sum_{a>0} X_{a0a0}$, or X_{0000} are linearly related to those above and so do not constitute independent RG directions.

Two RG-relevant directions is the right number for the Θ point. One relevant perturbation is automatically tuned to zero by taking the polymer to be long [43] and the other must be varied to reach the collapse point. Therefore, the above shows that the DS behavior is robust for nonintersecting models on the honeycomb lattice. (The generic form of such a model is given in Appendix A.) This conclusion is consistent with (and explains) early numerical transfer matrix calculations [7], which investigated several perturbations of the honeycomb Boltzmann weights and found that that the exponents remained the same to within numerical accuracy.

In order to infer that the DS universality class is robust to any local perturbations (with the exception of allowing the polymer to cross itself), we should check that there is no fine-tuning hidden in the choice of lattice. Fortunately, the lattice gauge theory representation makes clear that U(n) is retained so long as we do not allow the polymer to cross itself, regardless of the choice of lattice [44]; and so long as U(n) symmetry is retained, only the two relevant perturbations discussed above can appear in the action. Therefore, the DS universal behavior is generic so long as crossings are forbidden. Physically, only one parameter needs to be varied to reach a collapse transition in this universality class.

This resolves a longstanding theoretical question about this transition. One of the reasons why the stability of the DS fixed point has previously been a vexed question is that the traditional Coulomb gas [45] approach to loop models hides the SU(N) symmetry: This prevents us from being able to classify and count perturbations. (Other sources of confusion have included the assumption that the DS exponents are the same as those of de Gennes's tricritical O($n \to 0$) model, which we will argue is not the case, and the existence of various other solvable fixed points that were candidates for the θ point [5,46–48].) The σ model is the right formulation, as we have seen. However, even in this formulation it is easy to be misled, as we will see in a moment.

The potentially confusing point is related to the fact that the perturbation $\mathcal{S}$ does not appear in the above analysis. This corresponds to a crossing between a polymer and a background strand and effects a more drastic symmetry breaking than that in Eq. (21) (Sec. V). This perturbation does not appear when we perturb models that are truly in the DS universality class like that above, as we have seen, which is why the DS universal behavior is generic (in the absence of crossings). However, we will now see that it does appear when we perturb model T (the square lattice model with turns at every node). This additional perturbation means that model T is fine-tuned, as first argued by Blote and Nienhuis [5,18]. It also occupies a different position in the phase diagram from the true DS fixed point and strictly it should be regarded as a distinct universality class. The apparent paradox, that the two models have different numbers of relevant perturbations despite being described by the same CP^{N-1} field theory, will be resolved in the next section, where we discuss models without crossings on the square lattice. Then in Sec. VI we use the square lattice to introduce crossings, which is awkward to do on the honeycomb (though possible if we allow double occupancy of links [49]).

V. SQUARE LATTICE MODEL: PARADOX AND RESOLUTION

In the square lattice model of Eq. (8) (model T), the polymer is constrained to turn at each node. The gentlest perturbations of this model change the interactions while retaining this constraint. For this class of models, the story is the same as the previous section: The only relevant operators that arise are Q_{00} and $\mathcal{A}$ (Appendix B) and the universal behavior remains unchanged. However, an additional relevant perturbation arises if we relax the constraint of turning at every node [19].

Recall that this polymer model is related to a completely packed loop gas (Fig. 3). In the language of the loop gas, a nonturning node is a crossing between a polymer strand and a background strand. This is a vertex resembling Fig. 5 (right) where the two outgoing links have color index $a = 0$ (background) and the two incoming links have $a > 0$ (polymer), or vice versa. The new perturbation is denoted by $\mathcal{S}$,

$$\mathcal{S} = -(\mathcal{N} + \mathcal{N}^*), \quad \mathcal{N} \equiv \sum_{a=1}^{n} X_{aa00} = (Z_0^*)^2 \vec{Z}_\perp^T \vec{Z}_\perp. \tag{23}$$

The appearance of this perturbation can be confirmed directly using the lattice field theory representation of the loop gas (Appendix B). The operator $\mathcal{S}$ effects the symmetry breaking

$$\mathrm{SU}(N) \to \mathrm{O}(n), \tag{24}$$

where the remaining symmetry rotates $\vec{Z}_\perp$.

Both $\mathcal{S}$ and $\mathcal{A}$ derive from the four-leg tensor, but they are linearly independent operators. Therefore, for model T, the number of relevant directions is 3 when straight segments are allowed. This implies that model T describes a fine-tuned collapse point. That this model is fine-tuned was originally suggested by Blote and Nienhuis [5]: The above provides a precise field-theoretic version of their argument (from the decay of the appropriate correlator) that nonturning nodes should be a relevant perturbation. (The field theory formulation makes it clear that this perturbation is linearly independent of $\mathcal{A}$ at the fixed point.) Note that model T is more unstable than the honeycomb model, although they are both related to the CP^{N-1} field theory; we will discuss this below. We emphasize that, strictly speaking, model T is not in the DS universality class (Sec. V D).

The presence of a relevant additional perturbation in this model prompts various questions. First, we have already argued (Sec. IV) that all models in which the polymer does not cross itself have a U(n) symmetry. At first sight this is in conflict with Eq. (24), which says that the SU(N) symmetry of model T is broken all the way to O(n) when the model is perturbed. In fact, both statements are correct. The perturbed model T does have a U(n) symmetry, which is revealed by mapping it to a lattice field theory in a different way. However, this U(n) is not a subgroup of the SU(N) of the unperturbed model T! This subtlety arises because the two ways of mapping the model to field theory are not related by any local change of variable. In order to see this (Sec. V B) it will be convenient first to introduce a less peculiar square lattice model (Sec. V A). The latter also gives an explicit example of a model that has nonturning nodes and is truly in the DS universality class.

Second, the fact that the perturbation $\mathcal{S}$ appears in model T makes it surprising at first sight that it does not appear for models in the true DS universality class such as the honeycomb model. The result of Sec. IV is enough to show that it does not appear in those models, but we can still ask what it would mean to add it to the Lagrangian in that case. We discuss this briefly in Sec. V C: We find that for the true DS models, $\mathcal{S}$ corresponds to a nonlocal perturbation of the polymer Boltzmann weight. Thus, while the honeycomb model and model T are described by the same field theory, the mapping between operators in the field theory and in the polymer model works slightly differently in the two cases.

Third, it is natural to ask for a heuristic understanding of why model T represents a different universality class to the true DS models. Here the key player is the Ising variable of Refs. [5,18]. Model T occupies a different position in the phase diagram from the true DS fixed point: It represents a transition into a collapsed phase with an additional lattice dependent Ising order. Analogously, it is natural to ask how models with the same field theory description can be in different universality classes. We discuss this in Sec. V D. We also briefly discuss the possibilities for what fixed point model T flows to when it is perturbed with nonturning nodes.

A. A less peculiar square lattice model

This section introduces a square lattice model that is in the (true) DS universality class. This model does not have model T's peculiar feature of turning at every node. It lends itself to a different field theory mapping that sheds light on the above issues.

The mappings between polymer models and field theories in Sec. II D started by relating the former to a gas of oriented loops. We have seen two types of convention for doing this. For the honeycomb model, the loops were oriented by viewing them as cluster boundaries, while for the completely packed model the loops were oriented by assigning fixed directions to the links. The fact that we could consistently orient the loops in this way relied on fine-tuning in model T (the absence of nonturning nodes).

We may also consider loop gases on the square lattice that are not completely packed and associated polymer models. In fact, since the background loops are not physical degrees of freedom, we may be able to map a given polymer model to a loop gas (and then to a lattice field theory) in more than one way, and one mapping may reveal a symmetry which is hidden by the other.

For a specific polymer model (with nonturning nodes) that is demonstrably in the DS universality class, let us consider the natural square lattice analog of the honeycomb model. Again we begin with an $N = 1$ loop model in which the loops can be viewed as cluster boundaries (see Fig. 7). The only difference with the honeycomb case is that now two clusters can meet at a corner. In this case there are two possible ways to connect the cluster boundaries (similar to Fig. 4), which means that a given configuration of shaded faces can correspond to more than one loop configuration.

The loop gas partition function is

$$Z = \sum_{\substack{\text{loop}\\ \text{configs}}} \alpha^{-T} N^{\text{No. of loops}}. \tag{25}$$

Nodes may be visited twice but the loops do not cross. Here T denotes the number of twice-visited nodes and α is a constant that we take to be $\alpha = 1/2$ when $N = 1$. The loop gas then maps to a percolation problem in which we (i) color the faces black or white with equal probability and (ii) make random binary choices for how to connect the cluster boundaries at each twice-visited node. The fact that the weight of a percolation configuration is shared equally between the two ways of connecting up the clusters at twice-visited nodes gives $\alpha = 1/2$. (The nonstandard definition of clusters here means that this is different from conventional site percolation on the square lattice. Symmetry between black and white ensures that the present model is critical.)

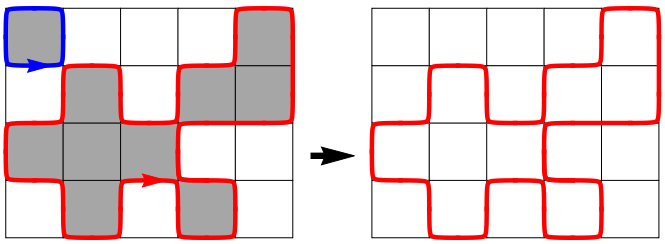

FIG. 7. A second square lattice loop gas (left) that maps to a polymer model (right). In this model the polymer can visit a site twice, but cannot cross. Nonturning nodes are allowed.

We can relate this loop gas to a polymer model in the usual way (Sec. II A). The precise polymer interactions, given in Appendix C, are cumbersome to write down, but are perfectly local. The relation with percolation ensures that the polymer is right at its collapse point and in the DS universality class.

We may also map this model to a lattice gauge theory in an identical manner to the honeycomb model. The continuum limit is again the CP^{N-1} model at $\Theta = \pi$. As for the honeycomb model, this lattice gauge theory representation can be generalized to allow an arbitrary local perturbation to the Boltzmann weight. This is explained in Appendix C. A convenient intermediate step is to first map the problem to a loop model on a modified lattice, in which each node is replaced by a cluster of trivalent nodes: This ensures that the conformation is specified uniquely by which links are visited, making it easy to write down the interactions in the lattice field theory language.

The conclusions about stability confirm what we already know from Sec. IV. So long as the polymer cannot cross itself, U(n) symmetry is retained and DS universal behavior is robust against (sufficiently weak) perturbations.

B. Hidden U(n) symmetry in noncrossing models

By a suitable (noninfinitesimal) deformation of the lattice gauge theory representation introduced for the model above, we can in principle describe any noncrossing polymer model on the square lattice while retaining U(n) symmetry (Appendix C). This includes model T perturbed by straight segments.

How do we reconcile this with the SU(N) $\to$ O(n) symmetry breaking that we found in Sec. V? Both results are correct: The symmetry depends on the way in which we map the polymer model to field theory or equivalently on the way in which we introduce the background loops. For model T, the advantage of the original representation (based on the completely packed loop model) is that it makes the SU(N) symmetry of the unperturbed model T manifest. The advantage of the alternative representation is that it makes the U(n) symmetry of the perturbed model manifest. However, this U(n) should not be regarded as a subgroup of the SU(N) of the unperturbed model T, since the two representations involve distinct sets of fields (not related by any local transformation). For this reason the alternative representation does not make the SU(N) symmetry of model T manifest. On the other hand, it does reveal another SU(N) symmetry at a different point in parameter space, namely, for the less peculiar model of the previous section. The common feature of the points at which an SU(N) symmetry exists is that they map to $N = 1$ loop gases. (However, differences between these loop gases lead to differences in the polymer models, which we touch on in Secs. V C and V D.)

Retaining U(n) symmetry is enough to ensure that models (without crossings) that are sufficiently close to the DS fixed point will flow to it. This includes, for example, any model that is sufficiently close to the less peculiar model [in which U(n) is enlarged to SU(N) microscopically]. This does not of course imply that all models with U(n) symmetry lie in the basin of attraction of the DS fixed point. Therefore, we cannot assume that when we perturb model T with nonturning nodes it will flow to the true DS fixed point. It may flow to a different fixed point that is also stable. We discuss this briefly in Sec. V D.

C. Absence of perturbation $\mathcal{S}$ at the (true) DS fixed point

As a concrete instance of the DS universality class let us take the less peculiar model on the square lattice (Sec. V A). When mapped to field theory appropriately, this is seen to have SU(N) symmetry. This is broken down to U(n) when the model is perturbed and we have seen explicitly that the perturbation $\mathcal{S}$ does not arise. However, what happens if we insist on adding this operator to the Lagrangian?

The perturbation $\mathcal{S}$ corresponds to a crossing between a polymer strand and a background strand. In the loop gas this is a perfectly local perturbation. However, it corresponds to a nonlocal perturbation of the polymer model. To see this, consider (for simplicity) a polymer loop in the shape of a large square with sides of length L. Let the weight of this configuration in the polymer partition function be $\mathcal{W}(\lambda_{\mathcal{S}})$, where $\lambda_{\mathcal{S}}$ is the weight associated with a crossing between polymer and background strands. We may easily check (using the relation between the loops and percolation) that at small $\lambda_{\mathcal{S}}$ and large L,

$$\mathcal{W}(\lambda_{\mathcal{S}}) = \mathcal{W}(0)\big[1 + O\big(\lambda_{\mathcal{S}}^2 L^2\big)\big]. \tag{26}$$

The leading correction is $O(\lambda_{\mathcal{S}}^2 L^2)$ for a simple reason: If a background strand enters the polymer loop, it must also leave (giving two $\lambda_{\mathcal{S}}$ insertions) and there are $O(L)$ choices for both the entry point and the exit point. However, we may easily check that a Taylor expansion of this form cannot arise if $\mathcal{W}(\lambda_{\mathcal{S}})$ is a local Boltzmann weight for the polymer, i.e., a function of the schematic form $\mathcal{W}(\lambda_{\mathcal{S}}) = \exp\sum_{\vec{r}} C_r(\lambda_{\mathcal{S}})$, where $C_{\vec{r}}$ is a local term in the Hamiltonian that depends on some finite region around position $\vec{r}$. Expanding this in $\lambda_{\mathcal{S}}$ gives

$$\mathcal{W}(\lambda_{\mathcal{S}}) = \mathcal{W}(0)\Bigg\{1 + \lambda_{\mathcal{S}} \sum_{\vec{r}} C'_{\vec{r}}(0) + \frac{\lambda_{\mathcal{S}}^2}{2}\Bigg[\Bigg(\sum_{\vec{r}} C'_{\vec{r}}(0)\Bigg)^2 + \sum_{\vec{r}} C''_{\vec{r}}(0)\Bigg] + \dots\Bigg\}. \tag{27}$$

Generically, the leading correction is $O(\lambda_{\mathcal{S}} L)$. We see that it vanishes only if the $O(\lambda_{\mathcal{S}}^2 L^2)$ term also vanishes, so an expansion of the form (26) is not possible for a local Hamiltonian.

D. Why is model T different?

Recall that for models on the square lattice, we may define an Ising variable associated with the polymer [18]. The following definition is equivalent to that of Ref. [18]. We consider a single polymer loop, which we take to be consistently oriented along its length. On each link we can then compare the polymer's orientation with the fixed link orientation defined by the L lattice (Fig. 3). We define the Ising-like variable σ_ℓ on link ℓ to be $+1$ if the two orientations agree and -1 if they disagree. As we go along the polymer, the domain walls in σ are precisely the nonturning nodes.

The role of σ is simplest in the phase in which the polymer is dense (but not necessarily completely dense), accessed by

increasing its length fugacity beyond the critical value. Since the polymer visits a finite fraction of the sites of the lattice we can define a coarse-grained Ising spin $\sigma(r)$, and the dense phase can be subdivided into two types, depending on whether σ is ordered or disordered [5,18]. The same is true of the collapsed phase (the collapsed polymer is essentially a bubble of the dense phase, surrounded by the vacuum).

For model T, σ_ℓ is perfectly ordered along the length of the polymer, while for models in the true DS universality class, σ_ℓ is disordered. Heuristically, this order in σ_ℓ is the reason that model T has an additional RG relevant direction, which corresponds to allowing σ to fluctuate. The order in σ also implies that model T exists in a different part of the phase diagram to the generic θ point [5,18]. Model T describes the transition between the extended phase and the Ising-ordered collapsed phase with $\langle\sigma\rangle \neq 0$ (which is what we access by perturbing model T with an additional attraction [50]). For the models in the DS universality class, however, infinitesimal perturbations will instead lead to the Ising-disordered collapsed phase with $\langle\sigma\rangle = 0$. The very existence of the Ising-ordered phase is of course a lattice artifact [5,18]. (Similar phases [51] are also seen in hierarchical lattice models [52].)

In what sense are the universal properties of model T different from those of the true DS point? It shares the same field theory description and many of the same exponents (the watermelon exponents for an even number of legs are the same). The correlations of the Ising order parameter σ are one difference. More importantly, the exponent γ governing the scaling of the partition function for an open chain is different for the two fixed points [36]. Both fixed points are described by the CP^{N-1} model, but in order to fully specify the universality class we need some additional information about how correlators in the CP^{N-1} model map to correlators for the polymer. This differs slightly for model T since the mapping arises from a completely packed loop gas. We have seen an example of this in Sec. V C, where an operator in CP^{N-1} is mapped to a local object for the polymer in one case but not the other. The interpretation of the polymer one-leg operator in terms of CP^{n-1} is also different in the two cases, reflecting the well-known fact that in model T a one-leg operator for the polymer corresponds to a two-leg operator in the loop gas.

As an aside, let us consider a simpler example of the fact that the same field theory can be compatible with two slightly different universality classes. These are the dense polymer phases with and without Ising order. Here order for σ has a more straightforward meaning than at the collapse point, since the polymer visits a finite fraction of the links on the lattice. This case is also simpler because we can stick with a single mapping from the polymer to field theory instead of worrying about two.

Consider the less peculiar polymer model of Sec. V A and its mapping to the CP^{N-1} model via the incompletely packed loop gas and lattice gauge theory. We increase the polymer's length fugacity (i.e., perturb with Q_{00}) so that we enter a dense polymer phase. The field Z_0 becomes massive and we can integrate it out [53]. This leaves us with the CP^{n-1} σ model at $n \to 0$, which is the expected description of a dense polymer [17,29]. Initially we are in the Ising-disordered dense phase $\langle\sigma\rangle = 0$.

By decreasing the weight of nonturning nodes, we may drive the transition into the Ising-ordered dense phase. In both phases, the fluctuations of σ are massive and decoupled from the CP^{n-1} sector. (The two sectors are decoupled even at the Ising transition [5].) We might think that the scaling of the watermelon correlators will be the same in the two phases, since the nontrivial CP^{n-1} sector has not undergone a phase transition [54], but this is not quite true. Consider the one-leg operator for the polymer. This acts both in the CP^{n-1} sector and in the Ising sector. In the Ising sector, the end point of an open chain should be viewed as a twist or disorder operator, i.e., the end point of a branch cut, for σ. This convention is necessary in order to ensure that the interactions between the σ values of different parts of the chain are effectively local: For example, two parts of the chain can only visit the same node if they have the same value of σ. (In the CP^{n-1} sector, we cannot write the one-leg operator simply as Z_a, since that is not gauge invariant, but one can argue from the lattice gauge theory that the one-leg operator can be incorporated as a twist defect [55].) When σ is disordered, the branch cut in σ costs only $O(1)$ free energy, so the scaling of the one-leg correlator is determined solely by the CP^{n-1} sector, giving a power law decay. However, when σ is ordered the branch cut costs a free energy proportional to its length. Therefore, we expect that in this phase the one-leg correlator scales exponentially with length and the two end points of an open chain are confined together.

Returning to the collapse transition in the regime where σ is playing a role, the nature of the RG flows between the various fixed points is not yet clear. (See Ref. [6] for a related discussion.) In particular, what universality class of collapse transition do we get when we slightly perturb model T with nonturning nodes? *A priori* there are two possible scenarios.

(i) We could flow from model T to the true DS universality class. This would be rather unusual, because it would be a flow from one fixed point described by CP^0 to another fixed point also described by CP^0, with the interpretation of the background loops changing during the flow. In this scenario, the perturbation would destroy the Ising order along the length of the polymer, but would leave the statistics of a large ring unchanged. The statistics of an open chain would change, since the exponent γ is different in the two cases. This scenario would leave the role of the branch 3 fixed point mentioned below somewhat mysterious, however.

(ii) We could flow from model T to a third universality class, denoted by U. Blöte and Nienhuis suggested that this scenario occurred and that U should be the branch 3 fixed point for which exact results are available [5,18,48]. This critical point has been revisited very recently by Vernier *et al.* and has been shown to have an unusual scaling limit [6]. In this scenario the presence of incipient Ising order then gives a natural explanation for why U is different from the generic DS behavior [5,6,18].

Note that we have already ruled out a third scenario, namely, that the flow is from the Blöte-Nienhuis fixed point to the fixed point of model T. The ISAT multicritical point, which allows crossings and is described by RP^{N-1} rather than CP^{N-1}

provides a simpler setting for investigating some of the issues of Ising ordering [12,56].

We emphasize that these questions about model T, while fascinating, are only indirectly relevant to our basic topic of the generic collapse behavior. From this point of view, the possibility of Ising order in the collapsed phase is a lattice artifact. The true DS fixed point is robust and the Ising ordering plays no role there. We now return to questions about generic models.

VI. MODELS WITH CROSSINGS

Since in a realistic situation polymers will not be *strictly* confined to two dimensions, we expect the chain to be able to cross itself, perhaps at some energy cost (Fig. 1, right). To understand how this affects the universal behavior and also to clarify the relevance of de Gennes's tricritical O($n \to 0$) model to 2D polymer collapse, we now perturb the square lattice models by allowing crossings. (Note that a crossing is not the same as a branching [57]: The polymers we consider are always topologically linear.)

Consider either of the two models on the square lattice. The rules for orienting the strands imply that at a four-leg vertex, the two incoming strands are opposite each other and the two outgoing strands are opposite each other (see, e.g., Fig. 5). Therefore, a crossing between two polymer strands (one of color index $a > 0$ and one of color index $b > 0$) corresponds to a four-leg vertex where the two incoming links are of color a and the two outgoing links are of color b (or vice versa). The corresponding perturbation is

$$\begin{aligned}\mathcal{C} &\equiv -\sum_{a,b=1}^{n} X_{aabb} \\ &= -|\vec{Z}_\perp^T \vec{Z}_\perp|^2 + 4(N+2)^{-1}|\vec{Z}_\perp|^2 + \text{const}, \end{aligned} \tag{28}$$

as we can check on the lattice (Appendix B). On its own, this operator gives the symmetry breaking

$$\mathrm{SU}(N) \to \mathrm{O}(n) \times \mathrm{U}(1), \tag{29}$$

where the U(1) is $\vec{Z}_\perp \to e^{i\theta}\vec{Z}_\perp$. (This is not a gauge transformation, since the phase multiplies only $\vec{Z}_\perp$ and not Z_0.) If we start with model T and make a fully generic perturbation (including $\mathcal{A}$, $\mathcal{S}$, and $\mathcal{C}$), then the symmetry is broken down to O(n) in the original representation.

This symmetry is what we would originally have expected from de Gennes. The resulting RG flow away from the DS fixed point, together with the fact that noncrossing models always have a higher symmetry (despite the subtlety discussed in Sec. V B) indicates that the DS exponents are unlikely to apply to models with crossings. That is, contrary to what is often assumed, we must allow for crossings in order to see the exponents of de Gennes's tricritical O(n) model. In particular, we do not expect the θ point of the ISAW to show the tricritical O($n \to 0$) exponents.

One point should be clarified. Just as we found in the case without crossings, it is again possible to choose a lattice field theory representation in which we avoid introducing the operator $\mathcal{S}$. Then, we in fact retain an O(n) × U(1) symmetry for generic models with crossings. However, we expect that this extra U(1) can be neglected when considering the generic collapse transition. That is, we expect the latter can be described by a Lagrangian for a real vector that transforms only under O(n). Symmetries of the Lagrangian are important because they encode information about microscopic constraints on the polymer configurations; here however, the U(1) does not appear to encode any additional constraints beyond those encoded in O(n). [Such a U(1) can always be included in a model of a single polymer 'for free'. The current associated with the U(1) has a simple interpretation. We decorate the polymer with an arrow indicating the direction of U(1) current flow, using the rule that the polymer's orientation flips whenever it crosses itself. Current is conserved because each crossing has two outgoing and two incoming strands [58].]

There is a special class of perturbations of model T that introduces crossings while preserving the equivalence between polymer and background loops. (All smart walk models, which have the feature that polymer configurations can be regarded as deterministic walks in a random environment, preserve the equivalence between the polymer and the background loops. This includes the models related to percolation and the collapse point of the interacting self-avoiding trail.) The equivalence is preserved if all the four-leg perturbations have exactly equal strength:

$$\mathcal{A} + \mathcal{S} + \mathcal{C} = -\sum_{a,b=0}^{m} X_{aabb}. \tag{30}$$

The symmetry breaking is then

$$\mathrm{SU}(N) \to \mathrm{SO}(N). \tag{31}$$

The RG flow then leads to the interacting self-avoiding trail (ISAT) fixed point, which is analytically tractable. Unfortunately, it is not the generic θ point for polymers with crossings. Viewed as a description of a polymer [59], the ISAT fixed point is extremely unstable: It has an infinite number of RG-relevant perturbations that break the symmetry from SO(N) to the generic O(n) [12]. Signs of this have been seen numerically [60].

Therefore, the generic critical exponents for models with crossings remain unknown. A natural model that does not appear fine-tuned has been studied numerically in Ref. [25]. It was conjectured in Ref. [25] that the critical exponents were those of the DS universality class. This would be surprising in view of the present results. Further numerical results would be valuable.

VII. OPERATORS IN THE CP^{N-1} MODEL

Our analysis of perturbations relied on the fact that all the symmetry-allowed operators that could appear in the action were components of Q and X. In order to confirm this we must now derive some features of the operator content of the σ model (about which there is currently limited knowledge). We will see that the correspondence with the loop gas implies some surprising things about operators in this field theory. The operators we need to consider in detail are those with the dimension of the four-leg operator (Sec. VII A) and the marginal operators (Sec. VII B) (Ref. [17] shows there are no

other relevant eigenvalues in the spectrum, apart from y_2). We will see that one of the marginal operators is an interesting odd-parity version of the two-leg operator.

A. Two types of four-leg operator

Consider the operators in the field theory that correspond to four-leg operators in the loop model. These are operators whose two-point function gives the probability that r and r' (or rather small regions around r and r') are joined by four strands of loop. At $N = 1$, they have scaling dimension $x_4 = 5/4$.

In the field theory, the obvious operator of this type is $X_{aa'bb'}$ described above: the traceless part of $Z_a Z_{a'} Z_b^* Z_{b'}^*$, which is invariant under parity. Indeed, it is straightforward to check that a lattice operator with the same symmetries as X allows us to write the four-leg watermelon correlator in the loop model. (Strictly speaking, the lattice operators cannot have the full symmetry of X, since complete invariance under spatial rotations only emerges in the continuum, but this will not be important in what follows.)

Surprisingly, the σ model also contains an odd-parity four-leg operator, which we denote by Y, with the same scaling dimension [19]. In the field theory, this operator has the following symmetry:

$$Y_{aa'bb'} = i\epsilon_{\mu\nu}(Z_a \overleftrightarrow{\nabla}_\mu Z_{a'})(Z_b^* \overleftrightarrow{\nabla}_\nu Z_{b'}^*) - \mathcal{T}, \tag{32}$$

with $\mathcal{T}$ denoting trace terms. Here $A\overleftrightarrow{\nabla}_\mu B = A\vec{\nabla}_\mu B - (\vec{\nabla}_\mu A)B$. Unlike X, this tensor changes sign under parity and also under the exchanges $a \leftrightarrow a'$ and $b \leftrightarrow b'$. It is easily checked to be gauge invariant. Once the trace terms are subtracted, Y forms an irreducible representation of dimension

$$d_Y = \frac{N^2(N+1)(N-3)}{4}. \tag{33}$$

From the point of view of field theory it is surprising that this operator, which has completely different symmetry properties and a different number of derivatives, has the same scaling dimension as X. This is in fact true for all $N \leqslant 2$, as we now argue on geometrical grounds.

Consider a component $X_{aa'bb'}$ of X, with all indices distinct. Graphically, this is a vertex with incoming strands of color a and a' and outgoing strands of color b and b'. Crucially, the no-crossing constraint means that the outgoing strands are opposite each other (Fig. 8). This leaves two possibilities for the ordering of the color indices as we go around the vertex

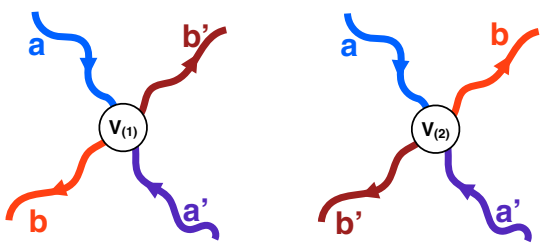

FIG. 8. Two types of four-leg vertices, distinguished by the ordering of the indices and related by parity (reflections). The even-parity and odd-parity four-leg operators correspond to the sum and difference, respectively, $V_{(1)} \pm V_{(2)}$.

counterclockwise, starting with a. Either the colors occur in the order a,b,a',b' or in the order a,b',a',b. We may in fact define two distinct operators corresponding to the two orderings, which we denote by $V_{(1)}$ and $V_{(2)}$. We may take each to be invariant under spatial rotations, but parity exchanges $V_{(1)}$ and $V_{(2)}$. The operator $X_{aa'bb'}$, which is invariant under parity, is then

$$X_{aa'bb'} = V_{(1)} + V_{(2)}. \tag{34}$$

However, there is also an odd-parity operator

$$Y_{aa'bb'} = i(V_{(1)} - V_{(2)}). \tag{35}$$

Note that Y also changes sign under either of the exchanges $a \leftrightarrow a'$ and $b \leftrightarrow b'$. These symmetry properties identify it with a component of the operator Y defined above (up to normalization).

Now consider the correlators $\langle V_{(i)}(r)V_{(j)}^*(r')\rangle$, where the conjugate operators $V_{(i)}^*$ are obtained by reversing the arrows on the strands. These correlators are sums over loop configurations in which the legs of corresponding color at the two vertices are joined. However, the key point is that, because of the no-crossing constraint, no such configurations are possible if $i = j$. We also have $\langle V_{(1)}(r)V_{(2)}^*(r')\rangle = \langle V_{(2)}(r)V_{(1)}^*(r')\rangle$. This implies

$$\langle X_{1234}(r)X_{1234}^*(r')\rangle = \langle Y_{1234}(r)Y_{1234}^*(r')\rangle. \tag{36}$$

Therefore, the scaling dimensions of X and Y are equal. This argument holds for any N, and generalizes immediately to the supersymmetric versions of the σ models (where we do not have to use the replicalike continuation from $N \geqslant 4$ to the desired value of N).

Note that the argument only shows that the two-point functions of X and Y are the same. More complex correlation functions will reveal the difference between the two operators.

The above argument applies for general N. In the special case $N = 1$ we may also see that there are additional operators with scaling dimension x_4 (i.e., beyond X) by an alternative argument. This is because the total multiplicity of each scaling dimension must vanish in the limit $N \to 1$ on general grounds [61]. The multiplicity of X is zero in this limit ($d_X = 0$), but we encounter the problem that there is another operator whose dimension x_W coincides with x_4 when $N \to 1$. This is simply the operator in the Θ term,

$$W = i\epsilon_{\mu\nu}\,\mathrm{tr}\,Q\vec{\nabla}_\mu Q\vec{\nabla}_\nu Q, \tag{37}$$

which in the percolation language drives the model off criticality. There must therefore be at least one more multiplet, whose multiplicity cancels that of W as $N \to 1$. This requirement is filled by Y, since $d_Y \to -1$ as $N \to 1$. The multiplicities of lattice operators in the spin chain formulation have also been discussed [62], reaching similar conclusions about the cancellation of multiplicities, but without clarifying the geometrical relation between X and Y or the role of parity symmetry.

References [31,32] revealed an enlarged symmetry algebra in completely packed loop models, related to quantum groups, which is independent of the phase the models are in but depends on the loops not crossing [31]. This implies larger degeneracies in the spectrum than expected from SU(N) alone. This must be the deeper explanation for the above

phenomenon. The above argument gives intuitive physical picture for this simple case.

We have found that at $N = 1$, there are three types of operators X, Y, and W, all with scaling dimension x_4. How do we know that there are not more? Fortunately, we can use the result of Read and Saleur [17] for the multiplicity of each scaling dimension in the supersymmetric (SUSY) σ model. This formula gives the total number of linearly independent operators with dimension x_4, without determining their symmetry properties. However, if we translate the above operators into the supersymmetric language (we find that the analogs of W and Y form an odd-parity indecomposable representation) and compute their multiplicities, we can check that the value for the multiplicity in Ref. [17] is saturated. This simple calculation is done in Appendix D. This shows that there are no other operators with dimension x_4 and gives an explicit identification of the supersymmetric operators contributing to the multiplicity formula.

Another unconventional feature of the CP^{N-1} model, related to the symmetry discussed in Sec. V B and presumably also a consequence of the extended symmetry of Ref. [31], is that the operator product expansions are more constrained than would be expected from symmetry. On geometrical grounds it is clear that the operator product expansion (OPE) of $\mathcal{S}$ with itself cannot generate $\mathcal{C}$, although SU(N) symmetry would allow this. Equivalently, perturbing the action with $\mathcal{S}$ does not generate $\mathcal{C}$ under the RG, consistent with the fact that models with crossings show different universal behavior from those without.

B. Marginal operators in the CP^{N-1} model

Having pinned down the relevant operators that can perturb the field theory for the polymer, we must also consider whether any marginal perturbations can appear. If present, such perturbations could destabilize the fixed point or give continuously varying exponents. However, we will argue that such perturbations are forbidden by symmetry. This also leads us to an operator that may be independently interesting.

The counting of multiplicities of Ref. [17] is a useful starting point. In the supersymmetric theory, the multiplicity of the scaling dimension $x = 2$ indicates that there are two marginal operators, each transforming in the adjoint [17]. This will also be true in the replica formalism. Let us write these so-far unknown operators as matrices A_{ab} and A'_{ab}. The question boils down to whether they are even parity or odd parity. If either operator (say, A) was even parity, we would have to worry about the possibility of A_{00} appearing in the action, just as Q_{00} can appear (Sec. IV). However, we argue here that A and A' are odd-parity operators. Therefore, spatial symmetry prevents them from appearing in the action. (One of them is also a total derivative in any case.) Our strategy is to exhibit two odd-parity marginal operators, which should therefore be identified with A and A'.

1. *Parity-odd two-leg operator related to winding angle*

First, we argue that there is an odd-parity analog of the two-leg operator, which we denote by Q^{odd}_{ab}, and that correlation functions involving this operator are related to winding angles of the critical curves. Recall that the usual

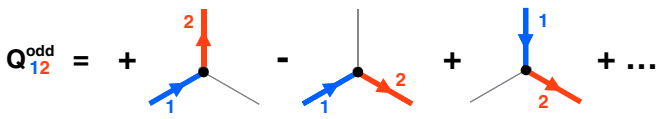

FIG. 9. Honeycomb lattice version of Q^{odd}_{12}. This emits an outgoing leg of color 2 and absorbs an incoming leg of color 1, like the two-leg operator, but also weights the configuration with a positive or negative sign according to the sign of the turning angle at the node.

two-leg operator Q_{ab} may be thought of as a vertex with an outgoing b strand and an incoming a strand [63]. We define Q^{odd}_{ab} similarly, except that we weight the vertex by a factor proportional to the signed angle through which the oriented strand turns at the vertex. This does not change the SU(N) symmetry properties of the operator; it remains in the adjoint, since it has one fundamental and one antifundamental index. However, it becomes manifestly odd parity, since reflections exchange clockwise and counterclockwise turns. For concreteness, we may take the operator to be defined at vertices of the honeycomb lattice, with left and right turns weighted by $+1$ and -1, respectively (see Fig. 9). (It is straightforward but not very illuminating to write down such operators in the lattice field theories of Sec. II D.) In the σ model, this operator has the symmetry of the traceless part of $i\epsilon_{\mu\nu}(Q\overleftarrow{\nabla}_\mu Q\overrightarrow{\nabla}_\nu Q)_{ab}$.

We must show that the scaling dimension of this operator x_{odd} is equal to 2. To see this, consider the ratio (say, on the honeycomb lattice)

$$\mathcal{R} = -\frac{\sum_{r,r'\ (r\neq r')} \left\langle Q_{12}(0)Q^{\mathrm{odd}}_{23}(r)Q_{34}(R)Q^{\mathrm{odd}}_{41}(r')\right\rangle}{\langle Q_{12}(0)Q_{21}(R)\rangle}. \tag{38}$$

The correlators in the numerator and denominator may both be written as sums over configurations with a loop passing through the sites at 0 and R. The denominator serves as a partition function for this restricted ensemble. For the correlator in the numerator, one arm of the loop (that from $0 \to R$) also passes through r, the other arm (from $R \to 0$) passes through r', and the configuration is weighted by the product of the turning angles at these two points. Altogether, $\mathcal{R}$ computes (minus) the expectation value of the product of the total turning angles of the two arms. Up to an $O(1)$ correction that is negligible at large R, this is just the square of the winding angle ψ for one of the arms (more precisely, the relevant winding angle is the sum of the winding angles about the two points 0 and R). This average is known to scale logarithmically as a result of scale invariance [64,65]:

$$\mathcal{R} \sim \langle \psi^2 \rangle \sim \ln R. \tag{39}$$

We compare this with the length scaling expected from the scaling dimensions of the operators [66] (κ stands for contact terms)

$$\mathcal{R} \sim -\frac{\int d^2r\, d^2r' \left\langle Q_{12}(0)Q^{\mathrm{odd}}_{23}(r)Q_{34}(R)Q^{\mathrm{odd}}_{41}(r')\right\rangle - \kappa}{\langle Q_{12}(0)Q_{21}(R)\rangle} \sim R^{2(2-x_{\mathrm{odd}})}. \tag{40}$$

A comparison with Eq. (40) indicates that $x_{\rm odd} = 2$, i.e., that $Q^{\rm odd}$ is marginal. Therefore, it accounts for one of the two marginal operators sought.

This argument is not specific to a particular value of N or even to the dense phase. By this reasoning, any conformally invariant fixed point for noncrossing loops should allow an odd-parity version of the two-leg operator, with dimension $x_{\rm odd} = 2$, in an appropriate field theory representation. This includes, for example, self-avoiding walks and Ising cluster boundaries.

2. *Effect of nonchirality of the currents*

Next, consider the conserved current $J_{\mu ab}$ associated with global SU(N) symmetry. Here a and b are SU(N) indices (J transforms in the adjoint) and μ is the spatial index. The current has length dimension -1 and satisfies $\vec{\nabla} \cdot \vec{J} = 0$ as a result of conservation. In a unitary conformal field theory, the current would also satisfy $\vec{\nabla} \times \vec{J} = 0$ [67]. In complex coordinates, this leads to J_z being purely holomorphic and $J_{\bar{z}}$ being purely antiholomorphic. However, it is known that this separation into holomorphic and antiholomorphic currents fails in the present nonunitary theory [17]. Equivalently, $\vec{\nabla} \times \vec{J}$ is nonzero as an operator (although it has a vanishing two-point function).

$\vec{\nabla} \times \vec{J}$ provides another marginal operator that is manifestly odd parity and transforms in the adjoint. It is distinct from the operator $Q^{\rm odd}$ defined above ($Q^{\rm odd}$ is not a total derivative, otherwise we could not use it to calculate the winding angle).

3. *Implication for polymers*

We infer that $Q^{\rm odd}$ and $\vec{\nabla} \times \vec{J}$ correspond to the only marginal scalar operators (that are local in the CP^{N-1} representation). This saturates the counting of states from Ref. [17]. It follows that there are no marginal perturbations allowed in the action for the polymer problem. This is also what we expect from numerical simulations, which do not see signs of the logarithmic drifts that would be expected for a marginally relevant or marginally irrelevant variable or the continuously varying exponents that would be expected for an exactly marginal one [11].

VIII. OUTLOOK

We have shown that the Duplantier-Saleur exponents for the θ point are generic for noncrossing polymers, as a result of symmetry enhancement under the RG flow. This resolves a longstanding question about the stability of the DS point for which previously there was only numerical evidence [11]. We have also argued that crossings induce a flow to a different universality class [and that the tricritical O(n) behavior will only be seen in this case]. Along the way we had to obtain a clearer picture of operators in the CP^{N-1} σ model. We also had to resolve some apparent paradoxes about the fine-tuned model T, which at first sight gives misleading conclusions about the robustness of the DS exponents and about the difference between models with and without crossings. The first of these issues is related to the fact that the same field theory may describe different models, but with a different relationship between polymer and field theory operators in each case. The second issue is related to the fact that the replicalike symmetry of a polymer model can be nontrivially dependent on the choice of field theory mapping.

Many exciting questions remain for the future. First, the full structure of the RG flows for a noncrossing polymer on the square lattice, in the regime where Ising order is playing a role [5], remains to be understood. Exciting progress has been made very recently on the conformal field theory of the branch 3 fixed point of Blöte and Nienhuis, which appears to be unconventional [6]. The flow away from model T may be to this fixed point [5,6]: It would be very interesting to have a heuristic understanding of this flow, from the point of view of an effective field theory obtained by perturbing the CP^{N-1} Lagrangian.

Another longstanding question concerns certain sequences of multicritical points found in supersymmetric theories and how to interpret them in terms of polymers [46,47].

Most importantly, models with crossings remain very little understood, despite the fact that a realistic model of a polymer living on a surface or in a quasi-2D geometry will likely include them. Historically such models have been neglected, perhaps because of the remarkable power of techniques like the Coulomb gas [45] and Schramm-Loewner evolution [68], which only work when crossings are forbidden. The present results motivate further examination of models with crossings. This will be necessary to understand polymer collapse in the fully generic situation and is likely to reveal novel aspects of 2D criticality [12,23,24].

Finally, there are interesting aspects of the CP^{N-1} field theory and its supersymmetric cousin [17] that deserve further study; for example it would be interesting to study the marginal operator $Q^{\rm odd}$ introduced here numerically.

Note added. Recently, work appeared on dilute loop models [69]; this addresses different questions from the present paper, but also considers a deformation of the lattice field theory for completely packed loop models [12,21,70]. Also, a pair of numerical studies of the phase diagrams of generalized square [71] and honeycomb lattice models [72] appeared. The results appear consistent with expectations from our analysis. The phase structure found in Ref. [71] seems to suggest that scenario (ii) in our Sec. V D is more likely than scenario (i) [73].

ACKNOWLEDGMENTS

I am very grateful to E. Bettelheim, J. Chalker, and J. Cardy for helpful discussions and comments. I also thank S. Vijay for useful discussions. I acknowledge the support of a fellowship from the Gordon and Betty Moore Foundation under the EPiQS initiative (Grant No. GBMF4303).

APPENDIX A: GENERIC PERTURBATIONS OF THE HONEYCOMB MODEL

We begin with the lattice field theory in Eq. (16), which maps to the polymer model $Z_{\rm polymer}$ in Eq. (2). We discuss how deforming the Boltzmann weight for the lattice field theory leads to a modified polymer model. First consider simply

inserting a factor y as follows:

$$Z(y) = \mathrm{Tr}\left\{ \prod_{\text{hexagons } H} \left(1 + \prod_{\langle ij\rangle \in H} U_{ij}\right) \times \prod_{\langle ij\rangle} [1 + U_{ij}(Z^*_{0i} Z_{0j} + y \vec{Z}^\dagger_{\perp i} \vec{Z}_{\perp j}) + \text{c.c.}] \right\}. \quad \text{(A1)}$$

In the graphical expansion, each segment of polymer loop ($\vec{Z}_\perp$ world line) now acquires a factor of y. We therefore obtain a polymer with a modified weight per unit length

$$Z_{\text{polymer}}(y) = \sum_{\substack{\text{polymer} \\ \text{configs}}} 2^{-H} y^{\text{length}}. \quad \text{(A2)}$$

Here H is the number of hexagons visited by the polymer. Making y smaller than one takes the model off criticality so that the polymer becomes of a finite typical size. Taking $y > 1$ will drive the model into the dense polymer (i.e., space-filling) phase, where the polymer's length scales with the total area of the lattice.

Varying y is a rather trivial perturbation to the Boltzmann weight. However, (A1) illustrates the basic point: Changing the polymer interactions induces local interactions in the lattice field theory, which break the symmetry from SU(N) to U(n). Next, we must check that any local perturbation to the polymer Boltzmann weight maps to a local perturbation in Eq. (A1). Let n_ℓ be the occupation number of the link ℓ in a given polymer configuration, i.e., $n_\ell = 1$ if the polymer passes through the link and $n_\ell = 0$ otherwise. The general perturbed partition function is

$$Z_{\text{polymer}}(y, J, K, \ldots) = \sum_{\substack{\text{polymer} \\ \text{configs}}} 2^{-H} \exp A, \quad \text{(A3)}$$

with

$$A = (\ln y) \sum_\ell n_\ell + \sum_{\ell,\ell'} J_{\ell,\ell'} n_\ell n_{\ell'} + \sum_{\ell,\ell',\ell''} K_{\ell,\ell',\ell''} n_\ell n_{\ell'} n_{\ell''} + \cdots.$$

Expanding the exponential in these couplings gives a sum of terms proportional to $n_{\ell_1} \cdots n_{\ell_k}$, where all the links can be taken distinct (since $n_\ell^2 = n_\ell$). Therefore, we must check that an insertion of n_ℓ corresponds to a local operator in the lattice field theory. This follows from the correspondence

$$n_{ij} \longrightarrow \frac{U_{ij}(\vec{Z}^\dagger_{\perp i} \vec{Z}_{\perp j} + \text{c.c.})}{1 + U_{ij}(\vec{Z}^\dagger_i \vec{Z}_j + \text{c.c.})}, \quad \text{(A4)}$$

which we may check by repeating the graphical expansion in the presence of n_{ij} insertions.

A crude effective action may be obtained by coarse graining Eq. (A1) or its perturbed version. We write $U_{ij} = e^{iA_{ij}}$ and expand the logarithm of the Boltzmann weight in A, in derivatives of $\vec{Z}$, and in the size of the perturbation (see, e.g., Refs. [12,20]). (For a crude picture of the perturbation terms, we may take U and $\vec{Z}$ as spatially constant; then the above formula is simply $n_{ij} \propto |\vec{Z}_\perp|^2$, so that, for example, $\exp \sum_{\ell,\ell'} J_{\ell,\ell'} n_\ell n_{\ell'}$ generates a quartic potential for $\vec{Z}_\perp$ at leading order in J.)

However, for our purposes all we need are the relevant operators that appear in the coarse-grained action, not the numerical values of the couplings. These operators are determined by symmetry and are given in Sec. IV.

Note that above we have not changed the allowed configurations for the polymer. Allowing configurations in which the polymer crosses itself (on the honeycomb lattice this can happen if, for example, we allow double occupancy of a link) introduces another relevant perturbation (see Sec. VI).

APPENDIX B: PERTURBATIONS OF THE L-LATTICE MODEL

For the square lattice model of Eq. (8) we will discuss a few illustrative deformations of the Boltzmann weight. Consider first the slightly generalized model

$$Z_{\text{polymer}} = \sum_{\substack{\text{polymer} \\ \text{configs}}} A^{\text{length}} B^{\text{No. of twice-visited nodes}}. \quad \text{(B1)}$$

This becomes model T when $A = 1/2$ and $B = 2$. To obtain this from the lattice magnet in Eq. (14),

$$Z_{\text{CPL}} = \mathrm{Tr} \prod_{\text{nodes}} e^{-S_{\text{node}}}, \qquad e^{-S_{\text{node}}} = (\vec{Z}^\dagger_o \vec{Z}_i)(\vec{Z}^\dagger_{o'} \vec{Z}_{i'}) + (\vec{Z}^\dagger_{o'} \vec{Z}_i)(\vec{Z}^\dagger_o \vec{Z}_{i'}), \quad \text{(B2)}$$

we note that each term in $e^{-S_{\text{node}}}$ of the form

$$W(\vec{Z}_i, \vec{Z}_{i'}, \vec{Z}_o, \vec{Z}_{o'}) \equiv (\vec{Z}^\dagger_o \vec{Z}_i)(\vec{Z}^\dagger_{o'} \vec{Z}_{i'}) \quad \text{(B3)}$$

can be expanded into terms that, depending on the values of the indices in the inner products, correspond either to (i) two segments of polymer passing through the node, (ii) one segment of polymer and one segment of background loop, or (iii) two segments of background loop. The weights of these possibilities can be adjusted by replacing the above with [12]

$$W(\vec{Z}_i, \vec{Z}_{i'}, \vec{Z}_o, \vec{Z}_{o'}) \to 2A^2 B(\vec{Z}^\dagger_{o\perp} \vec{Z}_{i\perp})(\vec{Z}^\dagger_{o'\perp} \vec{Z}_{i'\perp}) + 2A(Z^*_{o0} Z_{i0})(\vec{Z}^\dagger_{o'\perp} \vec{Z}_{i'\perp}) + 2A(\vec{Z}^\dagger_{o\perp} \vec{Z}_{i\perp})(Z^*_{o'0} Z_{i'0}) + (Z^*_{o0} Z_{i0})(Z^*_{o'0} Z_{i'0}). \quad \text{(B4)}$$

This is analogous to the perturbations discussed in Appendix A and breaks the symmetry down to U(n). The graphical expansion goes through straightforwardly and gives the desired modification to the polymer Boltzmann weight.

Next consider the introduction of straight segments for the polymer. This is achieved by the modification

$$e^{-S_{\text{node}}} \to e^{-S_{\text{node}}} + C R(\vec{Z}_i, \vec{Z}_{i'}, \vec{Z}_o, \vec{Z}_{o'}), \quad \text{(B5)}$$

where C is some weight and

$$R(\vec{Z}_i, \vec{Z}_{i'}, \vec{Z}_o, \vec{Z}_{o'}) = (Z^*_{o0} Z^*_{o'0})(\vec{Z}^T_{i\perp} \vec{Z}_{i'\perp}) + (\vec{Z}^\dagger_{o\perp} \vec{Z}^*_{o'\perp})\big(Z^T_{i0} Z_{i'0}\big). \quad \text{(B6)}$$

Note that this is a lattice analog of Eq. (23). This modification preserves the gauge invariance of the total Boltzmann weight. The color index of a strand is still preserved along its length,

but now the two outgoing links lie on one strand and the two incoming links lie on a different strand; the polymer and background segments cross at the node. The pattern of complex conjugation means that unitary symmetry is broken.

Finally, we allow nodes where the polymer crosses itself. This corresponds to adding to the node term a multiple of the expression

$$(\vec{Z}^{\dagger}_{o\perp}\vec{Z}^{*}_{o'\perp})(\vec{Z}^{T}_{i\perp}\vec{Z}_{i'\perp}). \tag{B7}$$

The corresponding additions to the continuum action follow on symmetry grounds. For a crude estimate of the couplings in the perturbed σ model, we can evaluate the Boltzmann weight with $\vec{Z}$ spatially constant and we see that the terms discussed in Secs. V and VI appear in the action (i.e., the logarithm of the Boltzmann weight) with the expected signs.

APPENDIX C: MORE DETAILS ON THE SECOND SQUARE LATTICE MODEL

1. Boltzmann weight for the associated polymer model

At $N = 1$, the model in Eq. (25) maps to a model for a polymer ring. Any loop drawn on the square lattice corresponds to an allowed polymer conformation so long as no edge is visited more than once, no site is visited more than twice, and the loop does not cross itself. The Boltzmann weight for a given polymer configuration is simply the probability of that loop appearing in the loop gas. This is easily evaluated using the mapping to a percolation problem

$$Z_{\mathrm{LP}} = \sum_{\substack{\text{polymer}\\ \text{configs}}} \left(\frac{1}{2}\right)^{N_F} \left(\frac{1}{2}\right)^{N_T} \prod_{k=1}^{4} \left(\frac{1}{2} + \frac{1}{2^k}\right)^{N_k}. \tag{C1}$$

Here N_F is the number of faces of the square lattice that contain at least one edge visited by the polymer, N_T is the number of sites that are visited twice by the polymer, and N_k is the number of faces that contain k sites visited by the polymer but no link visited by the polymer. This Boltzmann weight is somewhat cumbersome to write down (it has a simpler representation, described below) but it is perfectly local and the model has the advantage that, due to the mapping to percolation, we know that the statistics of the polymer ring are those of the DS point.

2. Lattice gauge theory

The loop gas in Eq. (25) maps to a lattice gauge theory identical to Eq. (16), modulo the substitution of the square for the honeycomb lattice (and square faces for hexagons). In order to consider the most general perturbations of the Boltzmann weight, however, it is convenient to map the loop gas to field theory in a slightly different way. First, we resolve the nodes as in Fig. 10, inserting a small diamond at each vertex so that the lattice becomes three-coordinated (the resulting lattice is known as the 4–8 lattice). Now we consider a straightforward percolation model in which we randomly color (all) the faces of this new lattice black or white with equal probability. This percolation problem maps to the previous one in a trivial way. Previously, there were two possibilities for how to connect the clusters when two of them met at a corner; now these two possibilities correspond to the two colorings of the diamond.

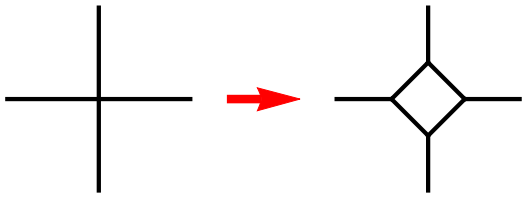

FIG. 10. It is convenient to expand the nodes of the square lattice so that the polymer becomes strictly self-avoiding not only on the links but also on the nodes.

Similarly, the loop gases are simply related at $N = 1$. The polymer partition function also takes a simple form if we regard the polymer as living on the new lattice; the weight is simply $(1/2)$ raised to the power of the number of faces visited. We emphasize that this is simply a different and more convenient representation of the same polymer model that we started with.

The loop gas on the new lattice again has a lattice gauge theory representation like Eq. (25):

$$Z = \mathrm{Tr} \prod_{\text{faces } F} \left(1 + \prod_{\langle ij\rangle \in F} U_{ij}\right) \prod_{\langle ij\rangle} (1 + U_{ij}\vec{Z}^{\dagger}_{i}\vec{Z}_{j} + \text{c.c.}). \tag{C2}$$

The product over faces now runs over both four-sided and eight-sided faces.

The reason for adopting this representation is that the polymer is now strictly self-avoiding (on sites as well as on links). This means that the configuration is completely determined by which links are occupied, which was not the case on the original square lattice (we needed to specify both which links were occupied and also how the strands were connected up at twice-visited nodes). This means that we can use the mapping between operators described in Appendix A to map any local perturbation of the polymer Boltzmann weight to a local perturbation in the lattice field theory.

APPENDIX D: OPERATORS WITH DIMENSION x_4 IN SUSY LANGUAGE

In the text we discussed three operators with dimension x_4, namely, the even-parity four-leg operator $X_{aa'bb'}$, the odd-parity four-leg operator $Y_{aa'bb'}$, and the odd-parity singlet operator in the Θ term W. The total multiplicity of these operators was $d_X + d_Y + 1$, which tends to zero in the limit $N \to 1$. This vanishing of the total multiplicity is required by general constraints on the spectrum of theories with central charge zero [61]. (In fact the multiplicities should vanish separately in the odd-parity and even-parity sectors, $d_X \to 0$ and $d_Y + 1 \to 0$.) This strongly suggests that X, Y, and W form the full set of operators with dimension x_4. To make absolutely sure, we translate these statements into the language of the SUSY σ model, where we can use the results of Ref. [17] for the counting of states.

In the SUSY model, the $\mathrm{CP}^{n_b-1|n_f}$ σ model, the spin $\vec{Z}$ is upgraded to a superspin $\Psi = (Z_1, \ldots, Z_{n_b}, \chi_1, \ldots, \chi_{n_f})$ with n_b bosonic and n_f fermionic components, which transforms under the superunitary group; see Ref. [17] for more information. The SUSY σ model describes the same physics as the replica σ model so long as $n_b - n_f = N$. The value

of $n_b + n_f$ is arbitrary; increasing $n_b + n_f$ gives a richer spectrum of operators but does not change the partition function or the mapping to the loop gas. The matrix Q_{ab} becomes $\Psi_a \Psi_b^\dagger - N^{-1}\delta_{ab}$, which is supertraceless. (To form the supertrace, indices are contracted up not with δ_{ab} but with $\eta_{ab} = \eta_a \delta_{ab}$, where η_a is $+1$ for bosonic and -1 for fermionic values of the index.) At $N = 1$, the multiplicity of the scaling dimension x_4 is [17]

$$d_{x_4} = 4n_f\left(2n_f^3 + 4n_f^2 + n_f - 1\right). \tag{D1}$$

We now translate our operators X, Y, and W into the SUSY language, calculate their total multiplicity, and confirm that at $N = 1$ it saturates d_{x_4}. This shows that there cannot be any other operators with dimension x_4. (In the SUSY representation, all multiplicities are of course positive, since there is no replicalike analytic continuation.)

First, consider the matrix $M_{aa'} = \Psi_a \Psi_{a'}$. This is symmetric in a,a' unless both of the indices are fermionic, in which case it is antisymmetric. Let us call this a "symmetric" tensor in quotation marks and a tensor that is antisymmetric except when both indices are fermionic an "antisymmetric" tensor. The number of independent components in $M_{aa'}$ is therefore

$$n_b(n_b + 1)/2 + n_b n_f + n_f(n_f - 1)/2. \tag{D2}$$

Next, consider $\widetilde{M}_{ab} = \Psi_a \overleftrightarrow{\nabla}_\mu \Psi_b$ (the spatial index will not play a role). This object is "antisymmetric." The number of components is obtained by exchanging $n_b \leftrightarrow n_f$ in the above formula.

The SUSY version of X is given by taking $\widetilde{X} = \Psi_a \Psi_{a'} \Psi_b^\dagger \Psi_{b'}^\dagger$ and subtracting appropriate terms to make it supertraceless (meaning that it vanishes when an index on one of the Ψs is contracted with an index on one of the $\Psi^\dagger$s) and therefore irreducible under the superunitary symmetry. Here such subtractions are possible for positive N, but see below.

The number of independent components in $\widetilde{X}$ is simply the square of the number of components in $M_{aa'}$ above. However, in making it traceless we remove $(n_b + n_f)^2$ components. This is simply the number of independent components that are left when we contract one pair of indices on $\widetilde{X}$ (all of these are set to zero when we make X supertraceless). Therefore, when $N = n_b - n_f = 1$, the dimension of the irreducible representation in which X transforms is

$$d_X = 4n_f^2(n_f + 1)^2. \tag{D3}$$

Next consider $\widetilde{Y}_{aa'bb'} = \epsilon_{\mu\nu}(\Psi_a \overleftrightarrow{\nabla}_\mu \Psi_{a'})(\Psi_b^\dagger \overleftrightarrow{\nabla}_\nu \Psi_{b'}^\dagger)$. The number of independent components is the square of the number in $\widetilde{M}$. When we perform subtractions to make $\widetilde{Y}$ supertraceless, naively we again remove $(n_b + n_f)^2$ components. This is correct for $N > 1$: In that case $\widetilde{Y}$ splits into an irreducible fully traceless object (the analog of Y in the replica theory), a singlet (the analog of W), and a two-index (adjoint) object, which is not of interest to us here. For $N = 1$, however, the invariant four-index tensor with the appropriate "antisymmetry" properties, denoted by $c_{aa'bb'}$, has vanishing trace, so subtracting $c_{aa'bb'}W$ does not change the trace of $\widetilde{Y}$. As a result, we expect that the resulting the four-index object forms an indecomposable representation that includes the singlet W. [For a simpler analog, consider the matrix $Q_{ab} = \Psi_a \Psi_b^\dagger - (n_b - n_f)^{-1}\delta_{ab}\Psi^\dagger\Psi$. When $n_b > n_f$, the subtraction ensures that the supertrace of Q vanishes. However, when $n_b = n_f$, the supertrace of the identity vanishes, so we cannot make Q traceless. This means that at $N = 0$, the supermatrix $\Psi\Psi^\dagger$ forms a single indecomposable representation of dimension $(n_b + n_f)^2$ that includes the singlet $\Psi^\dagger\Psi$.] The dimension of this representation is ($N = 1$)

$$d_{Y/W} = 4n_f(n_f + 1)\left(n_f^2 + n_f - 1\right). \tag{D4}$$

The notation indicates that this indecomposable representation in the SUSY model subsumes the analogs of both $Y_{aa'bb'}$ and W in the replica formulation. (Recall that Y was the parity odd four-leg operator and W was the singlet that appears in the Θ term.)

From the analysis in the main text, we know that the operators discussed above all have scaling dimension x_4. The scaling dimensions in the replica theory and the SUSY model are of course the same. (For example, this follows from the fact that they are related to the same correlators in the loop gas.) Adding up the multiplicities d_X and $d_{Y/W}$ gives perfect agreement with Eq. (D1) for d_{x_4}. This confirms that we have identified the complete set of operators with this scaling dimension and also explains where Eq. (D1) comes from physically.

[1] P.-G. de Gennes, Phys. Lett. **38**, 339 (1972).
[2] P.-G. de Gennes, J. Phys. (Paris) Lett. **36**, 55 (1975).
[3] B. Duplantier and H. Saleur, Phys. Rev. Lett. **59**, 539 (1987).
[4] A. Coniglio, N. Jan, I. Majid, and H. E. Stanley, Phys. Rev. B **35**, 3617 (1987).
[5] H. W. J. Blote and B. Nienhuis, J. Phys. A: Math. Gen. **22**, 1415 (1989).
[6] E. Vernier, J. L. Jacobsen, and H. Saleur, J. Stat. Mech. (2015) P09001.
[7] B. Duplantier and H. Saleur, Phys. Rev. Lett. **62**, 1368 (1989).
[8] F. Seno and A. L. Stella, J. Phys. (Paris) **49**, 739 (1988).
[9] T. Prellberg and A. L. Owczarek, J. Phys. A: Math. Gen. **27**, 1811 (1994).
[10] P. Grassberger and R. Hegger, J. Phys. (France) I **5**, 597 (1995).
[11] S. Caracciolo, M. Gherardio, M. Papinutto, and A. Pelissetto, J. Phys. A: Math. Theor. **44**, 115004 (2011).
[12] A. Nahum, P. Serna, A. M. Somoza, and M. Ortuño, Phys. Rev. B **87**, 184204 (2013).
[13] J. W. Lyklema, J. Phys. A: Math. Gen. **18**, L617 (1985).
[14] H. Meirovitch and H. A. Lim, Phys. Rev. A **38**, 1670 (1988).
[15] A. L. Owczarek and T. Prellberg, J. Stat. Phys. **79**, 951 (1995).
[16] D. P. Foster, J. Phys. A: Math. Theor. **42**, 372002 (2009).
[17] N. Read and H. Saleur, Nucl. Phys. B **613**, 409 (2001).
[18] H. W. J. Blöte, M. T. Batchelor, and B. Nienhuis, Physica A **251**, 95 (1998).
[19] A. Nahum, *Critical Phenomena in Loop Models* (Springer, Heidelberg, 2015).
[20] A. Nahum and J. T. Chalker, Phys. Rev. E **85**, 031141 (2012).

[21] A. Nahum, J. T. Chalker, P. Serna, M. Ortuño, and A. M. Somoza, Phys. Rev. Lett. **107**, 110601 (2011).

[22] A. Nahum, J. T. Chalker, P. Serna, M. Ortuño, and A. M. Somoza, Phys. Rev. B **88**, 134411 (2013).

[23] J. L. Jacobsen, N. Read, and H. Saleur, Phys. Rev. Lett. **90**, 090601 (2003).

[24] M. J. Martins, B. Nienhuis, and R. Rietman, Phys. Rev. Lett. **81**, 504 (1998).

[25] A. Bedini, A. L. Owczarek, and T. Prellberg, J. Phys. A: Math. Theor. **46**, 085001 (2013).

[26] I. Affleck, Nucl. Phys. B **257**, 397 (1985).

[27] E. Fradkin, *Field Theories of Condensed Matter Physics*, 2nd ed. (Cambridge University Press, Cambridge, 2013).

[28] I. Affleck, Phys. Rev. Lett. **66**, 2429 (1991).

[29] C. Candu, J. L. Jacobsen, N. Read, and H. Saleur, J. Phys. A **43**, 142001 (2010).

[30] J. Cardy, in *Fluctuating Geometries in Statistical Mechanics and Field Theory*, edited by F. David, P. Ginsparg, and J. Zinn-Justin (Elsevier, Amsterdam, 1998).

[31] N. Read and H. Saleur, Nucl. Phys. B **777**, 263 (2007).

[32] N. Read and H. Saleur, Nucl. Phys. B **777**, 316 (2007).

[33] E. Domany, D. Mukamel, B. Nienhuis, and A. Schwimmer, Nucl. Phys. B **190**, 279 (1981).

[34] B. Nienhuis, Phys. Rev. Lett. **49**, 1062 (1982).

[35] H. Saleur and B. Duplantier, Phys. Rev. Lett. **58**, 2325 (1987).

[36] R. M. Bradley, Phys. Rev. A **39**, 3738(R) (1989).

[37] This is associated with the fact that adjacent loops in Fig. 2, left, have opposite orientations or, in the 1+1D language, the fact that particles and antiparticles alternate in space [38].

[38] R. Shankar and G. Murthy, Phys. Rev. B **72**, 224414 (2005).

[39] A. Nahum, J. T. Chalker, P. Serna, M. Ortuño, and A. M. Somoza, Phys. Rev. X **5**, 041048 (2015).

[40] Θ changes sign under reflections, but since Θ is equivalent to $\Theta + 2\pi$ (up to a change of boundary conditions), $\Theta = \pi$ is compatible with spatial symmetry. A trick to obtain the value of Θ is as follows. Take a system with the topology of the disk (without dangling links at the boundary). Then, if Q is smoothly varying, it is possible to find a gauge where $\vec{Z}$ is smoothly varying and furthermore where the Boltzmann weight is real and positive for all the nodes in the interior. The imaginary part of the action then comes only from the Boltzmann weight for the boundary links. Labeling them by l, this is $\prod_l \vec{Z}^\dagger_{l+1}\vec{Z}_l$. In the continuum this becomes an exponential of the line integral of $\vec{Z}^\dagger\vec{\nabla}\vec{Z}$ along the boundary. This is precisely the Θ term in the action, which is a total derivative when written in terms of $\vec{Z}$ (so reduces to a boundary integral if we can choose a gauge where $\vec{Z}$ is smoothly varying). This term is pure imaginary. See Refs. [39,19] for a discussion of similar issues in a 3D loop model.

[41] The true scaling operators of the fixed point theory are not of course identical to simple expressions in the σ model fields, as different operators with the same symmetry properties can mix (for example, Q_{ab} could mix with $Q_{ab}\,\mathrm{tr}(\vec{\nabla}Q)^2$, etc.). However this is not important for the symmetry analysis.

[42] An operator defined on a link is not of course rotationally invariant: The leading contribution to this lattice operator is the rotationally invariant continuum operator, which we denote by Q.

[43] J. Cardy, *Scaling and Renormalization in Statistical Physics* (Cambridge University Press, Cambridge, 1996).

[44] One could even consider off-lattice (e.g., bead-and-spring-type) models. While the regularization of the appropriate gauge theories would be tedious to perform explicitly, it is clear that the same symmetry arguments would apply.

[45] For reviews see B. Nienhuis, in *Phase Transitions and Critical Phenomena*, edited by C. Domb and J. L. Lebowitz (Academic, Singapore, 1987), Vol. 11, Chap. 1; B. Duplantier, Phys. Rep. **184**, 229 (1989); J. L. Jacobsen, in *Polygons, Polyominoes, and Polyhedra*, edited by A. J. Guttmann (Springer, New York, 2009), Chap. 14.

[46] H. Saleur, Nucl. Phys. B **382**, 532 (1992).

[47] J. Cardy, J. Phys. A: Math. Gen. **34**, L665 (2001).

[48] S. O. Warnaar, M. T. Batchelor, and B. Nienhuis, J. Phys. A: Math. Gen. **25**, 3077 (1992).

[49] J. Krawczyk, T. Prellberg, A. L. Owczarek, and A. Rechnitzer, Phys. Rev. Lett. **96**, 240603 (2006).

[50] In the field theory obtained from the completely packed loop representation (but not in that obtained from the incompletely packed representation) a domain wall in the Ising order parameter is essentially a world line of Z_0 [12], so the presence or absence of Ising order in the dense phase is determined by whether this component is massive or condensed.

[51] P. Serra, J. F. Stilck, W. L. Cavalcanti, and K. D. Machado, J. Phys. A: Math. Gen. **37**, 8811 (2004).

[52] M. Kaufman, T. Berger, P. D. Gujrati, and D. Bowman, Phys. Rev. A **41**, 4371 (1990).

[53] This is easiest to see in the formulation of the $\mathrm{CP}^{N-1}\sigma$ model that includes an explicit fluctuating gauge field.

[54] A subtlety is that within the dense phase itself the exponents vary continuously as the σ model coupling is exactly marginal [17,29].

[55] We will discuss this in more detail elsewhere.

[56] When the polymer can cross itself, the definition given in the text for the Ising order parameter needs to be slightly modified: We take the orientation given to the polymer to flip every time it crosses itself.

[57] If included, four-leg branchings would induce the perturbation $\sum_{a>1} X_{aaaa}$, breaking the part of the symmetry group that mixes different components to permutations S_n.

[58] This rule for orienting the polymer is also that which is necessary in order to extend the definition of σ in Sec. V D to the case with crossings [56]. That is, the sign of σ is determined by whether the U(1) current is parallel or antiparallel to the orientation of the link in Fig. 3.

[59] While the interacting self-avoiding trail is fine-tuned when regarded as a model for a polymer, it can also be regarded as a model for a deterministic walk in a random environment, i.e., a Lorentz lattice gas, and in this interpretation it is not fine-tuned.

[60] A. Bedini, A. L. Owczarek, and T. Prellberg, J. Phys. A: Math. Theor. **47**, 145002 (2014).

[61] J. Cardy, arXiv:cond-mat/9911024; J. Phys. A: Math. Theor. **46**, 494001 (2013).

[62] R. Vasseur, Phys. Rev. B **92**, 014205 (2015).

[63] Strictly, this description applies when $a \neq b$. When $a = b$, the operator modifies the weight for a configuration with an a-flavored loop passing through the point.

[64] B. Duplantier and H. Saleur, Phys. Rev. Lett. **60**, 2343 (1988).

[65] T. Prellberg and B. Drossel, Physica A **249**, 337 (1998).

[66] Naively coarse graining Eq. (38) we might expect a contact term, with a nonuniversal coefficient, scaling like $\langle Q(0)Q(R)\rangle^{-1}\int d^2r\langle Q(0)X(r)Q(R)\rangle$. However, this would give a result of order $\mathcal{R}\sim R^{y_4}=R^{3/4}$ (which is the scaling of the number of close approaches of the two arms). Since this is much larger than the correct result, namely, $\ln R$, the coefficient of this term must be zero when the lattice expression is correctly coarse grained. Equivalently, the correct continuum limit should be $\lim_{\epsilon\to 0}[\int_{|r-r'|>\epsilon} d^2r\,d^2r'\langle Q_{12}(0)Q_{23}^{\rm odd}(r)Q_{34}(R)Q_{23}^{\rm odd}(r')\rangle - (8\pi c/3\epsilon^{3/4})\int d^2r\langle Q_{12}(0)X_{2431}(r)Q_{34}(R)\rangle]$, where c is the appropriate OPE coefficient.

[67] I. Affleck, Phys. Rev. Lett. **55**, 1355 (1985).

[68] For reviews see J. Cardy, Ann. Phys. (N.Y.) **318**, 81 (2005); I. Gruzberg, J. Phys. A: Math. Gen. **39**, 12601 (2006).

[69] E. Vernier, J. L. Jacobsen, and H. Saleur, J. Phys. A: Math. Theor. **49**, 064002 (2016).

[70] The deformation of Ref. [69] does not retain gauge invariance. In order to consider *infinitesimal* perturbations it is important to retain gauge invariance, since local operators in the conformal field theory are gauge invariant. For this, the approach and classification of perturbations provided in the present paper are appropriate.

[71] A. Bedini, A. L. Owczarek, and T. Prellberg, arXiv:1510.06686.

[72] A. Bedini, A. L. Owczarek, and T. Prellberg, arXiv:1510.06891.

[73] Since the true DS transition is into the collapsed phase without Ising order, scenario (i) in Sec. V D would imply that the two types of collapsed phase, $\langle\sigma\rangle=0$ and $\langle\sigma\rangle\neq 0$, both existed arbitrarily close to model T, i.e., that model T represented the point where the Ising phase transition line between them terminated.