Communications in Mathematical Physics manuscript No.
(will be inserted by the editor)

Entropy-driven phase transition in low-temperature antiferromagnetic Potts models

Roman Kotecký[1,2], **Alan D. Sokal**[3,4], **Jan M. Swart**[5]

[1] Department of Mathematics, University of Warwick, Coventry CV4 7AL, United Kingdom, E-mail: R.Kotecky@warwick.ac.uk
[2] Center for Theoretical Study, Charles University, Jilská 1, Prague, Czech Republic
[3] Department of Physics, New York University, 4 Washington Place, New York, NY 10003, USA, E-mail: sokal@nyu.edu
[4] Department of Mathematics, University College London, Gower Street, London WC1E 6BT, United Kingdom
[5] Institute of Information Theory and Automation of the ASCR (ÚTIA), Pod vodárenskou věží 4, 18208 Praha 8, Czech Republic, E-mail: swart@utia.cas.cz

The date of receipt and acceptance will be inserted by the editor

Abstract: We prove the existence of long-range order at sufficiently low temperatures, including zero temperature, for the three-state Potts antiferromagnet on a class of quasi-transitive plane quadrangulations, including the diced lattice. More precisely, we show the existence of (at least) three infinite-volume Gibbs measures, which exhibit spontaneous magnetization in the sense that vertices in one sublattice have a higher probability to be in one state than in either of the other two states. For the special case of the diced lattice, we give a good rigorous lower bound on this probability, based on computer-assisted calculations that are not available for the other lattices.

Key words. Antiferromagnetic Potts model, proper coloring, plane quadrangulation, phase transition, diced lattice.

MSC 2000 Subject Classification: Primary 82B20; Secondary 05C15, 05C63, 60K35.

Contents

1. Introduction and main results

1.1. Introduction. We are interested here in the three-state antiferromagnetic Potts model on a class of infinite plane quadrangulations. Recall that a graph embedded in the plane is called a quadrangulation if all its faces are quadrilaterals (i.e., have four vertices and four edges).[1] Some examples of infinite plane quadrangulations are drawn in Figure 1.1: these include the square lattice $\mathbb{Z}^2$ (with nearest-neighbor edges) and the so-called diced lattice.

---

[1] In this paper we restrict attention to *nondegenerate* quadrangulations, i.e. each face has four *distinct* vertices and four *distinct* edges. Some discussion of degenerate plane quadrangulations (in the case of finite graphs) can be found in [38].

On the square lattice, the three-state Potts antiferromagnet at zero temperature can be mapped into a special case of the six-vertex model that admits an exact (but nonrigorous) solution [5, section 8.13]. This model is therefore believed to be critical at zero temperature but disordered for any positive temperature.[2]

On the diced lattice, by contrast, a proof was outlined in [46] showing that the three-state Potts antiferromagnet has a phase transition at nonzero temperature and has long-range order at all sufficiently low temperatures (including zero temperature). In the present paper, we present the details of this proof and we extend the result to a large class of quasi-transitive quadrangulations, including some hyperbolic lattices.

To explain the class of lattices that we can cover, let us start by observing that a quadrangulation is a connected bipartite graph $G = (V, E)$, so that the vertex set has a canonical bipartition $V = V_0 \cup V_1$. We may view the two sublattices V_0 and V_1 as graphs in their own right by connecting vertices along the diagonals of the quadrilateral faces of the original lattice: this yields graphs $G_0 = (V_0, E_0)$ and $G_1 = (V_1, E_1)$ as shown in Figure 1.1. Note that G_0 and G_1 are duals of each other, i.e. each face of G_0 contains a unique vertex of G_1, and vice versa; and each edge of G_0 crosses a unique edge of G_1, and vice versa (see Figure 1.2). Conversely, given any dual pair of (finite or infinite) graphs $G_0 = (V_0, E_0)$ and $G_1 = (V_1, E_1)$ embedded in the plane, we can form a plane quadrangulation $G = (V, E)$ by setting $V = V_0 \cup V_1$ and placing an edge between each pair of vertices $v \in V_0$ and $w \in V_1$ where w lies in a face of G_0 that has v on its boundary (or equivalently vice versa). The main assumption that we will make in this paper is that one sublattice (say, G_0) is a triangulation. In particular, our proofs cover the lattices shown in (b)–(d) of Figure 1.1, but not the square lattice (a).

To explain the nature of the phase transition, note that ground states of the three-state Potts antiferromagnet are simply proper three-colorings of the lattice. On any bipartite lattice, we may construct special ground states by coloring one sublattice (say, V_0) in one color and using the other two colors to color the other sublattice in any possible way. Note that in this way, the second sublattice carries all the entropy. Of course, the special ground states in which the first sublattice uses *only* one color are atypical of the Gibbs measure, even at zero temperature. Nevertheless, the underlying idea applies more generally: there may be a preference for the first sublattice to be colored *mostly* in one color because this increases the freedom of choice of colors on the other sublattice. Otherwise put, integrating out the colors on the second sublattice may induce an effective ferromagnetic interaction on the first sublattice. If this effective interaction is strong enough, it may result in long-range order on the first sublattice. We call this an *entropy-driven phase transition*.[3] In [46] a proof was sketched along these lines that an entropy-driven transition indeed occurs on the diced lattice. Here we will present the details of this proof and extend it to a large class of plane quadrangulations in which one sublattice is a triangulation. The extension uses a variant of the Peierls argument that works whenever the Peierls sum is finite (even if it is not small), followed by a random-cluster argument.

In all the cases handled in this paper, there is a strong asymmetry between the two sublattices, so that it is entropically favorable to ferromagnetically order the triangulation (G_0) and place the entropy on the other sublattice (G_1). By contrast, in the square lattice, where no finite-temperature phase transition is believed to occur, the two sublattices are isomorphic. It is therefore natural to ask whether asymmetry is a necessary and/or sufficient condition for the existence of a finite-temperature phase transition. This is a subtle question, and we discuss it further in Section 1.3 below.

---

[2] See e.g. the discussion below formula (2.8) in [57]. See also [24, 64] for Monte Carlo data supporting these beliefs.

[3] The method for encoding such entropy costs in terms of certain Peierls contours was suggested already in [45], but in that paper it led to a proof of the transition only for some toy models including the three-state Potts antiferromagnet on the "decorated cubic lattice".

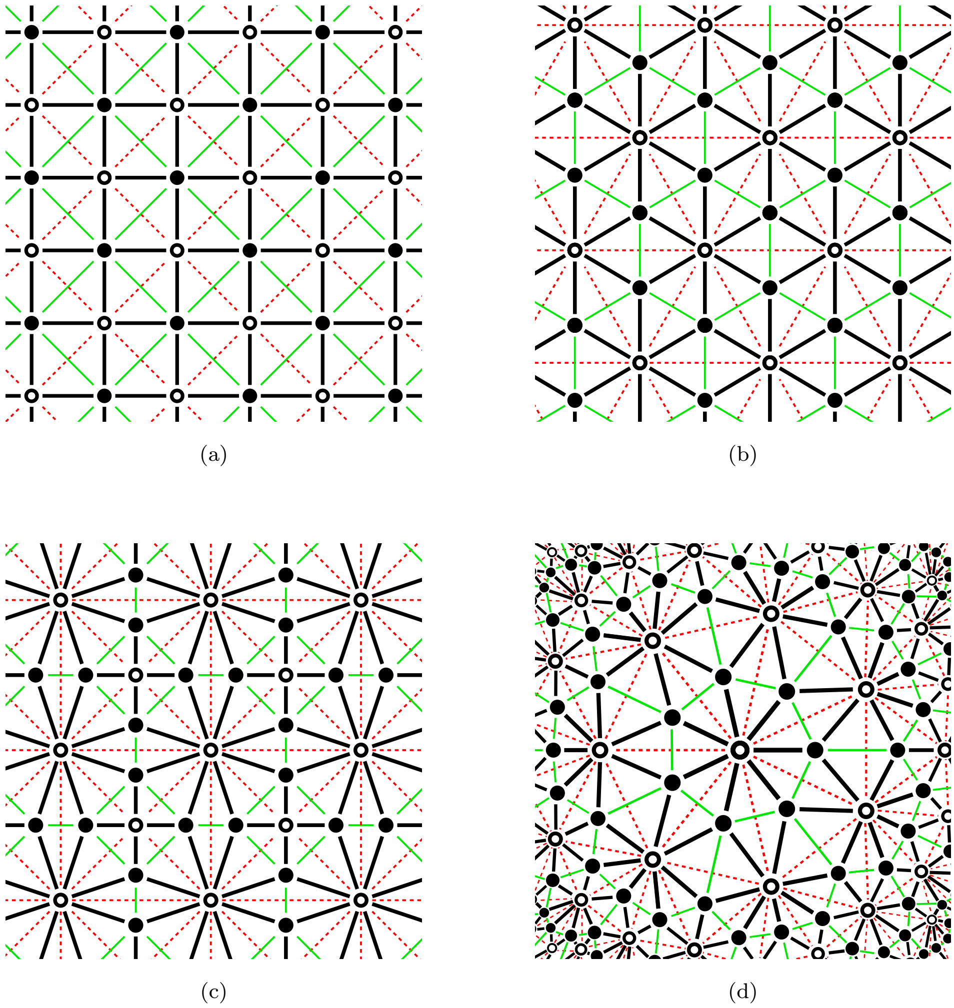

Fig. 1.1. Four quasi-transitive quadrangulations with their sublattices G_0 (open circles joined by dashed red edges) and G_1 (filled circles joined by solid green edges). In these examples, the sublattice G_0 is (a) the square lattice, (b) the triangular lattice, (c) the union-jack lattice, and (d) the hyperbolic lattice with Schläfli symbol $\{3, 7\}$. Note that in (b)–(d), G_0 is a triangulation. In (a) and (b), the quadrangulations G are, respectively, the square lattice and the diced lattice.

1.2. Statement of the results. Let us now formulate our results precisely. We first need to define more precisely the class of graphs we will be considering. We quickly review here the essential definitions; a more thorough summary of the needed theory of infinite graphs can be found in the Appendix.

A graph $G = (V, E)$ is called *locally finite* if every vertex has finitely many neighbors; in this paper we will consider only locally finite graphs. If G has at least $k + 1$ vertices, then G is called k*-connected* if one needs to remove at least k vertices to disconnect it. A graph G is said to have *one end* if after the removal of finitely many edges, there is exactly one infinite connected component; note that this implies in particular that G is infinite.

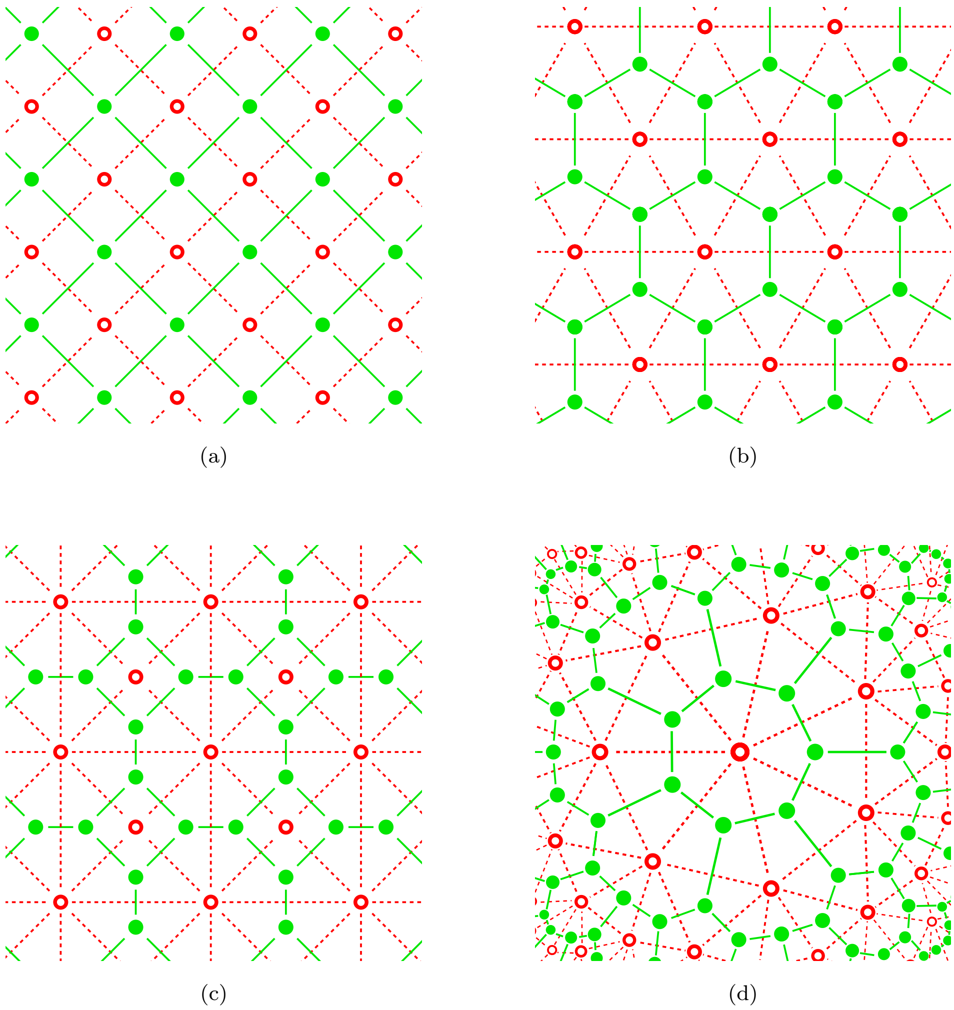

Fig. 1.2. The two sublattices G_0 and G_1 for each of the four lattices from Figure 1.1. The graph G_0 (red open circles and dashed edges) is dual to the graph G_1 (green filled circles and solid edges).

A graph is said to be *planar* if it can be drawn in the plane $\mathbb{R}^2$ with vertices represented by distinct points and edges represented by closed continuous arcs joining their endvertices, mutually disjoint except possibly at their endpoints. A *plane graph* is a planar graph with a given embedding in the plane. An embedding of a connected graph in the plane is called *edge-accumulation-point-free* (or *EAP-free* for short) if there are no points in the plane with the property that each neighborhood of the point intersects infinitely many edges. An EAP-free embedding divides the plane into connected open sets called *faces*. The boundary of a face is either a finite cycle or a two-way-infinite path.

Consider an EAP-free embedding of a 3-connected graph G. Then one may define a *dual graph* $G^* = (V^*, E^*)$ whose vertex set V^* is the set of faces of G and with edge set $E^* \simeq E$, where by definition two faces are linked by an edge $e^* \in E^*$ if and only if the corresponding edge $e \in E$ lies in the border of both faces. (Since planar embeddings of 3-connected graphs are unique up to isomorphism, this dual graph is likewise unique up to isomorphism.) Clearly G^* is locally finite if and only if every face of G is bounded by a finite cycle. In this case we can embed G^* in an EAP-free way in the plane such that each face of G contains exactly one vertex of G^* and each edge of G is crossed by exactly one edge of G^*, and vice versa; in particular, G is also the dual of G^*. We say that G is a *triangulation* (resp. *quadrangulation*) if every face of G is bounded by a triangle (resp. quadrilateral), or equivalently if each vertex of G^* has degree 3 (resp. 4). It turns out that triangulations and quadrangulations, defined in this way, are always graphs with at most one end.[4]

The final set of definitions we need concerns some form of "translation invariance" of our lattices. An *automorphism* of a graph $G = (V, E)$ is a bijection $g\colon V \to V$ that preserves the graph structure. Two vertices $u, v \in V$ are of the *same type* if there exists an automorphism that maps u into v. This relation partitions the vertex set V into equivalence classes called *types.* The graph G is called *vertex-transitive* if there is just one equivalence class, and *vertex-quasi-transitive* if there are finitely many equivalence classes. Edge-transitivity and (for plane graphs) face-transitivity are defined similarly. The corresponding forms of quasi-transitivity are all equivalent (see Lemma A.1 in the Appendix), which is why we simply talk about quasi-transitivity without specifying whether in the vertex-, edge- or face-sense.

We will study the three-state antiferromagnetic Potts model on plane quadrangulations G, constructed from mutually dual sublattices G_0 and G_1, such that G_0 is a locally finite 3-connected quasi-transitive triangulation with one end.[5] Note that quasi-transitivity refers only to the structure of G_0 (or G_1 or G) as an abstract graph (i.e., without reference to any embedding). It turns out [3, Theorem 4.2] (see also [66, Theorem 1]) that any locally finite 3-connected quasi-transitive planar graph with one end can be *periodically* embedded in either the Euclidean or hyperbolic plane, i.e., so that the automorphisms of G correspond to a discrete subgroup of the group of isometries of the embedding space. But we do not use this fact anywhere in this paper.

Let us now define the q-state Potts antiferromagnet on an arbitrary infinite graph $G = (V, E)$, for an arbitrary positive integer q. The state space is the set

$$S \;:=\; [q]^V \;=\; \bigl\{\sigma = (\sigma_v)_{v\in V}\colon\; \sigma_v \in [q]\ \forall v \in V\bigr\}\,, \tag{1.1}$$

where we have used the shorthand notation $[q] = \{1, 2, \ldots, q\}$. We also let

$$S_{\mathrm{g}} \;:=\; \bigl\{\sigma \in S\colon\; \sigma_u \neq \sigma_v\ \forall \{u, v\} \in E\bigr\} \tag{1.2}$$

denote the set of proper q-colorings of G. We sometimes use the terms "spin configuration" for $\sigma \in S$ and "ground-state configuration" for $\sigma \in S_{\mathrm{g}}$. For each finite subset $\Lambda \subset V$ we let

$$\partial\Lambda \;:=\; \bigl\{v \in V\backslash\Lambda\colon\; \{v, u\} \in E \text{ for some } u \in \Lambda\bigr\} \tag{1.3}$$

denote the external boundary of Λ. For any boundary condition $\tau\colon\; V \to [q]$ and any spin configuration $\sigma\colon\; \Lambda \to [q]$ on Λ, we define the Hamiltonian of σ under the boundary condition τ by

$$H_\Lambda(\sigma\,|\,\tau) \;:=\; \sum_{\substack{u,v\in\Lambda\\ \{u,v\}\in E}} \delta_{\sigma_u,\sigma_v} + \sum_{\substack{u\in\Lambda,\ v\in\partial\Lambda\\ \{u,v\}\in E}} \delta_{\sigma_u,\tau_v} \tag{1.4}$$

---

[4] A more general definition, which allows for multiple ends, is discussed in Section A.2 of the Appendix. Briefly, we say that a locally finite 3-connected graph G is a triangulation (resp. quadrangulation) if G has an *abstract dual* in which each vertex has degree 3 (resp. 4); see Section A.2 for the definition of abstract duals. With this more general definition, it can be shown that a triangulation (resp. quadrangulation) has an EAP-free embedding in the plane if and only if it is finite or has one end: see Proposition A.13 in Section A.3.

[5] The quasi-transitivity of G_0 implies that also G_1 is quasi-transitive: see Theorem A.3(v) in the Appendix. It is not hard to see that now also G must be quasi-transitive.

where $\delta_{\sigma_u,\sigma_v}$ is the Kronecker delta, i.e.

$$\delta_{ab} \;=\; \delta(a,b) \;=\; \begin{cases} 1 & \text{if } a = b \\ 0 & \text{if } a \neq b \end{cases} \tag{1.5}$$

For $\beta \in [0,\infty)$, we define the Gibbs measure in volume Λ with boundary condition τ at inverse temperature β:

$$\mu^\tau_{\Lambda,\beta}(\sigma) \;:=\; \frac{1}{Z^\tau_{\Lambda,\beta}} \exp[-\beta H_\Lambda(\sigma \,|\, \tau)] \;. \tag{1.6}$$

For $\beta = \infty$, we define

$$\mu^\tau_{\Lambda,\infty}(\sigma) \;:=\; \lim_{\beta\to\infty} \mu^\tau_{\Lambda,\beta}(\sigma) \;. \tag{1.7}$$

That is, $\mu^\tau_{\Lambda,\infty}$ is the uniform distribution on configurations σ that minimize $H_\Lambda(\sigma \,|\, \tau)$. [Note that for some τ this minimum energy might be strictly positive, i.e. there might not exist proper colorings of $\Lambda \cup \partial\Lambda$ that agree with τ on $\partial\Lambda$.] Of course, these definitions actually depend on τ only via the restriction $\tau_{\partial\Lambda} := (\tau_u)_{u\in\partial\Lambda}$ of τ to $\partial\Lambda$.

We then define infinite-volume Gibbs measures in the usual way through the Dobrushin–Lanford–Ruelle (DLR) conditions [29], i.e., we say that a probability measure μ on S is an infinite-volume Gibbs measure for the q-state antiferromagnetic Potts model at inverse temperature $\beta \in [0,\infty]$ if for each finite $\Lambda \subset V$ its conditional probabilities satisfy

$$\mu(\sigma_\Lambda |\, \sigma_{V\setminus\Lambda} = \tau_{V\setminus\Lambda}) \;=\; \mu^\tau_{\Lambda,\beta}(\sigma_\Lambda) \quad \text{for } \mu\text{-a.e. } \tau \;. \tag{1.8}$$

For the remainder of this paper we specialize to $q = 3$. Here is our main result:

Theorem 1.1. (Gibbs state multiplicity and positive magnetization) *Let $G = (V,E)$ be a quadrangulation of the plane, and let $G_0 = (V_0,E_0)$ and $G_1 = (V_1,E_1)$ be its sublattices, with edges drawn along the diagonals of quadrilaterals. Assume that G_0 is a locally finite 3-connected quasi-transitive triangulation with one end. Then there exist $\beta_0, C < \infty$ and $\epsilon > 0$ such that for each inverse temperature $\beta \in [\beta_0,\infty]$ and each $k \in \{1,2,3\}$, there exists an infinite-volume Gibbs measure $\mu_{k,\beta}$ for the 3-state Potts antiferromagnet on G satisfying:*

(a) For all $v_0 \in V_0$, we have $\mu_{k,\beta}(\sigma_{v_0} = k) \geq \frac{1}{3} + \epsilon$.
(b) For all $v_1 \in V_1$, we have $\mu_{k,\beta}(\sigma_{v_1} = k) \leq \frac{1}{3} - \epsilon$.
(c) For all $\{u,v\} \in E$, we have $\mu_{k,\beta}(\sigma_u = \sigma_v) \leq Ce^{-\beta}$.

In particular, for each inverse temperature $\beta \in [\beta_0,\infty]$, the 3-state Potts antiferromagnet on G has at least three distinct extremal infinite-volume Gibbs measures.

Remarks. 1. The bound (c) shows in particular that the zero-temperature Gibbs measure $\mu_{k,\infty}$ is supported on ground states.

2. Any subsequential limit as $\beta \to \infty$ of the measures $\mu_{k,\beta}$ with $\beta < \infty$ also satisfies the bounds (a)–(c). Therefore, there exist zero-temperature Gibbs measures with these properties that are limits of finite-temperature Gibbs measures with these properties.[6]

[6] To see that this is a nontrivial property, consider on the square lattice $\mathbb{Z}^2$ the configuration $\tau \in S$ of the 3-state Potts antiferromagnet defined by $\tau_{(i,j)} = 1 + (i + j \bmod 3)$. Then $\tau \in S_{\rm g}$ is a ground-state configuration such that its restriction to any row and any column, suitably shifted, is the sequence $(\ldots,1,2,3,1,2,3,\ldots)$. Since for any finite $\Lambda \subset \mathbb{Z}^2$ there is precisely one ground-state configuration that agrees with τ on $\mathbb{Z}^2\setminus\Lambda$, namely τ itself, we see that the Dirac measure δ_τ is a zero-temperature infinite-volume Gibbs measure. But this measure is not a limit of positive-temperature Gibbs measures: the reason can be traced to the fact that at $\beta < \infty$ with boundary condition τ in a sufficiently large volume Λ, there exists a configuration $\bar\tau$ [namely, $\bar\tau_{(i,j)} = 1 + (i + j \bmod 2)$] such that we can replace τ on the internal boundary of Λ by $\bar\tau$ at an energetic cost of order $|\partial\Lambda|$, while gaining a bulk entropic advantage (with the boundary condition $\bar\tau$, one can color Λ with 1 on one sublattice and arbitrarily 2 or 3 on the other sublattice).

3. We construct the infinite-volume Gibbs measures $\mu_{\beta,k}$ as subsequential limits of finite-volume Gibbs measures. We expect that there is no need to go to a subsequence and that our approximation procedure yields extremal infinite-volume Gibbs measures, but we have not proven either of these assertions.

In the special case where G is the diced lattice, we have a good explicit bound on the probabilities in Theorem 1.1(a,b):

Theorem 1.2. (Quantitative bound for the diced lattice) *Let $G = (V, E)$ be the diced lattice and let $G_0 = (V_0, E_0)$ and $G_1 = (V_1, E_1)$ be its triangular and hexagonal sublattices, respectively. Then there exists $C < \infty$ such that for each inverse temperature $\beta \in [0, \infty]$ and each $k \in \{1, 2, 3\}$, there exists an infinite-volume Gibbs measure $\mu_{k,\beta}$ for the 3-state Potts antiferromagnet on G satisfying:*

(a) For all $v_0 \in V_0$, we have $\mu_{k,\beta}(\sigma_{v_0} = k) \geq 0.90301 - Ce^{-\beta}$.
(b) For all $v_1 \in V_1$, we have $\mu_{k,\beta}(\sigma_{v_1} = k) \leq 0.14549 + Ce^{-\beta}$.
(c) For all $\{u, v\} \in E$, we have $\mu_{k,\beta}(\sigma_u = \sigma_v) \leq Ce^{-\beta}$.

The lower bound 0.90301 should be compared with the estimated zero-temperature value 0.957597 ± 0.000004 from Monte Carlo simulations [46].[7]

1.3. Discussion. The phase diagram of non-attractive (i.e., non-ferromagnetic) spin sytems is generally harder to predict than for attractive (ferromagnetic) spin sytems, and may sometimes depend subtly on the microscopic details of the model. In particular, this is true for the two-dimensional 3-state Potts antiferromagnet, for which we have shown that it has a phase transition at positive temperature on the diced lattice, while no such phase transition is believed to occur on the square lattice — even though both lattices are bipartite and are in fact plane quadrangulations.

The existence of a positive-temperature transition in the diced-lattice model was a surprise when it was first discovered [46], for the following reason: Some two-dimensional antiferromagnetic models at zero temperature have the property that they can be mapped exactly onto a "height" model (in general vector-valued) [64, 37]. In such cases one can argue heuristically that the height model must always be in either a "smooth" (ordered) or a "rough" (massless) phase; correspondingly, the underlying zero-temperature spin model should either be ordered or critical, never disordered. Experience teaches us (or at least seemed to teach us) that the most common case is criticality.[8] In particular, when the q-state zero-temperature Potts antiferromagnet on a two-dimensional periodic lattice admits a height representation, one ordinarily expects that model to have a zero-temperature critical point. This prediction is confirmed (at least non-rigorously) in most heretofore-studied cases: 2-state (Ising) triangular [7, 56], 3-state square-lattice [57, 42, 11, 64], 3-state kagome [35, 44], 4-state triangular [54], and 4-state on the line graph of the square lattice [43, 44]. Indeed, before the work of [46], no exceptions were known.

It was furthermore observed in [46] that the height mapping employed for the 3-state Potts antiferromagnet on the square lattice [64] carries over unchanged to any plane quadrangulation. One would therefore have expected the 3-state Potts antiferromagnet to have a zero-temperature critical point on every periodic plane quadrangulation. The example of the diced lattice showed that this is not the case; and the results of the present paper provide further counterexamples. Clearly, the mere existence of a height representation does *not* guarantee that the model will be

[7] The value $M_0 = 0.936395 \pm 0.000006$ reported in [46] is the spontaneous magnetization in the hypertetrahedral representation, i.e. $M_0 = \mu_{1,\infty}(\sigma_v = 1) - \frac{1}{2}\mu_{1,\infty}(\sigma_v \neq 1)$.

[8] Some exceptions known already at the time of [46] were the constrained square-lattice 4-state antiferromagnetic Potts model [11] and the triangular-lattice antiferromagnetic spin-s Ising model for large enough s [72], both of which appear to lie in a non-critical ordered phase at zero temperature.

critical. Indeed, criticality may well be an exception — corresponding to cases with an unusual degree of symmetry — rather than the generic case.

The mechanism behind all these transitions is what we have called an "entropy-driven phase transition": namely, ordering on one sublattice increases the entropy available to the other sublattice; or said in a different way, integrating out the spins on the second sublattice induces an effective ferromagnetic interaction on the first sublattice. If this effective interaction is strong enough, it may result in long-range order. Such a phase transition can therefore occur in principle in any antiferromagnetic model on any bipartite lattice[9]; whether it actually does occur is a quantitative question concerning the strength of the induced ferromagnetic interaction. Thus, such an entropy-driven phase transition is believed not to occur in the 3-state Potts antiferromagnet on the square lattice $\mathbb{Z}^2$; but Monte Carlo evidence [70,30,31] suggests that it does occur in this same model on the simple-cubic lattice $\mathbb{Z}^3$ and presumably also on $\mathbb{Z}^d$ for all $d \geq 3$; moreover, Peled [61] has recently proven this for all sufficiently large d and also for a "thickened" version of $\mathbb{Z}^2$.

From the point of view of the Peierls argument, the relevant issue is the strength of the entropic penalty for domain walls between differently-ordered regions, compared to the entropy associated to those domain walls. In order to successfully carry out the Peierls argument, one must consider all the relevant ordered phases, find an appropriate definition of Peierls contours separating spatial regions resembling those ordered phases, and prove that long Peierls contours γ are suppressed like $e^{-c|\gamma|}$ with a sufficiently large constant c.

The simplest situation arises when there is an asymmetry between the two sublattices, so that it is entropically more favorable for one of them (say, V_0) to be ferromagnetically ordered and for the other (V_1) to carry all the entropy. This situation is expected to occur, for instance, if V_1 has a higher density of points than V_0. The case treated in this paper, in which G_0 is a triangulation, achieves this in the strongest possible way: namely, for the Euclidean lattices in our class, it is easy to see using Euler's formula that the spatial densities of the sublattices V_0 and V_1 are in the proportion 1:2, which is the most extreme ratio achievable for two dual periodic Euclidean lattices.

In this asymmetric situation, one knows in advance which sublattice (V_0) is going to be ferromagnetically ordered (if the entropic effect is strong enough to produce any long-range order at all); therefore, for the 3-state Potts antiferromagnet on G, one expects at low temperature to have (at least) three distinct ordered phases, corresponding to the three possible choices for the color that dominates on V_0. One may therefore define Peierls contours just as one would for a ferromagnetic Potts model on G_0 (see Section 2.1 below for details), and then try to show that long Peierls contours are sufficiently suppressed, i.e. that it is sufficiently costly to create an interface between regions where one and another color are used on V_0. This is a quantitative problem, which is made difficult by the fact that (unlike in a ferromagnetic model) one does not have any parameter that can be varied to make the suppression of long contours as large as one wishes.

The situation is even more delicate for lattices, such as $\mathbb{Z}^d$, where the two sublattices play a symmetric role (in the sense that there exists an automorphism of G carrying one sublattice onto the other). Indeed, for models with symmetry between the sublattices, for every Gibbs measure where one sublattice is ordered (in the sense of being colored more often with one preferred color), there must obviously exist a corresponding Gibbs measure where the other sublattice is ordered. Therefore, the system has *two* "choices" to make: first, of the sublattice to be ordered, and then of the color in which it is ordered — which leads to a total of (at least) six distinct ordered phases for the 3-state model. For this sort of long-range order to occur, it must be sufficiently costly to create an interface between *any* pair of distinct ordered phases; in particular, it must be costly to create an interface between regions where one and the other sublattice are ordered (in whatever colors). To prove such a result will almost certainly require a different (and more subtle) definition of Peierls contour than is used in the asymmetric case.

---

[9] It can also occur in antiferromagnetic models on non-bipartite lattices: for instance, in the 4-state Potts antiferromagnets on the union-jack and bisected-hexagonal lattices [14], which are tripartite, and for which ordering on one sublattice increases the entropy available to the other *two* sublattices. However, we are concerned here for simplicity with the bipartite case.

The example of $\mathbb{Z}^d$ for d large shows that asymmetry is not necessary for the existence of a finite-temperature phase transition. But one can nevertheless say heuristically that asymmetry enhances the effect driving the transition, by increasing the strength of the effective ferromagnetic interaction on the favored sublattice (while of course decreasing it on the disfavored sublattice).

Everything said so far holds for an arbitrary bipartite lattice. But the case in which G is a plane quadrangulation is special, because G_0 and G_1 are not merely the two sublattices: they are a dual pair of plane graphs.[10] In particular, there is a symmetry between G_0 and G_1 if and only if G_0 is *self-dual*; and there may be special reasons, connected with the topology of the plane, that make this self-dual case special (e.g. critical at zero temperature). Now, it is well known that the square lattice is self-dual; what seems to be less well known[11] is that there exist many other examples of self-dual periodic plane graphs [2,59,60,66,71,65,**?**], including the "hextri" lattice [59, Fig. 1] [66, Fig. 16] [71, Fig. 1b], the "house" lattice [59, Fig. 2], and the martini-B lattice [65, Fig. 8]. Preliminary results [13] of Monte Carlo simulations on a variety of plane quadrangulations suggest that

(a) If G_0 is self-dual, then the 3-state Potts antiferromagnet on the associated quadrangulation G has a zero-temperature critical point; and
(b) If G_0 is not self-dual, then the 3-state Potts antiferromagnet on G has (always? usually?) a finite-temperature phase transition.

In other words, it seems that for plane quadrangulations — unlike for general bipartite lattices — asymmetry may be both necessary and sufficient for the existence of a finite-temperature phase transition. It would be very interesting to find a deeper theoretical explanation, and ultimately a proof, of this apparent fact. We conjecture that there is an exact duality mapping that explains why (a) is true. As for (b), one could argue for it heuristically as follows: Because the model at zero temperature has a height representation, it should be either critical or ordered. If the self-dual cases are critical, then the non-self-dual cases should be ordered, since asymmetry enhances the phase transition; and if the self-dual cases are ordered, then the non-self-dual cases should be even more strongly ordered. It goes without saying that this heuristic argument is extremely vague — no criterion for comparing lattices is given — and hence very far from suggesting a strategy of proof.

Entropy-driven phase transitions are also possible in the q-state Potts antiferromagnet for $q > 3$, but now one must consider the possibility of Gibbs measures associated to other partitions $[q] = Q_0 \cup Q_1$, in which the vertices in V_0 (resp. V_1) take predominantly colors from Q_0 (resp. Q_1). Depending on the size and shape of V_0 and V_1 and the value of q, such measures might be entropically favored. For instance, such ordering with $|Q_0| = |Q_1| = 2$ has been claimed to occur in the 4-state Potts antiferromagnet on the simple-cubic lattice $\mathbb{Z}^3$ [4,36]. Naive entropic considerations suggest that if the densities of the sublattices V_0 and V_1 are in the ratio $\alpha{:}1-\alpha$, then the dominant ordering would have $|Q_0| \approx \alpha q$. In general, one would expect to have $\binom{q}{|Q_0|}$ ordered phases in the asymmetric case, and $2\binom{q}{|Q_0|} = 2\binom{q}{\lfloor q/2 \rfloor}$ in the symmetric case. The cases with $|Q_0| > 1$ will require a different (and more subtle) definition of Peierls contour than the one used here for $|Q_0| = 1$.

The foregoing considerations are purely entropic; a more complicated phase diagram, involving tradeoffs between entropy and energy, can presumably be obtained by adding additional couplings into the Hamiltonian (1.4). Suppose, for instance, that in the 3-state Potts antiferromagnet on a plane quadrangulation G where G_0 is a triangulation, we add an explicit ferromagnetic interaction, of strength λ, between adjacent vertices in the sublattice G_1. Then for small λ we expect that the favored ordering at low temperature will be the same as for $\lambda = 0$, namely monocolor on V_0 and

---

10 For *any* connected bipartite graph $G = (V, E)$ with vertex bipartition $V = V_0 \cup V_1$, one can define graphs $G_0 = (V_0, E_0)$ and $G_1 = (V_1, E_1)$ by setting $E_0 = \big\{\{u, v\}\colon\ u, v \in V_0 \text{ and } d_G(u, v) = 2\big\}$ and likewise for E_1. But if G is non-planar, or is planar but not a quadrangulation, it is not clear whether these definitions will be useful.

11 Including to the authors until very recently.

bicolor on V_1; but for large positive λ the favored ordering will instead be monocolor on V_1 and bicolor on V_0. It is then natural to guess that for large β there is either a switchover between the two orderings at some particular value $\lambda_t(\beta)$, or else a pair of phase transitions $\lambda_{t1}(\beta) < \lambda_{t2}(\beta)$ with a disordered phase in-between [and possibly $\lambda_{t1}(\infty) = \lambda_{t2}(\infty)$]. It is an interesting open question to determine the correct qualitative phase diagram in the (λ, β)-plane and the order of the phase transition(s).

1.4. Some further open problems. Here are some further open problems suggested by our work:

1) Prove (or disprove) that

(a) the finite-volume measures $\mu^k_{\Lambda,\beta}$ used in the proof of Theorem 1.1 (see Section 2.3 below) converge as $\Lambda \uparrow V$ (i.e., there is no need to take a subsequence);
(b) the resulting infinite-volume Gibbs measures $\mu_{\beta,k}$ are *extremal* Gibbs measures; and
(c) $\mu_{\beta,k}$ are invariant with respect to the automorphism of the graph G.

2) Prove (or disprove) that for our lattices there are *no more than three* extremal translation-invariant Gibbs measures at small but strictly positive temperature. For this, one would need to control more general boundary conditions than the uniform colorings on V_0 that we have used here.

Please note that at *zero* temperature, there are in fact *more than three* extremal infinite-volume Gibbs measures on the diced lattice, since there exist ground-state configurations τ, similar to the example on $\mathbb{Z}^2$ sketched in footnote 6 above, such that for any finite $\Lambda \subset V$ there exists only one ground state that agrees with τ on $V \backslash \Lambda$ (namely, τ itself). The delta measure on such a ground state is therefore a zero-temperature Gibbs measure; but by the argument sketched at the end of footnote 6, this Gibbs measure is not a limit of positive-temperature Gibbs measures.

It is worth pointing out, however, that this latter argument makes essential use of the fact that the lattice is Euclidean (in particular, its isoperimetric constant is zero). This raises the question whether on hyperbolic lattices there might exist delta-measure zero-temperature Gibbs measures that *are* limits of positive-temperature Gibbs measures.

3) Extend these techniques to the q-state Potts antiferromagnet with $q > 3$ on suitable lattices. For instance, one might hope to prove the existence of an entropy-driven phase transition in the q-state Potts antiferromagnet on $\mathbb{Z}^d$ for suitable pairs (q, d), i.e., for $q <$ some $q_c(\mathbb{Z}^d)$. In this case it is not completely clear, even heuristically, how $q_c(\mathbb{Z}^d)$ should behave as $d \to \infty$. The example of the infinite Δ-regular tree, which has multiple Gibbs measures when $q \le \Delta$ [8] and a unique Gibbs measure when $q \ge \Delta + 1$ [41], suggests that we might have $q_c(\mathbb{Z}^d) \approx 2d$.

1.5. Plan of this paper. The remainder of this paper is organized as follows: In Section 2 we introduce the Peierls-contour representation of our model and sketch the main ideas underlying our proofs. In particular, we formulate the key steps in our proof as precise lemmas (Lemmas 2.1–2.5) that will be proven later, and we show how they together imply Theorems 1.1 and 1.2. In Section 3 we prove Lemmas 2.1 and 2.5 in the zero-temperature case $\beta = \infty$, using a Peierls argument. In Section 4 we extend these proofs to the low-temperature case $\beta \ge \beta_0$, and we also prove the technical Lemmas 2.2 and 2.4. In Section 5 we use a random-cluster argument to deduce positive magnetization (Lemma 2.3). In the Appendix we review the needed theory of infinite graphs.

2. Main structure of the proofs

2.1. Contour model. Our proofs of Theorems 1.1 and 1.2 are based on suitable bounds for finite-volume Gibbs measures, uniform in the system size and in the inverse temperature above a certain value. We will concentrate on finite-volume Gibbs measures with uniform 1 boundary conditions on the sublattice V_0; by symmetry, all statements immediately imply analogous results for boundary conditions 2 or 3. We will always employ finite sets $\Lambda \subset V$ whose external boundary lies entirely in the sublattice V_0, i.e. $\partial\Lambda \subset V_0$. We will also assume that $\Lambda \subset V$ is *simply connected,* by which we mean that both Λ and $V\backslash\Lambda$ are connected in G. Thus, let us fix any configuration τ that equals 1 on V_0, and let $\mu^1_{\Lambda,\beta}$ denote the finite-volume Gibbs measure in Λ with boundary condition τ and inverse temperature $\beta \in [0,\infty]$. (Since $\partial\Lambda \subset V_0$, this measure is the same for all configurations τ that equal 1 on V_0.)

Our proofs are based on a version of Peierls argument relying on a contour reformulation of the measure $\mu^1_{\Lambda,\beta}$. Our goal is to prove that the sublattice V_0 exhibits ferromagnetic order of a suitable kind. Therefore we will define Peierls contours just as one would for studying the ferromagnetic Potts model on G_0. Thus, for any configuration σ, we look only at the restriction of σ to V_0, and we define $E_0(\sigma)$ to be the set of "unsatisfied edges", i.e.

$$E_0(\sigma) \;:=\; \big\{\{u,v\} \in E_0 \colon\; u,v \in V_0 \cap (\Lambda \cup \partial\Lambda) \text{ and } \sigma_u \neq \sigma_v\big\}\;. \tag{2.1}$$

Letting

$$E_1(\sigma) \;:=\; \{e \in E_1 \colon\; e \text{ crosses some } f \in E_0(\sigma)\} \tag{2.2}$$

denote the edges in the dual graph G_1 that cross an edge in $E_0(\sigma)$, we see that edges in $E_1(\sigma)$ correspond to boundaries separating areas where the vertices of G_0 are uniformly colored in one of the colors $1,2,3$. Note that since $\partial\Lambda \subset V_0$ and $\partial\Lambda$ is uniformly colored, every edge of $E_1(\sigma)$ has *both* of its endpoints in Λ.

Since G_0 is a triangulation, each vertex of G_1 is of degree 3. If the three vertices of G_0 surrounding a vertex $v \in V_1$ are colored with three different colors, then one of these vertices must have the same color as v. This is clearly not possible for a ground state σ (i.e., a proper coloring), so at $\beta = \infty$ *at most two* different colors can surround any vertex $v \in V_1$. It follows that at zero temperature, either zero or two edges of $E_1(\sigma)$ emanate from the vertex v. Hence $E_1(\sigma)$ consists of a collection $\Gamma(\sigma)$ of disjoint simple circuits that we call contours.

At positive temperature, we define contours to be connected components of $E_1(\sigma)$, which can be much more complicated than a circuit. Nevertheless, we will show that at low temperatures, contours that are not simple circuits are rare.

2.2. The basic lemmas. Let us now sketch in broad lines the main ideas of our proofs, and formulate a number of precise lemmas, to be proven later, that together will imply our main results. We have seen that uniformly colored areas in the sublattice G_0 are separated by contours in the sublattice G_1, which at zero temperature are simple circuits. The number of different simple circuits of a given length L surrounding a given point is roughly of order α^L, where α is the connective constant of the lattice G_1. Since each vertex in G_1 has degree 3, a contour entering a vertex has two possible directions in which to continue. In view of this, it is easy to see that $\alpha \leq 2$. With a bit more work using quasi-transitivity, this can be improved to $\alpha < 2$. On the other hand, each vertex v in G_1 that lies on a contour is surrounded by vertices in G_0 of two different colors. At zero temperature, this means that there is only one color available for v, compared to two for a vertex in G_1 that does not lie on a contour. As a result, for each contour of length L we have to pay an entropic price 2^{-L}. In view of this, we will prove in Section 3 below that the expected number of contours surrounding a given site is of order $\sum_L \alpha^L 2^{-L}$, which is finite.

Note that this reasoning tells us that the Peierls sum is *finite*, but not necessarily that it is *small*. In a traditional Peierls argument (such as, for example, the proof of [50, Theorem IV.3.14]), one argues that if the Peierls contour sum is smaller than a certain model-dependent threshold (typically a number somewhat less than 1), then the model has spontaneous magnetization. This is indeed how we will prove Theorem 1.2 for the diced lattice. But for the general class of lattices in Theorem 1.1, all one can hope to prove is that the Peierls sum is finite; it need not be small. To handle this situation, we use a trick that we learned from [22, section 6a], where it is used for percolation. We observe that if $\Delta_0 \subset V_0$ is connected in G_0, then Δ_0 is uniformly colored in one color if and only if no contours cut through Δ_0. On the other hand, if Δ_0 is sufficiently large, then by the finiteness of the Peierls sum, Δ_0 is unlikely to be surrounded by a contour. It follows that, *conditional on Δ_0 being uniformly colored* (which is of course a rare event), it is much more likely for Δ_0 to be uniformly colored in the color 1 than in either of the other two colors.

More precisely, for each $k \in \{1, 2, 3\}$ and each finite set $\Delta_0 \subset V_0$, let $\mathcal{J}_{k,\Delta_0}$ denote the event that all sites in Δ_0 have the color k, and let $\mathcal{J}_{\Delta_0} = \bigcup_{k=1}^{3} \mathcal{J}_{k,\Delta_0}$ denote the event that all sites in Δ_0 are colored with the same color. By the arguments sketched above for the case of zero temperature, together with a careful estimate of non-simple contours at small positive temperatures, we are able to prove the following lemma:

Lemma 2.1. (Long-range dependence) *There exists $\beta_0 < \infty$ such that for each $\epsilon > 0$, there exists $M_\epsilon < \infty$ such that for every finite set $\Delta_0 \subset V_0$ that is connected in G_0 and satisfies $|\Delta_0| \geq M_\epsilon$, one has*

$$\mu^1_{\Lambda,\beta}(\mathcal{J}_{1,\Delta_0} \mid \mathcal{J}_{\Delta_0}) \;\geq\; 1 - \epsilon \tag{2.3}$$

uniformly for all $\beta \in [\beta_0, \infty]$ and all simply connected finite sets $\Lambda \supseteq \Delta_0$ such that $\partial\Lambda \subset V_0$.

In order for Lemma 2.1 to be of any use, we need to show that the event on which we are conditioning in (2.3) has positive probability, uniformly in the system size:

Lemma 2.2. (Uniformly colored sets) *Let $\Delta_0 \subset V_0$ be finite and connected in G_0. Then there exists a constant $\delta > 0$ such that*

$$\mu^1_{\Lambda,\beta}(\mathcal{J}_{\Delta_0}) \;\geq\; \delta \tag{2.4}$$

uniformly for all $0 \leq \beta \leq \infty$ and all finite and simply connected $\Lambda \supseteq \Delta_0$ such that $\partial\Lambda \subset V_0$.

Of course δ gets very small as Δ_0 gets large, but we do not care, as we will take Δ_0 to be large but *fixed*.

Let us note that Lemmas 2.1 and 2.2 are sufficient, by themselves, to prove the existence of at least three distinct infinite-volume Gibbs measures at all $\beta \in [\beta_0, \infty]$.[12] These infinite-volume Gibbs measures may or may not have spontaneous magnetization, but they do at least have long-range order of a special kind: namely, they assign unequal probabilities to the (rare) events $\mathcal{J}_{k,\Delta_0}$ ($k = 1, 2, 3$) for some large but finite set Δ_0. This part of the argument is quite general and applies to other models as well, as long as it can be shown that the Peierls sum is finite, even if it is not necessarily small.[13]

But for our particular model, we can actually do better and prove that there is spontaneous magnetization, thanks to the following lemma, which says that if a sufficiently "thick" block is more likely to be uniformly colored in one color than in the other two colors, then the same must be true for single sites within that block. Let us say that a set $\Delta \subset V$ is *thick*[14] if there exists a nonempty finite subset $\Delta_1 \subset V_1$ that is connected in G_1 and such that $\Delta = \{v \in V\colon\ d_G(v, \Delta_1) \leq 1\}$. Then Δ is connected in G, and we have $\Delta_1 = \Delta \cap V_1$; we write $\Delta_0 := \Delta \cap V_0$.

[12] To see this, just follow the proof of Theorem 1.1 given in Section 2.3 below and disregard all references to Lemma 2.3. The bound (2.10) and its analogues for $k = 2, 3$ survive to the (subsequential) infinite-volume limit and hence show that the Gibbs measures $\mu_{k,\beta}$ for $k = 1, 2, 3$ are distinct.

[13] For example, consider the Ising model on $\mathbb{Z}^2$, with the contours constructed as in [29, Lemma 6.14] by taking a circuit on the dual lattice that is the external boundary of the connected (by nearest-neighbor edges in $\mathbb{Z}^2$) region of constant spin containing origin. Then this argument gives the existence of at least two distinct infinite-volume Gibbs measures whenever $e^{-2\beta}\alpha_{\mathrm{SQ}} < 1$ (i.e. $\beta > \frac{1}{2}\log\alpha_{\mathrm{SQ}}$), where α_{SQ} is the connective constant for self-avoiding polygons (or equivalently, self-avoiding walks [52, Corollary 3.2.5]) on the square lattice, and β is the inverse temperature in the standard Ising normalization. (It is known that $\alpha_{\mathrm{SQ}} < 2.679192495$ [62]; the current best numerical estimate is $\alpha_{\mathrm{SQ}} \approx 2.63815853031(3)$ [39].

[14] Notice that the sublattices V_0 and V_1 enter the definition of thickness in an asymmetric way.

Lemma 2.3. (Positive magnetization) *Fix $\beta_0 > 0$ and let $\Delta \subset V$ be thick. Then there exists $\epsilon > 0$ such that for each $v_0 \in \Delta_0$,*

$$\mu^1_{\Lambda,\beta}(\sigma_{v_0} = 1) - \mu^1_{\Lambda,\beta}(\sigma_{v_0} = 2) \;\geq\; \epsilon\big[\mu^1_{\Lambda,\beta}(\mathcal{J}_{1,\Delta_0}) - \mu^1_{\Lambda,\beta}(\mathcal{J}_{2,\Delta_0})\big] \;, \tag{2.5}$$

and similarly, for each $v_1 \in \Delta_1$,

$$\mu^1_{\Lambda,\beta}(\sigma_{v_1} = 2) - \mu^1_{\Lambda,\beta}(\sigma_{v_1} = 1) \;\geq\; \epsilon\big[\mu^1_{\Lambda,\beta}(\mathcal{J}_{1,\Delta_0}) - \mu^1_{\Lambda,\beta}(\mathcal{J}_{2,\Delta_0})\big] \;, \tag{2.6}$$

uniformly for all $\beta \in [\beta_0, \infty]$ and all simply connected finite sets $\Lambda \supseteq \Delta$ such that $\partial\Lambda \subset V_0$.

The proof of Lemma 2.3 is not very complicated but is very much dependent on the specific properties of our Potts model. Inspired by the cluster algorithm introduced in [69,70], we condition on the position of the 3's and use the random-cluster representation for the Ising model of 1's and 2's on the remaining diluted lattice. We show that the difference of probabilities that Δ_0 is uniformly colored in the color 1 or in the color 2 equals the probability that Δ_0 is uniformly colored and there is a 1-2 random-cluster connection between Δ_0 and the boundary of Λ. Using this latter quantity, it is then easy to produce (by a finite-energy argument) a lower bound on the probability that there is a 1-2 random-cluster connection between a fixed lattice site v_0 and the boundary of Λ; and this, in turn, equals the magnetization.

The last main missing ingredient of Theorem 1.1 is the following lemma, which shows that improperly colored edges are rare when β is large; in particular it shows that any limit as $\beta \to \infty$ of the finite-temperature infinite-volume Gibbs measures that we will construct is concentrated on the set $S_{\rm g}$ of ground states.

Lemma 2.4. (Rarity of improperly colored edges) *There exists $C < \infty$ such that*

$$\mu^1_{\Lambda,\beta}\big(\sigma_u = \sigma_v\big) \;\leq\; Ce^{-\beta} \tag{2.7}$$

for all $\beta \in [0,\infty]$, all $\{u,v\} \in E$, and all finite and simply connected $\Lambda \ni u, v$ such that $\partial\Lambda \subset V_0$.

Finally, to prove Theorem 1.2 for the diced lattice, we need the following quantitative bound:

Lemma 2.5. (Explicit Peierls bound for the diced lattice) *If G is the diced lattice, then there exists $C < \infty$ such that*

$$\mu^1_{\Lambda,\beta}(\sigma_{v_0} = 1) \;\geq\; 0.90301 - Ce^{-\beta} \tag{2.8}$$

uniformly for all $\beta \in [0,\infty]$, all $v_0 \in V_0$, and all simply connected finite sets $\Lambda \ni v_0$ such that $\partial\Lambda \subset V_0$.

2.3. Proof of the main theorems, given the basic lemmas. Let us now show how to prove Theorems 1.1 and 1.2, given Lemmas 2.1–2.5.

Proof of Theorem 1.1. Fix $\epsilon > 0$, let β_0, M_ϵ and Δ_0 be as in Lemma 2.1, and let δ be as in Lemma 2.2. Since the colors 2 and 3 play a symmetric role under the measure $\mu^1_{\Lambda,\beta}$, we have

$$\mu^1_{\Lambda,\beta}(\mathcal{J}_{2,\Delta_0} \mid \mathcal{J}_{\Delta_0}) \;=\; \tfrac{1}{2}\big[1 - \mu^1_{\Lambda,\beta}(\mathcal{J}_{1,\Delta_0} \mid \mathcal{J}_{\Delta_0})\big] \tag{2.9}$$

and hence

$$\begin{aligned} \mu^1_{\Lambda,\beta}(\mathcal{J}_{1,\Delta_0}) \;-\; \mu^1_{\Lambda,\beta}(\mathcal{J}_{2,\Delta_0}) \;&=\; \tfrac{1}{2}\big[3\mu^1_{\Lambda,\beta}(\mathcal{J}_{1,\Delta_0} \mid \mathcal{J}_{\Delta_0}) - 1\big]\, \mu^1_{\Lambda,\beta}(\mathcal{J}_{\Delta_0}) \\ &\geq\; \tfrac{1}{2}[3(1-\epsilon) - 1]\delta \;=\; \tfrac{1}{2}(2 - 3\epsilon)\delta \;, \end{aligned} \tag{2.10}$$

which is positive for $\epsilon < 2/3$ (which we henceforth assume). Then, for any $v_0 \in V_0$, we may choose a thick set $\Delta \subset V$ such that $|\Delta_0| \geq M_\epsilon$ and $v_0 \in \Delta_0$. By (2.10) together with Lemma 2.3, there exists $\overline{\epsilon}(v_0) > 0$ such that

$$\mu^1_{\Lambda,\beta}(\sigma_{v_0} = 1) \;-\; \mu^1_{\Lambda,\beta}(\sigma_{v_0} = 2) \;\geq\; \overline{\epsilon}(v_0) \tag{2.11}$$

uniformly for all $\beta \in [\beta_0, \infty]$ and all finite and simply connected $\Lambda \supset \Delta_0$ such that $\partial\Lambda \subset V_0$. Since the measure $\mu^1_{\Lambda,\beta}$ treats the colors 2 and 3 symmetrically, it follows that

$$\mu^1_{\Lambda,\beta}(\sigma_{v_0} = 1) \;-\; \tfrac{1}{2}\bigl[1 - \mu^1_{\Lambda,\beta}(\sigma_{v_0} = 1)\bigr] \;\geq\; \overline{\epsilon}(v_0) \tag{2.12}$$

and hence

$$\mu^1_{\Lambda,\beta}(\sigma_{v_0} = 1) \;\geq\; \tfrac{1}{3} + \tfrac{2}{3}\overline{\epsilon}(v_0)\;. \tag{2.13}$$

Similarly, for any $v_1 \in V_1$, we may choose a thick set Δ such that $|\Delta_0| \geq M_\epsilon$ and $v_1 \in \Delta_1$. An analogous argument then shows that

$$\mu^1_{\Lambda,\beta}(\sigma_{v_1} = 1) \leq \tfrac{1}{3} - \tfrac{2}{3}\overline{\epsilon}(v_1)\;. \tag{2.14}$$

In this argument $\overline{\epsilon}$ depends on v_0 or v_1. But since all the quantities under study are invariant under graph automorphisms, $\overline{\epsilon}$ actually depends only on the type of v_0 or v_1. And since by quasi-transitivity there are only finitely many types, we may choose $\overline{\epsilon}$ such that (2.13) and (2.14) hold uniformly for all $v_0 \in V_0$ and $v_1 \in V_1$.

To construct the desired infinite-volume Gibbs measures, we use a compactness argument. For any $\beta \in [0, \infty]$ and finite $\Lambda \subset V$, let $\overline{\mu}_{\Lambda,\beta} := \mu^1_{\Lambda,\beta} \otimes \delta_{\tau_{V\setminus\Lambda}}$, i.e., if $\sigma \in \{1,2,3\}^V$ is distributed according to $\overline{\mu}_{\Lambda,\beta}$, then $(\sigma_v)_{v\in\Lambda}$ is distributed according to $\mu^1_{\Lambda,\beta}$ and $\sigma_v = \tau_v$ for all $v \in V\backslash\Lambda$. Choose finite and simply connected $\Lambda_n \uparrow V$ such that $\partial\Lambda_n \subset V_0$. Since $S = \{1,2,3\}^V$ is a compact space, the set of measures $\{\overline{\mu}_{\Lambda_n,\beta}\}$ is automatically tight. It follows from [29, Theorem 4.17] that each weak subsequential limit μ_β as $\Lambda_n \uparrow \Lambda$ is an infinite-volume Gibbs measure at inverse temperature β. Taking the limit $\Lambda_n \uparrow V$ in (2.13)/(2.14), we see that

$$\mu_\beta(\sigma_{v_0} = 1) \geq \tfrac{1}{3} + \overline{\epsilon} \qquad \text{for } v_0 \in V_0 \tag{2.15}$$

$$\mu_\beta(\sigma_{v_1} = 1) \leq \tfrac{1}{3} - \overline{\epsilon} \qquad \text{for } v_1 \in V_1 \tag{2.16}$$

Taking the limit $\Lambda_n \uparrow V$ in Lemma 2.4, we obtain

$$\mu_\beta\bigl(\sigma_u = \sigma_v\bigr) \;\leq\; Ce^{-\beta} \qquad \text{for } \{u,v\} \in E\;. \tag{2.17}$$

□

Proof of Theorem 1.2. (a) and (c) follow from the same arguments as in the proof of Theorem 1.1, but with the inequality (2.13) replaced by (2.8).

To prove (b), consider any $v_1 \in V_1 \cap \Lambda$ and let $w_1, w_2, w_3 \in V_0 \cap (\Lambda \cup \partial\Lambda)$ be its neighbors in G. Then the DLR equations for the volume $\{v_1\}$ imply that

$$\mu^1_{\Lambda,\beta}(\sigma_{v_1} = 1 | \sigma_{w_1} = \sigma_{w_2} = \sigma_{w_3} = 1) \;=\; \frac{e^{-3\beta}}{2 + e^{-3\beta}} \tag{2.18a}$$

$$\mu^1_{\Lambda,\beta}(\sigma_{v_1} = 1 | \sigma_{w_1} = \sigma_{w_2} = 1,\, \sigma_{w_3} \neq 1) \;=\; \frac{e^{-2\beta}}{1 + e^{-\beta} + e^{-2\beta}} \tag{2.18b}$$

(we call these the "good" cases). In the "bad" cases (i.e., those with two or three spins $\sigma_{w_i} \neq 1$) we will use only that the conditional probability is ≤ 1. On the other hand, using (a) we can bound the probability of the "bad" cases by

$$\begin{aligned} \mu^1_{\Lambda,\beta}(\text{two or three spins } \sigma_{w_i} \neq 1) &\leq \tfrac{1}{2}\, \mathbb{E}_{\mu^1_{\Lambda,\beta}}(\#\text{ spins } \sigma_{w_i} \neq 1) \\ &\leq \tfrac{1}{2} \times 3 \times (1 - 0.90301 + Ce^{-\beta}) \\ &\leq 0.14549 + C' e^{-\beta}\,, \end{aligned} \tag{2.19}$$

and we bound the probability of the "good" cases trivially by 1. Putting together (2.18a,b) and (2.19), we conclude that $\mu^1_{\Lambda,\beta}(\sigma_{v_1} = 1) \leq 0.14549 + C'e^{-\beta} + e^{-2\beta}$. □

3. The zero-temperature case

In the present section, we will prove Lemmas 2.1 and 2.5 in the zero-temperature case $\beta = \infty$. Then, in Section 4, we will show how our arguments can be adapted to cover the (more complicated) case of low positive temperatures; there we will also prove the technical Lemmas 2.2 and 2.4. The proof of Lemma 2.3 is postponed to Section 5.

3.1. Contour model for zero temperature. Let $\Lambda \subset V$ be finite and simply connected in G and such that $\partial\Lambda \subset V_0$. Recall that $\mu^1_{\Lambda,\infty}$ is the uniform distribution on the set $(S_{\rm g})^1_\Lambda$ of all proper 3-colorings of $\Lambda \cup \partial\Lambda$ that take the color 1 on $\partial\Lambda$, i.e. all configurations $\sigma \in \{1,2,3\}^{\Lambda\cup\partial\Lambda}$ such that $\sigma_u \neq \sigma_v$ for all $u, v \in \Lambda \cup \partial\Lambda$ with $\{u,v\} \in E$ and $\sigma_v = 1$ for all $v \in \partial\Lambda$. Since $\partial\Lambda \subset V_0$, the set $(S_{\rm g})^1_\Lambda$ is nonempty: for instance, it includes all configurations in which all sites of V_0 are colored 1 and all sites of V_1 are colored 2 or 3.

As explained in Section 2.1, for any $\sigma \in (S_{\rm g})^1_\Lambda$, we let $E_1(\sigma)$ be the collection of edges in G_1 that separate areas where the vertices of G_0 are uniformly colored in one of the colors $1, 2, 3$. And since at zero temperature at most two different colors on V_0 can meet at any vertex in V_1, the set $E_1(\sigma)$ consists of a collection $\Gamma(\sigma)$ of disjoint simple circuits that we call contours. [This is what makes the zero-temperature case so easy to handle. At positive temperature, a connected component of $E_1(\sigma)$ can be much more complicated than a circuit: see Section 4 below.]

In the zero-temperature case, therefore, we use the term *contour* to denote any simple circuit in G_1. We write $|\gamma|$ to denote the length of a contour γ, defined as the number of its edges (or equivalently the number of its vertices). For a collection Γ of disjoint contours, we write $|\Gamma| := \sum_{\gamma\in\Gamma} |\gamma|$ for the total length of the contours in Γ, and $\#\Gamma$ for the number of contours in Γ. Each contour γ divides V_0 into two connected (in the sense of G_0) components, of which one is infinite and the other is finite and simply connected. We call these the *exterior* $\mathrm{Ext}(\gamma)$ and *interior* $\mathrm{Int}(\gamma)$ of γ, respectively. We will say that a contour γ *surrounds* Δ_0 if $\Delta_0 \subseteq \mathrm{Int}(\gamma)$. We say that a contour lies in Λ if all its vertices are in $V_1 \cap \Lambda$. Note that if γ lies in Λ, then by our assumption that Λ is simply connected, we have $\mathrm{Int}(\gamma) \subseteq \Lambda$.

To each configuration $\sigma \in (S_{\rm g})^1_\Lambda$, there thus corresponds a unique collection $\Gamma(\sigma)$ of disjoint contours in Λ. Conversely, to each collection Γ of disjoint contours, there are $2^{\#\Gamma} 2^{|V_1\cap\Lambda| - |\Gamma|}$ distinct configurations $\sigma \in (S_{\rm g})^1_\Lambda$ that yield the collection $\Gamma(\sigma) = \Gamma$. Here the first and second factor are the number of restrictions $\sigma_{V_0\cap\Lambda}$ and $\sigma_{V_1\cap\Lambda}$, respectively, that are consistent with the specified collection of contours and the fixed boundary condition $\sigma_v = 1$ for $v \in \partial\Lambda$. To understand the first factor, observe that in passing through any contour (from outside to inside) we have 2 $(= q-1)$ independent alternatives for the choice of the color on V_0 just inside the contour. The second factor comes from the fact that there are either one $(= q-2)$ or two $(= q-1)$ colors available for a vertex in G_1, depending on whether this vertex lies on a contour or not. Notice that, given Γ, this latter number is independent of the configuration $\sigma_{V_0\cap\Lambda}$.

Let us therefore introduce the probability measure

$$\nu_\Lambda(\Gamma) \;=\; \frac{1}{Z_\Lambda} \, 2^{\#\Gamma} \, 2^{-|\Gamma|} \tag{3.1}$$

on the set of all collections Γ of disjoint contours in Λ, where $Z_\Lambda = \sum_\Gamma 2^{\#\Gamma} \, 2^{-|\Gamma|}$ is the normalizing constant. We have just shown that, under the probability measure $\mu^1_{\Lambda,\infty}$, the contour configuration $\Gamma(\sigma)$ is distributed according to ν_Λ.

The first step in any Peierls argument is to obtain an upper bound on the probability that Γ contains a given contour γ:

Lemma 3.1. (Bound on probability of a contour) *Let $\Lambda \subset V$ be a simply connected finite set such that $\partial\Lambda \subset V_0$. Then, for any contour γ in Λ,*

$$\nu_\Lambda(\{\Gamma\colon\; \gamma \in \Gamma\}) \;\le\; \frac{2^{1-|\gamma|}}{1 + 2^{1-|\gamma|}} \; . \tag{3.2}$$

Proof. We have

$$\begin{aligned}
\nu_\Lambda(\{\Gamma\colon\; \gamma \in \Gamma\}) \;&=\; \sum_{\Gamma \ni \gamma} \nu_\Lambda(\Gamma) \;=\; 2^{1-|\gamma|} \sum_{\Gamma \ni \gamma} \nu_\Lambda(\Gamma \backslash \{\gamma\}) \\
&\le\; 2^{1-|\gamma|} \sum_{\Gamma \not\ni \gamma} \nu_\Lambda(\Gamma) \;=\; 2^{1-|\gamma|} \bigl[1 - \nu_\Lambda(\{\Gamma\colon\; \gamma \in \Gamma\})\bigr] \; ,
\end{aligned} \tag{3.3}$$

which proves (3.2). □

Now let $\Delta_0 \subseteq \Lambda \cap V_0$ be connected in G_0. Let us say that a contour γ *cuts* Δ_0 if γ contains an edge that separates some pair of vertices $v, w \in \Delta_0$ that are adjacent in G_0. Then, obviously, the event that Δ_0 is uniformly colored corresponds to the event that no contour $\gamma \in \Gamma$ cuts Δ_0. Let $\nu_{\Lambda|\Delta_0}$ denote the measure ν_Λ from (3.1) conditioned on this event. Let $S_{\Delta_0}(\Gamma)$ denote the number of contours in a contour configuration Γ that surround Δ_0. We obtain a lower bound for the conditional probability in (2.3) by writing

$$\begin{aligned}
&\mu^1_{\Lambda,\infty}(\mathcal{J}_{1,\Delta_0} \,|\, \mathcal{J}_{\Delta_0}) \;\ge\; \nu_{\Lambda|\Delta_0}(\{\Gamma\colon\; S_{\Delta_0}(\Gamma) = 0\}) \\
&\quad \ge 1 - \sum_\Gamma \nu_{\Lambda|\Delta_0}(\Gamma) S_{\Delta_0}(\Gamma) \;=\; 1 - \sum_{\gamma\colon\; \mathrm{Int}(\gamma) \supseteq \Delta_0} \nu_{\Lambda|\Delta_0}(\{\Gamma\colon\; \gamma \in \Gamma\}) \; .
\end{aligned} \tag{3.4}$$

Then the probability that Γ contains a given contour γ is bounded under $\nu_{\Lambda|\Delta_0}$ in exactly the same way as it was bounded under ν_Λ [cf. (3.3)], yielding

$$\nu_{\Lambda|\Delta_0}(\{\Gamma\colon\; \gamma \in \Gamma\}) \;\le\; \frac{2^{1-|\gamma|}}{1 + 2^{1-|\gamma|}} \tag{3.5}$$

for every γ that surrounds Δ_0. Inserting this into (3.4) yields:

Lemma 3.2. (Peierls bound for zero temperature) *Let $\Lambda \subset V$ be a simply connected finite set such that $\partial\Lambda \subset V_0$, and let $\Delta_0 \subseteq \Lambda \cap V_0$ be connected in G_0. Then*

$$1 - \mu^1_{\Lambda,\infty}(\mathcal{J}_{1,\Delta_0} \,|\, \mathcal{J}_{\Delta_0}) \;\le\; \sum_{\gamma\colon\; \mathrm{Int}(\gamma) \supseteq \Delta_0} \frac{2^{1-|\gamma|}}{1 + 2^{1-|\gamma|}} \;=\; \sum_{L=3}^{\infty} N_{\Delta_0}(L) \, \frac{2^{1-L}}{1 + 2^{1-L}} \; , \tag{3.6}$$

where $N_{\Delta_0}(L)$ denotes the number of contours of length L surrounding Δ_0.

Our proofs of Lemmas 2.1 and 2.5 in the zero-temperature case will be based on the estimate (3.6) and suitable bounds on the numbers $N_{\Delta_0}(L)$.

Remarks. 1. In the special case that Δ_0 is a singleton, the event $\mathcal{J}_{\Delta_0}$ is trivially fulfilled and the conditional probability in (3.6) reduces to an unconditional probability.

2. The simpler but slightly weaker bound

$$1 - \mu^1_{\Lambda,\infty}\big(\mathcal{J}_{1,\Delta_0} \,\big|\, \mathcal{J}_{\Delta_0}\big) \;\le\; \sum_{L=3}^{\infty} N_{\Delta_0}(L)\, 2^{1-L} \tag{3.7}$$

is sufficient for nearly all purposes. Indeed, even for quantitative bounds the difference between (3.6) and (3.7) is very small: for instance, when G is the diced lattice and G_1 is the hexagonal lattice, we have $L \ge 6$, so one sees immediately that the difference between (3.6) and (3.7) cannot be more than about 3%. See also the proof of Lemma 2.5 for $\beta = \infty$ in Section 3.4 below.

3.2. Bounds on contours for zero temperature. The main ingredient in the proof of Lemma 2.1 will be a bound on the number of simple circuits in G_1 of a given length surrounding a given vertex in G_0. We start by bounding the number of self-avoiding paths in G_1, or more generally in quasi-transitive graphs of bounded degree. We then use this bound to obtain a bound on self-avoiding polygons, i.e. simple circuits.

Let $H = (V, E)$ be any graph. It will be convenient to view H as a directed graph, by introducing a pair of directed edges (one in each direction) corresponding to each edge of the undirected graph H. So let A be the set of directed edges of H, i.e., A is the set of all ordered pairs (v, w) of vertices such that $\{v, w\} \in E$. By definition, a *self-avoiding path* in G of length n is a finite sequence of vertices $v_0, \ldots, v_n \in V$, all different from each other, such that $(v_{k-1}, v_k) \in A$ for all $k = 1, \ldots, n$. We call (v_0, v_1) the *starting edge* and (v_{n-1}, v_n) the *final edge* of the path. For $n \ge 1$ and $a, b \in A$, we denote by $C_n(a, b)$ the number of self-avoiding paths in G of length n with starting edge a and final edge b. We then set

$$C_n(a) \;:=\; \sum_{b\in A} C_n(a,b) \quad \text{and} \quad C^*_n \;:=\; \sup_{a\in A} C_n(a)\,. \tag{3.8}$$

Lemma 3.3. (Exponential bound on self-avoiding paths) *Let $H = (V, E)$ be an infinite connected graph in which each vertex has degree at most k. Then the limit*

$$\alpha(H) \;:=\; \lim_{n\to\infty} (C^*_n)^{1/n} \tag{3.9}$$

exists and equals $\inf\limits_{n\ge 1} (C^*_{n+1})^{1/n}$; it satisfies $1 \le \alpha(H) \le k-1$. Furthermore, if H is quasi-transitive and is anything other than a tree in which every vertex has degree k, then $\alpha(H) < k - 1$.

Proof. For $m, n \ge 1$ and $a, c \in A$ we have

$$C_{m+n-1}(a,c) \;\le\; \sum_{b\in A} C_m(a,b)\, C_n(b,c) \tag{3.10}$$

because any self-avoiding path of length $m+n-1$ can be decomposed uniquely into a first m steps and a last n steps, each of which is a self-avoiding path, which overlap in a single directed edge (here called b). This implies the submultiplicativity

$$C^*_{m+n-1} \;\le\; C^*_m\, C^*_n\,. \tag{3.11}$$

Defining $D_n = C^*_{n+1}$ for $n \geq 0$, we see that $n \mapsto \log D_n$ is subadditive, which implies (see, e.g., [51, Theorem B.22]) that the limit

$$\alpha(H) \;:=\; \lim_{n\to\infty} D_n^{1/n} \;=\; \inf_{n\geq 1} D_n^{1/n} \;=\; \inf_{n\geq 1} \, (C^*_{n+1})^{1/n} \tag{3.12}$$

exists, with $0 \leq \alpha(H) < \infty$.

By Lemma A.2(a), there exists an infinite self-avoiding path $(v_0, v_1, v_2, \ldots)$; so taking $a = (v_0, v_1)$ we see that $C_n(a) \geq 1$ for all $n \geq 1$. Hence $\alpha(H) \geq 1$.

Since each $v \in V$ is of degree at most k, self-avoidance trivially implies that

$$C^*_{n+1} \;\leq\; (k-1)^n \;, \tag{3.13}$$

so that $\alpha(H) \leq k - 1$.

If H is anything other than a k-regular tree, then since H is connected, for each $a \in A$ there exists an integer m (depending only on the equivalence class of a under the automorphism group of H) such that $C_{m+1}(a) < (k-1)^m$: it suffices to walk to a vertex of degree $< k$ and then one step more, or else walk into and around a circuit. Using the submultiplicativity (3.11) together with (3.13), it follows that $C_{n+1}(a) < (k-1)^n$ for all $n \geq m$. If now H is (vertex-)quasi-transitive, then Lemma A.1 tells us that it is also directed-edge-quasi-transitive, i.e. there are finitely many equivalence classes of directed edges, so we can choose an m that works for all $a \in A$. It follows that $C^*_{n+1} < (k-1)^n$ for some n (in fact for all sufficiently large n), which shows that the infimum in (3.12) is strictly less than $k - 1$. $\Box$

Remark. Most of this proof can alternatively be carried out in terms of the more familiar vertex-to-vertex counts $c_n(u,v)$ for $n \geq 0$ and the corresponding quantities $c^*_n = \sup_{u\in V} \sum_{v\in V} c_n(u,v)$. (Since $C^*_n \leq c^*_n \leq kC^*_n$, the two counts have identical asymptotic growth.) Indeed for $m, n \geq 0$ we have

$$C_{m+n}(u,w) \;\leq\; \sum_{v\in V} C_m(u,v)\, C_n(v,w) \tag{3.14}$$

and hence $c^*_{m+n} \leq c^*_m \, c^*_n$, from which it follows that

$$\alpha(H) \;=\; \lim_{n\to\infty} (c^*_n)^{1/n} \;=\; \inf_{n\geq 1} \, (c^*_n)^{1/n} \tag{3.15}$$

exists. But it is more difficult in this framework to prove that $\alpha(H) < k - 1$, since the bound $c^*_n \leq k(k-1)^{n-1}$ has an extra factor $k/(k-1)$ that we must somehow overcome. It is for this reason that we found it convenient to work with directed edges instead of vertices.

It follows from (3.9) that for each $\epsilon > 0$ there exists $K_\epsilon < \infty$ such that

$$C^*_n \;\leq\; K_\epsilon \, [\alpha(H) + \epsilon]^n \qquad \text{for all } n \geq 0 \;. \tag{3.16}$$

Now let $G = (V,E)$ be as in Theorem 1.1 and let $G_0 = (V_0, E_0)$ and $G_1 = (V_1, E_1)$ be its sublattices. Recall from Section 3.1 that $N_{\Delta_0}(L)$ denotes the number of simple circuits of length L in G_1 surrounding a set $\Delta_0 \subset V_0$. Lemma 3.3 applied to $H = G_1$ implies the following bound on $N_{\Delta_0}(L)$:

Lemma 3.4. (Exponential bound on circuits surrounding a point) *We have $\alpha(G_1) < 2$. Moreover, for every $\epsilon > 0$, there exists a constant $C_\epsilon < \infty$ such that*

$$N_{\{v\}}(L) \;\leq\; C_\epsilon \, [\alpha(G_1) + \epsilon]^L \tag{3.17}$$

for all $v \in V_0$ and all $L \geq 1$.

PROOF. Since every vertex in G_1 has degree 3 and G_1 is not a tree (indeed, each vertex in G_0 is surrounded by a circuit in G_1), it follows from Lemma 3.3 that $\alpha(G_1) < 2$.

Since G is infinite, connected and locally finite, it is not hard to show [see Lemma A.2(a) in the Appendix] that for each $v \in V_0$ we can find an infinite self-avoiding path $\pi = (v_0, v_1, \ldots)$ in G_0 starting at $v_0 = v$ such that the graph distance (in G_0) of v_n to v is n. It is not hard to see that any simple circuit surrounding v must cross some edge of π. With a bit more work, we can get a quantitative bound on how far this edge can be from the starting point of π. Indeed, it follows from Proposition A.5 that there exists a constant $K < \infty$, depending only on the graph G_0, such that any simple circuit of length L surrounding v must cross one of the first N edges of π, where

$$N \;:=\; 1 + K + \tfrac{1}{2}(\tfrac{3}{2} - 1)L \;=\; 1 + K + L/4 \,. \tag{3.18}$$

So let γ be a simple circuit of length L surrounding v. Let (v_{k-1}, v_k) be the first edge of π that is crossed by γ, and let a be the corresponding (dual) edge in γ. We can view a as a directed edge by agreeing that we turn (v_{k-1}, v_k) anticlockwise to get a. Then we can specify γ completely by specifying the first edge of π to be crossed by γ and by specifying the self-avoiding path formed by the first $L-1$ edges of γ, starting with a. By (3.18), this yields the bound

$$N_{\{v\}}(L) \;\le\; (1 + K + L/4)\, C^*_{L-1} \,. \tag{3.19}$$

By (3.16), the claim follows: it suffices to absorb the factor $(1+K+L/4)$ into a change of the base of the exponential term. □

3.3. Long-range dependence for zero temperature. We are now ready to prove Lemma 2.1 for zero temperature.

PROOF OF LEMMA 2.1 FOR $\beta = \infty$. It follows from Proposition A.5 that for each $L_0 < \infty$, there exists $M < \infty$ such that each finite, G_0-connected set $\Delta_0 \subset V_0$ with $|\Delta_0| \ge M$ has the property that any simple circuit in G_1 surrounding Δ_0 must be of length at least L_0.

Then the weak Peierls bound (3.7) and Lemma 3.4 imply that for any $\epsilon > 0$ there exists $C_\epsilon < \infty$ such that for every finite and simply connected $\Lambda \supset \Delta_0$ with $\partial\Lambda \subset V_0$, we have

$$\mu^1_{\Lambda,\infty}\big(\mathcal{J}_{1,\Delta_0} \,\big|\mathcal{J}_{\Delta_0}\big) \;\ge\; 1 - C_\epsilon \sum_{L=L_0}^{\infty} 2^{1-L}[\alpha(G_1) + \epsilon]^L. \tag{3.20}$$

Since $\alpha(G_1) < 2$, by choosing first ϵ small enough and then L_0 large enough (and M appropriately), we can make the conditional probability in (3.20) as close to 1 as we wish, uniformly in Λ. □

3.4. Quantitative bound for the diced lattice. Let p_L denote the number of simple circuits (i.e., self-avoiding polygons) of length L in the hexagonal lattice, modulo translation. And let q_L denote the number of simple circuits of length L in the hexagonal lattice that surround a given vertex of the triangular lattice. We have the following bounds:

Lemma 3.5. (Supermultiplicativity of hexagonal-lattice polygons) *The number p_L of hexagonal-lattice self-avoiding polygons of length L, modulo translation, satisfies*

$$p_{L+M-2} \;\ge\; p_L\, p_M \,. \tag{3.21}$$

Corollary 3.6. (Bound on hexagonal-lattice circuits surrounding a point) *The number of simple circuits in the hexagonal lattice G_1 surrounding a given vertex in G_0 is bounded as*

$$q_L \;\le\; (L^2/36)\,(2+\sqrt{2})^{(L-2)/2}. \tag{3.22}$$

Proof of Lemma 3.5. We use concatenation: Consider two polygons γ_1 and γ_2 contributing to p_L and p_M, respectively. Let $(x, x+e_2)$ be the highest vertical edge of γ_1 in its rightmost column, and let $(y, y+e_2)$ be the lowest vertical edge of γ_2 in its leftmost column, where $e_1 := (1,0)$ and $e_2 := (0,1)$ denote the natural basisvectors of $\mathbb{R}^2$. Uniting the polygon γ_2 with γ_1 shifted by $y-x$ and erasing the edges $(y, y+e_2)$, we get a contour $\gamma = T_{y-x}(\gamma_1) \cup \gamma_2 \setminus (y, y+e_2)$ contributing to p_{L+M-2}. To complete the argument, we must show that different choices of γ_1 and/or γ_2 lead to a different γ (modulo translation), i.e., we can reconstruct γ_1 and γ_2 (modulo translation) from γ. To this aim, we observe that $(y, y+e_2)$ is the only vertical edge in its column that cuts the interior of γ. Also, if another column cuts the interior of γ in a single edge, then the contours γ_1' and γ_2' obtained by cutting γ at this edge into a left and right piece will have lengths different from L and M. Thus, for fixed L and M, each different (modulo translations) ordered pair (γ_1, γ_2) of polygons of lengths L and M yields a different (modulo translations) polygon of length $L+M-2$. $\square$

Proof of Corollary 3.6. The proof combines three ingredients. The first is the fact, conjectured in [55] and proven in [21], that the connective constant of the hexagonal lattice is exactly $\alpha = \sqrt{2+\sqrt{2}} \approx 1.847759$. The second ingredient is the isoperimetric inequality for the hexagonal lattice: the number of faces surrounded by a circuit of length L is at most $L^2/36$. The third ingredient is a bound on the number p_L of L-step hexagonal-lattice self-avoiding polygons modulo translation in terms of the connective constant α for self-avoiding walks on the hexagonal lattice, namely [46]

$$p_L \;\le\; \alpha^{L-2}\,. \tag{3.23}$$

Indeed, the supermultiplicativity $p_{L+M-2} \ge p_L p_M$ implies, by standard arguments, that $\alpha_{\rm SAP} = \lim_{L\to\infty} (p_L)^{1/L}$ exists and that $p_L \le (\alpha_{\rm SAP})^{L-2}$. On the other hand, since $p_L \le c_{L-1}/(2L)$ where c_n is the number of n-step self-avoiding paths starting at a given vertex, we manifestly have $\alpha_{\rm SAP} \le \alpha$. $\square$

Remarks. 1. The supermultiplicativity $p_{L+M-2} \ge p_L p_M$ for the hexagonal lattice is stronger than the inequality $p_{L+M} \ge p_L p_M$ that holds for the square lattice [52, Theorem 3.2.3]. As a consequence, we are able to prove $p_L \le \alpha^{L-2}$ rather than just $p_L \le \alpha^L$.

2. For self-avoiding paths and polygons on $\mathbb{Z}^d$ it is known [52, Corollary 3.2.5] that $\alpha_{\rm SAP} = \alpha$. The same presumably holds also for the hexagonal lattice and for other lattices periodically embedded in Euclidean space, but we are not aware of any proof. Since we need only an upper bound on $\alpha_{\rm SAP}$, we have refrained from addressing this question. Note also that $\alpha_{\rm SAP} < \alpha$ on hyperbolic lattices (with the possible exception of eight such lattices) [53], so the equality $\alpha_{\rm SAP} = \alpha$ is a somewhat delicate matter.

Proof of Lemma 2.5 for $\beta = \infty$. We use the explicit values of q_L for $L = 6, 8, \ldots, 140$ obtained by Jensen's computer-assisted enumerations [40][15] together with the bound (3.22) for even $L \ge 142$. From [40] we get

$$\sum_{L=6}^{140} q_L\, 2^{-L} \;=\; \frac{220742338993408811335836925197618724052 49}{2^{139}} \;<\; 0.03168\,. \tag{3.24}$$

On the other hand, we have

$$\sum_{\text{even } L \ge 142} (L^2/36)\, (2+\sqrt{2})^{(L-2)/2}\, 2^{-L} \;=\; \frac{(2+\sqrt{2})^{70}\, (2907 + 1531\sqrt{2})}{9 \cdot 2^{139}} \;<\; 0.01731\,. \tag{3.25}$$

[15] The relevant series is called there the “first area-weighted moment” for honeycomb-lattice polygons and is contained in the file `hcsapmom1.ser`.

Putting these together, we have

$$\sum_{L=6}^{\infty} q_L \, 2^{-L} \;<\; 0.04899 \; . \tag{3.26}$$

Inserting this into the weak Peierls bound (3.7) specialized to $\varDelta_0 = \{v\}$, we obtain

$$\mu^1_{\varLambda,\infty}(\sigma_v = 1) \;>\; 1 \,-\, 2(0.04899) \;=\; 0.90202 \; . \tag{3.27}$$

A slight improvement can be obtained by using (3.6) in place of (3.7): we have

$$\sum_{L=6}^{140} q_L \, \frac{2^{-L}}{1 + 2^{1-L}} \;<\; 0.03119 \; . \tag{3.28}$$

(The improvement in the tail sum $L \geq 142$ is of course utterly negligible.) The final result (3.27) is then improved from 0.90202 to 0.90301. $\square$

Remarks. 1. Jensen [40] conjectured, based on his enumerations for $L \leq 140$, that the large-L asymptotic asymptotic behavior of q_L is

$$q_L \;=\; \frac{1}{4\pi} \, (2 + \sqrt{2})^{L/2} \, L^{-1} \big[1 + o(1)\big] \; . \tag{3.29}$$

(At $L = 140$ the exact value for q_L is already within 0.4% of this asymptotic form.) Using this formula in place of the bound (3.22), we find for the tail

$$\sum_{\text{even } L \geq 142} q_L \, 2^{-L} \;\approx\; 4.7 \times 10^{-8} \;\; (\ll 0.01731) \; . \tag{3.30}$$

It follows that *if* we could know q_L exactly for all L, then our Peierls argument using (3.6) would be capable of proving a lower bound 0.93762 in (3.27). This should be compared with the actual zero-temperature value 0.957597 ± 0.000004 obtained by Monte Carlo simulations [46].[16]

2. When [46] was written, the exact result $\alpha = \sqrt{2 + \sqrt{2}} \approx 1.847759$ was not yet a rigorous theorem, so we used instead the bound $\alpha < 1.868832$ due to Alm and Parviainen [1]. Then, to get a sufficient final estimate, the additional factor α^{-2} from the improved bound (3.23) implied by stronger supermultiplicativity (see Remark 1 after the proof of Corollary 3.6) was crucial.

4. The positive-temperature case

In this section we extend the Peierls argument to positive temperature, allowing us to complete the proof of Lemmas 2.1 and 2.5. We also prove the technical Lemmas 2.2 and 2.4.

4.1. Contour model for positive temperature. As before, we consider a graph $G = (V, E)$ satisfying the conditions of Theorem 1.1 and take a finite and simply connected set $\varLambda \subset V$ such that $\partial\varLambda \subset V_0$. Our aim is to derive bounds on the probabilities of certain events under the finite-volume Gibbs measures $\mu^1_{\varLambda,\beta}$ which correspond to uniform color-1 boundary conditions on $\partial\varLambda$.

We recall from Section 2.1 that every color configuration σ on $\varLambda$ (2.1)/(2.2) a collection $E_1(\sigma)$ of edges in the sublattice G_1 that separate differently colored vertices in G_0 (or equivalently faces in G_1). Since $\partial\varLambda \subset V_0$ and $\partial\varLambda$ is uniformly colored (in color 1), each edge of $E_1(\sigma)$ has both its endvertices in $V_1 \cap \varLambda$. In general, we define contours to be connected components of $E_1(\sigma)$. [More precisely, we define contours to be the connected components of the graph $(V_1 \cap \varLambda,\, E_1(\sigma))$ other

[16] See footnote 7 above.

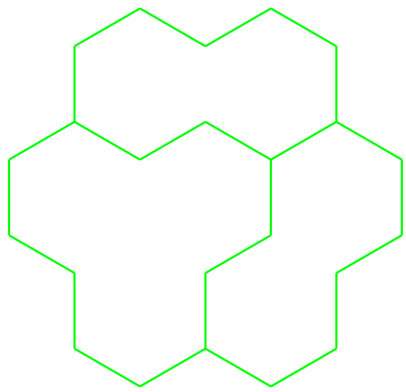

Fig. 4.1. A contour γ with $\chi(\gamma) = 0$.

than isolated vertices.] If σ is a ground state, then at each $v \in V_1 \cap \Lambda$, either zero or two edges of $E_1(\sigma)$ are incident, hence the connected components of $E_1(\sigma)$ are simple circuits in G_1. But for general color configurations σ, the connected components of $E_1(\sigma)$ may be more complicated. In particular, it is possible that three edges of $E_1(\sigma)$ are incident on a vertex $v \in V_1 \cap \Lambda$. Recall that a connected graph is called *bridgeless* (or *2-edge-connected*) if it contains no *bridges* (i.e., single edges the removal of which disconnects the graph). We observe that for any color configuration σ, the connected components of $E_1(\sigma)$ must be bridgeless, since otherwise there would be a uniformly colored region of G_0 that bounds such a bridge on both sides, contradicting the definition of $E_1(\sigma)$.

In view of this, in the positive-temperature model let us define a *contour* to be a finite connected bridgeless subgraph γ of G_1 containing at least one edge. Note that each vertex of such a contour has degree 2 or 3. It is easy to see that the number of vertices of degree 3 must be even. We let $|\gamma|$ denote the number of edges of γ, to which we will refer as the *length* of γ. We let $t(\gamma)$ be the number such that γ has $2t(\gamma)$ vertices of degree 3. Then γ has $|\gamma| - 3t(\gamma)$ vertices of degree 2. Moreover, γ divides V_0 into $2 + t(\gamma)$ connected components, of which one is infinite and the others are finite and simply connected. We say that a contour γ *surrounds* a set $\Delta_0 \subset V_0$, denoted as $\gamma \circlearrowleft \Delta_0$, if Δ_0 is contained in one of the finite components. We call the infinite component the *exterior* $\mathrm{Ext}(\gamma)$ of γ, and we refer to the union of all the finite components as the *interior* $\mathrm{Int}(\gamma)$ of γ. [Please note that saying that γ surrounds a set Δ_0 is stronger than saying that $\Delta_0 \subseteq \mathrm{Int}(\gamma)$, since "surrounding" is defined as Δ_0 being entirely contained in *one* of the finite components.] Given that the exterior of γ is colored in one particular color, we let $\chi(\gamma)$ denote the number of possible three-colorings of the connected components of $\mathrm{Int}(\gamma)$ in such a way that along each edge of γ, two different colors meet.[17] Please note that it is possible to have $\chi(\gamma) = 0$: see Figure 4.1. Obviously, such contours are "not allowed", and we shall soon see that their probability is zero.[18] Finally, let us observe that $\chi(\gamma) \le 2^{t(\gamma)+1}$.

We now claim that if σ is distributed according to $\mu^1_{\Lambda,\beta}$, and $\Gamma(\sigma)$ is the collection of connected components of $(V_1 \cap \Lambda,\, E_1(\sigma))$ other than isolated vertices, then $\Gamma(\sigma)$ is distributed according to

---

[17] Otherwise put, $\chi(\gamma)$ is $\frac{1}{3}$ times the number of proper 3-colorings of the dual graph γ^*.

[18] We could, if we wanted, redefine the term "contour" to include only those having $\chi(\gamma) > 0$. But there is little to be gained from complicating the definition in this way, since our counting of contours (Lemma 4.2 below) is too crude to distinguish between those having $\chi(\gamma) > 0$ or $\chi(\gamma) = 0$.

the law

$$\nu_{\Lambda,\beta}(\Gamma) \;=\; \frac{1}{Z_{\Lambda,\beta}} \prod_{\gamma\in\Gamma} \chi(\gamma)\, p_\beta^{|\gamma|}\, q_\beta^{t(\gamma)} \;, \tag{4.1}$$

where $Z_{\Lambda,\beta}$ is a normalizing constant and

$$p_\beta := \frac{1+e^{-\beta}+e^{-2\beta}}{2+e^{-3\beta}} \tag{4.2a}$$

$$q_\beta := \frac{9e^{-2\beta}(2+e^{-3\beta})}{(1+e^{-\beta}+e^{-2\beta})^3} \tag{4.2b}$$

To see this, note that there are $\prod_{\gamma\in\Gamma}\chi(\gamma)$ ways of coloring the sites in $V_0\cap\Lambda$ in a way that is consistent with Γ. Given a coloring of $V_0\cap\Lambda$, summing the probabilities of all possible colorings of $V_1\cap\Lambda$ yields for each site in $V_1\cap\Lambda$ a factor

$$1+1+e^{-3\beta}, \quad 1+e^{-\beta}+e^{-2\beta} \quad \text{or} \quad e^{-\beta}+e^{-\beta}+e^{-\beta} \tag{4.3}$$

depending on whether the site has neighbors with one, two or three different colors, respectively. These cases correspond, respectively, to sites not on a contour, sites of degree 2 on a contour, and sites of degree 3 on a contour. Absorbing the factor $1+1+e^{-3\beta}$ into the normalization contant $Z_{\Lambda,\beta}$, we get a factor $(1+e^{-\beta}+e^{-2\beta})/(2+e^{-3\beta})$ for each of the $|\gamma|-3t(\gamma)$ sites of degree 2, and a factor $3e^{-\beta}/(2+e^{-3\beta})$ for each of the $2t(\gamma)$ sites of degree 3. Putting this all together, we arrive at (4.1)/(4.2).

In the limit $\beta\to\infty$, we have $p_\beta\to\frac{1}{2}$ and $q_\beta\to 0$; in particular, the only countours that get nonzero weight in this limit are simple circuits, for which $\chi(\gamma)=2$. Then the contour law (4.1) reduces to (3.1), as expected.

More generally, it can be easily verified that p_β decreases monotonically from 1 to $\frac{1}{2}$ as β runs from 0 to ∞, and behaves for large β as $\frac{1}{2}+O(e^{-\beta})$; and that q_β decreases monotonically from 1 to 0 as β runs from 0 to ∞, and behaves for large β as $O(e^{-2\beta})$.

By the same arguments as in (3.4)–(3.5), and using $\chi(\gamma)\le 2^{t(\gamma)+1}$, we find:

Lemma 4.1. (Peierls bound for positive temperature) *Let $\Lambda\subset V$ be a simply connected finite set such that $\partial\Lambda\subset V_0$, and let $\Delta_0\subseteq\Lambda\cap V_0$ be connected in G_0. Then*

$$1-\mu^1_{\Lambda,\beta}(\mathcal{J}_{1,\Delta_0}\,|\,\mathcal{J}_{\Delta_0}) \;\le\; \sum_{T=0}^{\infty}\sum_{L=3}^{\infty} N_{\Delta_0}(L,T)\,\frac{2^{T+1}\,p_\beta^L\,q_\beta^T}{1+2^{T+1}\,p_\beta^L\,q_\beta^T} \;, \tag{4.4}$$

where $N_{\Delta_0}(L,T)$ denotes the number of contours γ surrounding Δ_0 satisfying $|\gamma|=L$ and $t(\gamma)=T$.

4.2. Bounds on contours for positive temperature. In this section, we prove Lemmas 2.1 and 2.5. We first need to generalize Lemma 3.4 to contours that are not simple circuits. Recall that $N_{\Delta_0}(L,T)$ denotes the number of contours γ surrounding Δ_0 satisfying $|\gamma|=L$ and $t(\gamma)=T$. Recall also from Lemma 3.3 that $\alpha(H)$ denotes the connective constant of a graph H as defined in (3.9).

Lemma 4.2. (Bound on number of contours) *We have $\alpha(G_1)<2$. Moreover, for every $\epsilon>0$ there exists a constant $C'_\epsilon<\infty$ such that*

$$N_{\{v\}}(L,T) \;\le\; (L^T/T!)^2\,(C'_\epsilon)^{T+1}\,[\alpha(G_1)+\epsilon]^L \tag{4.5}$$

for all $v\in V_0$ and all $L,T\ge 0$.

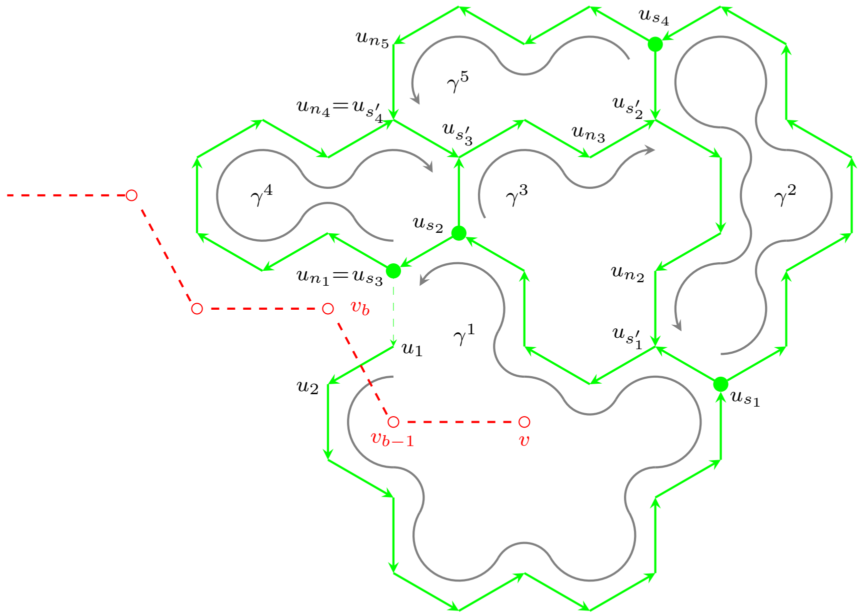

Fig. 4.2. A contour γ in the case when G is the diced lattice and G_1 is the hexagonal lattice. This contour γ contains 8 vertices of degree 3, hence $t(\gamma) = 4$. It is not hard to check that $\chi(\gamma) = 2$.

Proof. $\alpha(G_1) < 2$ was already proven in Lemma 3.4.

Now let γ be a contour surrounding $\{v\}$ such that $|\gamma| = L$ and $t(\gamma) = T$. We need a suitable way to encode γ. We begin, as in the proof of Lemma 3.4, by letting $\pi = (v_0, v_1, \ldots)$ be an infinite self-avoiding path in G_0 starting at $v_0 = v$ such that the graph distance (in G_0) of v_n to v is n. According to Proposition A.5 and formula (3.18), the contour γ intersects an edge (v_{b-1}, v_b) of π with $b \leq N := 1 + K + L/4$, where K is a constant depending only on the graph G_0. Thus, we can find some directed simple circuit $\gamma^* = (u_1, \ldots, u_{n_1}, u_1)$ contained in γ, such that (u_1, u_2) crosses the edge (v_{b-1}, v_b) in the anticlockwise direction (see Figure 4.2).

Let us write $\gamma^1 = (u_1, \ldots, u_{n_1})$, which is a self-avoiding path. If $T = 0$, then $\gamma = \gamma^*$ and our encoding is complete. Otherwise, let $s_1 := \min\{i \geq 1\colon\ u_i \text{ is of degree 3 in } \gamma\}$. Then we can find a self-avoiding path

$$\gamma^2 \;=\; (u_{s_1}, u_{n_1+1}, \ldots, u_{n_2}, u_{s'_1}) \tag{4.6}$$

in γ such that only the starting and ending points u_{s_1} and $u_{s'_1}$ are in γ^1. If $T = 1$, then $\gamma = \gamma^* \cup \gamma^2$ and we are done. Otherwise, let $s_2 := \min\{i > s_1\colon\ i \neq s'_1 \text{ and } u_i \text{ is of degree 3 in } \gamma\}$. Then we can find another self-avoiding path

$$\gamma^3 \;=\; (u_{s_2}, u_{n_2+1}, \ldots, u_{n_3}, u_{s'_2}) \tag{4.7}$$

in γ such that only the starting and ending points u_{s_2} and $u_{s'_2}$ are in $\gamma^1 \cup \gamma^2$. Continuing in this way, we see that we can code all the information needed to construct γ by specifying numbers

$$b \leq N, \quad 2 = n_0 < n_1 < \cdots < n_{T+1} = L - T \quad \text{and} \quad 0 < s_1 < \cdots < s_T < L - T \tag{4.8}$$

and self-avoiding paths $\gamma^1, \ldots, \gamma^{T+1}$ of lengths $n_1 - n_0 + 1, n_2 - n_1 + 1, \ldots, n_{T+1} - n_T + 1$ whose starting edges are uniquely determined by the information previously coded. By Lemma 3.3 and its consequence (3.16), for each $\epsilon > 0$ there exists $K_\epsilon < \infty$ such that the number of self-avoiding paths of length n with a specified starting edge is bounded from above by $K_\epsilon[\alpha(G_1) + \epsilon]^n$.

Therefore, there are at most

$$\prod_{i=1}^{T+1} K_\epsilon[\alpha(G_1) + \epsilon]^{n_i - n_{i-1}+1} \;=\; (K_\epsilon)^{T+1} \, [\alpha(G_1) + \epsilon]^{L-1} \tag{4.9}$$

different contours γ associated with given data $b, n_1, \ldots, n_T, s_1, \ldots, s_T$. Since there are $\binom{L-T-3}{T}$ and $\binom{L-T-1}{T}$ ways of choosing numbers $2 < n_1 < \cdots < n_T < L - T$ and $0 < s_1 < \cdots < s_T < L - T$, respectively, and since $b \le N = 1 + K + L/4$, summing over all ways to choose the numbers $b, n_1, \ldots, n_T, s_1, \ldots, s_T$ shows that the total number of contours γ surrounding v with given $|\gamma| = L$ and $t(\gamma) = T$ is bounded by

$$\begin{aligned} &\big(1 + K + L/4\big)\binom{L-T-3}{T}\binom{L-T-1}{T}(K_\epsilon)^{T+1}\,\big[\alpha(G_1) + \epsilon\big]^{L-1} \\ &\le\; (C'_\epsilon)^{T+1}\binom{L-T}{T}^2\big[\alpha(G_1) + 2\epsilon\big]^{L-1} \;\le\; (C'_\epsilon)^{T+1}\big(L^T/T!\big)^2\big[\alpha(G_1) + 2\epsilon\big]^{L}\;, \end{aligned} \tag{4.10}$$

where the factor $1 + K + L/4$ was absorbed into a change of base of the exponential term followed by the change of constant into C'_ϵ. □

4.3. Long-range dependence for positive temperature.

Proof of Lemmas 2.1 and 2.5 in the positive-temperature case. In the zero-temperature case, both lemmas have already been proven in Sections 3.3 and 3.4, respectively, by showing that for some sufficiently large Δ_0 (respectively for $\Delta_0 = \{v\}$) the right-hand side of (3.6) can be made sufficiently small. To generalize the two lemmas to small positive temperatures, it therefore suffices to show that the right-hand side of (4.4) converges to the right-hand side of (3.6) as $\beta \to \infty$ [for Lemma 2.5 we should also show that the error is $O(e^{-\beta})$]. In view of this, Lemmas 2.1 and 2.5 are consequences of the following lemma. □

Lemma 4.3. (Large-β behavior of the Peierls bound) *There exist $\beta_0, C < \infty$ such that*

$$0 \;\le\; \sum_{T=0}^{\infty}\sum_{L=3}^{\infty} N_{\Delta_0}(L,T)\,\frac{2^{T+1}\,p_\beta^L\,q_\beta^T}{1 + 2^{T+1}\,p_\beta^L\,q_\beta^T} \;-\; \sum_{L=3}^{\infty} N_{\Delta_0}(L)\,\frac{2^{1-L}}{1 + 2^{1-L}} \;\le\; Ce^{-\beta} \tag{4.11}$$

and

$$\sum_{T=1}^{\infty}\sum_{L=3}^{\infty} N_{\Delta_0}(L,T)\,\frac{2^{T+1}\,p_\beta^L\,q_\beta^T}{1 + 2^{T+1}\,p_\beta^L\,q_\beta^T} \;\le\; \sum_{T=1}^{\infty}\sum_{L=3}^{\infty} N_{\Delta_0}(L,T)\,2^{T+1}\,p_\beta^L\,q_\beta^T \;\le\; Ce^{-2\beta} \tag{4.12}$$

uniformly for $\beta \in [\beta_0, \infty]$ and for nonempty finite G_0-connected sets $\Delta_0 \subset V_0$.

PROOF. The lower bound in (4.11) is a trivial consequence of $p_\beta \ge \frac{1}{2}$ and $q_\beta \ge 0$. To prove the upper bounds, we split the double sum in (4.11) into its contributions $T = 0$ and $T \ge 1$ and bound them separately, using $p_\beta = \frac{1}{2} + O(e^{-\beta})$ and $q_\beta = O(e^{-2\beta})$.

$\boldsymbol{T = 0.}$ By Lemma 3.4, there exist $C < \infty$ and $\alpha < 2$ such that $N_{\Delta_0}(L) \le C\alpha^L$. The term $T = 0$ can therefore be bounded (using $p_\beta \ge \frac{1}{2}$) as

$$\sum_{L=3}^{\infty} N_{\Delta_0}(L)\, \frac{2p_\beta^L}{1+2p_\beta^L} \;\le\; \sum_{L=3}^{\infty} N_{\Delta_0}(L)\, \frac{2^{1-L}}{1+2^{1-L}} \;+\; 2C \sum_{L=3}^{\infty} \alpha^L\, \frac{p_\beta^L - (\frac{1}{2})^L}{1+2^{1-L}}\ . \tag{4.13}$$

Choosing β_0 large enough so that $\alpha\, p_{\beta_0} < 1$, it is easy to see, using $p_\beta = \frac{1}{2} + O(e^{-\beta})$, that the last term in (4.13) is $O(e^{-\beta})$.

$\boldsymbol{T \ge 1.}$ By Lemma 4.2, there exist $C < \infty$ and $\alpha < 2$ such that $N_{\Delta_0}(L,T) \le (L^T/T!)^2 C^{T+1}\alpha^L$. Therefore the terms $T \ge 1$ in (4.11) can be bounded as

$$\begin{aligned}
\sum_{T=1}^{\infty}\sum_{L=3}^{\infty} N_{\Delta_0}(L,T)\, \frac{2^{T+1}\, p_\beta^L\, q_\beta^T}{1+2^{T+1}\, p_\beta^L\, q_\beta^T} \;&\le\; \sum_{T=1}^{\infty}\sum_{L=3}^{\infty} N_{\Delta_0}(L,T)\, 2^{T+1}\, p_\beta^L\, q_\beta^T \\
&\le\; 2C \sum_{L=3}^{\infty} (\alpha p_\beta)^L \sum_{T=1}^{\infty} \frac{(2Cq_\beta)^T}{(T!)^2}\, L^{2T} \\
&\le\; 2C \sum_{L=3}^{\infty} (\alpha p_\beta)^L \sum_{T=1}^{\infty} \frac{(8Cq_\beta)^T}{(2T)!}\, L^{2T} \\
&=\; 16C^2 q_\beta \sum_{L=3}^{\infty} L^2\, (\alpha p_\beta)^L \sum_{T=0}^{\infty} \frac{(8Cq_\beta)^T}{(2T+2)!}\, L^{2T} \\
&\le\; 16C^2 q_\beta \sum_{L=3}^{\infty} L^2\, \big(\alpha p_\beta\, e^{\sqrt{8Cq_\beta}}\big)^L
\end{aligned} \tag{4.14}$$

where we used $\dfrac{(2T)!}{(T!)^2} = \dbinom{2T}{T} \le 2^{2T}$. Choosing β_0 large enough so that one has $\alpha\, p_{\beta_0}\, e^{\sqrt{8Cq_{\beta_0}}} < 1$, we see that (4.14) is $O(q_\beta) = O(e^{-2\beta})$, which proves (4.12) and completes the proof of (4.11). $\square$

The bound (4.12) from Lemma 4.3 has a useful corollary. Let us say that a contour γ is *simple* if it is a simple circuit, i.e. $t(\gamma) = 0$. For any contour configuration Γ and any $v \in V_0$, let $S_v^{\mathrm{t}}(\Gamma)$ denote the number of non-simple contours in Γ that surround $\{v\}$. We then have the following bound showing that non-simple contours are rare at low temperature:

Corollary 4.4. (Rarity of non-simple contours) *Let $\nu_{\Lambda,\beta}$ be the contour measure from (4.1). Then there exist $\beta_0, C < \infty$ such that*

$$\sum_{\Gamma} \nu_{\Lambda,\beta}(\Gamma)\, S_v^{\mathrm{t}}(\Gamma) \;\le\; Ce^{-2\beta} \tag{4.15}$$

uniformly for $\beta \in [\beta_0, \infty]$, for finite simply connected $\Lambda \subset V$ such that $\partial\Lambda \subset V_0$, and for $v \in \Lambda \cap V_0$.

PROOF. This is an immediate consequence of (4.12) together with the positive-temperature analogue of Lemma 3.1. $\square$

4.4. Proof of the technical lemmas. In this section we prove Lemmas 2.2 and 2.4.

Proof of Lemma 2.2. It is easy to show that for each $\beta_0 < \infty$ there exists an $\epsilon > 0$ such that $\mu^1_{\Lambda,\beta}(\mathcal{J}_{\Delta_0}) \geq \epsilon$, uniformly for all $0 \leq \beta \leq \beta_0$ and all finite and simply connected $\Lambda \supset \Delta_0$ such that $\partial\Lambda \subset V_0$. Indeed, this follows from a "finite energy" argument: given any configuration $\sigma \in \{1,2,3\}^\Lambda$, we can recolor the sites in Δ_0 in any color of our choice at an energetic cost of at worst $e^{-\beta|\partial\Delta_0|}$ and an entropic cost of at worst $3^{-|\Delta_0|}$. Note that here Δ_0 is fixed and finite, so the precise dependence of the costs on Δ_0 is irrelevant. The only difficulty is that the bound one obtains in this way is not uniform in β as $\beta \to \infty$. Therefore, to complete the proof, it suffices to show that there exists some $\beta_0 < \infty$ such that $\mu^1_{\Lambda,\beta}(\mathcal{J}_{\Delta_0})$ can be estimated from below uniformly in $\beta_0 \leq \beta \leq \infty$ and Λ.

In order to prove this, let $\Delta_0 \subset V_0$ be finite and G_0-connected, and let $\overline{\Delta_0}$ denote the union of Δ_0 with its boundary in G_0, i.e., $\overline{\Delta_0} := \Delta_0 \cup \{v \in V_0 : \exists u \in \Delta_0 \text{ s.t. } \{u,v\} \in E_0\}$. By Corollary 4.4, the probability that a non-simple contour intersects $\overline{\Delta_0}$ tends to zero as $\beta \to \infty$, uniformly in Λ. Thus, we may choose $\beta_0 < \infty$ such that

$$\nu_{\Lambda,\beta}\big(\{\Gamma\colon \nexists\gamma \in \Gamma \text{ s.t. } t(\gamma) \geq 1,\ \gamma \text{ intersects } \overline{\Delta_0}\}\big) \geq 1/2, \tag{4.16}$$

uniformly in Λ and $\beta_0 \leq \beta \leq \infty$. If all contours intersecting $\overline{\Delta_0}$ are simple, then we claim that we can change our contour configuration at a finite energetic cost *uniformly* in $\beta_0 \leq \beta \leq \infty$, so that no contour intersects Δ_0. To describe the algorithm of changing a contour configuration Γ into a configuration Γ' with no contour intersecting Δ_0, we first observe that relying on the fact that all contours intersecting $\overline{\Delta_0}$ are simple, we can color the vertices in $\Lambda \cap V_0$ in three colors in such a way that boundaries between different colors correspond to contours and only two different colors occur in $\overline{\Delta_0}$. (Note that this part of the argument uses that the contours intersecting $\overline{\Delta_0}$ are simple everywhere and not just that there are no triple points inside $\overline{\Delta_0}$.) Now we change our coloring by painting Δ_0 uniformly in one of these two colors, defining thus the new contour configuration Γ' (see Figure 4.3). Since in the construction of Γ' a two-color configuration in $\overline{\Delta_0}$ was changed using the same two colors *no new triple points* are introduced. This, together with (4.1) and a standard finite-energy argument proves our claim.

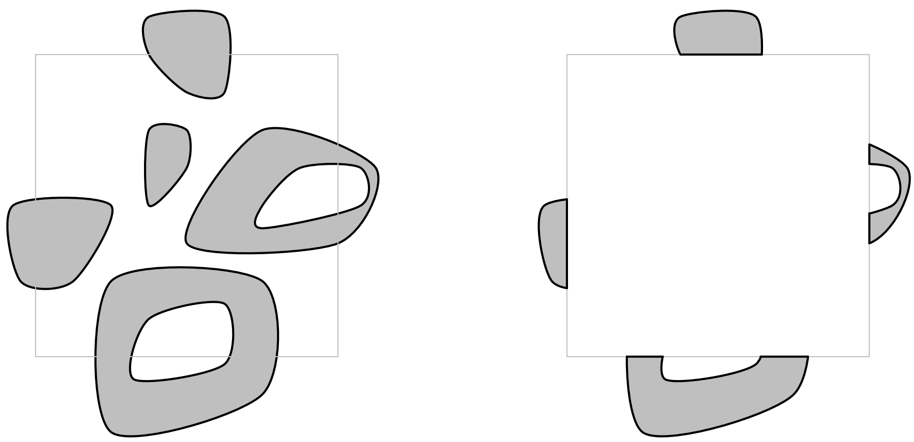

Fig. 4.3. Simple contours intersecting the square $\overline{\Delta_0}$ colored with two colors: 1 (white) and 2 (gray). After flipping all sites in Δ_0 to the color 1, no contour is intersecting Δ_0 and no triple point was created. The same would be true when flipping all sites in Δ_0 to the color 2.

For completeness, we write down this finite-energy argument in detail. Let Γ and $\Gamma' = \psi(\Gamma)$ denote the old and new contour configuration obtained by the procedure described above. We need to estimate the relative probability of Γ' with respect to Γ and the number of different configurations Γ that can be mapped onto the same Γ', $|\Psi^{-1}(\Gamma')|$. Let $|\Delta_0|$ be the number of sites in Δ_0, let M_{Δ_0} be the number of edges in E_1 that separate sites in Δ_0 from each other and let $M_{\partial\Delta_0}$ be the number of edges in E_1 that separate sites in Δ_0 from sites in $\overline{\Delta_0}\backslash\Delta_0$. Further, let

$$\chi(\Gamma) := \prod_{\gamma\in\Gamma} \chi(\gamma), \quad |\Gamma| := \sum_{\gamma\in\Gamma} |\gamma| \quad \text{and} \quad t(\Gamma) := \sum_{\gamma\in\Gamma} t(\gamma). \tag{4.17}$$

Since all contours we remove or alter are simple contours with $\chi(\gamma) = 2$ and we remove or alter no more than M_{Δ_0} contours from our configuration and add no more than $M_{\partial\Delta_0}$ edges, we have

$$\chi(\Gamma') \geq 2^{-M_{\Delta_0}}\chi(\Gamma) \quad \text{and} \quad |\Gamma'| \leq |\Gamma| + M_{\partial\Delta_0}, \tag{4.18}$$

while $t(\Gamma') = t(\Gamma)$, which by (4.1) implies that

$$\nu_{\Lambda,\beta}(\Gamma') \geq 2^{-M_{\Delta_0}} p_\beta^{M_{\partial\Delta_0}} \nu_{\Lambda,\beta}(\Gamma). \tag{4.19}$$

Moreover, since there are $2^{|\Delta_0|}$ ways of coloring the vertices in Δ_0 using only two colors, we see that there are at most $2^{|\Delta_0|}$ different contour configurations Γ in $\Psi^{-1}(\Gamma')$. Recall that $\mathcal{J}_{\Delta_0} = \{\Gamma\colon$ no contour in Γ intersects $\Delta_0\}$ corresponds to the event that Δ_0 is uniformly colored in one color. Let $\mathcal{S}_{\overline{\Delta_0}}$ be the event that all contours intersecting $\overline{\Delta_0}$ are simple. Then

$$\begin{aligned} \nu_{\Lambda,\beta}(\mathcal{J}_{\Delta_0}) &= \sum_{\Gamma'\in\mathcal{J}_{\Delta_0}} \nu_{\Lambda,\beta}(\Gamma') \geq 2^{-|\Delta_0|} \sum_{\Gamma\in\mathcal{S}_{\overline{\Delta_0}}} \nu_{\Lambda,\beta}(\Psi(\Gamma)) \\ &\geq 2^{-|\Delta_0|-M_{\Delta_0}} p_\beta^{M_{\partial\Delta_0}} \sum_{\Gamma\in\mathcal{S}_{\overline{\Delta_0}}} \nu_{\Lambda,\beta}(\Gamma) \geq 2^{-1-|\Delta_0|-M_{\Delta_0}} p_\beta^{M_{\partial\Delta_0}}, \end{aligned} \tag{4.20}$$

where we have used (4.16) in the last step. □

Proof of Lemma 2.4. Consider any $v \in V_1 \cap \Lambda$ and let $w_1, w_2, w_3 \in V_0$ be its neighbors in G. Then the DLR equations for the volume $\{v\}$ imply that

$$\begin{aligned} &\mu^1_{\Lambda,\beta}(\exists i \text{ with } \sigma_v = \sigma_{w_i} | \sigma_{w_1}, \sigma_{w_2}, \sigma_{w_3}) \\ &= \begin{cases} \dfrac{e^{-3\beta}}{2 + e^{-3\beta}} & \text{if } \sigma_{w_1} = \sigma_{w_2} = \sigma_{w_3} \\ \dfrac{e^{-\beta} + e^{-2\beta}}{1 + e^{-\beta} + e^{-2\beta}} & \text{if } |\{\sigma_{w_1}, \sigma_{w_2}, \sigma_{w_3}\}| = 2 \\ 1 & \text{if } |\{\sigma_{w_1}, \sigma_{w_2}, \sigma_{w_3}\}| = 3 \end{cases} \end{aligned} \tag{4.21}$$

Let $\mathcal{B} := \{\sigma\colon |\{\sigma_{w_1}, \sigma_{w_2}, \sigma_{w_3}\}| = 3\}$ be the ("bad") event that w_1, w_2, w_3 are colored in three different colors. It follows from Corollary 4.4 that

$$\mu^1_{\Lambda,\beta}(\mathcal{B}) \;\leq\; Ce^{-2\beta} \tag{4.22}$$

uniformly for $\beta \in [\beta_0, \infty]$ and for finite and simply connected $\Lambda \ni v$ such that $\partial\Lambda \subset V_0$; and by increasing C we can make this hold uniformly for $\beta \in [0, \infty]$. It then follows from (4.21) and (4.22) that

$$\mu^1_{\Lambda,\beta}(\exists i \text{ with } \sigma_v = \sigma_{w_i}) \;\leq\; C'e^{-\beta} \tag{4.23}$$

uniformly for $\beta \in [0, \infty]$ and for finite and simply connected $\Lambda \ni v$ such that $\partial\Lambda \subset V_0$. □

5. Positive magnetization

In this section we prove Lemma 2.3: by using this lemma we can improve the statement that sufficiently large blocks are more likely to be uniformly colored in the color 1 than in any other color, to the "positive magnetization" statements in Theorem 1.1(a,b), which say that single vertices in the sublattices V_0 and V_1 are colored with the color 1 with a probability that is strictly larger (resp. strictly smaller) than 1/3.

We fix an arbitrary $\beta_0 > 0$ throughout this section; our estimates will be uniform in $\beta \in [\beta_0, \infty]$. We will later also fix a finite G_0-connected set $\varDelta_0 \subset V_0$. As in all our proofs, we work with the finite-volume Gibbs measures $\mu^1_{\varLambda,\beta}$, where $\varLambda \subset V$ is finite and simply connected in G and satisfies $\partial\varLambda \subset V_0$. We aim to derive bounds that are uniform in such $\varLambda$ with $\varLambda \supseteq \varDelta_0$.

Unlike what was done in the preceding subsections, we will not make use of the contour description of $\mu^1_{\varLambda,\beta}$, nor will we integrate out one sublattice. Rather, we will work directly with the Potts antiferromagnet on our original quadrangulation $G = (V, E)$.

Note first that the measures $\mu^1_{\varLambda,\beta}$ are invariant under global interchange of the colors 2 and 3. In particular, we have $\mu^1_{\varLambda,\beta}(\sigma_v = 2) = \mu^1_{\varLambda,\beta}(\sigma_v = 3)$ for all $v \in \varLambda$. Thus, to show that $\mu^1_{\varLambda,\beta}(\sigma_v = 1) > 1/3$ (resp. $< 1/3$), we may equivalently show that $\mu^1_{\varLambda,\beta}(\sigma_v = 1) - \mu^1_{\varLambda,\beta}(\sigma_v = 2) > 0$ (resp. < 0). Because of the antiferromagnetic nature of our model, it is in fact already nontrivial to show that these quantities are nonnegative (resp. nonpositive) for $v \in V_0$ (resp. $v \in V_1$). This problem has been solved, however, in [24, Appendix A], where a first Griffiths inequality for antiferromagnetic Potts models on bipartite graphs is proven using ideas based on the cluster algorithm from [69, 70].

We will elaborate on these ideas. The main step will be to give a random-cluster representation for the law of the 1's and 2's conditional on the 3's. In this representation, we will see that for $v_0 \in V_0$, the difference between the probabilities that v_0 is colored 1 or colored 2 equals the probability that v_0 percolates, i.e., that v_0 is in the same random cluster of 1's and 2's as the boundary $\partial\varLambda$. Moreover, a similar statement holds for the probability that $\varDelta_0$ is uniformly colored in the color 1 minus the probability that it is uniformly colored in the color 2. Thus, by showing that both of these quantities are related to percolation of the 1's and 2's, we can prove that if one of them is strictly positive, then so must be the other. Note that conditioning on the positions of the 3's would not in general be a very useful thing to do when trying to prove statements about our model, since we have no *a priori* knowledge of the distribution of the 3's. Nevertheless, as we see here, it can be used to show that a certain statement that has already been proved is equivalent to another statement for which we have no direct control.

So let $G = (V, E)$ be our original quadrangulation, and let $E_\varLambda$ be the set of edges in E that have at least one endvertex in $\varLambda$. We define the measure $\rho^1_{\varLambda,\beta}$ on $\{1,2,3\}^\varLambda \times \{0,1\}^{E_\varLambda}$ so that the marginal distribution of $\rho^1_{\varLambda,\beta}(\sigma,\eta)$ on σ is the Gibbs measure $\mu^1_{\varLambda,\beta}$ and so that, conditional on σ, independently for each $e \in E_\varLambda$, one has $\eta_e = 1$ with probability $p := 1 - e^{-\beta}$ if $\sigma_u, \sigma_v \in \{1, 2\}$ and $\sigma_u \neq \sigma_v$, and $\eta_e = 0$ in all other cases. That is,

$$\begin{aligned}\rho^1_{\varLambda,\beta}(\sigma,\eta) := {} & \frac{1}{Z^1_{\varLambda,\beta}} \exp\bigl[-\beta H_\varLambda(\sigma \,|\, \tau)\bigr] \\ & \times \prod_{\{u,v\}\in E_\varLambda} \Bigl(1_{\mathcal{A}_{u,v}}\bigl[p 1_{\{\eta_{\{u,v\}}=1\}} + (1-p) 1_{\{\eta_{u,v}=0\}}\bigr] + 1_{\mathcal{A}^{\mathrm{c}}_{u,v}} 1_{\{\eta_{u,v}=0\}}\Bigr)\end{aligned} \tag{5.1}$$

where $\mathcal{A}_{u,v}$ is the event

$$\mathcal{A}_{u,v} \;:=\; \bigl\{\sigma_u, \sigma_v \in \{1,2\} \text{ and } \sigma_u \neq \sigma_v\bigr\}\;, \tag{5.2}$$

$\mathcal{A}^{\mathrm{c}}_{u,v}$ is its complement, τ is any spin configuration that assumes the value 1 on $\partial\varLambda$, $H_\varLambda(\sigma|\tau)$ is defined in (1.4), and $Z^1_{\varLambda,\beta}$ is the same normalizing constant as in (1.6).

Now let

$$\varLambda^{12} := \bigl\{v \in \varLambda \cup \partial\varLambda\colon\; \sigma_v \in \{1,2\}\bigr\} \tag{5.3a}$$

$$\varLambda^{3} := \bigl\{v \in \varLambda\colon\; \sigma_v = 3\bigr\} \tag{5.3b}$$

be the sets of vertices in $\Lambda \cup \partial\Lambda$ where σ assumes the values 1 or 2 (resp. 3), and set

$$E^{12} \;:=\; \{e \in E_\Lambda :\; \eta_e = 1\}\,. \tag{5.4}$$

Conditional on Λ^3, the spins $(\sigma_v)_{v\in\Lambda^{12}}$ are distributed as an antiferromagnetic Ising model, with 1 boundary conditions, on the diluted lattice Λ^{12}. Since Λ^{12} is bipartite and the boundary conditions lie entirely on the sublattice V_0, we may flip the spins on the other sublattice (i.e., on $\Lambda^{12} \cap V_1$) to obtain a ferromagnetic Ising model $(\sigma'_v)_{v\in\Lambda^{12}}$ on Λ^{12}. After this flipping, the conditional joint law of $(\sigma'_v)_{v\in\Lambda^{12}}$ and η given Λ^3 is just the standard coupling of this ferromagnetic Ising model and its corresponding random-cluster model on Λ^{12} (see [23] and [32, Section 1.4]). (Notice that for all edges $\{u,v\}$ such that $\{u,v\} \cap \Lambda^3 \neq \varnothing$, we have $\eta_{\{u,v\}} = 0$.) Returning to the original (unflipped) spins $(\sigma_v)_{v\in\Lambda^{12}}$, we see from [32, Theorem 1.13] that, conditional on Λ^3 and η, the connected components of the graph $G^{12} = (\Lambda^{12}, E^{12})$ are independently given proper 2-colorings (with colors 1 and 2) as follows: for any component not connected to the boundary $\partial\Lambda$, each of the two proper 2-colorings arises with probability 1/2; and any component connected to the boundary is given the unique proper 2-coloring that is compatible with the boundary conditions (namely, color 1 on V_0 and color 2 on V_1). In particular, for points $v_0 \in V_0 \cap \Lambda$ one has

$$\rho^1_{\Lambda,\beta}\big(\sigma_{v_0} = 1 \,\big|\, \Lambda^3, \eta\big) \;=\; \begin{cases} 1 & \text{if } v_0 \leftrightarrow_\eta \partial\Lambda \\ \frac{1}{2} & \text{if } v_0 \in \Lambda^{12} \text{ and } v_0 \not\leftrightarrow_\eta \partial\Lambda \\ 0 & \text{if } v_0 \in \Lambda^3 \end{cases} \tag{5.5}$$

$$\rho^1_{\Lambda,\beta}\big(\sigma_{v_0} = 2 \,\big|\, \Lambda^3, \eta\big) \;=\; \begin{cases} 0 & \text{if } v_0 \leftrightarrow_\eta \partial\Lambda \\ \frac{1}{2} & \text{if } v_0 \in \Lambda^{12} \text{ and } v_0 \not\leftrightarrow_\eta \partial\Lambda \\ 0 & \text{if } v_0 \in \Lambda^3 \end{cases} \tag{5.6}$$

where $v \leftrightarrow_\eta \partial\Lambda$ denotes the event that v is connected to $\partial\Lambda$ through a path of edges with $\eta_e = 1$ [note that $v_0 \in \Lambda^3$ implies $v_0 \not\leftrightarrow_\eta \partial\Lambda$]. For the unconditional law, it follows that

$$\rho^1_{\Lambda,\beta}(\sigma_{v_0} = 1) - \rho^1_{\Lambda,\beta}(\sigma_{v_0} = 2) \;=\; \rho^1_{\Lambda,\beta}\big(v_0 \leftrightarrow_\eta \partial\Lambda\big)\,. \tag{5.7}$$

For $v_1 \in V_1 \cap \Lambda$, one has similar equations with the roles of colors 1 and 2 interchanged, so that

$$\rho^1_{\Lambda,\beta}(\sigma_{v_1} = 2) - \rho^1_{\Lambda,\beta}(\sigma_{v_1} = 1) \;=\; \rho^1_{\Lambda,\beta}\big(v_1 \leftrightarrow_\eta \partial\Lambda\big)\,. \tag{5.8}$$

Now consider a finite set $\Delta_0 \subset V_0$, and recall that $\mathcal{J}_{k,\Delta_0}$ denotes the event that Δ_0 is uniformly colored in the color k, and that $\mathcal{J}_{\Delta_0} = \bigcup_{k=1}^3 \mathcal{J}_{k,\Delta_0}$ denotes the event that all sites in Δ_0 are uniformly colored in some color. Let $\Delta_0 \leftrightarrow_\eta \partial\Lambda$ denote the event that there is at least one site in Δ_0 that is connected to $\partial\Lambda$ through a path of edges with $\eta_e = 1$. Since

$$\rho^1_{\Lambda,\beta}\big(\mathcal{J}_{1,\Delta_0} \,\big|\, \Lambda^3, \eta\big) \;=\; \rho^1_{\Lambda,\beta}\big(\mathcal{J}_{2,\Delta_0} \,\big|\, \Lambda^3, \eta\big) \quad \text{a.s. on} \quad \Delta_0 \not\leftrightarrow_\eta \partial\Lambda \tag{5.9}$$

and

$$\rho^1_{\Lambda,\beta}\big(\mathcal{J}_{2,\Delta_0} \,\big|\, \Lambda^3, \eta\big) \;=\; 0 \quad \text{a.s. on} \quad \Delta_0 \leftrightarrow_\eta \partial\Lambda, \tag{5.10}$$

we see that

$$\rho^1_{\Lambda,\beta}(\mathcal{J}_{1,\Delta_0}) - \rho^1_{\Lambda,\beta}(\mathcal{J}_{2,\Delta_0}) \;=\; \rho^1_{\Lambda,\beta}\big(\mathcal{J}_{\Delta_0} \cap \{\Delta_0 \leftrightarrow_\eta \partial\Lambda\}\big)\,. \tag{5.11}$$

Now, recall that a finite set $\Delta \subset V$ is termed *thick* if there exists a nonempty finite subset $\Delta_1 \subset V_1$ that is connected in G_1 and such that $\Delta = \{v \in V :\; d_G(v, \Delta_1) \le 1\}$. We therefore fix some thick set $\Delta \subset V$ and define $\Delta_0 := \Delta \cap V_0$ (of course we have $\Delta_1 = \Delta \cap V_1$). Comparing (5.7)/(5.8)/(5.11) and noting that $\rho^1_{\Lambda,\beta}$ can be replaced by $\mu^1_{\Lambda,\beta}$ on the left-hand sides, we see that Lemma 2.3 is implied by the following claim:

Lemma 5.1. (Comparison lemma) *Fix $\beta_0 > 0$ and let $\Delta \subset V$ be thick. Then there exists an $\epsilon > 0$ such that*

$$\rho^1_{\Lambda,\beta}\big(v \nleftrightarrow_\eta \partial\Lambda \text{ for all } v \in \Delta\big) \;\geq\; \epsilon\, \rho^1_{\Lambda,\beta}\big(\mathcal{J}_{\Delta_0} \cap \{\Delta_0 \nleftrightarrow_\eta \partial\Lambda\}\big) \tag{5.12}$$

uniformly for all $\beta \in [\beta_0, \infty]$ and all simply connected finite sets $\Lambda \supseteq \Delta$ such that $\partial\Lambda \subset V_0$. In fact, we can choose $\epsilon = 3^{-|\Delta_1|}\,(1 - e^{-\beta_0})^{|E_\Delta|}$ where $E_\Delta = \big\{\{u, v\} \in E\colon\; u, v \in \Delta\big\}$.

Proof. The proof is by a finite-energy argument: that is, to each $(\sigma, \eta) \in \mathcal{J}_{\Delta_0} \cap \{\Delta_0 \nleftrightarrow_\eta \partial\Lambda\}$ we associate a $(\sigma'', \eta'') \in \{v \nleftrightarrow_\eta \partial\Lambda$ for all $v \in \Delta\}$; we then compute a lower bound on the ratio of $\rho^1_{\Lambda,\beta}(\sigma'', \eta'')$ to the total $\rho^1_{\Lambda,\beta}$-weight of the configurations (σ, η) that map onto it. The construction is in two steps $(\sigma, \eta) \mapsto (\sigma', \eta') \mapsto (\sigma'', \eta'')$. In the first step we recolor all spins $(\sigma_v)_{v\in\Delta_1}$ to $\sigma'_v = 2$ (leaving all other variables as is). In the second step we set all bond variables $(\eta_e)_{e\in E_\Delta}$ to $\eta''_e = 1$ (again leaving all other variables as is). Let us now compute a lower bound on the ratio of weights, as follows:

Since $(\sigma, \eta) \in \mathcal{J}_{\Delta_0} \cap \{\Delta_0 \nleftrightarrow_\eta \partial\Lambda\}$ and $\partial\Lambda$ is colored 1, it follows that $\sigma_v = 1$ for all $v \in \Delta_0$. Since Δ is thick, every vertex in Δ_1 has all its neighbors in Δ_0. Therefore we can recolor all sites in Δ_1 with the color 2 without increase in energy, i.e. $\rho^1_{\Lambda,\beta}(\sigma', \eta') \geq \rho^1_{\Lambda,\beta}(\sigma, \eta)$.[19] We lose a factor $3^{|\Delta_1|}$ because $3^{|\Delta_1|}$ configurations (σ, η) map onto the same configuration (σ', η').

We now have $\sigma'_u \neq \sigma'_v$ for all $\{u, v\} \in E_\Delta$. Therefore $\rho^1_{\Lambda,\beta}(\sigma'', \eta'')$ is precisely p^{E_Δ} times the total $\rho^1_{\Lambda,\beta}$-weight of the configurations (σ', η') that map onto it, where $p = 1 - e^{-\beta} \geq 1 - e^{-\beta_0}$. □

Remark. The ideas in this section — in particular, formulas (5.7), (5.8) and (5.11) — have an obvious generalization to the q-state Potts antiferromagnet for any $q \geq 2$ on any bipartite graph (not necessarily a plane quadrangulation), where we condition on $q-2$ colors and use the random-cluster representation for the remaining 2 colors. Indeed, as in [24, Appendix A], we may consider an even more general situation: Suppose that the vertex set V is partitioned as $V = V_0 \cup V_1$; then we can consider a Potts model with antiferromagnetic interactions on edges connecting V_0 to V_1 and ferromagnetic interactions on edges V_0–V_0 and V_1–V_1.

A. Some facts about infinite planar graphs

The purpose of this appendix is to collect some facts about infinite graphs, and in particular about infinite graphs embedded in the plane, that will be needed in the main part of the paper. An excellent general introduction to the theory of infinite graphs can be found in [16, Chapter 8]; but we shall require here some further facts that are scattered throughout the recent research literature, plus a few that appear to be new.

A.1. Basic facts and definitions. Recall that a graph is a pair $G = (V, E)$ consisting of a (not necessarily finite) vertex set V and edge set E. Unless mentioned otherwise, when we say "graph" we will always mean a *simple* graph, i.e., an undirected graph that has no loops or multiple edges. Thus, the elements of E (the edges) are unordered sets $\{v, w\}$ containing two distinct elements of V. Two vertices $v, w \in V$ are called *adjacent* if $\{v, w\} \in E$. An edge e containing a vertex v is said to be *incident* on v. The *degree* of a vertex $v \in V$ is the number of edges incident on it. We say that G is *finite* (resp. *countable*) if both V and E are finite (resp. countable). We say that G is *locally finite* (resp. *locally countable*) if every vertex has finite (resp. countable) degree.

A graph $G' = (V', E')$ such that $V' \subseteq V$ and $E' \subseteq E$ is called a *subgraph* of $G = (V, E)$; we also say that G *contains* G'. If E' contains all edges $\{v, w\} \in E$ with $v, w \in V'$, then G' is called

[19] In fact, (5.12) would still hold (with a worse ϵ) even if there were an energy cost associated to this operation, provided that this energy cost is uniformly bounded.

the subgraph of G *induced* by V'. Likewise, if $V' = \{v \in V\colon \exists w \in V \text{ s.t. } \{v,w\} \in E'\}$, then we call G' the subgraph *induced* by E'.

We will say that a graph G is *connected* if for every *proper* subset $W \subset V$ (the word "proper" means that $W \neq \varnothing, V$) there is an edge $\{v,w\} \in E$ with $v \in V \backslash W$, $w \in W$. Let us note that every locally countable, connected graph is countable.[20] A connected graph in which each vertex has degree ≤ 2 will be called a *generalized path*. The *length* of a generalized path is the number of its edges. Vertices of degree 2 are called *internal vertices* of the generalized path, while vertices of degree one or zero are called *endvertices*.[21] An infinite generalized path with one endvertex is called a one-way-infinite path or *ray*; an infinite generalized path without endvertices is called a two-way-infinite path or *double ray*; a finite generalized path without endvertices is called a *cycle*; and a finite generalized path with one or two endvertices is called a *path*. In particular, a graph consisting of a single vertex and no edges is a path of length zero.

Two vertices v, w in a graph G are *linked* by a path if G contains a path that has v and w as its endpoints. Then it is easy to see that a graph is connected (according to our definition above) if and only if every pair of vertices in G is linked by a path. The *graph distance* $d(v,w)$ between two vertices $v, w \in V$ is the length of a shortest path linking v and w if such a path exists, and ∞ otherwise. An edge $\{v,w\}$ in a path P is said to be a *final edge* of P if either v or w (or both) has degree one. The *graph distance* $d(e,f)$ between two edges $e, f \in E$ is defined as the minimal length minus one of a path that has e and f as final edges. Note, in particular, that $d(e,f) = 0$ iff $e = f$, and that $d(e,f) = 1$ iff $e \neq f$ but e and f share a vertex. It is easy to check that, whenever G is connected, the graph distance between vertices (resp. edges) defines a metric on V (resp. E).

Let $G = (V,E)$ be a connected graph. A set $B \subseteq E$ is called *separating* if $G \backslash B$ is not connected. A set $C \subseteq E$ is called a *cutset* if there exists a partition $\{V_1, V_2\}$ of V into two nonempty sets such that $C = E(V_1, V_2) := \{\{v_1, v_2\} \in E\colon\ v_1 \in V_1,\ v_2 \in V_2\}$. (Since G is assumed connected, $\varnothing$ is not a cutset.) A separating set (resp. cutset) is called *minimal* if it contains no proper subset that is a separating set (resp. a cutset). In fact, the two concepts are the same: each minimal separating set is also a minimal cutset and vice versa. Moreover, a set $C \subseteq E$ is a minimal cutset if and only if $(V, E \backslash C)$ has exactly two connected components.[22]

Two rays in an infinite graph G are said to be *end-equivalent* (or *equivalent* for short) if one (hence all) of the following equivalent conditions holds:

1) there exists a third ray whose intersection with both of them is infinite;
2) there are infinitely many disjoint paths in G joining the two rays;
3) for every finite set $S \subset V$, the two rays are eventually contained in the same connected component of $G \backslash S$.

It is easy to see that end-equivalence is an equivalence relation. The corresponding equivalence classes are termed the *ends* of the graph G.[23] It is not hard to see that a connected, locally finite graph $G = (V,E)$ has one end if and only if for every finite $E' \subseteq E$, the subgraph $(V, E \backslash E')$ has exactly one infinite component; or equivalently, if every finite minimal cutset divides V into two connected components, of which exactly one is of infinite size.

---

[20] Indeed, if $G = (V,E)$ is connected and W is the set of all vertices at finite graph distance from a given vertex, then connectedness implies $W = V$.

[21] Note that a vertex of degree zero can occur in a connected graph (and more particularly in a path) only if the graph has precisely one vertex.

[22] In the graph-theory literature, minimal cutsets are sometimes called *bonds*. Alas, in the statistical-mechanics literature, "bond" is often used as a synonym for "edge". To forestall confusion, we prefer to avoid the term "bond" altogether.

[23] This definition of the ends of a graph is due to Halin [33] in 1964 (see also Freudenthal [26] for the locally finite case). The notion of "end" of a topological space was introduced much earlier by Freudenthal [25], in 1931. If one identifies G with the "topological realization" defined below (at the beginning of Section A.3), then the two notions coincide for *locally finite* graphs but are different in general [18]. See [47] for an elementary introduction to the theory of ends of locally finite graphs; and see [49] for a survey treating both graphs and topological spaces.

For $k \geq 1$, a graph $G = (V, E)$ is called *k-connected* if $|V| \geq k+1$ and the subgraph induced by $V \backslash W$ is connected for all $W \subset V$ satisfying $|W| < k$. (Otherwise put, to disconnect G one must remove at least k vertices.) Two vertices v, w in G are said to be *k-edge-connected* if one needs to remove at least k edges to unlink them; a graph is called *k-edge-connected* if every pair of vertices in it is k-edge-connected. It is easy to see that k-edge-connectedness of vertices is an equivalence relation; the corresponding equivalence classes of vertices (and their induced subgraphs) are called the *k-edge-connected components* of the graph. Two paths are called *vertex-disjoint* (resp. *edge-disjoint*) if their sets of internal vertices (resp. edges) are disjoint. By Menger's theorem, two vertices are k-edge-connected if and only if they are linked by k edge-disjoint paths, and a graph is k-connected if and only if every two vertices are linked by k vertex-disjoint paths.

An *automorphism* of a graph $G = (V, E)$ is a bijection $g\colon\ V \to V$ such that $\{g(v), g(v')\} \in E$ if and only if $\{v, v'\} \in E$. We say that two vertices $v, w \in V$ are of the same *type*, denoted $v \sim w$, if there exists an automorphism g of G such that $g(v) = w$. Then $\sim$ is an equivalence relation that divides the vertex set V into equivalence classes called types. A graph is called *vertex-transitive* if there is only one type of vertex, and *vertex-quasi-transitive* if there are only finitely many types of vertices. Similarly, we say that two edges $\{v, v'\}$, $\{w, w'\} \in E$ are of the same type if there exists an automorphism g of G such that $\{g(v), g(v')\} = \{w, w'\}$; and we say that two directed edges (v, v'), $(w, w') \in V \times V$ with $\{v, v'\}$, $\{w, w'\} \in E$ are of the same type if there exists an automorphism g of G such that $g(v) = w$ and $g(v') = w'$. Edge- and directed-edge- transitivity or quasi-transitivity are then defined in the obvious way. We shall need the following fairly easy result:

Lemma A.1. *For a locally finite graph G, the following are equivalent:*

(a) G is vertex-quasi-transitive.
(b) G is edge-quasi-transitive.
(c) G is directed-edge-quasi-transitive.

PROOF. (b) $\Leftrightarrow$ (c): Obviously, if two directed edges (v, v') and (w, w') are of the same type, then the corresponding undirected edges $\{v, v'\}$ and $\{w, w'\}$ are also of the same type. This shows that there are as most as many types of edges as there are types of directed edges. Conversely, since there are only two ways to order a set with two elements, there are at most twice as many types of directed edges as there are types of edges.

(c) $\Rightarrow$ (a): If two directed edges (v, v') and (w, w') are of the same type, then obviously v and w are of the same type. Since all isolated vertices (i.e., vertices of degree zero) are of the same type, this shows that the number of types of vertices is at most the number of types of directed edges plus one.

(a) $\Rightarrow$ (c): Assume that there are m types of vertices and that these have degrees $d_1, \ldots, d_m$. Pick representatives $v_1, \ldots, v_m$ of these equivalence classes. For any directed edge (v, v'), there exists a $k \in \{1, \ldots, m\}$ and a graph automorphism that maps v to v_k. Since a graph automorphism preserves the graph structure, w' must be mapped into one of the d_k vertices adjacent to v_k. Thus, we have found $d_1 + \cdots + d_m$ directed edges such that each directed edge can be mapped into one of these by a graph automorphism. In particular, the number of types of directed edges is at most $d_1 + \cdots + d_m$. $\square$

In view of Lemma A.1, we usually talk about quasi-transitive graphs without specifying whether we mean in the vertex, edge or directed-edge sense.[24]

A *geodesic* in a graph G is a generalized path P such that for each pair of vertices v, w in P, the graph distance from v to w in P coincides with the graph distance from v to w in G. It is not hard to see that any path of minimal length linking two vertices v', w' is a geodesic. For completeness, we prove the following simple fact. Part (a) of this lemma can be found, for example, in [58, Prop. 1];

[24] The term "almost transitive" is also used as a synonym for "quasi-transitive".

it is a simple corollary of König's Infinity Lemma [16, Lemma 8.1.2]. We did not find a reference for part (b), but it is presumably well known.

Lemma A.2. (Infinite geodesics) *Let $G = (V, E)$ be a locally finite connected graph with infinite vertex set V. Then:*

(a) *For each $v \in V$, there exists a geodesic ray whose endpoint is v.*
(b) *If G is moreover quasi-transitive, then G contains a geodesic double ray.*

Proof. Since each vertex is of finite degree, the set of vertices at distance k from v is finite for each $k \geq 0$. Therefore, since V is infinite and G is connected, for each $n \geq 1$ we can find a vertex $v_n \in V$ at distance $d(v, v_n) = n$. Consider a path $(v_0^{(n)} = v, v_1^{(n)}, \ldots, v_{n-1}^{(n)}, v_n^{(n)} = v_n)$ and define the function $f_n \colon \mathbb{N} \to V$ by taking $f_n(k) = v_k^{(n)}$ for $k \leq n$ and $f_n(k) = v_n$ for all $k \geq n$. Since the set of points at distance k from v is finite for each $k \geq 0$, we may select a subsequence f_{n_m} that converges pointwise in the discrete topology. It is easy to see that the limit of such a subsequence is a geodesic ray starting at v, proving part (a) of the lemma.

To prove also part (b), we use that by quasi-transitivity, the geodesic ray constructed in part (a) must pass through at least one type of vertex infinitely often. It follows that for a vertex v of this type, for each $n \geq 0$ we can find a function $f_n \colon \mathbb{Z} \to V$ such that $f_n(0) = v$, $f_n(k) = f_n(-n)$ for all $k \leq -n$, the vertices $\{f_n(k) \colon k \geq -n\}$ are all different, and the graph induced by this set is a geodesic ray. (It suffices to employ the corresponding automorphism to shift the original geodesic ray by identifying a vertex of given type sufficiently far on it with the vertex v.) Now the statement follows from the same sort of compactness argument as used in the proof of part (a). $\square$

A.2. Duality. In the main part of this paper, we make extensive use of the fact that the two sublattices G_0 and G_1 are each other's dual in the sense of planar graph duality. Such duals may be defined abstractly, using only basic concepts of graph theory, without any reference to embeddings of graphs in the plane. In fact, this abstract notion of duality is sufficient for all our proofs, as we shall show in the present subsection. Nevertheless, in the next subsection we will complement this abstract theory by showing that sufficiently "nice" embeddings of a graph in the plane give rise to abstract duals, and conversely that every locally finite abstract dual arises in this way.

The basic theory of duals of infinite graphs was developed by Thomassen [67,68], which we largely follow here; see also Bruhn and Diestel [9,17,10] for a partially alternative approach.[25] Abstract duals can be defined for 2-connected graphs, but in this case the dual may have multiple edges; we therefore restrict ourselves for simplicity to the 3-connected case.

If $G = (V, E)$ is a 3-connected graph, then an *abstract dual* of G is a connected graph $G^\dagger = (V^\dagger, E^\dagger)$ together with a bijection $E \ni e \mapsto e^\dagger \in E^\dagger$ such that a finite set $C \subseteq E$ is a cycle in G if and only if $C^\dagger := \{e^\dagger \colon e \in E\}$ is a minimal cutset in $G^\dagger$. We stress that in this generality the term "dual" is something of a misnomer, since G need not be an abstract dual of $G^\dagger$; indeed, $G^\dagger$ need not have any abstract dual at all, even when G is planar and locally finite [67, p. 266]. However, if both G and $G^\dagger$ are locally finite, then the situation becomes particularly nice:

Theorem A.3. (Locally finite abstract duals) *Let $G = (V, E)$ be a locally finite 3-connected graph that has a locally finite abstract dual $G^\dagger = (V^\dagger, E^\dagger)$. Then:*

[25] The approach of Thomassen [67,68] is based on the study of *finite* cycles and minimal cutsets, as explained below. The alternative approach of Bruhn and Diestel [9,17,10] introduces (by topological means) a notion of *infinite* "cycles", and shows that this notion allows a somewhat cleaner duality theory. When $G^\dagger$ is locally finite, the two concepts of duality coincide [9, Lemma 4.7]. We are therefore entitled to use here the theorems from [9,17,10] under the added hypothesis that $G^\dagger$ is locally finite.

(i) G, with the inverse map $E^\dagger \ni e^\dagger \mapsto e \in E$, is an abstract dual of $G^\dagger$.
(ii) $G^\dagger$ is 3-connected.
(iii) $G^\dagger$ is, up to isomorphism, the only abstract dual of G.
(iv) If G has one end, then so does $G^\dagger$.
(v) If G is quasi-transitive, then so is $G^\dagger$.

Proof. Part (ii) follows from [68, Theorem 4.5]. Then parts (i) and (iii) follow from [67, Theorem 9.5]. Part (iv) follows from [10, Theorem 1.1], which states that there is a homeomorphism between the spaces of ends of G and $G^\dagger$ (considered as subspaces of the Freudenthal compactification, to be defined in the next subsection); so in particular G has one end if and only if $G^\dagger$ does. (Alternatively, this can be deduced from Proposition A.12(iii) below, using Lemma A.7 and Theorem A.10.)

We did not find a reference for part (v), but this is not hard to prove using some more results from [67]. We will prove the following, stronger statement. Let $g\colon V \to V$ be a graph automorphism of G and let $g(\{v,w\}) := \{g(v), g(w)\}$ also denote the induced map $g\colon E \to E$ on edges. Then there exists an automorphism $g^\dagger$ of $G^\dagger$ such that the induced map on edges satisfies $g^\dagger(e^\dagger) = g(e)^\dagger$. This shows that two edges in $G^\dagger$ are of the same type if the corresponding edges in G are of the same type. Since by part (i), duality is a symmetric relation, this "if" is an "if and only if". In particular, $G^\dagger$ is edge-quasi-transitive if and only if G is (with the same number of types of edges).

To prove the existence of $g^\dagger$, we need some definitions. Let $C = (V(C), E(C))$ be a cycle in the graph $G = (V, E)$. We say that C is an *induced cycle* if C is the subgraph of G induced by $V(C)$; equivalently, this says that C has no *diagonals*, i.e., there are no edges in $E \backslash E(C)$ that have both endvertices in $V(C)$. We say that C is a *separating cycle* if there are vertices v_1, v_2 in $V \backslash V(C)$ that are linked in G but not in the subgraph of G induced by $V \backslash V(C)$.

Now let $G = (V, E)$ be a locally finite 3-connected graph and let $G^\dagger = (V^\dagger, E^\dagger)$ be a locally finite abstract dual of G. Then [67, Theorem 9.5] says that there is a one-to-one correspondence between vertices of $G^\dagger$ and induced non-separating cycles of G. Indeed, for each $v^\dagger \in V^\dagger$, the set $C^\dagger$ of edges in $G^\dagger$ that are incident on $v^\dagger$ has the property that $C := \{e\colon e^\dagger \in C^\dagger\}$ is an induced non-separating cycle of G, and conversely, every induced non-separating cycle of G arises in this way.

Now let g be a graph automorphism of G. Since g maps induced non-separating cycles into induced non-separating cycles, there is a bijection $g^\dagger\colon V^\dagger \to V^\dagger$ that maps a vertex $v^\dagger$ into a vertex $w^\dagger$ of $G^\dagger$ if and only if g maps the associated induced non-separating cycles of G into each other. Since two vertices of $G^\dagger$ are adjacent if and only of the associated induced non-separating cycles of G share an edge, we see that $g^\dagger$ is a graph automorphism of $G^\dagger$ such that the induced map $g\colon E^\dagger \to E^\dagger$ on edges satisfies $g^\dagger(e^\dagger) = g(e)^\dagger$. $\square$

Let G and $G^\dagger$ be as in Theorem A.3, let $v^\dagger \in V^\dagger$ be a vertex in $G^\dagger$, and let $E^\dagger_{v^\dagger} := \{e^\dagger \in E^\dagger\colon e^\dagger \text{ is incident on } v^\dagger\}$. Then $E^\dagger_{v^\dagger}$ is a minimal cutset in $G^\dagger$, which is finite by virtue of the local finiteness of $G^\dagger$; hence $E_{v^\dagger} := \{e \in E\colon e^\dagger \in E^\dagger_{v^\dagger}\}$ is a cycle in G. We loosely call $v^\dagger$ a *face* of G and we call $E_{v^\dagger}$ the *boundary* of this face. Indeed, we will see in the next subsection that for a suitable embedding of G in the plane $\mathbb{R}^2$, $v^\dagger$ corresponds to a connected component of $\mathbb{R}^2 \backslash G$ and $E_{v^\dagger}$ corresponds to its boundary. In view of this, we abstractly define a *triangulation* (resp. *quadrangulation*) to be a locally finite 3-connected graph G that has an abstract dual $G^\dagger$ in which each vertex has degree 3 (resp. 4).

Now assume that G has one end. Then each finite minimal cutset C of G corresponds to a partition $\{V_1, V_2\}$ of V into two connected components, of which exactly one is infinite. Let V_1, V_2 denote the finite and infinite component, respectively. Since by Theorem A.3(i), G is an abstract dual of $G^\dagger$, the set $C^\dagger := \{e^\dagger\colon e \in C\}$ is a cycle in $G^\dagger$, and each cycle in $G^\dagger$ arises in this way. We call $\mathrm{Int}(C^\dagger) := V_1$ and $\mathrm{Ext}(C^\dagger) := V_2$ the *interior* and *exterior* of $C^\dagger$, respectively. We say that $C^\dagger$ *surrounds* a vertex $v \in V$ if $v \in \mathrm{Int}(C^\dagger)$.

An essential ingredient in our proofs in this paper is an upper bound (Lemma 3.4) for certain quasi-transitive triangulations G on the number of cycles in $G^\dagger$ of a given length surrounding a given vertex v in G. To derive this bound, we need some simple graph-theoretic facts.

Lemma A.4. (Distances in a graph and its dual) *Let $G = (V,E)$ be a 3-connected graph. Assume that each vertex in G has degree at most $d_{\max}$ and that G has a locally finite abstract dual $G^\dagger = (V^\dagger, E^\dagger)$. Then for all $e, f \in E$ we have*

$$d(e^\dagger, f^\dagger) \;\le\; (\tfrac{1}{2}d_{\max} - 1)d(e,f) + 1 \,, \tag{A.1}$$

where $d(e^\dagger, f^\dagger)$ denotes the distance between $e^\dagger$ and $f^\dagger$ in the dual graph $G^\dagger$.

Note that $d_{\max} \ge 3$ since G is 3-connected.

Proof of Lemma A.4. If $e = f$ (hence $e^\dagger = f^\dagger$), the statement is trivial, so consider the case $e \ne f$. Let P be a path of minimal length that has e and f as final edges. For each vertex v of G, let $E_v^\dagger := \{e^\dagger\colon\ e \text{ is incident on } v\}$ denote the boundary of the corresponding face of $G^\dagger$. Note that $E_v^\dagger$ is a cycle whose length is the degree of v. If some $E_v^\dagger$ and $E_w^\dagger$ share an edge $e^\dagger$, then e connects v and w. For internal vertices v, w of P, by the minimality of P, this is possible only if v and w are adjacent in P. Let $v_1, \ldots, v_n$ be the internal vertices of P, where $n = d(e,f)$, and let $d_1, \ldots, d_n$ denote their degrees in G. Then the symmetric difference $D := E_{v_1}^\dagger \,\triangle \cdots \triangle\, E_{v_n}^\dagger$ consists of exactly $\sum_{k=1}^n d_k - 2(n-1)$ edges which form a cycle in $G^\dagger$ containing $e^\dagger$ and $f^\dagger$. It follows that $G^\dagger$ contains two paths $P_1^\dagger, P_2^\dagger$ which have $e^\dagger$ and $f^\dagger$ as their final edges and are otherwise edge-disjoint, and whose respective lengths k_1, k_2 satisfy $k_1 + k_2 - 2 = \sum_{k=1}^n d_k - 2(n-1)$. We conclude from this that $\min(k_1, k_2) \le (k_1 + k_2)/2 \le \frac{1}{2}(n d_{\max} - 2n + 4)$ and hence $d(e^\dagger, f^\dagger) \le (\frac{1}{2}d_{\max} - 1)n + 1$. $\square$

Now let $G = (V,E)$ be a locally finite 3-connected quasi-transitive graph that has a locally finite abstract dual $G^\dagger = (V^\dagger, E^\dagger)$. By Theorem A.3(v), $G^\dagger$ is also quasi-transitive, hence of bounded degree; let $d_{\max}^\dagger$ denote the maximal degree of a vertex in $G^\dagger$. By Lemma A.2(b), G contains at least one geodesic double ray; let V_0 be the set of all vertices $v \in V$ that lie on some geodesic double ray. Then, by quasi-transitivity, the maximal distance of any vertex in G to the set V_0,

$$K \;:=\; \sup_{w \in V} \inf_{v \in V_0} d(v,w) \,, \tag{A.2}$$

is finite. Note, finally, that by Lemma A.2(a), at each $v \in V$ there starts at least one geodesic ray.

Proposition A.5. (Distance to a surrounding cycle) *Let $G, G^\dagger$, $d_{\max}^\dagger$ and K be as above. Assume that G has one end. Let $v \in V$, let R be a geodesic ray in G starting at v, and let $C^\dagger$ be a cycle in $G^\dagger$ of length L surrounding v. Then $C^\dagger$ must cross one of the first N edges of R, where*

$$N \;:=\; 1 + K + \tfrac{1}{2}(\tfrac{1}{2}d_{\max}^\dagger - 1)\, L. \tag{A.3}$$

Proof. As discussed above, the cycle $C^\dagger$ corresponds to a minimal cutset C of G, which divides G into two connected components, one of which is finite and the other of which is infinite; moreover, v is contained by hypothesis in the finite component. Since R is an infinite ray starting at v, it must use somewhere an edge of C; let f be the first such edge, and let $f^\dagger \in C^\dagger$ be the corresponding dual edge.

Let w be the point in V_0 that is closest to v, let P be a path of minimal length $\overline{K} \le K$ linking v and w, and let D be a geodesic double ray containing w. Write $D = R_1 \cup R_2$ where R_1, R_2 are geodesic rays starting at w, and observe that $R_1' := P \cup R_1$ and $R_2' := P \cup R_2$ are rays starting at v. The cycle $C^\dagger$ must cross some edge in R_1' and some edge in R_2'. Let e_1, e_2 be the first edges (counting from v) in R_1', R_2' crossed by $C^\dagger$. We distinguish two cases: I. $e_1 \ne e_2$ and II. $e_1 = e_2$.

In case I, e_1 and e_2 lie on D and are the $(\overline{K}+N_1)$-th and $(\overline{K}+N_2)$-th edge of the rays R'_1 and R'_2, say. Then $C^\dagger$ is the union of two paths $P_1^\dagger, P_2^\dagger$, each of which has $e_1^\dagger$ and $e_2^\dagger$ as their final edges, and which are disjoint except for their overlap at $e_1^\dagger$ and $e_2^\dagger$. Let L_1, L_2 denote the lengths of these paths in $G^\dagger$, where $L_1 + L_2 - 2 = L$. Without loss of generality we may assume that $f^\dagger$ lies on $P_1^\dagger$ and is the M_1-th edge of $P_1^\dagger$ starting from e_1 and the M_2-th edge starting from e_2, where $M_1 + M_2 - 1 = L_1$. By Lemma A.4 (applied with the roles of G and $G^\dagger$ reversed), we have

$$\min(L_1, L_2) \;\geq\; d(e_1^\dagger, e_2^\dagger) + 1 \;\geq\; \frac{1}{c}\big(d(e_1, e_2) - 1\big) + 1 \;=\; \frac{1}{c}\big(N_1 + N_2 - 2\big) + 1 \;, \tag{A.4}$$

where we have abbreviated $c := \frac{1}{2} d_{\max}^\dagger - 1$. Therefore

$$M_1 + M_2 \;=\; L_1 + 1 \;=\; L + 3 - L_2 \;\leq\; L + 2 - \frac{1}{c}\big(N_1 + N_2 - 2\big) \;. \tag{A.5}$$

Using Lemma A.4 again, there exists a path in G of length at most $c(M_1 - 1) + 1$ that has e_1 and f as its final edges, and another path of length at most $c(M_2 - 1) + 1$ that has e_2 and f as its final edges. Combining these paths with the pieces of R'_1 and R'_2 leading up to e_1 and e_2, respectively, we find two paths in G starting at v and with f as their final edge, with lengths of at most

$$\overline{K} + N_1 + c(M_1 - 1) \quad \text{and} \quad \overline{K} + N_2 + c(M_2 - 1) \;, \tag{A.6}$$

respectively. By (A.5), it follows that the average length of these two paths is at most

$$\begin{aligned} &\tfrac{1}{2}\big(2\overline{K} + N_1 + N_2 + c(M_1 + M_2 - 2)\big) \\ &\quad \leq\; \tfrac{1}{2}\Big[2K + N_1 + N_2 + c\big(L - \frac{1}{c}(N_1 + N_2 - 2)\big)\Big] \;=\; 1 + K + \tfrac{1}{2} cL \;. \end{aligned} \tag{A.7}$$

Taking the shorter of these two paths, we have found a path of length at most $1 + K + \frac{1}{2}cL$ starting at v and having f as its final edge. But since R is a geodesic ray, the distance from v to f along R must be at most this.

In case II, $e := e_1 = e_2$ must lie on P; it is the first edge of P (counting from v) that is crossed by some edge in $C^\dagger$. Then $C^\dagger$ contains two paths with lengths L_1, L_2 satisfying $L_1 + L_2 - 2 = L$ that have $e^\dagger$ and $f^\dagger$ as their final edges. It follows that $d(e^\dagger, f^\dagger) \leq \frac{1}{2}L$. Therefore, by Lemma A.4, $d(e, f) \leq \frac{1}{2}cL + 1$, which means that we can find a path in G of length at most $\frac{1}{2}cL + 2$ that has e and f as its final edges. Combining this path with the piece of P leading from v to e, we again find a path of length at most $1 + K + \frac{1}{2}cL$ starting at v and having f as its final edge. $\square$

A.3. Planar embeddings. In this section we collect some results that show that sufficiently "nice" embeddings of a graph in the plane give rise to a geometric dual that is also an abstract dual, and conversely that graphs having a locally finite abstract dual have "nice" embeddings in the plane such that the geometric dual coincides with the abstract dual.

Embeddings of finite (planar) graphs in the plane are treated in almost any elementary book on graph theory, but it is more difficult to find a good reference for infinite planar graphs. Some articles that we have found useful are [6, 9, 19, 20, 48, 63, 67, 68].

Each graph G gives rise to a topological space — which, by a slight abuse of notation, we shall continue to call G — that is defined by first assigning a disjoint copy of $[0, 1]$ to each edge of G and a point to each vertex of G and then identifying the endpoints of intervals with the endvertices of the corresponding edges.[26] We equip G with the quotient topology arising from this identification:

[26] More precisely, given a graph $G = (V, E)$, we start from the set $(E \times [0, 1]) \cup (V \times \{2\})$, and then for each edge $e = xy$ we identify $(e, 0)$ with $(x, 2)$ and $(e, 1)$ with $(y, 2)$.

thus, a neighborhood base of an inner point on an edge is formed by the open intervals on the edge containing that point, while a neighborhood base of a vertex x is formed by the unions of half-open intervals $[x, z)$ containing x, one interval being taken from every edge $[x, y]$ incident on x.[27] Such a *topological realization* of the graph G is compact if and only if G is finite, and locally compact if and only if G is locally finite. It is metrizable if and only if G is locally finite.[28]

An *embedding* of a graph G in the plane is a continuous injective map $\phi\colon\ G \to \mathbb{R}^2$.[29] A graph G that can be embedded in the plane is called *planar*. A *plane graph* is a pair (G, ϕ) where G is a graph and ϕ is an embedding of G in the plane. We (topologically) identify the sphere $\mathbb{S}$ with the one-point compactification $\mathbb{R}^2 \cup \{\infty\}$ of the plane $\mathbb{R}^2$. Embeddings of graphs in the sphere are defined analogously to embeddings in the plane; a graph can be embedded in the sphere if and only if it can be embedded in the plane.

If G is a finite graph, then $\phi(G)$, being the continuous image of a compact set, is a closed subset of $\mathbb{R}^2$. Moreover, ϕ is necessarily a homeomorphism to its image, i.e., the inverse map $\phi^{-1}\colon\ \phi(G) \to G$ is also continuous. Both statements fail in general when G is infinite. Mainly for these reasons, not much can be said in general about embeddings of infinite graphs; one needs extra conditions to proceed.

So let G be a graph with no isolated vertices, and let $\phi\colon\ G \to \mathbb{S}$ be an embedding of G in the sphere. Following an idea of [48], we say that ϕ is *self-accumulation-free* if no point $z \in \phi(G)$ is an accumulation point of edges that do not contain z.[30] We say that ϕ is *pointed* if the image under ϕ of each ray in G converges to a point in $\mathbb{S}$.[31] (Clearly, equivalent rays must converge to the same point. Inequivalent rays may or may not converge to the same point.)

Recall that a compactification $\overline{F}$ of a topological space F is a compact topological space $\overline{F}$ such that $F \subseteq \overline{F}$ is dense. We will always require that $\overline{F}$ be Hausdorff. As in [63], we say that a compactification $\overline{G}$ of a connected, locally finite (not necessarily planar) graph G is *pointed* if each ray in G converges to some point in $\overline{G} \backslash G$.[32] (Clearly, equivalent rays must converge to the same point. Inequivalent rays may or may not converge to the same point.) If G is a graph and $\overline{G}$ is any compactification of G, then an *embedding* of $\overline{G}$ in the sphere is a continuous injective map $\overline{\phi}\colon\ \overline{G} \to \mathbb{S}$; since $\overline{G}$ is compact, it is necessarily a homeomorphism to its image.

---

[27] With this topology, G is a 1-dimensional CW-complex [34, pp. 5–6 and pp. 519 ff.].

[28] Indeed, if x is a vertex of infinite degree, then x does not have a countable neighborhood basis in the topology we have given G. (It is possible to equip G with a different topology, which is always metrizable, in such a way that a neighborhood base of a vertex x is formed by the unions of half-open intervals $[x, z)$ using *the same* distance $\epsilon = d(x, z)$ for each incident edge [27]. However, we shall not use this topology.)

[29] Traditionally an embedding is defined as a drawing in which the vertices are represented by distinct points and the edges are represented by closed continuous arcs joining their endvertices, mutually disjoint except possibly at their endpoints. Given an embedding in this sense, pasting together the continuous mappings from $[0, 1]$ to $\mathbb{R}^2$ corresponding to the individual edges always yields a continuous map $\phi\colon\ G \to \mathbb{R}^2$ (with the topology we have given G); and the converse is trivial. Therefore, our definition of embedding is equivalent to the traditional one.

[30] More precisely, if $x = \phi^{-1}(z)$, we let $E(x) \subseteq G$ be the union of all edges containing x; we then require that z is not contained in the closure of $\phi(G \setminus E(x))$. Krön [48] uses the term "accumulation-free", but we prefer the term "self-accumulation-free" in order to emphasize that only accumulation points *on the graph itself* are forbidden. In this way we clearly distinguish this concept from the standard concepts "VAP-free" and "EAP-free" to be introduced later, which forbid accumulation points *everywhere in the finite plane* $\mathbb{R}^2 = \mathbb{S} \backslash \{\infty\}$.

[31] Observe that a ray $R \subseteq G$ is homeomorphic to $[0, \infty)$ with its usual topology; what we are requiring here is that $\lim\limits_{x \to +\infty\ (x \in R)} \phi(x)$ exists in $\mathbb{S}$. Since $\mathbb{S}$ is compact, $\phi(x)$ must necessarily have *at least one* limit point as $x \to +\infty$; what we are requiring here is that it have *exactly one* limit point.

[32] Note that in *any* compactification $\overline{G}$ of G, the limit points of a ray in G cannot lie in G, because the topology of $\overline{G}$ extends that of G (and rays have no limit points in G). On the other hand, since $\overline{G}$ is compact, each ray in G necessarily has *at least one* limit point in $\overline{G}$. So what we are requiring here is that each ray should have *exactly one* limit point in $\overline{G}$.

Lemma A.6. (Embeddings of compactifications) *Let (G, ϕ) be a locally finite plane graph with no isolated vertices. Then the following conditions are equivalent:*

(i) ϕ is self-accumulation-free.
(ii) ϕ is a homeomorphism to its image.
(iii) ϕ can be extended to an embedding $\overline{\phi}\colon\ \overline{G} \to \mathbb{S}$ in the sphere of some compactification $\overline{G}$ of G. [The extension is of course unique, since G is dense in $\overline{G}$.]

Moreover, under these conditions, the embedding ϕ determines the compactification $\overline{G}$ uniquely (up to trivial renamings of points in $\overline{G}\backslash G$); and $\overline{G}$ is a pointed compactification if and only if ϕ is a pointed embedding.

Proof. If ϕ is not a homeomorphism to its image, then we can find $x_n, x \in G$ such that $\phi(x_n) \to \phi(x)$ but $x_n \not\to x$. Since ϕ is continuous and injective, the sequence (x_n) cannot have any accumulation points in G other than x; so by passing to a subsequence way may assume that it has no accumulation points at all. Since any finite collection of edges is compact, the sequence (x_n) must visit infinitely many edges. Since G is locally finite, only finitely many of these edges can contain x. It follows that ϕ is not self-accumulation-free.

Conversely, if ϕ is not self-accumulation-free, then there exists $z = \phi(x) \in \phi(G)$ and a sequence (x_n) belonging to edges not containing x, such that $\phi(x_n) \to \phi(x)$. Clearly $x_n \not\to x$, so ϕ is not a homeomorphism to its image.

If (iii) holds, then $\overline{\phi}$ is a homeomorphism to its image, hence so is its restriction ϕ to G. Conversely, if ϕ is a homeomorphism to its image, then we may topologically identify G with its image $\phi(G)$. Then the closure $\overline{\phi(G)}$ of $\phi(G)$ in the sphere $\mathbb{S}$ is a compactification of $\phi(G)$ such that the identity map from $\phi(G)$ to itself can be continuously extended to $\overline{\phi(G)}$; moreover, $\overline{\phi(G)}$ is (up to renaming) the only compactification of $\phi(G)$ with this property.

Clearly, if $\overline{\phi(G)}$ is a pointed compactification, then ϕ is a pointed embedding. Conversely, if ϕ is a pointed embedding, then each ray in G converges to some point $z \in \mathbb{S}$; but since ϕ is self-accumulation-free, we have $z \notin \phi(G)$, hence $\overline{\phi(G)}$ is a pointed compactification. □

Remark. If G is not locally finite, then ϕ can *never* be a homeomorphism of G (with the topology we have given it) to its image. For if x is a vertex of infinite degree, then it is easy to choose points x_n lying on infinitely many distinct edges such that $\phi(x_n) \to \phi(x)$; but $x_n \not\to x$.

Let $\phi\colon\ G \to \mathbb{S}$ be any embedding of a graph G in the sphere. Following [67,68], let us define a *vertex accumulation point* (resp. *edge accumulation point*) of an embedding ϕ to be a point in $\mathbb{S}$ such that each of its open neighborhoods contains infinitely many vertices (resp. intersects infinitely many edges). We abbreviate "vertex accumulation point" and "edge accumulation point" by VAP and EAP, respectively. Let us also say that a point $x \in \mathbb{S}$ is an *endpoint* of (G, ϕ) if there exists a ray in G such that its image under ϕ converges to x. The next lemma says that for self-accumulation-free pointed embeddings of 2-connected graphs, all these concepts coincide:

Lemma A.7. (Endpoints of plane graphs) *Let G be a locally finite 2-connected graph, and let $\phi\colon\ G \to \mathbb{S}$ be a self-accumulation-free pointed embedding of G. Then the following four sets are equal:*

(i) The set of vertex accumulation points of (G, ϕ).
(ii) The set of edge accumulation points of (G, ϕ).
(iii) The set of endpoints of (G, ϕ).
(iv) $\overline{\phi(G)}\backslash\phi(G)$.

Proof. We will prove the inclusions (iii)⊆(i)⊆(iv)⊆(iii) and the same with (i) replaced by (ii). Clearly, every endpoint is also a vertex and edge accumulation point. Since ϕ is self-accumulation-free and locally finite, each vertex or edge accumulation point lies in $\overline{\phi(G)}\backslash\phi(G)$. For each point

$x \in \overline{\phi(G)} \setminus \phi(G)$, we can find a sequence of points $x_n \in \phi(G)$, all lying on different edges, such that $x_n \to x$. Since G is 2-connected, it can be shown that there exists a ray R in G whose image under ϕ passes through infinitely many of these points. [33] Since ϕ is pointed, it follows that x is the limit of the ϕ-image of R. □

Let $\overline{G}$ be a pointed compactification of a connected locally finite (not necessarily planar) graph G. Then, by definition, each ray in G converges to some limit in $\overline{G} \setminus G$; and conversely, if G is 2-connected, then the argument proving (iv)⊆(iii) in Lemma A.7 shows more generally [63, p. 4591] that each point in $\overline{G} \setminus G$ is the limit point of some ray in G.

Each connected locally finite graph G has a unique (up to renaming) pointed compactification $\overline{G}$ such that each point in $\overline{G} \setminus G$ is the limit point of some ray in G and moreover two nonequivalent rays always converge to different limit points: this is the *Freudenthal compactification* $\mathcal{F}(G)$, see [16, Section 8.5] or [63, Section 7].[34] Clearly, $\mathcal{F}(G) \setminus G$ is in bijection with the space of ends of G. If G is 2-connected, then, in a sense, $\mathcal{F}(G)$ is the "largest" pointed compactification of G (compare the remarks on [63, p. 4594]). At the other end of the scale, every connected, locally finite, infinite graph G has a smallest pointed compactification, namely the one-point compactification $G^\bullet = G \cup \{\infty\}$ in which every infinite sequence of distinct vertices or edges (hence in particular every ray) converges to the single point ∞.

By Lemma A.6, an embedding ϕ of a connected locally finite graph G is self-accumulation-free and pointed if and only if ϕ extends to an embedding $\overline{\phi}$ of a pointed compactification $\overline{G}$ of G. Richter and Thomassen [63, Theorems 1 and 13] have proven the following key result concerning the existence and uniqueness of such embeddings:

Theorem A.8. (Embeddings of 3-connected planar graphs) *Let G be a 3-connected locally finite planar graph. Then:*

(a) There exists an embedding of the Freudenthal compactification $\mathcal{F}(G)$ in the sphere $\mathbb{S}$.
(b) If $\overline{G}$ is any pointed compactification of G and $\overline{\phi}_1, \overline{\phi}_2$ are embeddings of $\overline{G}$ in the sphere $\mathbb{S}$, then there exists a homeomorphism $h\colon\ \mathbb{S} \to \mathbb{S}$ such that $\overline{\phi}_2 = h \circ \overline{\phi}_1$.

The unique (up to homeomorphism) embedding of the Freudenthal compactification $\mathcal{F}(G)$ in the sphere $\mathbb{S}$ will be called the *Freudenthal embedding*.

Henceforth we assume that G is a 2-connected locally finite graph and that $\phi\colon\ G \to \mathbb{S}$ is a self-accumulation-free pointed embedding of G in the sphere. By Lemma A.6, ϕ extends uniquely to an embedding $\overline{\phi}\colon\ \overline{G} \to \mathbb{S}$ of an essentially unique pointed compactification $\overline{G}$ of G, and we have $\overline{\phi}(\overline{G}) = \overline{\phi(G)}$. The connected components of the remaining open set $\mathbb{S} \setminus \overline{\phi}(\overline{G})$ are called the *faces* of the plane graph (G, ϕ) [and also of $(\overline{G}, \overline{\phi})$]. For any face f, we let $\partial f := \overline{f} \setminus f$ denote its (topological) boundary. If X is any topological space, we say that a subset $C \subseteq X$ is a *circle* if it is homeomorphic to the unit circle $\mathbb{S}^1$. We need the following fundamental facts about the boundaries of faces:

Theorem A.9. (Boundaries of faces) *Let G be a 2-connected locally finite (planar) graph, and let ϕ be a self-accumulation-free pointed embedding of G in the sphere $\mathbb{S}$. Then each face f of (G, ϕ) is bounded by a circle, and exactly one of the following three possibilities holds:*

---

[33] Given an infinite set $B \subseteq E$, we can use local finiteness to extract an infinite subset $B' \subseteq B$ that is pairwise vertex-disjoint; then by [63, Proposition 8], G has a ray that uses infinitely many of the edges in B'. Given an infinite subset $T \subseteq V$, we can obviously choose an infinite set B of edges containing all the vertices in T, and then proceed as before.

[34] The Freudenthal compactification is a general construction of point-set topology: it is defined for locally compact Hausdorff spaces, or more generally for completely regular rim-compact spaces (see e.g. [15,49]). For *locally finite* graphs — which are the only ones we are concerned with — this topological definition coincides with the graph-theoretic definition given here or in [16, Section 8.5].

(i) ∂f *is the* ϕ*-image of a (finite*[35]*) cycle in* G *[hence* $\partial f \cap [\overline{\phi(G)} \backslash \phi(G)] = \varnothing$*].*
(ii) $\partial f \cap \phi(G)$ *is the disjoint union of a nonempty countable collection of double rays, and* $\partial f \cap [\overline{\phi(G)} \backslash \phi(G)] \neq \varnothing$*.*
(iii) $\partial f \cap \phi(G) = \varnothing$ *[hence* $\partial f \cap [\overline{\phi(G)} \backslash \phi(G)] = \partial f \simeq \mathbb{S}^1$*].*

Moreover, when $\overline{G}$ *is the Freudenthal compactification* $\mathcal{F}(G)$*, the set* $\partial f \cap \phi(G)$ *is dense in* ∂f*; in particular, case (iii) cannot occur.*

Proof. By Lemma A.6, ϕ extends uniquely to an embedding $\overline{\phi}$ of an (essentially unique) pointed compactification $\overline{G}$ of G. Richter and Thomassen [63, Proposition 3 and Theorem 7] have proven the fundamental result that every face of $(\overline{G}, \overline{\phi})$ is bounded by a circle. By Lemma A.6, such a circle is the image of a circle C in $\overline{G}$ under the homeomorphism $\overline{\phi}$. We shall therefore prove the more general result that circles in $\overline{G}$ have properties analogous to those stated in the theorem (here G need not be planar). We first claim that

(a) Whenever C contains an inner point of an edge $e \in G$, it contains the entire edge e.
(b) Whenever C contains a vertex $x \in G$, it contains precisely two edges of G that are incident on x.

For the Freudenthal compactification, (a) and (b) are proven in [20, Lemma 2.3]. But since statements (a) and (b) concern only a neighborhood of a point of G, they hold for any Hausdorff compactification of G (since G is locally compact). For the Freudenthal compactification, it is moreover proven in [19, Lemma 4.3] that $C \cap G$ is dense in C.

By (a,b), the union of all edges belonging to C is a 2-regular graph $C \cap G$, so each connected component of $C \cap G$ is either a cycle or a double ray. Since $C \cap G$ is moreover homeomorphic to a subset of $\mathbb{S}^1$, we conclude that it is either empty [case (iii)], a cycle [case (i)], or a nonempty disjoint union of double rays [case (ii)]. In the latter case, there can be only countably many double rays (since G is countable), and C must also contain some points of $\overline{G} \backslash G$ (in fact, at least as many as there are double rays) in order to have the topology of a circle. □

Examples. 1. The lattices shown in Figure 1.1 have one end, which is mapped to ∞, and the boundary of every face is a cycle in G. As we will see in Proposition A.16 below, quasi-transitivity and the fact that there is only one end imply that every face is bounded by a cycle.

2. Consider $\mathbb{N} \times \mathbb{Z}$ with nearest-neighbor edges, with its usual embedding in the plane. This graph has one end, which is mapped to ∞. It has a face that is bounded by a double ray together with $\{\infty\} = \overline{\phi(G)} \backslash \phi(G)$.

3. The graph in the top left of Figure A.1 has uncountably many ends, which are mapped by the embedding ϕ onto the unit circle. This embedding has one face that is bounded by endpoints alone (namely, the exterior face), as well as faces that are bounded by cycles and faces that are bounded by a double ray together with one endpoint. This graph is constructed in the following way. Consider a 3-regular tree with origin $\varnothing$. Each vertex of G is labeled by a finite (possibly empty) sequence $\ell_0 \ell_1 \ell_2 \cdots \ell_n$ where $\ell_0 \in \{0, 1, 2\}$ and $\ell_i \in \{0, 1\}$ for $i \geq 1$; and each ray emanating from $\varnothing$ is labeled by an infinite sequence $\ell_0 \ell_1 \ell_2 \cdots$ of the same type. The tree can now be embedded in the unit disk such that the ray $\ell_0 \ell_1 \ell_2 \cdots$ converges to the point

$$2\pi \cdot \frac{1}{3}\Big(\ell_0 + \sum_{i=1}^{\infty} 2^{-i} \ell_i\Big) \tag{A.8}$$

[35] Of course, cycles in the standard graph-theoretic meaning of the word are by definition finite. We will always use the word in this standard sense; we shall not make explicit use of the "infinite cycles" introduced (by topological means) by Diestel and collaborators [19, 20, 9, 17, 10].

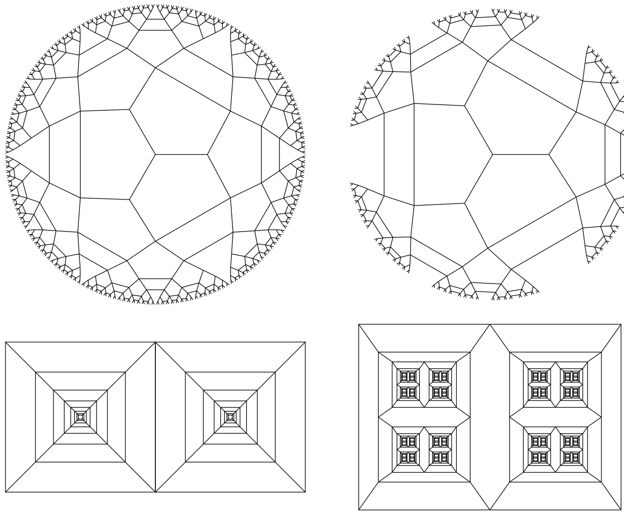

Fig. A.1. Four examples of self-accumulation-free pointed embeddings of 3-connected graphs, demonstrating some of the different possibilities for the set of endpoints and the boundaries of faces.

on the unit circle. We then make the tree into a 3-connected graph by adding edges. For each vertex $\ell \neq \varnothing$, we draw edges between the points $\ell 01$ and $\ell 10$ as well as between $\ell 011$ and $\ell 100$. Moreover, around the origin, we draw the edges

$$\{01,\, 10\},\ \{011,\, 100\},\ \{11,\, 20\},\ \{111,\, 200\},\ \{21,\, 00\},\ \{211,\, 000\}\ . \tag{A.9}$$

This embedding is not a Freudenthal embedding, since two nonequivalent rays (e.g. $0111\cdots$ and $1000\cdots$) may converge to the same point. However, we can easily modify this example to obtain a Freudenthal embedding of the same graph, by replacing the binary expansion $2^{-i}\ell_i$ in (A.8) by ternary expansion $2 \cdot 3^{-i}\ell_i$, as drawn in the top right of Figure A.1. In this latter example, the set of ends is mapped onto a Cantor subset of the unit circle. There is now a face f of (G, ϕ) that contains the exterior of the unit disc but also extends inside the unit disc; its boundary consists of a countably infinite collection of disjoint double rays together with all of $\overline{\phi(G)} \backslash \phi(G)$. Note that here $\partial f \cap \phi(G)$ is dense in ∂f.

4. The graphs in the bottom row of Figure A.1 are examples of Freudenthal embeddings of graphs with more than one end, in which each face is bounded by a cycle. The left graph has two ends (and can easily be modified to produce any finite or countable number of ends), while the right graph has uncountably many ends.

Let us now say that faces f_1 and f_2 *border* each other in an edge e if e lies in the boundary of both faces, with f_1 on one side of e and f_2 on the other. Since G is 2-connected, it is not hard to see that no face borders itself; rather, each edge e lies in the common boundary of precisely two faces, which border each other in e. If G is moreover 3-connected, then two faces border each other in at most one edge.[36] So let us assume henceforth that G is 3-connected. For each edge e of G,

[36] In fact, it is easy to see that 2-edge-connectedness and 3-edge-connectedness are sufficient, respectively, for these statements.

we let $e^* := \{f_1, f_2\}$ denote the pair of faces of G that border each other in e. We then define the *geometric dual* $G^* = (V^*, E^*)$ of the plane graph (G, ϕ) to be the graph whose vertex set V^* is the set of all faces of (G, ϕ) that have at least one edge in their boundary [i.e., fall in case (i) or (ii) of Theorem A.9] and whose edge set is $E^* := \{e^* \colon\ e \in E\}$. Note that the edge sets E and E^* are in bijection under $e \mapsto e^*$; for any subset $B \subseteq E$ we denote by $B^* = \{e^* \colon\ e \in B\}$ the corresponding subset of E^*.

We can now prove that geometric duals, in this generality, are always abstract duals:

Theorem A.10. (Geometric duals are abstract duals) *Let G be a locally finite 3-connected graph, and let $\phi\colon\ G \to \mathbb{S}$ be a self-accumulation-free pointed embedding of G in the sphere. Let G^* be the geometric dual associated with this embedding. Then:*

(i) G^* is connected.
(ii) G^* is an abstract dual of G.
(iii) G^* is locally finite if and only if every face of G that has at least one edge in its boundary is bounded by a cycle.

Proof. For each vertex v of G, let E_v denote the set of edges incident on v. Since E_v is finite (say, of cardinality n) and ϕ is self-accumulation-free, basic topological considerations (based on the Jordan curve theorem) imply that $\phi(v)$ is surrounded by exactly n faces that border each other pairwise (in cyclic order) in the edges of E_v, and that E_v^* is a cycle in $G^* = (V^*, E^*)$.

(i) Consider any pair $v^*, w^* \in V^*$. By definition, v^* and w^* are faces of (G, ϕ) that contain at least one edge in their boundary; so let v and w, respectively, be any endvertex of any such edge. Since G is connected, there exists a path $v = v_1, v_2, \ldots, v_k = w$ in G. Since each cycle $E_{v_i}^*$ is connected, and $E_{v_i}^* \cap E_{v_{i+1}}^* = \{\{v_i, v_{i+1}\}^*\} \neq \varnothing$, the set $\bigcup_{i=1}^k E_{v_i}^*$ is connected; moreover, v^* belongs to $E_{v_1}^*$ and w^* belongs to $E_{v_k}^*$. So there exists a path in $\bigcup_{i=1}^k E_{v_i}^*$ from v^* to w^*.

(ii) Consider any cycle C in G. Then $\phi(C)$ is a circle in $\phi(G)$ that (by the Jordan curve theorem) partitions V^* into two sets (call them V_1^* and V_2^*) corresponding to faces lying in the two components of $\mathbb{S}\backslash\phi(C)$. Since C consists exactly of those edges $e \in E$ that lie on the boundary of some pair of faces $v_1^* \in V_1^*$ and $v_2^* \in V_2^*$, it follows that V_1^* and V_2^* are both nonempty, so that the set C^* is a cutset of G^*.

For each $v \in C$, the set $E_v \cap C$ consists of precisely two edges (call them e and f) such that e borders $v_1^* \in V_1^*$ and $v_2^* \in V_2^*$ and f borders $w_1^* \in V_1^*$ and $w_2^* \in V_2^*$. Then $E_v^*\backslash C^*$ is the disjoint union of two paths P_1^*, P_2^* in G^* that link v_1^* and w_1^* in V_1^*, and v_2^* and w_2^* in V_2^*, respectively. Doing this for all $v \in C$, it follows that the set of all vertices in V_1^* (resp. V_2^*) that are incident to a an edge of the cutset C^* is connected in V_1^* (resp. V_2^*). Therefore C^* is a minimal cutset of G^*.

Conversely, if B is a finite subset of E such that B^* is a cutset of G^*, then for each $v \in V$, the cutset B^* must intersect E_v^* (which is a cycle in G^*) an even number of times (that is, $B^* \cap E_v^*$ has even cardinality). It follows that each vertex of G is incident on an even number of edges of B, and hence B contains a cycle C. But we have just proven that C^* is a cutset in G^*. Therefore, if B^* is a minimal cutset of G^*, we necessarily have $B = C$.

(iii) By Theorem A.9, for every face f of (G, ϕ) that has at least one edge in its boundary (i.e., belongs to V^*), its boundary ∂f must either be a cycle or contain a double ray. Clearly, f has finite degree in G^* in case ∂f is a cycle, and has countably infinite degree in case ∂f contains a double ray. □

In particular, whenever G is a locally finite 3-connected graph and $\phi\colon\ G \to \mathbb{S}$ is a self-accumulation-free pointed embedding such that every face of G is bounded by a cycle, the geometric dual G^* is a locally finite abstract dual of G, so that all the nice properties stated in Theorem A.3 ensue.

Conversely, let us now show that if a locally finite 3-connected graph G has a locally finite abstract dual, then G has a unique self-accumulation-free pointed embedding, and the geometric dual associated with this embedding coincides with the abstract dual:

Theorem A.11. (Embeddings of graphs with locally finite abstract duals) *Let $G = (V, E)$ be a locally finite 3-connected graph that has a locally finite abstract dual $G^\dagger = (V^\dagger, E^\dagger)$. Then:*

(i) G is planar.
(ii) In the Freudenthal embedding of G, each face is bounded by a cycle, and the geometric dual G^* of this embedding coincides with the (unique) abstract dual $G^\dagger$.
(iii) The Freudenthal embedding is (up to homeomorphism) the only self-accumulation-free pointed embedding of G in the sphere.

Proof. (i) is proven in [67, Theorem 9.3]. As a consequence of (i), Theorem A.8 guarantees the existence and uniqueness (up to homeomorphism) of the Freudenthal embedding.

As a preliminary to (ii) and (iii), consider any self-accumulation-free pointed embedding of G in the sphere, and let G^* be the geometric dual associated with this embedding. By Theorem A.10, G^* is an abstract dual of G. But by Theorem A.3(iii), G has a unique abstract dual $G^\dagger$. Hence $G^\dagger = G^*$.

(ii) For the Freudenthal embedding, Theorem A.9 guarantees that every face has at least one edge in its boundary. And since $G^\dagger = G^*$ is locally finite, Theorem A.10(iii) implies that every face is bounded by a cycle.

(iii) Let R_1, R_2 be inequivalent rays in G. Then there exists a finite set of edges B in G such that the tails of R_1 and R_2 lie in different connected components of $G\backslash B$. Making B smaller if necessary, we may assume without loss of generality that B is a minimal set with this property. It is not hard to see that B must then be a minimal cutset. Since G is an abstract dual of $G^\dagger$ [by Theorem A.3(i)], $B^\dagger$ is a cycle in $G^\dagger$. Since $G^\dagger = G^*$, B^* is a cycle in G^*. By (ii), each of the vertices of the cycle B^* is a face of (G, ϕ) whose boundary is a cycle in G. The sum modulo 2 of all these cycles forms a pair of disjoint cycles in G, each of which separates the tail of $\phi(R_1)$ from the tail of $\phi(R_2)$. As a result, these tails necessarily converge to different endpoints. Since this holds for any pair of inequivalent rays, we conclude that $\overline{G}$ must be the Freudenthal compactification, and ϕ the Freudenthal embedding. □

Until now we have defined dual graphs — whether in the abstract or geometric sense — simply as abstract graphs without a given embedding in the plane. However, given an embedding of a planar graph, there is a natural way to embed its dual. Let G be a locally finite 3-connected graph, let $\phi\colon\ G \to \mathbb{S}$ be a self-accumulation-free pointed embedding of G in which each face is bounded by a cycle, and let G^* be the geometric dual of (G, ϕ). If e is an edge of G, we denote by $\mathring{e}$ the interior of the edge e in the topological realization of G (namely, the edge without its endvertices). We then say that an embedding ϕ^* of G^* is a *dual embedding* to (G, ϕ) if

(i) For each $v^* \in V^*$, we have $\phi^*(v^*) \in v^*$.
(ii) For each $e \in E$, the arc $\phi^*(e^*)$ intersects $\phi(G)$ in a single point, which lies on $\phi(\mathring{e})$.

Less formally, (i) says that each vertex of G^* is represented by a point lying in the correponding face of (G, ϕ), and (ii) says that two such points that lie in faces that border each other in an edge e are linked by a dual edge e^* that crosses e in a single interior point and is otherwise disjoint from $\phi(G)$. We then have:

Proposition A.12. (Dual embeddings) *Let G be a locally finite 3-connected graph, and let $\phi\colon\ G \to \mathbb{S}$ be a self-accumulation-free pointed embedding of G such that each face of (G, ϕ) is bounded by a cycle. Let G^* be the geometric dual associated with this embedding. Then there exists a dual embedding ϕ^* of G^*. Moreover, for each such ϕ^*:*

(i) ϕ^* is self-accumulation-free and pointed, and each face of (G^*, ϕ^*) is bounded by a cycle.
(ii) G is a geometric dual of (G^*, ϕ^*) and ϕ is a dual embedding of G.
(iii) The sets of endpoints $\overline{\phi(G)}\backslash\phi(G)$ and $\overline{\phi^*(G^*)}\backslash\phi^*(G^*)$ coincide.

If (G, ϕ) and (G^*, ϕ^*) are as in Proposition A.12, then we say that they form a *geometric dual pair*.

Proof of Proposition A.12. To prove the existence of a dual embedding ϕ^* of G^*, we begin by choosing a point $\phi^*(v^*) \in v^*$ for each face v^* of (G, ϕ), and a point x on $\phi(\mathring{e})$ for each edge $e \in E(G)$. Let $x_1, \ldots, x_n$ be the points chosen on the edges of the cycle that bounds v^*. It follows from basic topological considerations[37] that we can connect $\phi^*(v^*)$ to each of the points $x_1, \ldots, x_n$ by continuous arcs that are disjoint except at their common endpoint $\phi^*(v^*)$ and that lie entirely in v^* except for their endpoints $x_1, \ldots, x_n$. Now if $e \in E(G)$ lies on the boundary of faces v_1^* and v_2^* and x is the chosen point on $\phi(\mathring{e})$, then the concatenation of the arcs from $\phi^*(v_1^*)$ to x and from x to $\phi^*(v_2^*)$ is an arc from $\phi^*(v_1^*)$ to $\phi^*(v_2^*)$ that we take as our definition of $\phi^*(e^*)$.

To see that ϕ^* is self-accumulation-free, it suffices to show that for each point $x \in G^*$ we can remove a finite set of edges E_0^* from G^* such that $\phi^*(x) \notin \overline{\phi^*(G^* \backslash E_0^*)}$. If $x = v^*$ is a vertex of G^*, then it suffices to remove the set E_0^* of all edges incident on v^*. Then $\phi^*(G^* \backslash E_0^*)$ is contained in the complement of the face v^* and hence its closure cannot contain $\phi^*(x)$. If x lies on an edge e^* connecting two vertices v_1^* and v_2^* of G^*, then we remove the set E_0^* of all edges incident on v_1^* or v_2^*. Then $\phi^*(G^* \backslash E_0^*)$ is contained in the complement of the open set formed by the union of the faces v_1^* and v_2^* and $\phi(\mathring{e})$, while $\phi^*(x)$ lies inside this open set.

To see that ϕ^* is pointed, let R^* be a ray in G^* consisting of consecutive vertices $v_0^*, v_1^*, \ldots$ connected by edges $e_1^*, e_2^*, \ldots$. Then $v_0^*, v_1^*, \ldots$ are faces of (G, ϕ) which by assumption are bounded by cycles $C_0, C_1, \ldots$ such that $C_{k-1} \cap C_k = e_k$. Adding these cycles modulo 2 yields a double ray consisting of two rays R_1, R_2 in $\phi(G)$. Moreover, we can construct a third ray R_3 that lies in the union of the cycles $C_0, C_1, \ldots$, passes through each of the edges $e_1, e_2, \ldots$, and has an infinite intersection with both R_1 and R_2. In particular, the rays R_1, R_2 and R_3 are all equivalent and thus, as ϕ is pointed, they all converge to the same point $x \in \mathbb{S}$. It follows that for each open ball around x, there exists an n such that the sum modulo 2 of the cycles $C_n, C_{n+1}, \ldots$ lies in this ball. But this is a double ray which together with the point x forms a circle in $\mathbb{S}$ that contains the tail of the ray R^* in its interior. Any open ball around x thus eventually contains the tail of R^*, which proves that R^* converges to x.

This proves not only that (G^*, ϕ^*) is pointed, but also that every point in $\mathbb{S}$ that is the limit of some ray in (G^*, ϕ^*) is also the limit of some ray in (G, ϕ). It follows that the set of endpoints of (G^*, ϕ^*) is contained in the set of endpoints of (G, ϕ): that is, $\overline{\phi^*(G^*)} \backslash \phi^*(G^*) \subseteq \overline{\phi(G)} \backslash \phi(G)$.

To complete the proof, it now suffices to show that each face of (G^*, ϕ^*) is bounded by a cycle [completing the proof of (i)] and contains a single vertex of (G, ϕ) [which implies (ii)]. Then (ii) together with what we already got implies the reverse inclusion $\overline{\phi(G)} \backslash \phi(G) \subseteq \overline{\phi^*(G^*)} \backslash \phi^*(G^*)$, so we conclude that (iii) holds.

Since every face of (G, ϕ) is bounded by a cycle, the set $\overline{\phi(G)} \backslash \phi(G)$ does not separate $\mathbb{S}$, so the same is true for $\overline{\phi^*(G^*)} \backslash \phi^*(G^*)$ which is contained in it. It follows that each face f of (G^*, ϕ^*) has at least one edge in its boundary and hence by Theorem A.9 is either a cycle C^* or contains a double ray R^*. For each vertex v^* in C^* or R^* we have that $\phi^*(v^*)$ lies inside a cycle C of G that crosses two consecutive edges of C^* or R^*. This cycle C can have only one vertex in f since otherwise there would be an edge of (G, ϕ) that lies entirely in f, which contradicts our construction of ϕ^*. It follows that all edges in (G, ϕ) that cross an edge in C^* or R^* are incident to one and the same vertex v of G with $\phi(v) \in f$. Since G is locally finite, we can rule out the double ray so we conclude that f is bounded by a cycle C^* and contains a single vertex of (G, ϕ). $\square$

Now recall the definition of vertex and edge accumulation points (just before Lemma A.7). Following [67,68], let us say that an embedding $\phi\colon\ G \to \mathbb{R}^2 \cup \{\infty\} \cong \mathbb{S}$ is *VAP-free* (resp. *EAP-free*) if $\phi(G) \subset \mathbb{R}^2$ and ϕ has no VAPs (resp. EAPs) in the finite plane $\mathbb{R}^2$. Note that if G has at most

---

[37] Indeed, this can be proven by repeated application of the Jordan curve theorem and the following basic topological fact: Let $C \subset \mathbb{S}$ be a circle, which by the Jordan curve theorem divides $\mathbb{S}$ into two connected open sets U, V; let x, z be two distinct points on C, and let $y \in U$. Then there exists a continuous arc from x to z that lies, except for its endpoints, entirely in U and passes through y.

finitely many isolated vertices (in particular, if G is connected), then every VAP is also an EAP. For such graphs, EAP-free embeddings are automatically self-accumulation-free and pointed.[38]

An EAP-free embedding ϕ of an infinite, connected, locally finite graph G can be extended to an embedding $\overline{\phi}$ of the one-point compactification $G^\bullet$ of G by setting $\overline{\phi}(\infty) = \infty$. Conversely, if the one-point compactification $G^\bullet$ of G can be embedded in the sphere $\mathbb{S} \cong \mathbb{R}^2 \cup \{\infty\}$, then without loss of generality we may assume $\overline{\phi}(\infty) = \infty$, yielding an EAP-free embedding of G. We then have:

Proposition A.13. (Graphs with one end) *Let G be a locally finite 3-connected planar graph.*

(i) If G has at most one end, then it has an EAP-free embedding, which is unique up to homeomorphism and coincides with the Freudenthal embedding.

If in addition G has a locally finite abstract dual $G^\dagger$, then:

(ii) G has an EAP-free embedding if and only if *it has at most one end; and in this case G and $G^\dagger$ can be represented as a geometric dual pair (G, ϕ), (G^*, ϕ^*) such that both (G, ϕ) and (G^*, ϕ^*) are EAP-free.*

Proof. (i) If G is finite, then it obviously has an EAP-free embedding. If G is infinite with one end, then it has an EAP-free embedding since its one-point compactification coincides with its Freudenthal compactification, which by Theorem A.8(a) can be embedded in the sphere. And by Theorem A.8(b) this embedding is unique up to homeomorphism.

(ii) If G has a locally finite abstract dual $G^\dagger$, then Theorem A.11(iii) and Lemma A.6 tell us that the Freudenthal compactification is the *only* pointed compactification of G that is embeddable in the sphere; therefore, such a graph has an EAP-free embedding *if and only if* it has at most one end. (This statement can also be found in [68, Theorem 5.9].) By Proposition A.12, G and $G^\dagger$ can be represented as a geometric dual pair (G, ϕ), (G^*, ϕ^*); and Proposition A.12(iii) then implies in particular that if (G, ϕ) is EAP-free, then so is (G^*, ϕ^*). □

A.4. Special embeddings. In this section we discuss embeddings with special "nice" properties: notably, straight-line embeddings and periodic embeddings.

We begin by citing a result of Thomassen [67] on straight-line embeddings with convex faces:

Proposition A.14. (Convex embeddings) *Let G be a locally finite, 3-connected graph with one end. Then there exists an EAP-free embedding of G in which every edge is a straight line segment and every face is convex. If in addition G has a locally finite abstract dual $G^\dagger$, then there exists an EAP-free embedding of G in which every face is a convex polygon.*

Proof. By Proposition A.13(i), G has an EAP-free embedding in the plane. Then [67, Theorems 7.4 and 8.6] imply that G has an EAP-free embedding in which every edge is a straight line segment and every face is convex. If in addition G has a locally finite abstract dual $G^\dagger$, then Theorem A.11 guarantees that every face is bounded by a cycle. □

Using Proposition A.14, we can deduce that G and $G^\dagger$ can be *jointly* embedded as a geometric dual pair such that edges are represented by straight-line segments in both graphs:

Proposition A.15. (Straight-line embedding of dual pair) *Let G be a locally finite, 3-connected graph with one end. Assume that G has a locally finite abstract dual $G^\dagger$. Then there exist EAP-free embeddings ϕ and $\phi^\dagger$ of G and $G^\dagger$ in the (Euclidean) plane such that (G, ϕ) and $(G^\dagger, \phi^\dagger)$ form a geometric dual pair and each edge of (G, ϕ) and $(G^\dagger, \phi^\dagger)$ is a straight line segment.*

[38] Since $\phi(G) \subset \mathbb{R}^2$ and ϕ has no EAPs in $\mathbb{R}^2$, it must be self-accumulation-free. Since every ray has at least one accumulation point in $\mathbb{R}^2 \cup \{\infty\}$ and no accumulation points in $\mathbb{R}^2$, all rays converge to ∞. It seems to us that the concepts "self-accumulation-free" and "pointed" can now largely replace the more restricted notion of being EAP-free.

PROOF. By Theorem A.11 and Proposition A.12, we may embed G and $G^\dagger$ as a geometric dual pair. Now we can also draw a graph H in the plane such that the vertex set of H is the union of the vertex sets of G and $G^\dagger$ and two vertices $v \in G$ and $w \in G^\dagger$ are joined by an edge in H when v and w are the endpoints of an edge in G and its dual edge in $G^\dagger$; thus, H is a quadrangulation, with G and $G^\dagger$ being its sublattices whose edges connect opposing vertices of the quadrilaterals of H. Then obviously H is planar, has an EAP-free embedding in the plane, and has a locally finite geometric (and hence abstract) dual. Since two faces of H never have the property that two non-adjacent vertices of H both lie on the boundary of both faces, we see that H is 3-connected. By Proposition A.13(ii), H has one end, so applying Proposition A.14 to H we see that H has an EAP-free embedding such that each face is a convex polygon (with four corners). Connecting opposite corners of these faces by straight line segments, we obtain the required straight-line embeddings of (G, ϕ) and $(G^\dagger, \phi^\dagger)$ as a geometric dual pair. $\square$

Remark. We do not know whether G and $G^\dagger$ can be jointly embedded as a geometric dual pair such that edges are represented by straight line segments in both graphs *and* faces are convex in both graphs.

We next turn our attention to embeddings of quasi-transitive graphs. Our main graphs of interest — namely, the graph G and its sublattices G_0 and G_1 arising in Theorem 1.1 — are all locally finite, 3-connected, quasi-transitive, planar graphs with one end that have a locally finite abstract dual $G^\dagger$. Remarkably, this latter condition turns out to be superfluous:

Proposition A.16. (Quasi-transitive graphs with one end) *Let G be a locally finite, 3-connected, quasi-transitive, planar graph with one end. Then:*

(i) G has a locally finite abstract dual $G^\dagger$.
(ii) The Freudenthal compactification is the only pointed compactification of G.
(iii) The Freudenthal embedding of G is EAP-free (when the endpoint is taken to map to ∞), and each face is bounded by a cycle.

PROOF. (ii) is trivial because G has one end (and is 2-connected). To prove (iii) and (i), consider the Freudenthal embedding of G. By Proposition A.13(i), it is EAP-free. By Theorem A.9, no face is bounded by endpoints alone; and since G has one end, a face boundary containing infinitely many edges must consist of a single double ray together with the one endpoint. But quasi-transitivity then implies (see [6, Theorem 2.3] or [48, Theorem 8(1)]) that no such infinite face can exist, i.e., every face is bounded by a cycle. From Theorem A.10 we then conclude that G has a locally finite abstract dual $G^\dagger$. $\square$

Quasi-transitivity is essential here, as the example of $\mathbb{N} \times \mathbb{Z}$ (see Example 2 after Theorem A.9) shows. The assumption that G has one end is also essential: let H be the graph with vertex set $\mathbb{Z}$ and edges $\{n, n+1\}$ and $\{2n, 2n+2\}$ for all $n \in \mathbb{Z}$, let G be the ladder graph $\{0,1\} \times H$, and let ϕ be the obvious embedding. Then G is quasi-transitive and 3-connected and has two ends; ϕ is EAP-free and maps both ends to the point ∞ (i.e., this is a non-Freudenthal embedding); and (G, ϕ) has a pair of faces that are each bounded by a double ray together with the point ∞.

When planar graphs are quasi-transitive, it is natural to ask if they can be embedded in a periodic way in the plane. This is not true if one restricts oneself to the Euclidean plane, but remarkably, it turns out to be correct if one also considers the hyperbolic plane. The following result has been proven in [3, Theorem 4.2] (see also [66, Theorem 1]):

Theorem A.17. (Periodic embeddings) *Every locally finite, 3-connected, quasi-transitive planar graph G with one end can be embedded in the Euclidean plane $\mathbb{R}^2$ or hyperbolic plane $\mathbb{H}^2$ such that every automorphism of G corresponds to an isometry of $\mathbb{R}^2$ or $\mathbb{H}^2$, respectively.*

We remark that in Theorem A.17, we do not know if the embedding can be chosen in such a way that, *moreover*, edges are represented by straight line segments in $\mathbb{R}^2$ or $\mathbb{H}^2$. Note also that Propositions A.14 and A.15 talk about embeddings such that edges are straight line segments in the Euclidean geometry. We are not aware of results about straight-line embeddings in the hyperbolic geometry.

A.5. Some examples. Finally, let us describe a method for creating examples of graphs satisfying the assumptions of Theorem 1.1 — i.e., locally finite, 3-connected, quasi-transitive triangulations with one end — and their duals. All our examples come naturally with a periodic straight-line embedding in the Euclidean or hyperbolic plane.

Let $p, q \geq 3$ be integers and let ABC be a triangle whose angles (in anticlockwise order) at the corners A, B, C are π/p, π/q, and $\pi/2$, respectively. Such a triangle can be constructed in either the sphere, the Euclidean plane, or the hyperbolic plane, depending on whether $1/p + 1/q + 1/2$ is larger than, equal to, or less than 1, respectively. By reflecting the triangle ABC in one of its edges and continuing this process, we can cover the whole space alternately by copies of ABC and its mirror image [12, Section 2]. This yields a planar graph with vertices of types A, B and C that are of degree $2p$, $2q$ and 4, respectively. In particular, each vertex of type C is adjacent to two vertices of types A and B each, in alternating order. We may view the A and B sublattices as planar graphs in their own right by erasing the vertices of type C and viewing the four edges emanating from C as two straight edges crossing each other in C, where one connects two A's and the other connects two B's. This yields two regular tesselations that are geometric duals of each other. In the tesselation formed by the A vertices, each vertex has degree p and each face is a regular polygon with q edges. This regular tesselation is denoted by the *Schläfli symbol* $\{q, p\}$ [12]. Likewise, the dual B lattice has the Schläfli symbol $\{p, q\}$.

In particular, the tesselations with Schläfli symbol $\{3, p\}$ (with $p \geq 3$) are regular triangulations of the sphere, the Euclidean plane, or the hyperbolic plane, depending on whether p is less than, equal to, or larger than 6, respectively. It is easy to see that $\{3, p\}$, as a graph, is 3-connected and vertex-transitive. It is finite for $p \leq 5$ and infinite for $p \geq 6$. In particular, Theorem 1.1 applies when $G_0 = \{3, p\}$ with $p \geq 6$. The case $p = 6$, which is the only Euclidean tesselation in this class, yields G_0 = triangular lattice, G_1 = hexagonal lattice, and G = diced lattice. The cases $p > 6$ yield hyperbolic tesselations. The graphs $\{3, 6\}$ and $\{3, 7\}$ and their duals are drawn in Figure 1.2(b,d).

The (dual) tesselations with Schläfli symbol $\{p, 3\}$ are planar Cayley graphs in which every vertex has degree three. A full classification of graphs with these properties can be found in [28]. In particular, [28, Table 1, items 12–19] lists those that are 3-connected and have at most one end. Note that all these graphs are vertex-transitive.

More general examples of quasi-transitive triangulations satisfying the assumptions of Theorem 1.1 can be contructed by starting with any regular tesselation and dividing the basic polygon into triangles in some suitable way, so that the resulting graph is 3-connected. It would go to far to attempt here a full classification of the class of tesselations covered by Theorem 1.1.

Acknowledgement. We wish to thank Youjin Deng, Jesper Jacobsen, and Jesús Salas for many helpful conversations over the course of this work. We also thank Agelos Georgakopoulos, Sebastian Müller, Bruce Richter, Jan Seidler, and Carsten Thomassen for helpful discussions and correspondence concerning infinite graphs and planar topology; Neal Madras and Gordon Slade for correspondence concerning self-avoiding walks; and Youjin Deng, Kun Chen and Yuan Huang for sharing with us their preliminary Monte Carlo data.

The research of R.K. and J.M.S. was supported in part by the grants GAČR 201-09-1931 and P201/12/2613. The research of A.D.S. was supported in part by U.S. National Science Foundation grant PHY–0424082.

References

1. S.E. Alm and R. Parviainen, Bounds for the connective constant of the hexagonal lattice, J. Phys. A: Math. Gen. **37**, 549–560 (2004).

2. J. Ashley, B. Grünbaum, G.C. Shephard and W. Stromquist, Self-duality groups and ranks of self-dualities, in *Applied Geometry and Discrete Mathematics*, DIMACS Series in Discrete Mathematics and Theoretical Computer Science, vol. 4 (American Mathematical Society, Providence RI, 1991), pp. 11–50.
3. L. Babai, The growth rate of vertex-transitive planar graphs, in *Proceedings of the Eighth Annual ACM-SIAM Symposium on Discrete Algorithms* (Association for Computing Machinery, New York, 1997), pp. 564–573.
4. J.R. Banavar, G.S. Grest and D. Jasnow, Antiferromagnetic Potts and Ashkin–Teller models in three dimensions, Phys. Rev. B **25**, 4639–4650 (1982).
5. R.J. Baxter, *Exactly Solved Models in Statistical Mechanics* (Academic Press, London, 1982).
6. C.P. Bonnington, W. Imrich and M.E. Watkins, Separating double rays in locally finite planar graphs, Discrete Math. **145**, 61–72 (1995).
7. H.W.J. Blöte and H.J. Hilhorst, Roughening transitions and the zero-temperature triangular Ising antiferromagnet, J. Phys. A: Math. Gen. **15**, L631–L637 (1982).
8. G.R. Brightwell and P. Winkler, Random colorings of a Cayley tree, in *Contemporary Combinatorics*, edited by B. Bollobás (Bolyai Mathematical Society, Budapest and Springer-Verlag, Berlin, 2002), pp. 247–276.
9. H. Bruhn and R. Diestel, Duality in infinite graphs, J. Combin. Theory. **B 96**, 225–239 (2006).
10. H. Bruhn and M. Stein, Duality of ends, Combin. Probab. Comput. **19**, 47–60 (2010).
11. J.K. Burton Jr. and C.L. Henley, A constrained Potts antiferromagnet model with an interface representation, J. Phys. A: Math. Gen. **30**, 8385–8413 (1997), cond-mat/9708171.
12. H.S.M. Coxeter, Regular compound tessellations of the hyperbolic plane, Proc. Roy. Soc. London A **278**, 147–167 (1964).
13. Y. Deng, K. Chen and Y. Huang, private communication (April 2012).
14. Y. Deng, Y. Huang, J.L. Jacobsen, J. Salas and A.D. Sokal, Finite-temperature phase transition in a class of four-state Potts antiferromagnets, Phys. Rev. Lett. **107**, 150601 (2011), arXiv:1108.1743.
15. R.F. Dickman and R.A. McCoy, The Freudenthal compactification, Dissertationes Math. (Rozprawy Mat.) **262** (1988), 35 pp.
16. R. Diestel, *Graph Theory*, 4th ed. (Springer-Verlag, New York, 2010). Also available on-line at `http://diestel-graph-theory.com/`
17. R. Diestel, Locally finite graphs with ends: a topological approach, arXiv:0912.4213v3. An earlier version of this survey was published in Discrete Math. **311**, 1423–1447 (2011) and **310**, 2750–2765 (2010).
18. R. Diestel and D. Kühn, Graph-theoretical versus topological ends of graphs, J. Combin. Theory B **87**, 197–206 (2003).
19. R. Diestel and D. Kühn, On Infinite Cycles I, Combinatorica **24**, 69–89 (2004).
20. R. Diestel and D. Kühn, On Infinite Cycles II, Combinatorica **24**, 91–116 (2004).
21. H. Duminil-Copin and S. Smirnov, The connective constant of the honeycomb lattice equals $\sqrt{2+\sqrt{2}}$, Ann. Math. **175**, 1653–1665 (2012), arXiv:1007.0575v2.
22. R. Durrett, *Lecture Notes on Particle Systems and Percolation* (Wadsworth, Pacific Grove CA, 1988).
23. R.G. Edwards and A.D. Sokal, Generalization of the Fortuin–Kasteleyn–Swendsen–Wang representation and Monte Carlo algorithm, Phys. Rev. D **38**, 2009–2012 (1988).
24. S.J. Ferreira and A.D. Sokal, Antiferromagnetic Potts models on the square lattice: A high-precision Monte Carlo study, J. Statist. Phys. **96**, 461–530 (1999).
25. H. Freudenthal, Über die Enden topologischer Räume und Gruppen, Math. Z. **33**, 692–713 (1931).
26. H. Freudenthal, Über die Enden diskreter Räume und Gruppen, Comment. Math. Helv. **17**, 1–38 (1945).
27. A. Georgakopoulos, Graph topologies induced by edge lengths, Discrete Math. **311**, 1523–1542 (2011), arXiv:0903.1744.
28. A. Georgakopoulos, The planar cubic Cayley graphs, preprint (February 2011), arXiv:1102.2087v2.
29. H.-O. Georgii, *Gibbs Measures and Phase Transitions* (de Gruyter, Berlin–New York, 1988).
30. A.P. Gottlob and M. Hasenbusch, Three-state anti-ferromagnetic Potts model in three dimensions: Universality and critical amplitudes, Physica A **210**, 217–236 (1994).
31. A.P. Gottlob and M. Hasenbusch, The *XY* model and the three-state antiferromagnetic Potts model in three dimensions: Critical properties from fluctuating boundary conditions, J. Statist. Phys. **77**, 919–930 (1994).
32. G.R. Grimmett, *The Random-Cluster Model* (Springer-Verlag, Berlin, 2006).
33. R. Halin, Über unendliche Wege in Graphen, Math. Ann. **157**, 125–137 (1964).
34. A. Hatcher, *Algebraic Topology* (Cambridge University Press, Cambridge, 2002). Also available on-line at `http://www.math.cornell.edu/~hatcher/AT/ATpage.html`
35. D.A. Huse and A.D. Rutenberg, Classical antiferromagnets on the Kagomé lattice, Phys. Rev. B **45**, 7536–7539 (1992).
36. M. Itakura, Monte Carlo simulation of the antiferromagnetic four-state Potts model on simple cubic and body-centered-cubic lattices, Phys. Rev. B **60**, 6558–6565 (1999).
37. J.L. Jacobsen, Conformal field theory applied to loop models, in *Polygons, Polyominoes and Polycubes*, edited by A.J. Guttmann, Lecture Notes in Physics #775 (Springer, Dordrecht, 2009), Chapter 14, pp. 347–424.
38. J.L. Jacobsen and A.D. Sokal, Mapping of a graph-homomorphism (RSOS) model onto a multivariate Tutte polynomial (Potts model), in preparation.
39. I. Jensen, A parallel algorithm for the enumeration of self-avoiding polygons on the square lattice, J. Phys. A: Math. Gen. **36**, 5731–5745 (2003), cond-mat/0301468.
40. I. Jensen, Honeycomb lattice polygons and walks as a test of series analysis techniques, J. Phys.: Conf. Ser. **42**, 163–178 (2006); and `http://www.ms.unimelb.edu.au/~iwan/polygons/Polygons_ser.html`

41. J. Jonasson, Uniqueness of uniform random colorings of regular trees, Statist. Probab. Lett. **57**, 243–248 (2002).
42. J. Kolafa, Monte Carlo study of the three-state square Potts antiferromagnet, J. Phys. A: Math. Gen. **17**, L777–L781 (1984).
43. J. Kondev and C.L. Henley, Four-coloring model on the square lattice: A critical ground state, Phys. Rev. B **52**, 6628–6639 (1995).
44. J. Kondev and C.L. Henley, Kac-Moody symmetries of critical ground states, Nucl. Phys. B **464**, 540–575 (1996).
45. R. Kotecký, Long-range order for antiferromagnetic Potts models, Phys. Rev. B **31**, 3088–3092 (1985).
46. R. Kotecký, J. Salas and A.D. Sokal, Phase transition in the three-state Potts antiferromagnet on the diced lattice, Phys. Rev. Lett. **101**, 030601 (2008), arXiv:0802.2270.
47. B. Krön, Introduction to ends of graphs, preprint (July 2005), `http://homepage.univie.ac.at/bernhard.kroen/ ends_22_07_2005.pdf`
48. B. Krön, Infinite faces and ends of almost transitive plane graphs, preprint (December 2006), Hamburger Beiträge zur Mathematik, Heft 257, `http://preprint.math.uni-hamburg.de/public/hbm.html`
49. B. Krön and E. Teufl, Ends — group theoretical and topological aspects, preprint (April 2009), `http://homepage.univie.ac.at/bernhard.kroen/kroen_teufl_ends.pdf`
50. T.M. Liggett, *Interacting Particle Systems* (Springer-Verlag, New York, 1985).
51. T.M. Liggett, *Stochastic Interacting Systems: Contact, Voter and Exclusion Processes* (Springer-Verlag, Berlin, 1999).
52. N. Madras and G. Slade, *The Self-Avoiding Walk* (Birkhäuser Boston, Boston MA, 1993).
53. N. Madras and C.C. Wu, Self-avoiding walks on hyperbolic graphs, Combin. Probab. Comput. **14**, 523–548 (2005).
54. C. Moore and M.E.J. Newman, Height representation, critical exponents, and ergodicity in the four-state triangular Potts antiferromagnet, J. Statist. Phys. **99**, 629–660 (2000), cond-mat/9902295.
55. B. Nienhuis, Exact critical point and critical exponents of $O(n)$ models in two dimensions, Phys. Rev. Lett. **49**, 1062–1065 (1982).
56. B. Nienhuis, H.J. Hilhorst and H.W.J. Blöte, Triangular SOS models and cubic-crystal shapes, J. Phys. A: Math. Gen. **17**, 3559–3581 (1984).
57. M.P.M. den Nijs, M.P. Nightingale, and M. Schick, Critical fan in the antiferromagnetic three-state Potts model, Phys. Rev. B **26**, 2490–2500 (1982).
58. S. Northshield, Geodesics and bounded harmonic functions on infinite planar graphs, Proc. Amer. Math. Soc. **113**, 229–233 (1991).
59. M. O'Keeffe, Self-dual plane nets in crystal chemistry, Austral. J. Chem. **45**, 1489–1498 (1992).
60. M. O'Keeffe and B.G. Hyde, *Crystal Structures I. Patterns and Symmetry* (Mineralogical Society of America, Washington DC, 1996), Section 5.3.7. Available on-line at `http://www.public.asu.edu/~rosebudx/okeeffe.htm`
61. R. Peled, High-dimensional Lipschitz functions are typically flat, preprint (May 2010), arXiv:1005.4636.
62. A. Pönitz and P. Tittman, Improved upper bounds for self-avoiding walks in $\mathbb{Z}^d$, Electron. J. Combin. **7**, Research Paper 21 (2000).
63. R.B. Richter and C. Thomassen, 3-connected planar spaces uniquely embed in the sphere, Trans. Am. Math. Soc. **354**, 4585–4595 (2002).
64. J. Salas and A.D. Sokal, The three-state square-lattice Potts antiferromagnet at zero temperature, J. Statist. Phys. **92**, 729–753 (1998), cond-mat/9801079.
65. C.R. Scullard, Exact site percolation thresholds using a site-to-bond transformation and the star-triangle transformation, Phys. Rev. E **73**, 016107 (2006).
66. B. Servatius and H. Servatius, Symmetry, automorphisms, and self-duality of infinite planar graphs and tilings, in *Proceedings of the International Scientific Conference on Mathematics* (Žilina, 30 June – 3 July 1998), edited by V. Bálint (University of Žilina, Žilina, 1998), pp. 83–116. Available on-line at `http://users.wpi.edu/~bservat/self5.html`
67. C. Thomassen, Planarity and duality of finite and infinite graphs, J. Combin. Theory B **29**, 244–271 (1980).
68. C. Thomassen, Duality of infinite graphs, J. Combin. Theory B **33**, 137–160 (1982).
69. J.-S. Wang, R.H. Swendsen and R. Kotecký, Antiferromagnetic Potts models, Phys. Rev. Lett. **63**, 109–112 (1989).
70. J.-S. Wang, R.H. Swendsen and R. Kotecký, Three-state antiferromagnetic Potts models: A Monte Carlo study, Phys. Rev. B **42**, 2465–2474 (1990).
71. J.C. Wierman, Construction of infinite self-dual graphs, in *Proceedings of the 5th Hawaii International Conference on Statistics, Mathematics and Related Fields* (2006).
72. C. Zeng and C.L. Henley, Zero-temperature phase transitions of an antiferromagnetic Ising model of general spin on a triangular lattice, Phys. Rev. B **55**, 14935–14947 (1997), cond-mat/9609007.
73. R.M. Ziff, C.R. Scullard, J.C. Wierman and M.R.A. Sedlock, The critical manifolds of inhomogeneous bond percolation on bow-tie and checkerboard lattices, J. Phys. A: Math. Theor. **45**, 494005 (2012), arXiv:1210.6609.