ON THE WIDOM-ROWLINSON OCCUPANCY FRACTION IN REGULAR GRAPHS

EMMA COHEN, WILL PERKINS, PRASAD TETALI

Abstract. We consider the Widom-Rowlinson model of two types of interacting particles on d-regular graphs. We prove a tight upper bound on the occupancy fraction: the expected fraction of vertices occupied by a particle under a random configuration from the model. The upper bound is achieved uniquely by unions of complete graphs on $d+1$ vertices, K_{d+1}'s. As a corollary we find that K_{d+1} also maximizes the normalized partition function of the Widom-Rowlinson model over the class of d-regular graphs. A special case of this shows that the normalized number of homomorphisms from any d-regular graph G to the graph H_{WR}, a path on three vertices with a self-loop on each vertex, is maximized by K_{d+1}. This proves a conjecture of Galvin.

1. The Widom-Rowlinson Model

A Widom-Rowlinson assignment or configuration on a graph G is a map $\chi : V(G) \to \{0, 1, 2\}$ so that 1 and 2 are not assigned to neighboring vertices, or in other words, a graph homomorphism from G to the graph H_{WR} consisting of a path on 3 vertices with a self-loop on each vertex (the middle vertex represents the label 0). Call the set of all such assignments $\Omega(G)$. The Widom-Rowlinson model on G is a probability distribution over $\Omega(G)$ parameterized by $\lambda \in (0, \infty)$, given by:

$$\Pr[\chi] = \frac{\lambda^{X_1(\chi)+X_2(\chi)}}{P_G(\lambda)},$$

where $X_i(\chi)$ are the number of vertices colored i under χ, and

$$P_G(\lambda) = \sum_{\chi \in \Omega(G)} \lambda^{X_1(\chi)+X_2(\chi)}$$

is the *partition function*. Evaluating $P_G(\lambda)$ at $\lambda = 1$ counts the number of homomorphisms from G to H_{WR}. We think of vertices assigned 1 and 2 as "colored" and those assigned 0 as "uncolored" (see Figure 1).

The Widom-Rowlinson model was introduced by Widom and Rowlinson in 1970 [13], as a model of two types of interacting particles with a hard-core exclusion between particles

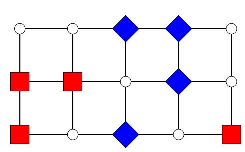

Figure 1. A configuration for the Widom-Rowlinson model on a grid. Vertices mapping to 1 and 2 are shown as squares and diamonds, respectively (corresponding to Figure 2).

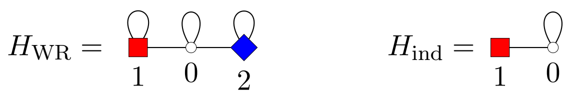

FIGURE 2. The target graphs for the Widom-Rowlinson model and the hard-core model.

of different types: color 1 and 2 represent particles of each type and color 0 represents an unoccupied site. The model has been studied both on lattices [9] and in the continuum [11, 2] and is known to exhibit a phase transition in both cases.

The Widom-Rowlinson model is one case of a general random model: that of choosing a random homomorphism from a large graph G to a fixed graph H. In the Widom-Rowlinson case, we take $H = H_{\mathrm{WR}}$. Another notable case is H_{ind}, an edge between two vertices, one of which has a self loop (see Figure 2). Homomorphisms from G to H_{ind} are exactly the independent sets of G, and the partition function of the hard-core model is the sum of $\lambda^{|I|}$ over all independent sets I. An overview of the connections between statistical physics models with hard constraints, graph homomorphisms, and combinatorics can be found in [1].

For every such model, there is an associated extremal problem. Denote by $\hom(G, H)$ the number of homomorphisms from G to H. Then we can ask which graph G from a class of graphs $\mathcal{G}$ maximizes $\hom(G, H)$, or if we wish to compare graphs on different numbers of vertices, ask which graph maximizes the scaled quantity $\hom(G, H)^{1/|V(G)|}$.

Kahn [8] proved that for any d-regular, bipartite graph G,

$$\hom(G, H_{\mathrm{ind}}) \leq \hom(K_{d,d}, H_{\mathrm{ind}})^{|V(G)|/2d}, \tag{1}$$

where $K_{d,d}$ is the complete d-regular bipartite graph. Equality holds in (1) if G is $K_{d,d}$ or a union of $K_{d,d}$'s. In other words, unions of $K_{d,d}$'s maximize the total number of independent sets over all d-regular, bipartite graphs on a fixed number of vertices.

In a broad generalization of Kahn's result, Galvin and Tetali [7] showed that in fact, (1) holds for all d-regular, bipartite G and *all* target graphs H (including, for example, H_{WR}). And using a cloning construction and a limiting argument, they showed that in fact the partition function of such models (a weighted count of homomorphisms) is maximized by $K_{d,d}$; for example, for a d-regular, bipartite G,

$$P_G(\lambda) \leq P_{K_{d,d}}(\lambda)^{|V(G)|/2d}, \tag{2}$$

where $P_G(\lambda)$ is the Widom-Rowlinson partition function defined above or the independence polynomial of a graph. Note that the case $\lambda = 1$ is the counting result.

There is no such sweeping statement for the class of all d-regular graphs with the bipartiteness restriction removed. In [14] and [15], Zhao showed that the bipartiteness restriction on G in (1) and (2) can be removed for some class of graphs H, including H_{ind}. But such an extension is not possible for all graphs H; for example, K_{d+1} has more homomorphisms to H_{WR} than does $K_{d,d}$ (after normalizing for the different numbers of vertices). In fact Galvin conjectured the following:

Conjecture 1 (Galvin [5, 6])**.** *Let G be a any d-regular graph. Then*

$$\hom(G, H_{WR}) \leq \hom(K_{d+1}, H_{WR})^{|V(G)|/(d+1)}.$$

The more general Conjecture 1.1 of [5] that the maximizing G for any H is either $K_{d,d}$ or K_{d+1} has been disproved by Sernau [12].

The above theorems of Kahn and Galvin and Tetali are based on the *entropy method* (see [10] and [6] for a survey), but in this context bipartiteness seems essential for the effectiveness of the method. We will approach the problem differently, using the *occupancy method* of [3].

We first define the *occupancy fraction* $\alpha_G(\lambda)$ to be the expected fraction of vertices which receive a (nonzero) color in the Widom-Rowlinson model:

$$\alpha_G(\lambda) = \frac{\mathbb{E}[X_1 + X_2]}{|V(G)|},$$

where X_i is the number of vertices colored i by the random assignment χ. A calculation shows that $\alpha_G(\lambda)$ is in fact the scaled logarithmic derivative of the partition function:

$$\alpha_G(\lambda) = \frac{\lambda}{|V(G)|} \cdot \frac{P'_G(\lambda)}{P_G(\lambda)} = \frac{\lambda \cdot (\log P_G(\lambda))'}{|V(G)|}. \tag{3}$$

Our main result is that for any λ, $\alpha_G(\lambda)$ is maximized over all d-regular graphs G by K_{d+1}.

Theorem 1. *Let G be any d-regular graph and $\lambda > 0$. Then*

$$\alpha_G(\lambda) \le \alpha_{K_{d+1}}(\lambda)$$

with equality if and only if G is a union of K_{d+1}'s.

We will prove this by introducing local constraints on random configurations induced by the Widom-Rowlinson model on a d-regular graph G, then solving a linear programming relaxation of the optimization problem over all d-regular graphs.

Theorem 1 implies maximality of the normalized partition function:

Corollary 1. *Let G be a d-regular graph and $\lambda > 0$. Then*

$$\frac{1}{|V(G)|} \log P_G(\lambda) \le \frac{1}{d+1} \log P_{K_{d+1}}(\lambda),$$

or equivalently,

$$P_G(\lambda) \le P_{K_{d+1}}(\lambda)^{|V(G)|/(d+1)},$$

with equality if and only if G is a union of K_{d+1}'s.

The quantity $\frac{1}{|V(G)|} \log P_G(\lambda)$ is known in statistical physics as the *free energy per unit volume*. Corollary 1 follows from Theorem 1 as follows: $\frac{1}{|V(G)|} \log P_G(0) = 0$ for any G, and so

$$\begin{aligned}\frac{1}{|V(G)|} \log P_G(\lambda) &= \frac{1}{|V(G)|} \int_0^\lambda (\log P_G(t))' \, dt \\ &\le \frac{1}{d+1} \int_0^\lambda \left(\log P_{K_{d+1}}(t)\right)' \, dt = \frac{1}{d+1} \log P_{K_{d+1}}(\lambda)\end{aligned}$$

where the inequality follows from Theorem 1 and (3). Exponentiating both sides gives Corollary 1.

By taking $\lambda = 1$ in Corollary 1, we get the counting result:

Corollary 2. *For all d-regular G,*

$$\hom(G, H_{WR}) \leq \hom(K_{d+1}, H_{WR})^{|V(G)|/(d+1)}$$

with equality if and only if G is a union of K_{d+1}'s.

This proves Conjecture 1.

Discussion and related work. The method we use is more probabilistic than the entropy method in the sense that Theorem 1 gives information about an observable of the model; in some statistical physics models, the analogue of $\alpha_G(\lambda)$ would be called the *mean magnetization*. We also work directly in the statistical physics model, instead of counting homomorphisms.

Davies, Jenssen, Perkins, and Roberts [3] applied the occupancy method to two central models in statistical physics: the hard-core model of a random independent set described above, and the monomer-dimer model of a randomly chosen matching from a graph G. In both cases they showed that $K_{d,d}$ maximizes the occupancy fraction over all d-regular graphs. In the case of independent sets this gives a strengthening of the results of Kahn, Galvin and Tetali, and Zhao, while for matchings, it was not known previously that unions of $K_{d,d}$ maximizes the partition function or the total number of matchings.

The idea of calculating the log partition function by integrating a partial derivative is not new of course; see for example, the interpolation scheme of Dembo, Montanari, and Sun [4] in the context of Gibbs distributions on locally tree-like graphs. The method is powerful because it reduces the computation of a very global quantity, $P_G(\lambda)$, to that of a locally estimable quantity, $\alpha_G(\lambda)$.

Some partial results towards the Widom-Rowlinson counting problem were obtained by Galvin [5], who showed that a graph with more homomorphisms than a union of K_{d+1}'s must be close in a specific sense to a union of K_{d+1}'s.

2. Proof of Theorem 1

2.1. **Preliminaries.** To prove Theorem 1, we will use the following experiment: for a d-regular graph G, we first draw a random χ from the Widom-Rowlinson model, then select a vertex v uniformly at random from $V(G)$. We then write our objective function, the occupancy fraction, in terms of local probabilities with respect to this experiment, and add constraints on the local probabilities that must hold for all G. We then relax the optimization problem to all distributions satisfying the local constraints, and optimize using linear programming.

Fix d and λ. Define a *configuration with boundary conditions* $C = (H, \mathcal{L})$ to be a graph H on d vertices with family of lists $\mathrm{Ł} = \{L_u\}_{u \in H}$, where each $L_u \subseteq \{1, 2\}$ is a set of allowed colors for the vertex u. Here H represents the neighborhood structure of a vertex $v \in V(G)$ and the color lists L_u represent the colors permitted to neighbors of v, given an assignment χ on the vertices outside of $N(v) \cup \{v\}$. (See Figure 3.) Denote by $\mathcal{C}$ the set of all possible configurations with boundary conditions in any d-regular graph.

We now pick the assignment χ at random from the Widom-Rowlinson model on a fixed d-regular graph G, pick a vertex v uniformly at random from $V(G)$, and consider the probability distribution induced on $\mathcal{C}$.

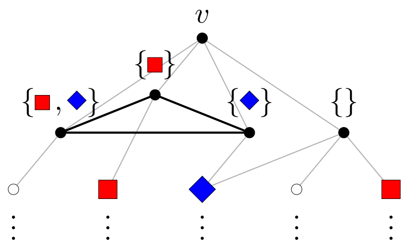

FIGURE 3. An example configuration with boundary conditions based on a coloring χ. The graph H consists of the four neighbors of v along with the black edges, and the list L_u is shown above each vertex u of H. The colors assigned by χ to v and its neighbors are immaterial and so are not shown.

For example, if $G = K_{d+1}$ then with probability 1 the random configuration C is $H = K_d$ with $L_u = \{1, 2\}$ for all $u \in V(H)$. If $G = K_{d,d}$ then H is always d isolated vertices and the color lists can be any (possibly empty) subset of $\{1, 2\}$, but the lists must be the same for all $u \in V(H)$.

For a configuration $C = (H, Ł)$, define

$$\alpha_i^v(C) = \Pr[\chi(v) = i \mid C]$$
$$\alpha_i^u(C) = \frac{1}{d} \sum_{u \in V(H)} \Pr[\chi(u) = i \mid C],$$

where the probability is over the Widom-Rowlinson model on G given the boundary conditions $\mathcal{L}$. Note that the spatial Markov property of the model means that these probabilities are "local" in the sense that they can be computed knowing only C. Let $\alpha^v(C) = \alpha_1^v(C) + \alpha_2^v(C)$ and $\alpha^u(C) = \alpha_1^u(C) + \alpha_2^u(C)$. Then we have

$$\alpha_G(\lambda) = \frac{1}{|V(G)|} \sum_{v \in V(G)} \Pr[\chi(v) \in \{1, 2\}] = \mathbb{E}_C[\alpha^v(C)] \tag{4}$$
$$= \frac{1}{d} \frac{1}{|V(G)|} \sum_{v \in V(G)} \sum_{u \sim v} \Pr[\chi(u) \in \{1, 2\}] = \mathbb{E}_C[\alpha^u(C)]\,,$$

where the expectations are over the probability distribution induced on $\mathcal{C}$ by our experiment of drawing χ from the model and v uniformly at random from $V(G)$, and the last sum is over all neighbors of v in G. Equality of the two expressions for α follows since sampling a uniform neighbor or a uniform vertex in a regular graph is equivalent to sampling a uniform vertex. We will show that this expectation is maximized when the graph G is K_{d+1}.

We can in fact write explicit formulae for $\alpha^v(C)$ and $\alpha^u(C)$. For a configuration $C = (H, Ł)$, let $P_C^{(0)}(\lambda)$ be the total weight of colorings of H satisfying the boundary conditions given by the lists Ł (corresponding to the partition function for the neighborhood of v conditioned on $\chi(v) = 0$). Also, write $P_C^{(i)}(\lambda)$ for the total weight of colorings of H satisfying the boundary conditions and using only color i and 0 (corresponding to the partition functions for the neighborhood of v conditioned on $\chi(v) = i$). Finally, let $P_C^{(12)}(\lambda) = P_C^{(1)}(\lambda) + P_C^{(2)}(\lambda)$ and let

$$P_C(\lambda) = P_C^{(0)}(\lambda) + \lambda P_C^{(12)}(\lambda)$$

be the partition function of $N(v) \cup \{v\}$ conditioned on the boundary conditions given by C. Note that if Ł has a_1 lists containing 1 and a_2 lists containing 2, then $P_C^{(i)}(\lambda) = (1+\lambda)^{a_i}$.

Now we can write

$$\alpha^v(C) = \frac{\lambda P_C^{(12)}}{P_C} \qquad \text{and} \qquad \alpha^u(C) = \frac{\lambda\Big((P_C^{(0)})' + \lambda (P_C^{(12)})'\Big)}{d\, P_C}, \tag{5}$$

where P' is the derivative of P in λ. We will suppress the dependence of the partition functions on λ from now on.

For $G = K_{d+1}$, we have

$$P_{K_{d+1}} = 2(1+\lambda)^{d+1} - 1$$
$$\alpha_{K_{d+1}}(\lambda) = \frac{2\lambda(1+\lambda)^d}{2(1+\lambda)^{d+1} - 1}.$$

If $G = K_{d+1}$ then the only possible configuration is $C_{K_{d+1}}$, the complete neighborhood K_d with full boundary lists, so we also have $\alpha^u(K_d) = \alpha^v(K_d) = \alpha_{K_{d+1}}(\lambda)$ (we can also compute these directly). Since this quantity will arise frequently, we will use the notation $\alpha_K = \alpha_{K_{d+1}}(\lambda)$.

2.2. **A linear programming relaxation.** Now let $q : \mathcal{C} \to [0,1]$ denote a probability distribution over the set of all possible configurations. Then we set up the following optimization problem over the variables $q(C)$, $C \in \mathcal{C}$.

$$\begin{aligned} \alpha^* = \max \sum_{C \in \mathcal{C}} q(C) \alpha^v(C) \qquad & \text{s.t.} \\ \sum_{C \in \mathcal{C}} q(C) = 1 & \\ \sum_{C \in \mathcal{C}} q(C)[\alpha^v(C) - \alpha^u(C)] = 0 & \\ q(C) \geq 0 \quad \forall C \in \mathcal{C}. & \end{aligned} \tag{6}$$

Note that this linear program is indeed a relaxation of our optimization problem of maximizing $\alpha_G(\lambda)$ over all d-regular graphs: any such graph induces a probability distribution on $\mathcal{C}$, and as we have seen above in (4), the constraint asserting the equality $\mathbb{E}\alpha^v(C) = \mathbb{E}\alpha^u(C)$ must hold in all d-regular graphs.

We will show that for any $\lambda > 0$ the unique optimal solution of this linear program is $q(C_{K_{d+1}}) = 1$, where $C_{K_{d+1}}$ is the configuration induced by K_{d+1}: $H = K_d$ and $L_u = \{1,2\}$ for all $u \in H$.

The dual of the above linear program is

$$\begin{aligned} \alpha^* = \min \Lambda_p \qquad & \text{s.t.} \\ \Lambda_p + \Lambda_c(\alpha^v(C) - \alpha^u(C)) \geq \alpha^v(C) \qquad & \forall C \in \mathcal{C}, \end{aligned}$$

with decision variables Λ_p and Λ_c.

To show that the optimum is attained by $C_{K_{d+1}}$, we must find a feasible solution to the dual program with $\Lambda_p = \alpha_K = \frac{2\lambda(1+\lambda)^d}{2(1+\lambda)^{d+1}-1}$. Note that with $\Lambda_p = \alpha_K$ the constraint for $C_{K_{d+1}}$ holds with equality for any choice of Λ_c. In other words, it suffices to find some convex combination of the two local estimates α^u and α^v which is maximized by $C_{K_{d+1}}$ over all $C \in \mathcal{C}$.

Let C_0 be a configuration with $L_u = \emptyset$ for all $u \in H$ (in which case the edges of H are immaterial, and so abusing notation we will refer to any one of these configurations as C_0). We find a candidate Λ_c by solving the constraint corresponding to C_0 with equality:

$$\begin{aligned}\alpha_K &= \Lambda_c(\alpha^u(C_0) - \alpha^v(C_0)) + \alpha^v(C_0) \\ &= (1-\Lambda_c)\frac{2\lambda}{1+2\lambda}\,.\end{aligned}$$

This gives

$$\Lambda_c = 1 - \frac{\alpha_K}{2\lambda}(1+2\lambda) = \frac{\alpha_K}{2\lambda}\frac{(1+\lambda)^d - 1}{(1+\lambda)^d}\,.$$

With this choice of Λ_c, the general dual constraint is

$$\alpha_K \geq \frac{\alpha_K}{2\lambda}\frac{(1+\lambda)^d-1}{(1+\lambda)^d}\alpha^u(C) + \frac{\alpha_K}{2\lambda}(1+2\lambda)\alpha^v(C)\,.$$

Using (5), this becomes

$$\frac{(P_C^{(0)})' + \lambda(P_C^{(12)})'}{2P_C^{(0)} - P_C^{(12)}} \leq \frac{d(1+\lambda)^d}{(1+\lambda)^d - 1}\,. \tag{7}$$

From this point on we may assume that C has some non-empty color list, since otherwise the configuration is equivalent to C_0 and the constraint holds with equality by our choice of Λ_c. This assumption tells us, among other things, that $(P_C^{(0)})' > 0$ and $2P_C^{(0)} - P_C^{(12)} > 0$.

Our goal is now to show that (7) holds for all C. We consider the two terms separately.

Claim 1. *For any $C \neq C_0$,*

$$\frac{\lambda(P_C^{(12)})'}{2P_C^{(0)} - P_C^{(12)}} \leq \frac{d\lambda(1+\lambda)^{d-1}}{(1+\lambda)^d - 1}\,,$$

with equality if and only if the lists L_u are all equal and C has no dichromatic colorings.

Proof. Since the partition function $P_C^{(0)}$ is at least the total weight $P_C^{(1)} + P_C^{(2)} - 1$ of monochromatic colorings (with equality when C has no dichromatic colorings), we have

$$\frac{(P_C^{(12)})'}{2P_C^{(0)} - P_C^{(12)}} \leq \frac{(P_C^{(12)})'}{P_C^{(12)} - 2} = \frac{a_1(1+\lambda)^{a_1-1} + a_2(1+\lambda)^{a_2-1}}{(1+\lambda)^{a_1} + (1+\lambda)^{a_2} - 2} \tag{8}$$

(where, as above, a_i is the number of vertices in H allowed color i under the given boundary conditions), and so we need to show that

$$\frac{a_1(1+\lambda)^{a_1-1} + a_2(1+\lambda)^{a_2-1}}{(1+\lambda)^{a_1} + (1+\lambda)^{a_2} - 2} \leq \frac{d(1+\lambda)^{d-1}}{(1+\lambda)^d - 1}\,.$$

In general, to show that $(a+b)/(c+d) \leq t$ it suffices to show that $a/c \leq t$ and $b/d \leq t$. Thus it is enough to show that

$$\frac{a(1+\lambda)^{a-1}}{(1+\lambda)^a - 1} \leq \frac{d(1+\lambda)^{d-1}}{(1+\lambda)^d - 1} \tag{9}$$

whenever $1 \leq a \leq d$. (Note that if either $a_1 = 0$ or $a_2 = 0$ then (8) reduces to (9), and if both $a_1, a_2 = 0$ then the configuration is C_0). Indeed, it is not hard to check via calculus that the left hand side of (9) is increasing with a. This completes the proof of the inequality in Claim 1.

We have equality in this final step when $a_1 = a_2 = d$ or when one is 0 and the other is d. So we have equality overall whenever the lists are all equal and there are no dichromatic colorings (recall that we are assuming C has some non-empty coloring list). □

Claim 2. *For any* $C \neq C_0$,

$$\frac{(P_C^{(0)})'}{2P_C^{(0)} - P_C^{(12)}} \leq \frac{d(1+\lambda)^{d-1}}{(1+\lambda)^d - 1},$$

with equality if and only if the lists L_u *are all equal and* C *has no dichromatic colorings.*

Proof. We can write

$$\frac{\lambda(P_C^{(0)})'}{2P_C^{(0)} - P_C^{(12)}} = \frac{\lambda(P_C^{(0)})'}{P_C^{(0)}} \cdot \frac{P_C^{(0)}}{(P_C^{(0)} - P_C^{(1)}) + (P_C^{(0)} - P_C^{(2)})}$$
$$= \frac{\mathbb{E}_C[X_1] + \mathbb{E}_C[X_2]}{\Pr_C[X_1 > 0] + \Pr_C[X_2 > 0]},$$

where now X_i is the number of vertices colored i in a random coloring chosen from the Widom-Rowlinson model on C. Noting that $\mathbb{E}_C[X_1] = 0$ whenever $\Pr_C[X_1 > 0] = 0$, it suffices as above to show that whenever color 1 is permitted anywhere in C,

$$\frac{\mathbb{E}_C[X_1]}{\Pr_C[X_1 > 0]} = \mathbb{E}_C[X_1 \mid X_1 > 0] \leq \frac{\lambda d(1+\lambda)^{d-1}}{(1+\lambda)^d - 1} = \mathbb{E}_{K_d}[X_1 \mid X_1 > 0], \tag{10}$$

and similarly for X_2, but this will follow by symmetry.

We can decompose the expectation as

$$\mathbb{E}_C[X_1 \mid X_1 > 0] = \sum_{S \subseteq V(H)} \Pr_C[\chi^{-1}(2) = S \mid X_1 > 0] \cdot \mathbb{E}_C[X_1 \mid X_1 > 0 \wedge \chi^{-1}(2) = S].$$

The partition function restricted to colorings satisfying $X_1 > 0$ and $\chi^{-1}(2) = S$ is just $P_S(\lambda) = \lambda^{|S|}((1+\lambda)^{a_S} - 1)$, where a_S is the number of vertices in $H \setminus S$ which are allowed color 1 and are not adjacent to any vertex of S. The conditional expectation is then

$$\mathbb{E}_C[X_1 \mid X_1 > 0 \wedge \chi^{-1}(2) = S] = \frac{a_S \lambda(1+\lambda)^{a_S - 1}}{(1+\lambda)^{a_S} - 1} \leq \frac{d\lambda(1+\lambda)^{d-1}}{(1+\lambda)^d - 1}$$

with equality precisely when S is empty and 1 is available for every vertex. That is,

$$\mathbb{E}_C[X_1 \mid X_1 > 0] \leq \sum_{S \subseteq V(H)} \Pr_C[\chi^{-1}(2) = S \mid X_1 > 0] \cdot \frac{d\lambda(1+\lambda)^{d-1}}{(1+\lambda)^d - 1} = \frac{\lambda d(1+\lambda)^{d-1}}{(1+\lambda)^d - 1},$$

as desired. We have equality in (10) when $\Pr_C[a_S = d \mid X_1 > 0] = 1$, which holds for the configurations where 1 is available to every vertex but which have no dichromatic colorings. That is, for equality to hold in the claim C must have no dichromatic colorings, and any color which is available to some vertex u must be available to every vertex (so the lists must be identical). □

Adding the inequalities in Claims 2 and 1 shows that (7) holds for all C, proving optimality of K_{d+1}.

2.3. Uniqueness.

Lemma 1. *The distribution induced by K_{d+1} is the unique optimum of the LP relaxation* (6)*.*

Proof. Complementary slackness for our dual solution says that any optimal primal solution is supported only on configurations C with identical boundary lists and no dichromatic colorings. These fall into three categories:

- **Case 0:** $L_u = \emptyset$ for all u. In this case the edges of H are immaterial, as none of H can be colored. This is the configuration C_0 above.
- **Case 1:** $L_u = \{i\}$ for all u (for $i = 1$ or 2). The edges of H are again immaterial, as every coloring of H with only color i is allowed. Call this configuration C_1.
- **Case 2:** $L_u = \{1, 2\}$ for all u. In this case the prohibition on dichromatic colorings requires that $C = C_{K_{d+1}}$.

We can calculate $\alpha^v(C)$ and $\alpha^u(C)$ for each case. For Case 0 we have

$$\alpha^v(C_0) = \frac{2\lambda}{1+2\lambda} \qquad \text{and} \qquad \alpha^u(C_0) = 0.$$

For Case 1 we have

$$\alpha^v(C_1) = \frac{\lambda + \lambda(1+\lambda)^d}{\lambda + (1+\lambda)^{d+1}} \qquad \text{and} \qquad \alpha^u(C_1) = \frac{\lambda(1+\lambda)^d}{\lambda + (1+\lambda)^{d+1}}.$$

And of course, for Case 2 we have

$$\alpha^v(K_d) = \alpha^u(K_d) = \alpha_K.$$

In both Case 0 and Case 1 we have $\alpha^u < \alpha^v$, so the only convex combination q of the three cases giving $\sum_C q(C)\alpha^u(C) = \sum_C q(C)\alpha^v(C)$ (as is required for feasibility) is the one which puts all of the weight on $C_{K_{d+1}}$. □

3. Distinct activities

It is also natural to consider a weighted version of the Widom-Rowlinson model with distinct activities λ_1, λ_2 for the two colors, so that the configuration χ is chosen with probability proportional to $\lambda_1^{X_1}\lambda_2^{X_2}$, and where the partition function $P_G(\lambda_1, \lambda_2)$ is again the normalizing factor. We can ask which graphs maximize $P(\lambda_1, \lambda_2)^{1/|V(G)|}$. We conjecture

Conjecture 2. *For any $\lambda_1, \lambda_2 > 0$, the* weighted occupancy fraction

$$\overline{\alpha}_G(\lambda_1, \lambda_2) = \frac{\lambda_2 \alpha_G^1(\lambda_1, \lambda_2) + \lambda_1 \alpha_G^2(\lambda_1, \lambda_2)}{\lambda_1 + \lambda_2}$$

is maximized over all d-regular graphs by K_{d+1}.

In fact, Conjecture 2 implies the following conjecture on the maximality of the partition function:

Conjecture 3. *For any $\lambda_1, \lambda_2 > 0$, and any d-regular graph G,*

$$P_G(\lambda_1, \lambda_2) \le P_{K_{d+1}}(\lambda_1, \lambda_2)^{|V(G)|/(d+1)}. \tag{11}$$

To see this, assume $\lambda_1 \geq \lambda_2$, and let $F_G(x) = \frac{1}{n} \log P_G(\lambda_1 - \lambda_2 + x, x)$. We have

$$\frac{1}{n} \log P_G(\lambda_1, \lambda_2) = F_G(\lambda_2) = F_G(0) + \int_0^{\lambda_2} \frac{dF_G}{dx}(x)\, dx$$

$F_G(0) = \frac{1}{n} \log P_G(\lambda_1 - \lambda_2, 0) = \log(1 + \lambda_1 - \lambda_2)$ for all graphs G, and so if we can show that for all $0 \leq x \leq \lambda_2$, $\frac{dF_G}{dx}(x)$ is maximized when $G = K_{d+1}$, then we obtain (the log of) inequality (11). We compute:

$$\begin{aligned}
\frac{dF_G}{dx}(x) &= \frac{1}{n} \frac{\frac{d}{dx} P_G(\lambda_1 - \lambda_2 + x, x)}{P_G(\lambda_1 - \lambda_2 + x, x)} \\
&= \frac{1}{n} \frac{\sum_\chi \frac{xX_1 + (\lambda_1 - \lambda_2 + x)X_2}{x(\lambda_1 - \lambda_2 + x)} (\lambda_1 - \lambda_2 + x)^{X_1} \cdot x^{X_2}}{P_G(\lambda_1 - \lambda_2 + x, x)} \\
&= \frac{1}{x(\lambda_1 - \lambda_2 + x)} \frac{1}{n} \frac{\sum_\chi (xX_1 + (\lambda_1 - \lambda_2 + x)X_2)(\lambda_1 - \lambda_2 + x)^{X_1} \cdot x^{X_2}}{P_G(\lambda_1 - \lambda_2 + x, x)} \\
&= \frac{1}{x(\lambda_1 - \lambda_2 + x)} \left[x\alpha_G^{(1)}(\lambda_1 - \lambda_2 + x, x) + (\lambda_1 - \lambda_2 + x)\alpha_G^{(2)}(\lambda_1 - \lambda_2 + x, x) \right]
\end{aligned}$$

Conjecture 2 implies that this is maximized by K_{d+1}.

References

[1] G. Brightwell and P. Winkler. Hard constraints and the bethe lattice: adventures at the interface of combinatorics and statistical physics. In *Proc. Int'l. Congress of Mathematicians*, volume III, pages 605–624, 2002.

[2] J. Chayes, L. Chayes, and R. Koteckỳ. The analysis of the widom-rowlinson model by stochastic geometric methods. *Communications in Mathematical Physics*, 172(3):551–569, 1995.

[3] E. Davies, M. Jenssen, W. Perkins, and B. Roberts. Independent sets, matchings, and occupancy fractions. *arXiv preprint arXiv:1508.04675*, 2015.

[4] A. Dembo, A. Montanari, N. Sun, et al. Factor models on locally tree-like graphs. *The Annals of Probability*, 41(6):4162–4213, 2013.

[5] D. Galvin. Maximizing h-colorings of a regular graph. *Journal of Graph Theory*, 73(1):66–84, 2013.

[6] D. Galvin. Three tutorial lectures on entropy and counting. *arXiv preprint arXiv:1406.7872*, 2014.

[7] D. Galvin and P. Tetali. On weighted graph homomorphisms. *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, 63:97–104, 2004.

[8] J. Kahn. An entropy approach to the hard-core model on bipartite graphs. *Combinatorics, Probability and Computing*, 10(03):219–237, 2001.

[9] J. Lebowitz and G. Gallavotti. Phase transitions in binary lattice gases. *Journal of Mathematical Physics*, 12(7):1129–1133, 1971.

[10] J. Radhakrishnan. Entropy and counting. *Computational Mathematics, Modelling and Algorithms*, page 146, 2003.

[11] D. Ruelle. Existence of a phase transition in a continuous classical system. *Physical Review Letters*, 27(16):1040, 1971.

[12] L. Sernau. Graph operations and upper bounds on graph homomorphism counts. *arXiv preprint arXiv:1510.01833*, 2015.

[13] B. Widom and J. S. Rowlinson. New model for the study of liquid–vapor phase transitions. *The Journal of Chemical Physics*, 52(4):1670–1684, 1970.

[14] Y. Zhao. The number of independent sets in a regular graph. *Combinatorics, Probability and Computing*, 19(02):315–320, 2010.

[15] Y. Zhao. The bipartite swapping trick on graph homomorphisms. *SIAM Journal on Discrete Mathematics*, 25(2):660–680, 2011.

School of Mathematics, Georgia Institute of Technology
E-mail address: ecohen32@gatech.edu

School of Mathematics, University of Birmingham
E-mail address: math@willperkins.org

School of Mathematics, Georgia Institute of Technology
E-mail address: tetali@math.gatech.edu