Extremal Regular Graphs: Independent Sets and Graph Homomorphisms

Yufei Zhao

Abstract. This survey concerns regular graphs that are extremal with respect to the number of independent sets and, more generally, graph homomorphisms. More precisely, in the family of of d-regular graphs, which graph G maximizes/minimizes the quantity $i(G)^{1/v(G)}$, the number of independent sets in G normalized exponentially by the size of G? What if $i(G)$ is replaced by some other graph parameter? We review existing techniques, highlight some exciting recent developments, and discuss open problems and conjectures for future research.

1. INDEPENDENT SETS. An *independent set* in a graph is a subset of vertices with no two adjacent. Many combinatorial problems can be reformulated in terms of independent sets by setting up a graph where edges represent forbidden relations.

A graph is d*-regular* if all vertices have degree d. In the family of d-regular graphs of the same size, which graph has the largest number of independent sets? This question was initially raised by Andrew Granville in connection to combinatorial number theory and appeared first in print in a paper by Alon [**1**], who speculated that, at least when n is divisible by $2d$, the maximum is attained by a disjoint union of complete bipartite graphs $K_{d,d}$. Some ten years later, Kahn [**23**] arrived at the same conjecture while studying a problem arising from statistical physics. Using a beautiful entropy argument, Kahn proved the conjecture under the additional assumption that the graph is already bipartite. While attending Joe Gallian's Research Experience for Undergraduates in Duluth in 2009, the author showed that the bipartite assumption can be dropped [**30**]. The precise theorem is stated below. We write $I(G)$ to denote the set of independent sets in G and $i(G) := |I(G)|$ for its cardinality. See Figure 1.

Theorem 1.1 ([23] for bipartite *G*; [30] for general *G*). *If G is a bipartite n-vertex d-regular graph, then*

$$i(G) \leq i(K_{d,d})^{n/(2d)} = (2^{d+1} - 1)^{n/(2d)}.$$

Equality occurs when n is divisible by $2d$ and G is a disjoint union of $K_{d,d}$'s. We do not concern ourselves here with what happens when n is not divisible by $2d$, as the extremal graphs are likely dependent on number-theoretic conditions, and we do not know a clean set of examples. Alternatively, the problem can phrased as maximizing $i(G)^{1/v(G)}$ over the set of d-regular bipartite graphs G, where $v(G)$ denotes the number of vertices of G. The above theorem says that this maximum is attained at $G = K_{d,d}$. Note that $i(G)^{1/v(G)}$ remains unchanged if G is replaced by a disjoint union of copies of G.

We provide an exposition of this theorem as well as a discussion of subsequent developments. Notably, Davies, Jenssen, Perkins, and Roberts [**13**] recently gave a new proof by introducing a powerful new technique, which has already had a number of surprising new consequences [**14, 15, 26**]. The results have been partially extended to graph homomorphisms, though many intriguing open problems remain. We also

http://dx.doi.org/10.4169/amer.math.monthly.124.9.827
MSC: Primary 05C35, Secondary 05C69; 05C60

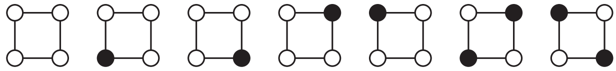

Figure 1. The independent sets of a 4-cycle: $i(C_4) = 7$.

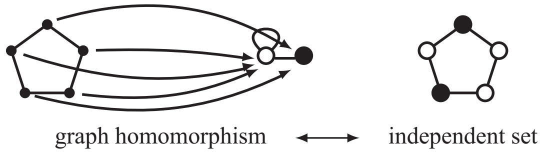

Figure 2. Homomorphisms from G to  correspond to independent sets of G.

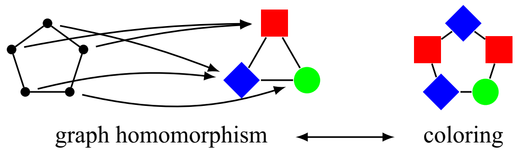

Figure 3. Homomorphisms from G to K_q correspond to proper colorings of vertices of G with q colors.

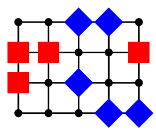

Figure 4. A configuration for the Widom–Rowlinson model on a grid, corresponding to a homomorphism to , where vertices of the grid that are mapped to the first vertex in 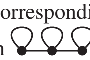 are marked  and those mapped to the third vertex are marked 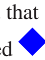.

discuss some recent work on the subject done by Luke Sernau [**28**] as an undergraduate student at Notre Dame.

2. GRAPH HOMOMORPHISMS.

Given two graphs G and H, a *graph homomorphism* from G to H is a map of vertex sets $\phi\colon V(G) \to V(H)$ that sends every edge of G to an edge of H, i.e., $\phi(u)\phi(v) \in E(H)$ whenever $uv \in E(G)$. Here, $V(G)$ denotes the vertex set of G and $E(G)$ the edge set. Denote the set of graph homomorphisms from G to H by

$$\mathrm{Hom}(G, H) := \{\phi\colon V(G) \to V(H) : \phi(u)\phi(v) \in E(H)\ \forall uv \in E(G)\}.$$

We use lowercase letters for cardinalities: $v(G) := |V(G)|$, $e(G) := |E(G)|$, and $\hom(G, H) := |\mathrm{Hom}(G, H)|$.

We usually use the letter G for the source graph and H for the target graph. It will be useful to allow the target graph H to have loops (but not multiple edges), and we shall refer to such graphs as *loop-graphs*. The source graph G is usually simple (without loops). By *graph*, we usually mean a simple graph.

Graph homomorphisms generalize the notion of independent sets. They are equivalent to labeling the vertices of G subject to certain constraints encoded by H.

Example 2.1 (Independent sets). Homomorphisms from G to  correspond bijectively to independent sets in G. Indeed, a map of vertices from G to 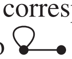 is a homomorphism if and only if the preimage of the nonlooped vertex in 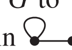 forms an independent set in G. So $\hom(G,$ $) = i(G)$. See Figure 2. In the statistical physics literature,[1] independent sets correspond to *hard-core models*. For example, they can be used to represent configurations of nonoverlapping spheres ("hard" spheres) on a grid.

Example 2.2 (Graph colorings). When the target graph is the complete graph K_q on q vertices, a graph homomorphism from G to K_q corresponds to a coloring of the vertices of G with q colors so that no two adjacent vertices of G receive the same color. Such colorings are called *proper q-colorings*. See Figure 3. Thus, $\hom(G, K_q)$ is the number of proper q-colorings of G. For a fixed G, the quantity $\hom(G, K_q)$ is a polynomial function in q, and it is called the *chromatic polynomial* of G, a classic object in graph theory.

Example 2.3 (Widom–Rowlinson model). A homomorphism from G to  corresponds to a partial coloring of the vertices of G with red or blue, allowing vertices to be left uncolored, such that no red vertex is adjacent to a blue vertex. Such a coloring is known as a *Widom–Rowlinson configuation*. See Figure 4.

As graph homomorphisms generalize independent sets, one may wonder whether Theorem 1.1 generalizes to graph homomorphisms. It turns out, perhaps surprisingly, that the bipartite case of Theorem 1.1, concerning the number of independent sets in a regular bipartite graph, always extends to graph homomorphisms.

Theorem 2.4 (Galvin and Tetali [21]). *Let G be a bipartite d-regular graph and H a loop-graph. Then*

$$\hom(G, H)^{1/v(G)} \leq \hom(K_{d,d}, H)^{1/(2d)}.$$

Can the bipartite hypothesis above be dropped as in Theorem 1.1? The answer is no. Indeed, with $H =$  being two disjoint loops, $\hom(G,$ $) = 2^{c(G)}$ where $c(G)$ is the number of connected components of G. In this case, $\hom(G,$ $)^{1/v(G)}$ is maximized when the sizes of the components of G are as small as possible (among d-regular graphs), i.e., when $G = K_{d+1}$.

The central problem of interest for the rest of this article is stated below. It has been solved for certain targets H, but it is open in general. The analogous minimization problem is also interesting and will be discussed in Section 8.

Problem 2.5. *Fix a loop-graph H and a positive integer d. Determine the supremum of* $\hom(G, H)^{1/v(G)}$ *taken over all d-regular graphs G.*

We have already seen two cases where Problem 2.5 has been solved: When $H =$ , the maximum is attained by $G = K_{d,d}$ (Theorem 1.1), and when $H =$ , the maximum is attained by $G = K_{d+1}$. The latter example can be extended to H being a disjoint union of complete loop-graphs. Another easy case is H bipartite, as $\hom(G, H) = 0$ unless G is bipartite, so the maximizer is $K_{d,d}$ by Theorem 2.4.

In his undergraduate senior project, the author extended Theorem 1.1 to solve Problem 2.5 for a certain family of graphs H called *bipartite swapping targets*. As a special case, we define a *loop-threshold graph* to be a loop-graph whose vertices can be ordered so that its adjacency matrix has the property that whenever an entry is 1, all entries to the left of it and above it are 1 as well. An example of a loop-threshold graph, along with its adjacency matrix, is shown below.

[1] See [5] for the connection between the combinatorics of graph homomorphisms and Gibbs measures in statistical physics.

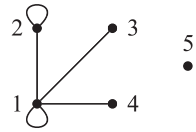

$$\begin{pmatrix} 1 & 1 & 1 & 1 & 0 \\ 1 & 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix}.$$

Loop-threshold graphs generalize  from Example 2.1. The following result was obtained by extending the proof method of Theorem 1.1.

Theorem 2.6 ([31]). *Let G be a d-regular graph G and H a loop-threshold graph. Then*

$$\hom(G, H)^{1/v(G)} \le \hom(K_{d,d}, H)^{1/(2d)}.$$

Sernau [**28**] recently extended the above results to an even larger family of H (see Section 5). The most interesting open case of Problem 2.5 is $H = K_q$, concerning the number of proper q-colorings of vertices of G (Example 2.2).

Conjecture 2.7. *For every d-regular graph G and integer $q \ge 3$,*

$$\hom(G, K_q)^{1/v(G)} \le \hom(K_{d,d}, K_q)^{1/(2d)}.$$

The conjecture was recently solved for $d = 3$ by Davies, Jenssen, Perkins, and Roberts [**15**] using a novel method they developed earlier. We will discuss the method in Section 7. The conjecture remains open for all $d \ge 4$ and $q \ge 3$. The above inequality is known to hold if q is sufficiently large as a function of G [**31**] (the current best bound is $q > 2\binom{v(G)d/2}{4}$ [**19**]).

The first nontrivial case of Problem 2.5 where the maximizing G is not $K_{d,d}$ was obtained recently by Cohen, Perkins, and Tetali [**8**].

Theorem 2.8 (Cohen, Perkins, and Tetali [8]). *For any d-regular graph G, we have*

$\hom(G,$  $)^{1/v(G)} \le \hom(K_{d+1},$  $)^{1/(d+1)}.$

Theorem 2.8 was initially proved [**8**] using the occupancy fraction method, which will be discussed in Section 7. Subsequently, a much shorter proof was given in [**7**] (also see Sernau [**28**]).[2] These methods can be used to prove that K_{d+1} is the maximizer for a large family of target loop-graphs H (see Section 5).

There are weighted generalizations of these problems and results. However, for clarity, we defer discussing the weighted versions until Section 7, where we will see that introducing weights leads to a powerful new differential method for proving the unweighted results.

We conclude this section with some open problems. Galvin [**19**] conjectured that in Problem 2.5, the maximizing G is always either $K_{d,d}$ or K_{d+1}, as with all the cases we have seen so far. However, this conjecture was recently disproved by Sernau [**28**][3] (see Section 6). As it stands, there does not seem to be a clean conjecture concerning the solution to Problem 2.5 on determining the maximizing G. Sernau suggested the possibility that there is a finite list of maximizing G for every d.

[2]Sernau also tackled Theorem 2.8, obtaining an approximate result in a version of [**28**] that predated [**8**] and [**7**]. After the appearance of [**7**], Sernau corrected an error (identified by Cohen) in [**28**], and the corrected version turned out to include Theorem 2.8 as a special case.

[3]A similar counterexample was independently found by Pat Devlin.

Conjecture 2.9. *For every $d \geq 3$, there exists a finite set $\mathcal{G}_d$ of d-regular graphs such that for every loop-graph H and every d-regular graph G one has*

$$\hom(G, H)^{1/v(G)} \leq \max_{G' \in \mathcal{G}_d} \hom(G', H)^{1/v(G')}.$$

It has been speculated that the maximizing G perhaps always has between $d + 1$ and $2d$ vertices (corresponding to K_{d+1} and $K_{d,d}$, respectively). Sernau suggested the possibility that for a fixed H, the maximizer is always one of $K_{d,d}$ and K_{d+1} as long as d is large enough.

Conjecture 2.10. *Let H be a fixed loop-graph. There is some d_H such that for all $d \geq d_H$ and d-regular graphs G,*

$$\hom(G, H)^{1/v(G)} \leq \max\{\hom(K_{d+1}, H)^{1/(d+1)}, \hom(K_{2d}, H)^{1/(2d)}\}.$$

We do not know if the supremum in Problem 2.5 can always be attained.

Question 2.11. Fix $d \geq 3$ and a loop-graph H. Is the supremum of $\hom(G, H)^{1/v(G)}$ over all d-regular graphs G always attained by some G?

It could be the case that the supremum is the limit coming from a sequence of graphs G of increasing size instead of a single graph G on finitely many vertices. This is indeed the case if we wish to *minimize* hom(G, 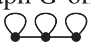)$^{1/v(G)}$ over d-regular graphs G. Csikvári [**9**] recently showed that the infimum of hom(G, )$^{1/v(G)}$ is given by a limit of d-regular graphs G with increasing girth (i.e., G locally looks like a d-regular tree at every vertex).

3. PROJECTION INEQUALITIES. The original proof of the bipartite case of Theorem 1.1 (as well as Theorem 2.4) uses a beautiful entropy argument, with a key input being Shearer's entropy inequality [**6**]. We will not cover the entropy arguments as they would lead us too far astray. See Galvin's lecture notes [**20**] for a nice exposition of the entropy method for counting problems. The first nonentropy proof of these two theorems was given in [**25**] using a variant of Hölder's inequality, which we describe in this section. We begin our discussion with a classical projection inequality. See Friedgut's MONTHLY article [**17**] concerning how the projection inequalities relate to entropy.

Let P_{xy} denote the projection operator from $\mathbb{R}^3$ onto the xy-plane. Similarly, define P_{xz} and P_{yz}. Let S be a body in $\mathbb{R}^3$ such that each of the three projections $P_{xz}(S)$, $P_{yz}(S)$, and $P_{xz}(S)$ has area 1. What is the maximum volume of S? (This is not as obvious as it may first appear. Note that we are projecting onto the 2-D coordinate planes as opposed to the 1-D axes.)

The answer is 1, attained when S is an axes-parallel cube of side-length 1. Indeed, equivalently (by rescaling), we have

$$\operatorname{vol}(S)^2 \leq \operatorname{area}(P_{xy}(S)) \operatorname{area}(P_{xz}(S)) \operatorname{area}(P_{yz}(S)). \tag{1}$$

Such results were first obtained by Loomis and Whitney [**24**]. More generally, for any functions $f, g, h \colon \mathbb{R}^2 \to \mathbb{R}$ (assuming integrability conditions)

$$\left(\int_{\mathbb{R}^3} f(x, y) g(x, z) h(y, z)\, dx dy dz\right)^2$$

$$\leq \left(\int_{\mathbb{R}^2} f(x,y)^2\,dxdy\right)\left(\int_{\mathbb{R}^2} g(x,z)^2\,dxdz\right)\left(\int_{\mathbb{R}^2} h(y,z)^2\,dydz\right). \quad (2)$$

To see how (2) implies (1), take f, g, h to be the indicator functions of the projections of S onto the three coordinates planes, and observe that $1_S(x, y, z) \leq f(x, y)g(x, z)h(y, z)$.

Let us prove (2). In fact, x, y, z can vary over any measurable space instead of $\mathbb{R}$. In our application, the domains will be discrete, i.e., the integral will be a sum. It suffices to prove the inequality when f, g, h are nonnegative. The proof is via three simple applications of the Cauchy–Schwarz inequality, to the variables x, y, z, one at a time in that order:

$$\begin{aligned}
&\int f(x,y)g(x,z)h(y,z)\,dxdydz \\
&\leq \int \left(\int f(x,y)^2\,dx\right)^{1/2}\left(\int g(x,z)^2\,dx\right)^{1/2} h(y,z)\,dydz \\
&\leq \int \left(\int f(x,y)^2\,dxdy\right)^{1/2}\left(\int g(x,z)^2\,dx\right)^{1/2}\left(\int h(y,z)^2\,dy\right)^{1/2} dz \\
&\leq \left(\int f(x,y)^2\,dxdy\right)^{1/2}\left(\int g(x,z)^2\,dxdz\right)^{1/2}\left(\int h(y,z)^2\,dydz\right)^{1/2} \\
&= \|f\|_2\|g\|_2\|h\|_2,
\end{aligned}$$

where

$$\|f\|_p := \left(\int |f|^p\right)^{1/p}$$

is the L^p norm. This proves (2). This inequality strengthens Hölder's inequality since a direct application of Hölder's inequality would yield

$$\int fgh \leq \|f\|_3\|g\|_3\|h\|_3. \quad (3)$$

What we have shown is that whenever each of the variables x, y, z appears in the argument of exactly two of the three functions f, g, and h, then the L^3 norms on the right-hand side of (3) can be sharpened to L^2 norms (we always have $\|f\|_2 \leq \|f\|_3$ by convexity).

The above proof easily generalizes to prove the following more general result [**16**] (also see [**25**, Theorem 3.1]). It is also related to the Brascamp–Lieb inequality [**3**].

Theorem 3.1. *Let $A_1, \ldots, A_m$ be subsets of $[n] := \{1, 2, \ldots, n\}$ such that each $i \in [n]$ appears in exactly d of the sets A_j. Let Ω_i be a measure space for each $i \in [n]$. For each j, let $f_j \colon \prod_{i\in A_j} \Omega_i \to \mathbb{R}$ be measurable functions. Let P_j denote the projection of $\mathbb{R}^n$ onto the coordinates indexed by A_j. Then*

$$\int_{\Omega_1\times\cdots\times\Omega_n} f_1(P_1(\mathbf{x}))\cdots f_m(P_m(\mathbf{x}))\,d\mathbf{x} \leq \|f_1\|_d\cdots\|f_m\|_d.$$

Using this inequality, we now prove Theorem 2.4.

Proof of Theorem 2.4. [**25**] Let $V(G) = U \cup W$ be a bipartition of G. Since G is d-regular, $|U| = |W| = v(G)/2$. For any $z_1, \ldots, z_d \in V(H)$, let

$$f(z_1, \ldots, z_d) := |\{z \in V(H) : z_1 z, \ldots, z_d z \in E(H)\}|$$

denote the size of the common neighborhood of $z_1, \ldots, z_d$ in H.

For any $\phi\colon U \to V(H)$, the number of ways to extend ϕ to a graph homomorphism from G to H can be determined by noting that, for each $w \in W$, there are exactly $f(\phi(u) : u \in N(w))$ choices for its image $\phi(w)$, independent of the choices for other vertices in W. Therefore,

$$\hom(G, H) = \sum_{\phi\colon U \to V(H)} \prod_{w \in W} f(\phi(u) : u \in N(w)).$$

Since G is d-regular, every $u \in U$ is contained in $N(w)$ for exactly d different $w \in W$. Therefore, by applying Theorem 3.1 with the counting measure on $V(H)$, we find that

$$\hom(G, H) \le \|f\|_d^{|W|}.$$

Note that

$$\|f\|_d^d = \sum_{z_1, \ldots, z_d \in V(H)} f(z_1, \ldots, z_d)^d = \hom(K_{d,d}, H).$$

Therefore,

$$\hom(G, H) \le \hom(K_{d,d}, H)^{|W|/d} = \hom(K_{d,d}, H)^{v(G)/(2d)}.$$ ■

4. A BIPARTITE SWAPPING TRICK. In the previous section, we proved Theorem 1.1 for bipartite graphs G. Now we use it to deduce Theorem 1.1 for nonbipartite graphs. The proof follows [**30, 31**]. The idea is to transform G into a bipartite graph, namely the *bipartite double cover* $G \times K_2$, with vertex set $V(G) \times \{0, 1\}$. The vertices of $G \times K_2$ are labeled v_i for $v \in V(G)$ and $i \in \{0, 1\}$. Its edges are $u_0 v_1$ for all $uv \in E(G)$. See Figure 5. This construction is a special case of the graph tensor product, which we define in the next section. Note that $G \times K_2$ is always a bipartite graph. The following key lemma shows that $G \times K_2$ always has at least as many independent sets as two disjoint copies of G.

Lemma 4.1 ([30]). *Let G be any graph (not necessarily regular). Then*

$$i(G)^2 \le i(G \times K_2).$$

Since $G \times K_2$ is bipartite and d-regular, the bipartite case of Theorem 1.1 implies

$$i(G)^2 \le i(G \times K_2) \le (2^{d+1} - 1)^{n/d}$$

so that Theorem 1.1 follows immediately. See Figure 5 for an illustration of the following proof.

Proof of Lemma 4.1. Let $2G$ denote a disjoint union of two copies of G. Label its vertices by v_i with $v \in V$ and $i \in \{0, 1\}$ so that its edges are $u_i v_i$ with $uv \in E(G)$ and $i \in \{0, 1\}$. We will give an injection $\phi\colon I(2G) \to I(G \times K_2)$. Recall that $I(G)$ is the set of independent sets of G. The injection would imply $i(G)^2 = i(2G) \le i(G \times K_2)$ as desired.

Fix an arbitrary order on all subsets of $V(G)$. Let S be an independent set of $2G$. Let

$$E_{\text{bad}}(S) := \{uv \in E(G) : u_0, v_1 \in S\}.$$

Note that $E_{\text{bad}}(S)$ is a bipartite subgraph of G since each edge of E_{bad} has exactly one endpoint in $\{v \in V(G) : v_0 \in S\}$ but not both (or else S would not be independent). Let A denote the first subset (in the previously fixed ordering) of $V(G)$ such that all edges in $E_{\text{bad}}(S)$ have one vertex in A and the other outside A. Define $\phi(S)$ to be the subset of $V(G) \times \{0, 1\}$ obtained by "swapping" the pairs in A, i.e., for all $v \in A$, $v_i \in \phi(S)$ if and only if $v_{1-i} \in S$ for each $i \in \{0, 1\}$, and for all $v \notin A$, $v_i \in \phi(S)$ if and only if $v_i \in S$ for each $i \in \{0, 1\}$. It is not hard to verify that $\phi(S)$ is an independent set in $G \times K_2$. The swapping procedure fixes the "bad" edges.

It remains to verify that ϕ is an injection. Let $S \in I(2G)$ and $T = \phi(S)$. We show how to recover S from T. Set

$$E'_{\text{bad}}(T) = \{uv \in E(G) : u_i, v_i \in T \text{ for some } i \in \{0, 1\}\}$$

so that $E_{\text{bad}}(S) = E'_{\text{bad}}(T)$. We find A as earlier and then swap the pairs of A back. (Remark: It follows that $T \in I(G \times K_2)$ lies in the image of ϕ if and only if $E'_{\text{bad}}(T)$ is bipartite.) ∎

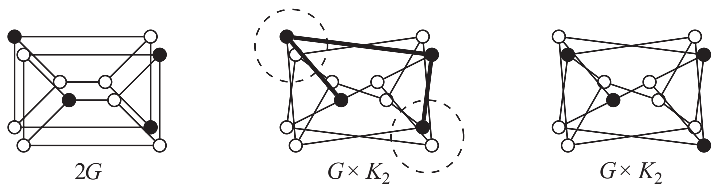

Figure 5. The bipartite swapping trick in the proof of Lemma 4.1: swapping the circled pairs of vertices (denoted A in the proof) fixes the bad edges (bolded), transforming an independent set of $2G$ into an independent set of $G \times K_2$.

The above method was used to drop the G bipartite hypothesis in Theorem 2.4 when H belongs to a certain class of loop-graphs called *bipartite swapping targets*, which includes the threshold graphs in Theorem 2.6 as a special case. See [**31**] for details.

We do not know how to extend the method to $H = K_q$, corresponding to the number of proper q-colorings. Nonetheless, the analogous strengthening of Conjecture 2.7 is conjectured to hold (the inequality was proved for q sufficiently large as a function of G).

Conjecture 4.2 ([31]). *For every graph G and every $q \geq 3$,*

$$\hom(G, K_q)^2 \leq \hom(G \times K_2, K_q).$$

5. GRAPH PRODUCTS AND POWERS.

We define several graph operations.

- *Tensor product* $G \times H$: Its vertices are $V(G) \times V(H)$, with (u, v) and $(u', v') \in V(G) \times V(H)$ adjacent in $G \times H$ if $uu' \in E(G)$ and $vv' \in E(H)$. This construction is also known as the *categorical product*.
- *Exponentiation* H^G: Its vertices are maps $f : V(G) \to V(H)$ (not necessarily homomorphisms), where f and f' are adjacent if $f(u)f'(v) \in E(H)$ whenever $uv \in E(G)$.
- G°: Same as G except that every vertex now has a loop.

- $\ell(H)$: Subgraph of H induced by its looped vertices, or equivalently, delete all nonlooped vertices from H.

Example 5.1. For any loop-graph H, the graph H^{K_2} has vertex set $V(H) \times V(H)$, with (u, v) and $(u', v') \in V(H) \times V(H)$ adjacent if and only if $uv', u'v \in E(H)$. In particular, if $H_{\text{ind}} =$ , then $H_{\text{ind}}^{K_2} =$ 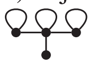.

Here are a few easy yet key facts relating the above operations with graph homomorphisms. The proofs are left as exercises for the readers:

$$\hom(G, H_1 \times H_2) = \hom(G, H_1)\hom(G, H_2), \tag{4}$$

$$\hom(G \times G', H) = \hom(G, H^{G'}), \tag{5}$$

$$\hom(G^\circ, H) = \hom(G, \ell(H)). \tag{6}$$

Clique as maximizer. Now we prove Theorem 2.8 concerning the Widom–Rowlinson model. Recall it says that for any d-regular graph G, we have

$\hom\Big(G,$  $\Big)^{1/v(G)} \le \hom\Big(K_{d+1},$  $\Big)^{1/(d+1)}.$

Proof of Theorem 2.8. [**7**] We have  $= \ell(H_{\text{ind}}^{K_2})$ (Example 5.1). For any graph G,

$\hom\Big(G,$  $\Big) = \hom\Big(G, \ell(H_{\text{ind}}^{K_2})\Big) = \hom\Big(G^\circ, H_{\text{ind}}^{K_2}\Big) = \hom\Big(G^\circ \times K_2,$  $\Big).$ (7)

When G is d-regular, $G^\circ \times K_2$ is a $(d + 1)$-regular bipartite graph, so Theorem 2.4 (or Theorem 1.1) implies that the above quantity is at most $\hom\Big(K_{d+1,d+1},$  $\Big)^{v(G)/(d+1)}$. Since $K_{d+1,d+1} = K_{d+1}^\circ \times K_2$, we have by (7),

$\hom(G,$  $)^{1/v(G)} = \hom(G^\circ \times K_2,$  $)^{1/v(G)}$

$\le \hom(K_{d+1}^\circ \times K_2,$  $)^{1/(d+1)}$

$= \hom(K_{d+1},$  $)^{1/(d+1)}.$ ■

The above proof exploits the connection (7) between the hard-core model (independent sets) and the Widom–Rowlinson model. This relationship had been previously observed in [**4**]. More generally, the above proof extends to give the following result.

Definition 5.2. The *extended line graph* $\widetilde{H}$ of a graph H has $V(\widetilde{H}) = E(H)$ and two edges e and f of H are adjacent in $\widetilde{H}$ if (1) $e = f$, or (2) e and f share a common vertex, or (3) e and f are opposite edges of a 4-cycle in G.

Theorem 5.3 ([7]). *Let $\widetilde{H}$ be the extended line graph of a bipartite graph H. For any d-regular graph G,*

$$\hom(G, \widetilde{H})^{1/v(G)} \le \hom(K_{d+1}, \widetilde{H})^{n/(d+1)}.$$

For a simple graph H, let H° denote H with a loop added at every vertex. Let P_k denote the path of k vertices and C_k the cycle with k vertices.

Example 5.4. One has $\widetilde{P}_{k+1} = P_k^\circ$ for all k. Also, $\widetilde{C}_k = C_k^\circ$ for all $k \neq 4$.

Corollary 5.5 ([7]). *Let $H = C_k^\circ$ with even $k \geq 6$ or $H = P_k^\circ$ for any $k \geq 1$. For any d-regular graph G,*

$$\hom(G, H)^{1/v(G)} \leq \hom(K_{d+1}, H)^{1/v(G)}.$$

The following related result was established by Sernau using similar techniques.

Theorem 5.6 ([28]). *Let $H = \ell(A^B)$ where A is any loop-graph and B is a bipartite graph. For any d-regular graph G,*

$$\hom(G, H)^{1/v(G)} \leq \hom(K_{d+1}, H)^{1/(d+1)}.$$

Closure under tensor products. Sernau [**28**] observed that, for any d, if $H = H_1$ and $H = H_2$ both have the property that $G = K_{d,d}$ maximizes $\hom(G, H)^{1/v(G)}$ over all d-regular graphs G, then $H = H_1 \times H_2$ has the same property by (4). In other words, the set of H such that $G = K_{d,d}$ is the maximizer in Problem 2.5 is closed under tensor products. This observation enlarges the set of such H previously obtained in [**31**].

6. NEITHER COMPLETE BIPARTITE NOR CLIQUE. In all cases of Problem 2.5 that we have considered, the maximizing G is always either $K_{d,d}$ or K_{d+1}. It was conjectured [**19**] that one of $K_{d,d}$ and K_{d+1} always maximizes $\hom(G, H)^{1/v(G)}$ for every H. However, Sernau [**28**] showed that this is false (a similar construction was independently found by Pat Devlin).

Let $d \geq 4$, and let G be a d-regular graph with $v(G) < 2d$ other than K_{d+1}. Brooks's theorem tells us that G is d-colorable so that $\hom(G, K_d) > 0$. It follows that for this G,

$$\begin{aligned}\hom(G, kK_d)^{1/v(G)} &= k^{1/v(G)} \hom(G, K_d)^{1/v(G)} \\ &> k^{1/(2d)} \hom(K_{d,d}, K_d)^{1/(2d)} = \hom(K_{d,d}, kK_d)^{1/(2d)}\end{aligned}$$

for sufficiently large k (as a function of d) since $v(G) < 2d$. Also,

$$\hom(G, kK_d)^{1/v(G)} > 0 = \hom(K_{d+1}, kK_d)^{1/(d+1)}.$$

Therefore, neither $G = K_{d,d}$ nor $G = K_{d+1}$ maximize $\hom(G, kK_d)^{1/v(G)}$ over all d-regular graphs G. For $d = 3$, Csikvári [**10**] found a counterexample using a similar construction.

In general, we do not know which graphs G (other than K_{d+1} and $K_{d,d}$) can arise as maximizers for $\hom(G, H)^{1/v(G)}$ in Problem 2.5. See the end of Section 2 for some open questions and conjectures.

7. OCCUPANCY FRACTION. The original proof of the bipartite case of Theorem 1.1 used the entropy method [**23**]. The proof in Section 3, following [**25**], used a variant of the Hölder's inequality and is related to the original entropy method proof. Recently, an elegant new proof of the result was found [**13**] using a novel method, which was unrelated to previous proofs. We discuss this new technique in this section. It will be necessary to introduce weighted versions of the problems.

The *independence polynomial* of a graph G is defined by

$$P_G(\lambda) := \sum_{I \in I(G)} \lambda^{|I|}.$$

Recall that $I(G)$ is the set of independent sets of G. In particular, $P_G(1) = i(G)$. Theorem 1.1 extends to this weighted version of the number of independent sets.

Theorem 7.1 ([21] for bipartite *G*; [30] for general *G*). *If G is a d-regular graph and $\lambda \geq 0$, then*

$$P_G(\lambda)^{1/v(G)} \leq P_{K_{d,d}}(\lambda)^{1/(2d)}.$$

The *hard-core model* with *fugacity* λ on G is defined as the probability distribution on independent sets of G where an independent set I is chosen with probability proportional to $\lambda^{|I|}$, i.e., with probability

$$\Pr_{\lambda}[I] = \frac{\lambda^{|I|}}{P_G(\lambda)}.$$

The *occupancy fraction* of I is the fraction of vertices of G occupied by I. The expected occupancy fraction of a random independent set from the hard-core model is

$$\alpha_G(\lambda) := \frac{1}{v(G)} \sum_{I \in I(G)} |I| \cdot \Pr_{\lambda}[I] = \frac{\sum_{I \in I(G)} |I| \lambda^{|I|}}{v(G) P_G(\lambda)} = \frac{\lambda P'_G(\lambda)}{v(G) P_G(\lambda)}.$$

It turns out that $K_{d,d}$ maximizes the occupancy fraction among all d-regular graphs.

Theorem 7.2 (Davies, Jenssen, Perkins, and Roberts [13]). *For all d-regular graphs G and all $\lambda \geq 0$, we have*

$$\alpha_G(\lambda) \leq \alpha_{K_{d,d}}(\lambda) = \frac{\lambda(1+\lambda)^{d-1}}{2(1+\lambda)^d - 1}. \tag{8}$$

Since the expected occupancy fraction is proportional to the logarithmic derivative of $P_G(\lambda)^{1/v(G)}$, the inequality for the expected occupancy fraction implies the corresponding inequality for the independence polynomial. Indeed, Theorem 7.2 implies Theorem 7.1 (and hence Theorem 1.1) since

$$\frac{1}{v(G)} \int_0^{\lambda} \frac{\overline{\alpha}_G(t)}{t}\, dt = \frac{1}{v(G)} \int_0^{\lambda} \frac{P'_G(t)}{P_G(t)}\, dt = \frac{\log P_G(\lambda)}{v(G)}.$$

We reproduce here two proofs of Theorem 7.2. They are both based on the following idea, introduced in [**13**] for this problem. We draw a random independent set I from the hard-core model and look at the neighborhood of a uniform random vertex $v \in V(G)$. The expected occupancy fraction is then the probability that $v \in I$. We then analyze how the neighborhood of v should look in relation to I. Since the graph is regular, a uniform random neighbor of v is uniformly distributed in $V(G)$. By finding an appropriate set of constraints on the probabilities of seeing various neighborhood configurations of v, we can bound the probability that $v \in I$.

The first proof is given below under the additional simplifying assumption that G is triangle-free (which includes all bipartite graphs and much more). See [**13**] for how to extend this proof to all regular graphs.

Proof of Theorem 7.2 for triangle-free G. Let I be an independent set of G drawn according to the hard-core model with fugacity λ. For each $v \in V(G)$, let p_v denote the probability that $v \in I$. We say $v \in V(G)$ is *uncovered* if none of the neighbors of v is in I, i.e., $N(v) \cap I = \emptyset$. If $v \in I$, then v is necessarily uncovered. Conversely, conditioned on v being uncovered, one has $v \in I$ with probability $\lambda/(1+\lambda)$. So the probability that v is uncovered is $p_v(1+\lambda)/\lambda$.

Let U_v denote the set of uncovered neighbors of v. Since G is triangle-free, U_v is an independent set. Conditioned on U_v being the uncovered neighbors of v, the probability that v is uncovered, which is equivalent to $U_v \cap I = \emptyset$, is exactly $(1+\lambda)^{-|U_v|}$. Hence,

$$\frac{1+\lambda}{\lambda} p_v = \mathbb{E}[(1+\lambda)^{-|U_v|}] \le 1 - \frac{\mathbb{E}[|U_v|]}{d}\left(1-(1+\lambda)^{-d}\right), \tag{9}$$

where the inequality follows from $0 \le |U_v| \le d$ and the convexity of the function $x \mapsto (1+\lambda)^{-x}$ so that $(1+\lambda)^{-x} \le 1 - \frac{x}{d}(1-(1+\lambda)^{-d})$ for all $0 \le x \le d$ by linear interpolation.

If v is chosen from $V(G)$ uniformly at random, then $\mathbb{E}[p_v] = \alpha_G(\lambda)$ is the expected occupancy fraction. Similarly, $\mathbb{E}[|U_v|]/d$ is the probability that a random vertex is uncovered (here we use again that G is d-regular), which equals $\mathbb{E}[p_v]\frac{1+\lambda}{\lambda} = \alpha_G(\lambda)\frac{1+\lambda}{\lambda}$. Setting into (9), we obtain

$$\frac{1+\lambda}{\lambda}\alpha_G(\lambda) \le 1 - \alpha_G(\lambda)\frac{1+\lambda}{\lambda}\left(1-(1+\lambda)^{-d}\right).$$

Rearranging gives us (8). ∎

In [**13**], Theorem 7.2 was proved for all d-regular graphs G by considering all graphs on d vertices that could be induced by the neighborhood of a vertex in G and using a linear program to constrain the probability distribution of the neighborhood profile of a random vertex. When G is triangle-free, the neighborhood of a vertex is always an independent set, which significantly simplifies the situation. The following conjecture extends Theorem 2.4 to triangle-free graphs.

Conjecture 7.3 ([7]). *Let G be a triangle-free d-regular graph and H a loop-graph. Then*

$$\hom(G,H)^{1/v(G)} \le \hom(K_{d,d},H)^{1/(2d)}.$$

Next, we give an alternative proof of Theorem 7.2 due to Perkins [**27**], based on a similar idea. In the following proof, we do not need to assume that G is triangle-free. We introduce an additional constraint, which allows us to obtain the result more quickly. This simplification seems to be somewhat specific to independent sets.

Second proof of Theorem 7.2. Let I be an independent set of G drawn according to the hard-core model with fugacity λ, and let v be a uniform random vertex in G. Let $Y = |I \cap N(v)|$ denote the number of neighbors of v in I (not including v itself). Let $p_k = \mathbb{P}(Y = k)$. Since $Y \in \{0, 1, \ldots, d\}$,

$$p_0 + p_1 + \cdots + p_d = 1. \tag{10}$$

However, not all vectors of probabilities $(p_0, \ldots, p_d)$ are feasible. The art of the method is in finding additional constraints on the probability distribution.

As in the previous proof, since v is uncovered if and only if $Y = 0$, we have

$$\alpha_G(\lambda) = \mathbb{P}(v \in I) = \frac{\lambda}{1+\lambda}\mathbb{P}(Y=0) = \frac{\lambda}{1+\lambda}p_0.$$

On the other hand, since G is d-regular, a uniform random neighbor of v is also uniformly distributed in $V(G)$, so we have

$$\alpha_G(\lambda) = \frac{1}{d}\mathbb{E}[Y] = \frac{1}{d}(p_1 + 2p_2 + \cdots + dp_d).$$

Comparing the previous two relations, we obtain

$$\frac{\lambda}{1+\lambda}p_0 = \frac{1}{d}(p_1 + 2p_2 + \cdots + dp_d). \tag{11}$$

Now, let us compare the probability that v has k versus $k-1$ neighbors in I. In an event where exactly k neighbors of v are occupied, we can remove any of the occupied neighbors from I and obtain another independent set where v has exactly $k-1$ neighbors. There are k ways to remove an element, but we overcount by a factor of at most $d-k+1$. Also factoring in the weight multiplier, we obtain the inequality

$$(d-k+1)\lambda p_{k-1} \geq kp_k, \qquad \text{for } 2 \leq k \leq d. \tag{12}$$

The constraints (10), (11), and (12) form a linear program with variables $p_0, \ldots, p_d$. Next, we show that these linear constraints together imply $p_0 \leq \frac{(1+\lambda)^d}{2(1+\lambda)^d - 1}$, which gives the desired bound on $\alpha_G(\lambda) = \frac{\lambda}{1+\lambda}p_0$. Equality is attained for the probability distribution $(p_0, \ldots, p_d)$ arising from $G = K_{d,d}$.

To prove this claim, first we show that if $(p_0, \ldots, p_d)$ achieves the maximum of value of p_0 while satisfying the constraints (10), (11), and (12), then every inequality in (12) must be an equality. Indeed, if we have $(d-k+1)\lambda p_{k-1} > kp_k$ for some k, then by increasing p_0 by ϵ, decreasing p_{k-1} by $(\frac{d\lambda}{1+\lambda} + k)\epsilon$, increasing p_k by $(\frac{d\lambda}{1+\lambda} + k - 1)\epsilon$, and leaving all other p_i's fixed, we can maintain all constraints and increase p_0, provided $\epsilon > 0$ is sufficiently small. Thus, in the maximizing solution, equality occurs in (12) for all $2 \leq k \leq d$. It can be checked that the vector $(p_0, \ldots, p_d)$ arising from $G = K_{d,d}$ satisfies all the equality constraints, and it is the unique solution since we have a linear system of equations with full rank. ∎

Remark. Conjecture 2.7 about the number of colorings was recently proved [**15**] for 3-regular graphs using an extension of the above method.

8. ON THE MINIMUM NUMBER OF INDEPENDENT SETS AND HOMOMORPHISMS.

Independent sets. Having explored the maximum number of independent sets in a regular graph, let us turn to the natural opposite question. Which d-regular graph has the minimum number of independent sets? It turns out that the answer is a disjoint union of cliques.

Theorem 8.1 (Cutler and Radcliffe [12]). *For a d-regular graph G,*

$$i(G)^{1/v(G)} \geq i(K_{d+1})^{1/(d+1)} = (d+2)^{1/(d+1)}.$$

In fact, a stronger result holds: A disjoint union of K_{d+1}'s minimizes the number of independent sets of every fixed size. We write aG for a disjoint union a copies of G. Let $i_t(G)$ denote the number of independent sets of G of size t.

Theorem 8.2 ([12]). *Let a and d be positive integers. Let G be a d-regular graph with $a(d+1)$ vertices. Then $i_t(G) \geq i_t(aK_{d+1})$ for every $0 \leq t \leq a(d+1)$.*

Proof. Let us compare the number of sequences of t vertices that form an independent set in G and aK_{d+1}. In aK_{d+1}, we have $a(d+1)$ choices for the first vertex. Once the first vertex has been chosen, there are exactly $(a-1)(d+1)$ choices for the second vertex. More generally, for $1 \leq j \leq a$, once the first $j-1$ vertices have been chosen, there are exactly $(a+1-j)(d+1)$ choices for the jth vertex.

On the other hand, in G, after the first $j-1$ vertices have been chosen, the union of these $j-1$ vertices along with their neighborhoods has cardinality at most $(j-1)(d+1)$, so there are at least $(a+1-j)(d+1)$ choices for the jth vertex, at least as many compared to aK_{d+1}. ■

Proof of Theorem 8.1. Theorem 8.2 implies that $i(G)^{1/v(G)} \geq i(K_{d+1})^{1/(d+1)}$ whenever $v(G)$ is divisible by $d+1$. When $v(G)$ is not divisible by $d+1$, we can apply the same inequality to a disjoint union of $(d+1)$ copies of G to obtain $i(G)^{1/v(G)} = i((d+1)G)^{1/((d+1)v(G))} \geq i(K_{d+1})^{1/(d+1)}$. ■

The situation changes significantly if we require G to be bipartite. In this case, the problem was solved very recently by Csikvári [**9**], who showed that the infimum of $i(G)^{1/v(G)}$ over d-regular bipartite graphs G is obtained by taking a sequence of G with increasing girth, i.e., G is locally tree-like. The limit of $i(G)^{1/v(G)}$ for a sequence of bipartite d-regular graphs G of increasing girth was determined by Sly and Sun [**29**] using sophisticated (rigorous) methods from statistical physics.

Colorings. Here is the infimum of $\hom(G, K_q)^{1/v(G)}$ over d-regular graphs G due to Csikvári [**10**].

Theorem 8.3. *For a d-regular graph G and any $q \geq 2$,*

$$\hom(G, K_q)^{1/v(G)} \geq \hom(K_{d+1}, K_q)^{1/(d+1)}.$$

Proof. Assume $q \geq d+1$ since otherwise the right-hand side is zero. Let σ be a random permutation of $V(G)$. For each $u \in V(G)$, let d_u^σ denote the number of neighbors of u that appears before u in the permutation σ. By coloring the vertices in the order of σ, there are at least $q - d_u^\sigma$ choices for the color of vertex u, so

$$\hom(G, K_q) \geq \prod_{u \in V(G)} (q - d_u^\sigma).$$

Taking the logarithm of both sides, we find that

$$\frac{1}{v(G)} \log \hom(G, K_q) \geq \frac{1}{v(G)} \sum_{u \in V(G)} \log(q - d_u^\sigma). \tag{13}$$

For each $u \in V(G)$, the random variable d_u^σ is uniformly distributed in $\{0, 1, \ldots, d\}$ since the ordering of $u \cup N(u)$ under σ is uniform. Therefore, the expected value of the right-hand side of (13) is

$$\frac{1}{d+1}(\log q + \log(q-1) + \cdots + \log(q-d)) = \frac{1}{d+1}\log \hom(K_{d+1}, K_q),$$

which proves the theorem. ■

What is the infimum of $\hom(G, K_q)^{1/v(G)}$ over *bipartite* d-regular graphs G? The following inequality was proved by Csikvári and Lin [**11**]. For $q \geq d+1$, the constant in the inequality is best possible as it is the limit for any sequence of d-regular graphs with increasing girth [**2**].

Theorem 8.4 ([11]). *For any d-regular bipartite graph G and any $q \geq 2$,*

$$\hom(G, K_q)^{1/v(G)} \geq q(1-1/q)^{d/2}.$$

Widom–Rowlinson model. In the previous two cases, for independent sets and colorings, the minimizing G is K_{d+1}, and if we restrict to bipartite G, the "minimizing" G is locally tree-like. For the Widom–Rowlinson model, we saw in Theorem 2.8 that the quantity hom(G, )$^{1/v(G)}$ is maximized, over d-regular graphs G, by $G = K_{d+1}$. Csikvári [**9**] recently showed that hom(G, )$^{1/v(G)}$ is minimized, over d-regular graphs G, by a sequence of graphs G with increasing girth, even without the bipartite assumption on G.

9. RELATED RESULTS AND FURTHER QUESTIONS.

Independent sets of fixed size. We saw in Theorem 1.1 that in the family of d-regular graphs on n vertices, a disjoint union of $K_{d,d}$'s maximizes the number of independent sets. It is conjectured that the latter maximizes the number of independent sets of every fixed size. Let $i_t(G)$ denote the number of independent sets of size t in G. Recall that kG denotes a disjoint union of k copies of G. See [**13**, Section 8] for the current best bounds on this problem.

Conjecture 9.1 ([23]). *If G is a d-regular graph with $2ad$ vertices, then $i_t(G) \leq i_t(aK_{d,d})$ for every t.*

Graphs with given degree profile. Kahn [**23**] made the following conjecture extending Theorem 1.1 to irregular graphs. We write d_u for the degree of vertex $u \in V(G)$.

Conjecture 9.2 ([23]). *For any graph G,*

$$i(G) \leq \prod_{uv \in E(G)} i(K_{d_u,d_v})^{1/d_u d_v} = \prod_{uv \in E(G)} (2^{d_u} + d^{d_v} - 1)^{1/(d_u d_v)}.$$

By the bipartite reduction in Section 4, it suffices to prove the conjecture for bipartite graphs G. Galvin and Zhao [**22**] proved Conjecture 9.2 for all G with maximum degree at most 5. The following conjecture, due to Galvin [**18**],[4] extends Theorem 2.4 and the bipartite case of Conjecture 9.2.

Conjecture 9.3. *For any bipartite graph G and loop-graph H,*

$$\hom(G, H) \leq \prod_{uv \in E(G)} \hom(K_{d_u,d_v}, H)^{1/(d_u d_v)}.$$

[4] A bipartite assumption on G is missing in [**18**, Conjecture 1.5].

Graphs with additional local constraints. We saw in Theorem 1.1 and Theorem 8.1 that the maximum and minimum of $i(G)^{1/v(G)}$ among d-regular graphs G are attained by $K_{d,d}$ and K_{d+1}, respectively. What if we impose additional "local" constraints to disallow $K_{d,d}$ and K_{d+1}? For example, consider the following.

- What is the infimum of $i(G)^{1/v(G)}$ among d-regular triangle-free graphs G?
- What is the supremum of $i(G)^{1/v(G)}$ among d-regular graphs G that do not contain any cycles of length 4?

These two questions were recently answered by Perarnau and Perkins [**26**].

Theorem 9.4. *(a) Among 3-regular triangle-free graphs G, the quantity $i(G)^{1/v(G)}$ is minimized when G is the Petersen graph.*

(b) Among 3-regular graphs G without cycles of length 4, the quantity $i(G)^{1/v(G)}$ is maximized when G is the Heawood graph.

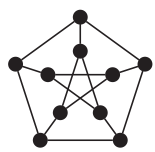

Petersen graph

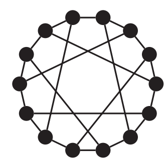

Heawood graph

Theorem 9.4 was proved using the occupancy method discussed in Section 7. The following general problem is very much open.

Problem 9.5. *Let $d \geq 3$ be an integer and $\mathcal{F}$ be a finite list of graphs. Determine the infimum and supremum of $i(G)^{1/v(G)}$ among d-regular graphs G that do not contain any element of $\mathcal{F}$ as an induced subgraph.*

We pose the following (fairly bold) conjecture that the extrema are always attained by finite graphs. It would be interesting to know which graphs can arise as extremal graphs in this manner.

Conjecture 9.6 (Local constraints imply bounded extrema). *Let $d \geq 3$ be an integer and $\mathcal{F}$ be a finite list of graphs. Let $\mathcal{G}_d(\mathcal{F})$ denote the set of finite d-regular graphs that do not contain any element of $\mathcal{F}$ as an induced subgraph. Then there exist $G_{\min}, G_{\max} \in \mathcal{G}_d(\mathcal{F})$ such that for all $G \in \mathcal{G}_d(\mathcal{F})$,*

$$i(G_{\min})^{1/v(G_{\min})} \leq i(G)^{1/v(G)} \leq i(G_{\max})^{1/v(G_{\max})}.$$

ACKNOWLEDGMENTS. The author is grateful to Joe Gallian for the REU opportunity in 2009 where he began working on this problem (resulting in [**30**]). He thanks Péter Csikvári, David Galvin, Joonkyung Lee, Will Perkins, and Prasad Tetali for carefully reading a draft of this paper and providing helpful comments, as well as the anonymous reviewers for suggestions that improved the exposition of the paper.

REFERENCES

1. N. Alon, Independent sets in regular graphs and sum-free subsets of finite groups, *Israel J. Math.* **73** (1991) 247–256.
2. A. Bandyopadhyay, D. Gamarnik, Counting without sampling: Asymptotics of the log-partition function for certain statistical physics models, *Random Structures Algorithms* **33** (2008) 452–479.

3. H. J. Brascamp, E. H. Lieb, Best constants in Young's inequality, its converse, and its generalization to more than three functions, *Advances Math.* **20** (1976) 151–173.
4. G. R. Brightwell, O. Häggström, P. Winkler, Nonmonotonic behavior in hard-core and Widom-Rowlinson models, *J. Statist. Phys.* **94** (1999) 415–435.
5. G. R. Brightwell, P. Winkler, Graph homomorphisms and phase transitions, *J. Combin. Theory Ser. B* **77** (1999) 221–262.
6. F. R. K. Chung, R. L. Graham, P. Frankl, J. B. Shearer, Some intersection theorems for ordered sets and graphs, *J. Combin. Theory Ser. A* **43** (1986) 23–37.
7. E. Cohen, P. Csikvári, W. Perkins, P. Tetali, The Widom–Rowlinson model, the hard-core model and the extremality of the complete graph, *European J. Combin.* **62** (2017) 70–76.
8. E. Cohen, W. Perkins, P. Tetali, On the Widom–Rowlinson occupancy fraction in regular graphs, *Combin. Probab. Comput.* **26** (2017) 183–194.
9. P. Csikvári, Extremal regular graphs: The case of the infinite regular tree (2016), arXiv:1612.01295.
10. ———, personal communication.
11. P. Csikvári, Z. Lin, Sidorenko's conjecture, colorings and independent sets, *Electron. J. Combin.* **24** (2017) P1.2.
12. J. Cutler, A. J. Radcliffe, The maximum number of complete subgraphs in a graph with given maximum degree, *J. Combin. Theory Ser. B* **104** (2014) 60–71.
13. E. Davies, M. Jenssen, W. Perkins, B. Roberts, Independent sets, matchings, and occupancy fractions (2015), arXiv:1508.04675.
14. ———, On the average size of independent sets in triangle-free graphs, *Proc. Amer. Math. Soc.*, to appear.
15. ———, Extremes of the internal energy of the Potts model on cubic graphs, *Random Structures Algorithms*, to appear.
16. H. Finner, A generalization of Hölder's inequality and some probability inequalities, *Ann. Probab.* **20** (1992) 1893–1901.
17. E. Friedgut, Hypergraphs, entropy, and inequalities, *Amer. Math. Monthly* **111** (2004) 749–760.
18. D. Galvin, Bounding the partition function of spin-systems, *Electron. J. Combin.* **13** (2006) 11.
19. ———, Maximizing H-colorings of a regular graph, *J. Graph Theory* **73** (2013) 66–84.
20. ———, Three tutorial lectures on entropy and counting (2014), arXiv:1406.7872.
21. D. Galvin, P. Tetali, On weighted graph homomorphisms, in *Graphs, Morphisms and Statistical Physics*. DIMACS Ser. Discrete Math. Theoret. Comput. Sci, Vol. 63, American Mathematical Society, Providence, RI, 2004. 97–104.
22. D. Galvin, Y. Zhao, The number of independent sets in a graph with small maximum degree, *Graphs Combin.* **27** (2011) 177–186.
23. J. Kahn, An entropy approach to the hard-core model on bipartite graphs, *Combin. Probab. Comput.* **10** (2001) 219–237.
24. L. H. Loomis, H. Whitney, An inequality related to the isoperimetric inequality, *Bull. Amer. Math. Soc* **55** (1949) 961–962.
25. E. Lubetzky, Y. Zhao, On replica symmetry of large deviations in random graphs, *Random Structures Algorithms* **47** (2015) 109–146.
26. G. Perarnau, W. Perkins, Counting independent sets in cubic graphs of given girth, preprint (2016), arXiv:1610.08496.
27. W. Perkins, personal communication.
28. L. Sernau, Graph operations and upper bounds on graph homomorphism counts, *J. Graph Theory*, to appear.
29. A. Sly, N. Sun, Counting in two-spin models on d-regular graphs, *Ann. Probab.* **42** (2014) 2383–2416.
30. Y. Zhao, The number of independent sets in a regular graph, *Combin. Probab. Comput.* **19** (2010) 315–320.
31. ———, The bipartite swapping trick on graph homomorphisms, *SIAM J. Discrete Math.* **25** (2011) 660–680.

YUFEI ZHAO received his B.Sc. and Ph.D. from MIT in 2010 and 2015 respectively, and his M.A.St. from Cambridge University in 2011. He is currently an Assistant Professor of Mathematics at MIT, and was previously the Esmée Fairbairn Junior Research Fellow in Mathematics at New College, Oxford. His mathematical research interests include extremal, probabilistic, and additive combinatorics.
Department of Mathematics, MIT, Cambridge, MA 02139, USA
yufeiz@mit.edu