

Counting proper colourings in 4-regular graphs via the Potts model

Ewan Davies

Korteweg-de Vries Institute for Mathematics
University of Amsterdam
1090 GE Amsterdam, The Netherlands
e.s.d.davies@uva.nl

Submitted: Mar 23, 2018; Accepted: Sep 19, 2018; Published: Oct 5, 2018

© The author. Released under the CC BY-ND license (International 4.0).

Abstract

We give tight upper and lower bounds on the internal energy per particle in the antiferromagnetic q -state Potts model on 4-regular graphs, for $q \geq 5$. This proves the first case of a conjecture of the author, Perkins, Jenssen and Roberts, and implies tight bounds on the antiferromagnetic Potts partition function.

The zero-temperature limit gives upper and lower bounds on the number of proper q -colourings of 4-regular graphs, which almost proves the case $d = 4$ of a conjecture of Galvin and Tetali. For any $q \geq 5$ we prove that the number of proper q -colourings of a 4-regular graph is maximised by a union of $K_{4,4}$'s.

Mathematics Subject Classifications: 05C15, 82B20, 90C35

1 Introduction

The Potts model is a probabilistic model on assignments of interacting colours to the vertices of a graph. A temperature parameter controls how strongly the colours interact, and in the zero-temperature limit valid assignments of colours are proper q -colourings of the graph. For a graph G we call a map $\sigma : V(G) \rightarrow [q]$ a q -colouring of G , and we let $m(\sigma)$ denote the number of monochromatic edges of G under σ . The usual graph-theoretic notion of a *proper colouring* is then such a map σ with zero monochromatic edges. The q -state Potts model partition function is

$$Z_G^q(\beta) = \sum_{\sigma: V(G) \rightarrow [q]} e^{-\beta m(\sigma)},$$

where β is the *inverse temperature* and for $\beta > 0$ we say the model is *antiferromagnetic*. For our purposes, this will mean that increasing β biases the Potts model towards colourings with few monochromatic edges.

Using this partition function, the Potts model [21] (with no external field) is a random q -colouring σ of G , chosen according to the distribution

$$\sigma \mapsto \frac{e^{-\beta m(\sigma)}}{Z_G^q(\beta)}.$$

Note that for $q = 2$ this is known as the *Ising model* [13]. See [19, 24] for more details on the Potts model.

There is a strong link between the Potts model and extremal combinatorics, as in the limit $\beta \rightarrow +\infty$, the model converges to the uniform distribution on proper q -colourings of a graph, and $Z_G^q(\beta)$ tends to $c^q(G)$, the number of proper q -colourings of G . Extremal questions on the number of q -colourings have received a lot of attention; see e.g. [14, 15] for examples in dense graphs and [25] for a survey (which also covers related problems) on regular graphs. A conjecture of Galvin and Tetali [10] states that over all d -regular graphs G and $q \geq 3$, we have

$$\frac{1}{|V(G)|} \log c^q(G) \leq \frac{1}{2d} \log c^q(K_{d,d}). \quad (1)$$

Note that the case $q = 2$ follows easily from the observation that $K_{d,d}$ is the smallest d -regular graph which is bipartite (that is, which admits a proper 2-colouring).

In this paper we obtain the above inequality via the Potts model, and a natural quantity to consider is the *free energy per particle* $F_G^q(\beta) := \frac{1}{|V(G)|} \log Z_G^q(\beta)$. The normalisation is convenient if one wishes to compare the number of q -colourings in graphs on different numbers of vertices. Inequality (1) is the zero-temperature version of

$$F_G^q(\beta) \leq F_{K_{d,d}}^q(\beta), \quad (2)$$

since in the limit $\beta \rightarrow +\infty$, inequality (2) becomes exactly (1).

For general d , Galvin [9, 11] showed that (2) holds for all β when G is bipartite, and (1) holds for all d when $q \geq 2 \binom{d+V(G)/2}{4}$. Recently, for $d = 3$ both (1) and (2) were shown to hold for all q and $\beta > 0$ by the author, Perkins, Jenssen, and Roberts [5], using somewhat different techniques than in previous work.

The main difference between the techniques of [5] and previous work is the focus on an *observable* (expectation over the Potts model) rather than the partition function itself. The expected number of monochromatic edges indeed an observable, and we scale by the number of vertices to obtain the *internal energy per particle* $U_G^q(\beta)$, where

$$\begin{aligned} U_G^q(\beta) &= \frac{1}{|V(G)|} \mathbb{E}[m(\sigma)] \\ &= \frac{1}{|V(G)|} \frac{\sum_{\sigma: V(G) \rightarrow [q]} m(\sigma) e^{-\beta m(\sigma)}}{Z_G^q(\beta)} \\ &= -\frac{1}{|V(G)|} \frac{1}{Z_G^q(\beta)} \frac{\partial}{\partial \beta} Z_G^q(\beta) \\ &= -\frac{1}{|V(G)|} \frac{\partial}{\partial \beta} \log Z_G^q(\beta) = -\frac{\partial}{\partial \beta} F_G^q(\beta). \end{aligned}$$

When $\beta = 0$ there are no interactions between colours and $Z_G^q(0) = q^{|V(G)|}$ for all G . Then to obtain the free energy per particle we can integrate the internal energy per particle,

$$F_G^q(\beta) = \log q - \int_0^\beta U_G^q(t) dt. \quad (3)$$

In [5], tight bounds on $U_G^q(\beta)$ were given for 3-regular graphs for all $\beta > 0$, leading to (2) and (1) for $d = 3$ and all $q, \beta > 0$. Here we extend the methods to $d = 4$ under the additional restriction that $q \geq 5$, almost settling the $d = 4$ case of (1), the conjecture of Galvin and Tetali, as only the cases $q \in \{3, 4\}$ remain.

Theorem 1. *For any 4-regular graph G , any $q \geq 5$, and any $\beta > 0$,*

$$U_{K_{4,4}}^q(\beta) \leq U_G^q(\beta) \leq U_{K_5}^q(\beta).$$

Equalities hold respectively if and only if G is a union of $K_{4,4}$'s or a union of K_5 's. As a corollary via (3), we have

$$F_{K_5}^q(\beta) \leq F_G^q(\beta) \leq F_{K_{4,4}}^q(\beta).$$

This result gives (1) by the following simple argument (also given in [5]).

Lemma 2. *Fix positive integers d and q . If for all d -regular graphs G , and all $\beta > 0$, we have $U_G^q(\beta) \geq U_{K_{d,d}}^q(\beta)$, then*

$$\frac{1}{|V(G)|} \log c^q(G) \leq \frac{1}{2d} \log c^q(K_{d,d}).$$

Proof. Let G be any d -regular graph. If $c^q(G) = 0$ then the conclusion clearly holds. Otherwise, we take logarithms and write

$$\begin{aligned} \frac{1}{|V(G)|} \log c^q(G) &= \lim_{\beta \rightarrow +\infty} \frac{1}{|V(G)|} \log Z_G^q(\beta) \\ &= \log q - \int_0^\infty U_G^q(\beta) d\beta \\ &\leq \log q - \int_0^\infty U_{K_{d,d}}^q(\beta) d\beta \\ &= \frac{1}{2d} \log c^q(K_{d,d}). \end{aligned}$$

This completes the proof. □

Similarly, Theorem 1 also implies the lower bound

$$\frac{1}{|V(G)|} \log c^q(G) \geq \frac{1}{5} \log c^q(K_5)$$

for 4-regular graphs G when $q \geq 5$, but the analogous bound was proved for all d and $q \geq d + 1$ by Bezáková, Štefankovič, Vazirani, and Vigoda [1].

Bounding the number of colourings of a graph via the internal energy per particle of the Potts model has another advantage that was highlighted in [7]. The main advance of [7] is a general method for turning proofs of the kind given below (and similar ones cited throughout) into stability results and tight bounds on the individual coefficients of the partition function. In this context, the most interesting result of these methods is a stability version of the fact that unions of $K_{4,4}$ maximise the number of proper q -colourings (see Lemma 2). For more details, and more general results about the coefficients of the Potts model partition function, one should combine the methods of [7] with our Theorem 1.

To state the stability result we require a little notation. When $2d$ divides n , let $H_{d,n}$ be the graph consisting of $n/(2d)$ copies of $K_{d,d}$, and let $\delta_{\circ}(G, H)$ denote the sampling distance between the bounded degree graphs G and H (see [7, 16]). This is a natural measure of how similar G and H are, and for 4-regular G (in lieu of greater detail) we can loosely interpret $\delta_{\circ}(G, K_{4,4})$ as the proportion of vertices in G that are contained in a $K_{4,4}$.

Corollary 3. *For all $q \geq 5$, there exists $\kappa(q) > 0$ such that for any n divisible by 8 and any 4-regular graph G on n vertices, we have*

$$\frac{1}{n} \log c^q(G) \leq \frac{1}{n} \log c^q(H_{4,n}) - \kappa(q) \cdot \delta_{\circ}(G, K_{4,4})$$

Finally, we suggest that it is not entirely surprising that the extremal problems we discuss are harder for smaller q , as in general much more is known about the Potts model for large q . The methods of Galvin [9] apply when $q \geq 2 \binom{d|V(G)|/2}{4}$, and for example, one can efficiently approximate the number of proper q -colourings when $q \geq 11d/6$ [23] with a suitable randomised algorithm, but this is impossible when $q < d$ unless $\text{NP} = \text{RP}$, as it is NP -hard to determine even whether a single proper q -colouring exists (when $q \geq 3$) [12].

The rest of this paper is organised as follows. In Section 2 we describe the method by which we bound the internal energy of the Potts model, noting that symmetry in the colours reduces the problem to a finite calculation. In section 3 we describe a practical implementation of the method by computer. This section contains the main advance of the paper: an application of efficient graph isomorphism and generation algorithms to the problem of generating the information necessary to carry out the method of Section 2. Finally, in Section 4 we give some concluding remarks.

The computational techniques described in this paper are implemented in Cython¹ code using libraries from the Sage² computer algebra system, version 8.0. The source code is included with the paper as a collection of ancillary files.

2 The method

The general form of the method used here has been applied to problems of a similar nature; in [4, 6, 20] bounds on the number of independent sets or matchings in certain

¹<http://cython.org>

²<https://www.sagemath.org>

classes of graphs are obtained via bounds on an observable of the hard-core or monomer-dimer models, and a similar problem for the Widom–Rowlinson model was solved in [3]. One can also obtain tight bounds on the individual coefficients of the partition function in certain cases [7]. We now describe the method for the Potts model, which is essentially the same in this paper and in [5]. We leave the regularity of the graph d as a parameter, but will only address the case $d = 4$ in the proof.

2.1 Outline

Fix a d -regular graph G , positive integer q , and $\beta > 0$. To bound $U_G^q(\beta)$, we consider the following experiment. Let v be a vertex of G chosen uniformly at random, and independently sample a colouring $\sigma : V(G) \rightarrow [q]$ from the Potts model on G . For each neighbour u of v , record the number of its *external neighbours* (neighbours outside $v \cup N(v)$) receiving each colour and record also any edges within $N(v)$. This gives a *local view* of σ from v . Note that although we have sampled a colouring σ of the whole graph, the colours of v and its neighbours are not revealed. The graph structure amongst the external neighbours is not required, and so formally the local view is a graph on $d + 1$ vertices such that for d of the vertices we have a multiset of externally observed colours, see Figure 1.




Figure 1: A local view where the multisets of observed colours are shown below each neighbour of v .

An important part of the method is that, conditioned on the local view, the distribution of colourings of v and its neighbours can be determined. This is a fundamental feature of the Potts model (and more generally any *Gibbs measure*) known as the spatial Markov property, which can easily be derived from the definition of the distribution and partition function. In our case, for any v and the random colouring σ , conditioned on the colours of the external neighbours of v , the distribution of the colours on $\{v\} \cup N(v)$ is independent of the colours of any vertex at distance greater than 2 from v .

For fixed d and q there is a finite set of possible local views which we denote $\mathcal{L}_{d,q}$. By the above experiment, each d -regular graph G and inverse temperature β induces a probability distribution on $\mathcal{L}_{d,q}$. Not all probability distributions on $\mathcal{L}_{d,q}$ can arise from a graph; there are certain consistency conditions that must hold. Here we consider a family of conditions identified in [5] but not needed for the case $d = 3$. Let $\mathcal{S}_{d,q}$ be the set of q -partitions of d ; that is, partitions of the form $d = a_1 + a_2 + \dots + a_k$ where $a_i \geq 1$ are positive integers and $k \leq q$. Let $x \in V(G)$, and note that any q -colouring of $N(x)$ induces

a q -partition of d which we denote $H(x)$ and compute by counting the frequencies of each colour on $N(x)$. For example, if $N(x)$ is assigned the multiset of colours $[1, 2, 4, 4]$, then $H(x) = 2 + 1 + 1$. The family of consistency conditions that we use to prove the lower bound in Theorem 1 is that for every $S \in \mathcal{S}_{d,q}$, the probability that for v as in the local view, $\mathbb{P}(H(v) = S)$ is the same as the average over $u \in N(v)$ of $\mathbb{P}(H(u) = S)$. This holds for in any regular graph since both v and a uniformly random neighbour are distributed uniformly over $V(G)$. Note that because $a_i \geq 1$ in any partition $d = a_1 + a_2 + \dots + a_k$, we must have $k \leq 4$, and hence $\mathcal{S}_{4,4} = \mathcal{S}_{4,q}$ for any $q \geq 5$. This simply means that our set of consistency conditions is bounded in size independent of q , which is important for our calculations. These consistency conditions are the first way in which the method of this paper and the method of [5] differ. For the $d = 3$ case addressed in [5], simpler consistency conditions suffice which we do not discuss here.

It turns out that $U_G^q(\beta)$ and the probabilities one must compute for these consistency conditions are linear functions of the probabilities that each local view appears in the random experiment. Then to maximise or minimise $U_G^q(\beta)$ over all d -regular graphs, we consider a linear programming relaxation of the problem, optimising over all probability distributions on $\mathcal{L}_{d,q}$ that satisfy the above consistency conditions. For each fixed d , q , and β , this is a finite linear program with $|\mathcal{L}_{d,q}|$ variables, but we wish to solve an infinite family of these programs indexed by q and β , hence we cannot yet reach for a computer.

To prove Theorem 1 we show, separately for the minimum and the maximum, that the distributions corresponding to $K_{4,4}$ and K_5 respectively are feasible in the dual linear program. In [5] it was conjectured that when $d \geq 4$ and $q \geq d + 1$, the linear program above is indeed minimised and maximised by distributions corresponding to $K_{d,d}$ and K_{d+1} respectively, hence this paper confirms that conjecture for $d = 4$.

2.2 Linear programs

In this section we describe the linear programs which are used to prove Theorem 1. Recall the process for sampling a local view: draw a colouring $\sigma : V(G) \rightarrow [q]$ according to the q -colour Potts model with inverse temperature β and independently, uniformly at random choose a vertex $v \in V(G)$. The random local view L consists of the induced subgraph of G on $\{v\} \cup N(v)$, together with, for each $u \in N(v)$, the multiset of colours that appears in $N(u) \setminus (\{v\} \cup N(v))$.

To make the program of finite size independent of q we consider equivalence classes of the local views under permutations of the colours. There are at most $d(d - 1)$ vertices whose colours we record to form the local view, and all of the quantities we need to calculate are invariant under colour permutations, hence for $d = 4$ we need only consider 12 colours in the linear program. Let $\hat{\mathcal{L}}_{d,q}$ be a subset of the possible local views $\mathcal{L}_{d,q}$ that contains exactly one representative from each equivalence class. The construction of $\hat{\mathcal{L}}_{4,q}$, which contains 3529 local views, is one of the computational problems that required a nontrivial implementation to be achievable in a short time. See Section 3 for details.

Suppose that the local view L arises from selecting the colouring σ and the vertex v . We refer to the coloured external neighbours at distance two from v as the *boundary*, and write V_L for the set of uncoloured vertices in L . The colouring σ induces a *local colouring*

$\chi : V_L \rightarrow [q]$ that, by the spatial Markov property of the Potts model, is distributed according to the Potts model on L (which includes some fixed colours on the boundary). For $\chi : V_L \rightarrow [q]$, write $m(\chi)$ for the total number of monochromatic edges in L (including any monochromatic edges between V_L and the boundary), and write $m_v(\chi)$ for the number of monochromatic edges in L incident to v . Then, with the local partition function defined as

$$Z_L^q(\beta) := \sum_{\chi:V_L \rightarrow [q]} e^{-\beta m(\chi)},$$

the random local colouring $\chi : V_L \rightarrow [q]$ is distributed according to

$$\chi \mapsto \frac{e^{-\beta m(\chi)}}{Z_L^q(\beta)}.$$

We can now write $U_G^q(\beta)$ as an expectation over the random local view L . Each edge of G is incident to exactly two vertices, hence

$$\begin{aligned} U_G^q(\beta) &= \frac{1}{|V(G)|} \mathbb{E}_\sigma[m(\sigma)] \\ &= \frac{1}{2|V(G)|} \sum_{v \in V(G)} \sum_{u \in N(v)} \mathbb{P}_\sigma(uv \text{ monochromatic}) \\ &= \frac{1}{2} \mathbb{E}_L \left[\sum_{u \in N(v)} \mathbb{P}_\chi(uv \text{ monochromatic} | L) \right] \\ &= \frac{1}{2} \mathbb{E}_{L,\chi} [m_v(\chi)], \end{aligned}$$

where the subscripts show which random variables are involved in the expectations and probabilities.

We define the vector $\vec{c} \in \mathbb{R}^{\hat{\mathcal{L}}_{4,q}}$ with components

$$\vec{c}_L := \frac{1}{2Z_L^q} \sum_{\chi:V_L \rightarrow [q]} m_v(\chi) e^{-\beta m(\chi)}, \quad (4)$$

which is the objective for both our linear programs.

To constrain the linear programs, recall that we consider the q -partitions $H(v)$ and $H(u)$ that appear in the local view. For $S \in \mathcal{S}_{4,q}$, we have

$$\mathbb{E}_{L,\chi} [\mathbf{1}_{\{H(v)=S\}}] = \mathbb{E}_{L,\chi} \left[\frac{1}{d} \sum_{u \in N(v)} \mathbf{1}_{\{H(u)=S\}} \right],$$

and for use in the linear program we define vectors $\vec{\gamma}^S \in \mathbb{R}^{\hat{\mathcal{L}}_{4,q}}$ for $S \in \mathcal{S}_{4,q}$ with components

$$\vec{\gamma}_L^S := \sum_{\chi:V_L \rightarrow [q]} \left(\mathbf{1}_{\{H(v)=S\}} - \frac{1}{d} \sum_{u \in N(v)} \mathbf{1}_{\{H(u)=S\}} \right) e^{-\beta m(\chi)}, \quad (5)$$

which give constraints for the linear program.

For $L \in \hat{\mathcal{L}}_{4,q}$ and $S \in \mathcal{S}_{4,q}$, the quantities \vec{c}_L and $\vec{\gamma}_L^S$ in (4) and (5) involve sums over local colourings $\chi : V_L \rightarrow [q]$. The fact that these sums grow in size with q is a problem for the general case $q \geq 5$ of Theorem 1, but again we can solve this by considering permutations of the colours. By permuting colours we only need consider L whose boundary colours form an initial segment of $\{1, 2, \dots\}$, and we write q_L for the number of colours used on the boundary of L . Since $|V_L| = d + 1$ we only need $q_L + d + 1$ colours to represent each equivalence class of local colourings χ . That is, for any quantity $f(\chi)$ which is invariant under permutations of the colours that are not used on the boundary of L , we have

$$\sum_{\chi: V_L \rightarrow [q]} f(\chi) = \sum_{\chi \in \Omega_L} z(\chi) f(\chi),$$

where Ω_L is the set of maps $\{\chi : V_L \rightarrow [q_L + d + 1]\}$ which use an initial segment (possibly empty) of the ‘extra colours’ $\{q_L + 1, \dots, q_L + d + 1\}$, and

$$z(\chi) = \binom{q - q_L}{\ell(\chi) - q_L},$$

is the size of the equivalence class of a map χ that uses largest colour $\ell(\chi)$. This idea turns the calculations giving \vec{c} and $\vec{\gamma}^S$ into finite sums whose size is bounded independently of q .

Let $\vec{p} \in \mathbb{R}^{\hat{\mathcal{L}}_{4,q}}$ be the variables of the primal programs, where \vec{p}_L corresponds to the probability that a local view isomorphic to L is sampled in the random process. In standard form, the primal linear programs in we wish to solve are

$$\left. \begin{aligned} U^{\min} &= \min \\ U^{\max} &= \max \end{aligned} \right\} \vec{p} \cdot \vec{c} \quad \text{s.t.} \quad A\vec{p} = \vec{b}, \quad \vec{p} \geq \vec{0}, \quad (6)$$

where A is the $|\hat{\mathcal{L}}_{4,q}| \times (1 + |\mathcal{S}_{4,q}|)$ matrix with first row all ones, and subsequent rows given by $\vec{\gamma}^S$ for $S \in \mathcal{S}_{4,q}$, and \vec{b} is the vector $(1, 0, \dots, 0)$ with exactly $|\mathcal{S}_{4,q}|$ zeros. Each row of A together with corresponding entry from \vec{b} corresponds to a constraint: the first row gives $\sum_L \vec{p}_L = 1$, and subsequent rows give $\sum_L \vec{p}_L \vec{\gamma}_L^S = 0$. For concreteness, we consider the rows of A and entries of \vec{b} to be indexed by the set $\{1\} \cup \mathcal{S}_{4,q}$ with 1 corresponding to the probability distribution constraint $\sum_L \vec{p}_L = 1$ being first.

2.3 Minimising

We solve the minimisation program by considering the dual. Linear programming duality [2] states that

$$U^{\min} = \max \vec{\Delta} \cdot \vec{b} \quad \text{s.t.} \quad A^T \vec{\Delta} \leq \vec{c}, \quad \vec{\Delta} \geq \vec{0},$$

where $\vec{\Delta}$ is a vector of dual variables. Let \vec{p}^* correspond to the distribution of the random local view in $K_{d,d}$, and $U^* = U_{K_{4,4}}^q(\beta)$. We know that U^* is feasible for the primal program,

which gives $U^{\min} \leq U^*$, and if we find a vector of dual variables $\vec{\Delta}^*$ which are feasible for the dual program such that $\vec{\Delta}_1^* = U^*$, then the dual program shows $U^{\min} \geq \vec{\Delta}^* \cdot \vec{b} = U^*$ and hence $U^{\min} = U^*$ as required. Then it suffices to find a dual-feasible vector $\vec{\Delta}^*$ with $\vec{\Delta}_1^* = U^*$.

It turns out we can solve for the dual variables uniquely provided we choose the minimal number of constraints. In general for an optimum supported on k local views, we can set all but k dual variables to zero and solve the subsystem of the dual constraints corresponding to the k local views with equality to obtain the remaining dual variables.

In $K_{d,d}$ every vertex v has exactly $d - 1$ vertices at distance 2 which form the external neighbours of any local view from v , and each $u \in N(v)$ is adjacent to every one of these external neighbours, hence the local views which arise in $K_{d,d}$ with positive probability have the same multiset for each $u \in N(v)$. Up to isomorphism there are $|\mathcal{S}_{d-1,q}|$ such local views. With $d = 4$ and any $q \geq 5$ we have $1 + |\mathcal{S}_{4,q}| = 6$, but only need $|\mathcal{S}_{3,q}| = 3$ constraints. Rather frustratingly there is no set of three constraints that work for the entire range of q , we must consider $q = 5$ and $q \geq 6$ separately. In either case we select the first constraint corresponding to $\sum_L \vec{p}_L = 1$, and two constraints corresponding to q -partitions in $\mathcal{S}_{4,q}$. For $q = 5$ we use the partitions 4 and $2 + 1 + 1$, and for $q \geq 6$ we use $2 + 1 + 1$ and $1 + 1 + 1 + 1$. In both cases we solve for the corresponding dual variables and verify dual feasibility.

In each case, we have a set \mathcal{L}^* of three local views which arise in $K_{4,4}$, the probability distribution constraint, and two q -partitions $S_1, S_2 \in \mathcal{S}_{4,q}$ giving the remaining constraints. We set $\vec{\Delta}_1^* = U^*$, $\vec{\Delta}_S^* = 0$ for $S \in \mathcal{S}_{4,q} \setminus \{S_1, S_2\}$, and find $\vec{\Delta}_{S_1}^*, \vec{\Delta}_{S_2}^*$ by solving the dual constraints corresponding to $L \in \mathcal{L}^*$ with equality. That is, we take the rows of $A^T \vec{\Delta}^* \leq \vec{c}$ corresponding to $L \in \mathcal{L}^*$, and solve the system of equations given by

$$U^* + \vec{\Delta}_{S_1}^* \vec{\gamma}_L^{S_1} + \vec{\Delta}_{S_2}^* \vec{\gamma}_L^{S_2} = \vec{c}_L \quad \forall L \in \mathcal{L}^* .$$

Note that we actually have three equations for the two unknowns $\vec{\Delta}_{S_1}^*$ and $\vec{\Delta}_{S_2}^*$, but the system is indeed consistent.

After the vector $\vec{\Delta}^*$ has been found, dual feasibility reduces to the statement that the entries of the *slack vector* $\vec{c} - A^T \vec{\Delta}^*$ are non-negative for all L . In fact we show that the optimising distribution \vec{p}^* is unique by complementary slackness: we verify that the slack is strictly positive for all $L \notin \mathcal{L}^*$. Since unions of $K_{4,4}$'s are the only graphs that have distribution \vec{p}^* , these graphs are the only graphs attaining equality in the corresponding bounds in Theorem 1 and (1).

We verify this on a computer by substituting $e^\beta = 1 + t$ and either $q = 5$ or $q = r + 6$, and showing that for some ‘magic factor’ $M > 0$ which is independent of L , the quantity

$$M \cdot \left(\vec{c}_L - U^* - \vec{\Delta}_{S_1}^* \vec{\gamma}_L^{S_1} - \vec{\Delta}_{S_2}^* \vec{\gamma}_L^{S_2} \right) , \tag{7}$$

is a polynomial with positive coefficients in t or both t and r , for all L which do not occur in $K_{4,4}$. Note that $\beta > 0$ corresponds to the range $t > 0$, and $q \geq 6$ corresponds to $r \geq 0$. The magic factor depends on the case $q = 5$ or $q \geq 6$, and was found by hand. For $q = 5$

we take

$$M = 6t^2(t+1)^{31}(2t^2+6t+5)(16t^{10}+176t^9+888t^8+2676t^7+5309t^6+7260t^5+6996t^4+4760t^3+2224t^2+660t+100),$$

and for $q \geq 6$ we have

$$\begin{aligned} M = & 2(r+3)(r+4)(t+1)^{35}(r^4t^9+9r^4t^8+36r^4t^7+84r^4t^6+126r^4t^5+ \\ & 126r^4t^4+84r^4t^3+36r^4t^2+9r^4t+r^4+11r^3t^9+105r^3t^8+ \\ & 447r^3t^7+1117r^3t^6+1809r^3t^5+1971r^3t^4+1445r^3t^3+ \\ & 687r^3t^2+192r^3t+24r^3+39r^2t^9+411r^2t^8+1926r^2t^7+ \\ & 5286r^2t^6+9393r^2t^5+11241r^2t^4+9090r^2t^3+4806r^2t^2+ \\ & 1512r^2t+216r^2+51rt^9+645rt^8+3477rt^7+10715rt^6+ \\ & 21096rt^5+27816rt^4+24812rt^3+14580rt^2+5184rt+ \\ & 864r+18t^9+342t^8+2250t^7+7954t^6+17508t^5+ \\ & 25428t^4+24888t^3+16200t^2+6480t+1296). \end{aligned}$$

Given M it is easy to use a computer algebra system to verify the assertion that (7) is a polynomial in t with non-negative coefficients, and is zero if and only if L arises in $K_{4,4}$. Indeed, for each L the verification is straightforward by hand, but requires lengthy calculations more suited to a computer.

2.4 Maximising

Note that K_5 itself is a valid local view that has no boundary, hence the optimum in the maximisation program is supported on a single local view. This means we wish to use a single constraint in the program, namely that $\sum_L p_L = 1$. It suffices to verify that the largest entry of \vec{c} corresponds to $L = K_5$, which we do by showing that for another ‘magic factor’ $M' = 2 * (t+1)^{25}$, and substitutions $e^b = 1+t$, $q = r+5$, the quantity

$$M' \cdot (\vec{c}_{K_5} - \vec{c}_L)$$

is a polynomial in t and r with positive coefficients for all $L \neq K_5$. Again, this is simple to verify with a computer.

3 Generating local views

The method of the previous section reduces an optimisation problem over the infinite set of 4-regular graphs to the verification of a finite number of inequalities, which can be done by computer. The main difficulty in establishing Theorem 1 is now to enumerate all local views so that this verification can be performed.

3.1 Representing local views as a simple graph

Formally, the local view is a graph on $d + 1$ vertices where an identified vertex v has degree d and for the other vertices we record a multiset of observed colours. For ease of use with standard graph isomorphism algorithms, we represent the entire local view as a simple graph with no coloured vertices. A local view already gives us the graph structure on $\{v\} \cup N(v)$, and we add to this graph a set B of pendant boundary vertices to the neighbours of v . Each neighbour u of v is degree 4 in the original graph, so we add enough pendant vertices such that this holds in the representation of the local view. Now we can represent the multiset of observed colours by colouring the pendant boundary vertices, but we wish to have a simple graph representation that uses no colours. So to represent the colours of the boundary we add a set C of auxiliary *colour vertices* and an edge from each $b \in B$ to exactly one $c \in C$ to represent b being coloured by colour c . Note that with this definition, any local view in $\hat{\mathcal{L}}_{4,q}$ can be represented by with at most $|B| \leq 12$ colour vertices. See Figure 2 for the simple graph representation of the local view pictured in Figure 1. We denote by L_{simp} this representation of $L \in \hat{\mathcal{L}}_{q,d}$.




Figure 2: A representation of the same local view as in Figure 1, but where colours on boundary vertices are indicated by edges to auxiliary colour vertices. There are 10 boundary vertices, hence we draw 10 colour vertices, 7 of which are unused.

We will enumerate the simple representations of local views as follows.

1. Start with a copy of $K_{1,d}$ which is the vertex v and its d neighbours.
2. Add any collection of edges inside $N(v)$, and add pendant boundary vertices to each $u \in N(v)$ such that each $u \in N(v)$ has degree d .
3. Add a set of colour vertices C of size $|B|$, and colour the boundary vertices by adding, for each $b \in B$, exactly one edge to a vertex in C .

Provided the possible edges inside $N(v)$ and the possible colourings are enumerated exhaustively in steps (2) and (3), this process gives at least one representative of every

possible local view and is easy to implement on a computer. Note that by virtue of being represented on a computer we are forced to consider a labelling of the vertices, so a naive implementation of step (2) involves generating every *labelled graph* on d vertices. Similarly, a naive implementation of step (3) will consider all functions from the boundary B to the colour vertices C . This is highly inefficient³ since the quantities we are interested in are independent of the labelling of $N(v)$ and invariant under permutations of the colours. Exploiting these symmetries drastically reduces the computation required for steps (2) and (3).

Using the simple graph representation, the generation of $\hat{\mathcal{L}}_{d,q}$ for any d is now an instance of a well-studied general problem: we want a single labelled representative of each isomorphism class of some unlabelled combinatorial objects (the graphs L_{simp} for $L \in \hat{\mathcal{L}}_{d,q}$) under the action of some symmetry group (permuting the neighbours of v and permuting colours). In this paper we use a canonical generation algorithm of McKay [18] to achieve this efficiently, and an advantage of the simple graph representation is that all the symmetries we must consider arise as graph automorphisms, facilitating the use of pre-existing routines for such problems.

3.2 Generating graphs

We first give an efficient version of step (2) of the local view generation algorithm. We want to generate a single labelled version of each graph on d vertices in order to enumerate the possible underlying graphs of the local views.

There is a natural ‘construction tree’ for this process. The root is the empty labelled graph on d vertices, and the children of any node X are the labelled graphs obtained from X by adding one edge. Every leaf in the tree is therefore the complete labelled graph on d vertices. McKay’s method works by generating only those X that have a ‘canonical construction path’ in this tree. A key ingredient is a *canonical labelling function* that, given a labelled graph, outputs an isomorphic labelled graph, and which is constant on equivalence classes of labelled graphs. Such a function can be used to identify a single, canonical representative of an unlabelled graph G : given an arbitrary initial labelling of G , the function produces a fixed labelling that depended only on G and not the initial labelling. These functions are well-studied, since any such function solves the *graph isomorphism problem*: testing whether two graphs are isomorphic reduces to checking equality of their corresponding canonical labellings.

For completeness, we describe a basic example of a canonical labelling function due to Read [22], and independently Faradžev [8], here. For practical purposes we use a more efficient algorithm of McKay [17] which exploits much more graph structure, specifically the implementation supplied with Sage 8.0. The Read–Faradžev function relates each labelled graph with vertex set $[n]$ to the binary sequence $a(G)$ of length $\binom{n}{2}$, where $a(G)_i = 1$ if the i -th pair in the lexicographic ordering of $\binom{[n]}{2}$ is an edge of G , and $a(G)_i = 0$

³For example, when $d = 4$ each triangle-free local view has 12 boundary vertices to colour with up to 12 colours. There are $12^{12} = 8\,916\,100\,448\,256$ such colourings, but in fact only 1636 isomorphism classes of triangle-free local view.

otherwise. This gives an ordering \prec where $G \prec H$ when $a(G) < a(H)$ in the lexicographic ordering of binary sequences. The canonical labelling of G is then the labelling which is maximal in the ordering \prec , and can be computed by checking every permutation of the vertices of G .

Armed with a canonical labelling function, we give in Figure 3 a short description of McKay's algorithm for canonical graph generation. When $\text{SCAN}(\overline{K}_d)$ is called, the algorithm explores a subtree of the construction tree for labelled graphs on d vertices and outputs exactly one labelled copy of each graph on d vertices.

```

procedure SCAN( $X$ )
  yield  $X$ 
  if  $X = K_d$  then return
  let  $U^*$  contain one representative of each isomorphism class
    of non-edges of  $X$  under the action of  $\text{Aut}(X)$ 
  for  $e \in U^*$  do
    let  $Y = X + e$ , and  $Y^*$  be the canonical labelling of  $Y$ 
    let  $y^*$  be the last edge of  $Y^*$ , and  $y$  be its relabelling in  $Y$ 
    let  $m(Y)$  be the orbit of  $y$  under the action of  $\text{Aut}(Y)$ 
    if  $e \in m(Y)$  then SCAN( $Y$ )

```

Figure 3: A canonical augmentation algorithm which yields a single labelled representative of all unlabelled graphs on d vertices.

As one can see from the algorithm, an explored node X is output (yielded) and then a set of representatives U^* of possible ‘next edges’ is computed. This can be done by computing the orbit of every possible next edge under the action of $\text{Aut}(X)$, and picking the least edge of each orbit in the lexicographic order on $\binom{[d]}{2}$. Here is the first use of symmetry to prune the construction tree: we only consider one representative of each equivalence class of ‘next edge’ e . The next step is to ensure that we only recursively explore the subtree rooted at $Y = X + e$ when e is, up to isomorphism, the ‘last edge’ of Y . Here we take the ‘last edge’ of Y to be the last edge of the canonical labelling Y^* of Y in the lexicographic order on $\binom{[d]}{2}$. Thus, we observe that the algorithm explores a subtree of the construction tree. We omit a more detailed analysis of the algorithm that would constitute a full proof that that exactly one labelled copy of each unlabelled graph is generated, see [18] for details.

3.3 Colouring graphs

The same algorithm can be used to generate colourings by restricting the edges one adds to the graph. For local views, the idea is to start with the necessary graph structure on $\{v\} \cup N(v) \cup B$ with the colour vertices C added as isolated vertices. Then colouring a vertex $b \in B$ with colour $c \in C$ corresponds to adding an edge between b and c in the graph. For brevity we say a vertex in B is coloured to mean it has exactly one edge from it to a vertex in C .

This again gives a construction tree: the root is a graph of the form L_{simp} but with no edges from B to C , and the children of each node are the graphs obtained by adding exactly one edge from an uncoloured boundary vertex to a colour vertex. The algorithm in Figure 3 can be slightly modified to output produce exactly one labelled representative of every possible colouring of the initial graph structure. Here we write $\text{Aut}(X)$ to mean the subgroup of graph automorphisms of the local view representation X that fixes $\{v\}$, $N(v)$, B , and C as sets.

```

procedure COLOUR( $X$ )
  if all vertices in  $B$  are coloured then
    yield  $X$ 
  return
  let  $U^*$  contain one representative of each isomorphism class
    of non-edges from uncoloured boundary vertices to colour
    vertices under the action of the subgroup of  $\text{Aut}(X)$ 
  for  $e \in U^*$  do
    let  $Y = X + e$ , and  $Y^*$  be the canonical labelling of  $Y$ 
    let  $y^*$  be the last edge of  $Y^*$ , and  $y$  be its relabelling in  $Y$ 
    let  $m(Y)$  be the orbit of  $y$  under the action of  $\text{Aut}(Y)$ 
    if  $e \in m(Y)$  then COLOUR( $Y$ )

```

Figure 4: A canonical augmentation algorithm for colouring the boundary of local views.

The changes compared to Figure 3 are that yielding X is only done when all boundary vertices are coloured, and that the edges we consider adding to X are only those from uncoloured boundary vertices to C .

3.4 Notes

In this section we justify some features of the above computational methods that might seem peculiar to the experienced user of the algorithms of McKay.

Our representation of local view seems wasteful in that to represent the externally observed colours for each $u \in N(v)$ we could omit the pendant boundary vertices and record edges from u to auxiliary colour vertices. To obtain this representation from the one given above one simply contracts the pendant edges from $u \in N(v)$ to its adjacent boundary vertices. This would be a more compact representation, but we would require in some cases multiple edges from the neighbours of v to auxiliary colour vertices. For ease of use with existing implementations of algorithms and data structures that do not allow for multiple edges between two vertices in a graph, we chose the representation given above.

In addition, McKay's work [17] on isomorphisms naturally allows for coloured vertices to be considered, and can be used to find automorphisms of labelled graphs that preserves the colour classes as sets. The reason we do not use this facility directly here is that in the algorithms of McKay one does not allow automorphisms that exchange entire colour

classes. The trick we use here with auxiliary colour vertices that can be exchanged (amongst themselves) is a standard way of getting around this problem. When calling McKay's algorithm as in Figure 4, we use four McKay-style non-interchangeable colour classes $\{v\}$, $N(v)$, B , and C , so that the group $\text{Aut}(X)$ computed there fixes these sets, as needed for our application.

4 Conclusions

This paper demonstrates that the approach conjectured in [5] works for $d = 4$, and gives reasonably efficient implementations of computationally expensive tasks that are required for the proof. The fact that the minimisation program used here (and first discussed in [5]) has optimum given by $K_{4,4}$ as desired is some evidence that the approach may work in general for $d \geq 3$ and $q \geq d + 1$.

We leave for further work the tasks of trying $d \geq 5$ with this method, noting that the computational techniques described here make no essential use of $d = 4$ and can be applied to larger d immediately, albeit slowly. For completeness, we repeat the relevant conjecture from [5] here, noting that it is open for $d \geq 5$.

Conjecture 4 (Davies, Jenssen, Perkins, Roberts). The minimisation program (6) (see Section 2.2) has optimum given by $K_{d,d}$ for $d \geq 3$, $q \geq d + 1$ and $\beta > 0$, which shows that $U_{K_{d,d}}^q(\beta) \leq U_G^q(\beta)$ for all d -regular graphs G .

For the proof of the case $d = 4$ of this conjecture we had to solve several difficulties which were not present in the case $d = 3$. Firstly, the computations required to solve the problem with this method require substantial computing power, and were only feasible due to the careful choice of efficient algorithms such as SCAN to generate $\hat{\mathcal{L}}_{d,q}$. We are able to compute that (for $q \geq 20$) $\hat{\mathcal{L}}_{5,q}$ contains approximately 7 million local views, compared to only 3529 in $\hat{\mathcal{L}}_{4,q}$. It seems that even computing the coefficients for the linear program for $d = 5$ is computationally infeasible today. Secondly, there is not a canonical choice of the minimal set of the $\mathcal{S}_{4,q}$ constraints that gives the correct result for all $q \geq 5$, and no apparent reason why some subsets (of the correct size) work, and some do not for $q = 5$ or for $q \geq 6$. It took a substantial amount of experimentation to find the proofs presented here, and for larger values of d such experiments would require even more expensive computation.

Fortunately, there is still hope that these computations might be avoidable. In [4] a general solution for an analogous problem on matchings was given, with an inductive description (in d) of the dual variables and proof that the $K_{d,d}$ yields a dual feasible solution. For such a proof here, one would need to find a description of the minimal set of constraints to use, solve for the dual variables, and show dual feasibility in general. These appear to be challenging tasks, but now that solutions exist for $d \in \{3, 4\}$ perhaps it will be possible to spot the beginnings of a pattern.

It would be interesting to settle the 4-regular case of the conjecture of Galvin and Tetali [10] in full; one still has to prove that (1) holds for all 4-regular G when $q \in \{3, 4\}$ (though this is known for bipartite graphs [10]). The linear program given here does not

give $U^{\min} = U_{K_{4,4}}^q(\beta)$ for such q and large β , hence one must try a stronger program, or perhaps a different approach. There are natural ways to strengthen the methods given here, such as adding more constraints (if they exist), or taking larger local views: e.g. to include vertices at distance at most 2 from v with the colours of vertices at exactly distance 3 revealed to form the boundary.

Acknowledgements

The author would like to thank Will Perkins and Matthew Jenssen for their careful reading of a draft of this article, and an anonymous referee for comments.

The research leading to these results has received funding from the European Research Council under the European Union’s Seventh Framework Programme (FP7/2007-2013) / ERC grant agreement № 339109.

References

- [1] I. Bezáková, D. Štefankovič, V. V. Vazirani, and E. Vigoda. Accelerating simulated annealing for the permanent and combinatorial counting problems. *SIAM Journal on Computing*, 37(5):1429–1454, 2008. doi:10.1137/050644033.
- [2] S. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge University Press, 2004.
- [3] E. Cohen, W. Perkins, and P. Tetali. On the Widom–Rowlinson occupancy fraction in regular graphs. *Combinatorics, Probability and Computing*, 26(2):183–194, 2017. doi:10.1017/S0963548316000249.
- [4] E. Davies, M. Jenssen, W. Perkins, and B. Roberts. Independent sets, matchings, and occupancy fractions. *Journal of the London Mathematical Society*, 96(1):47–66, 2017. doi:10.1112/jlms.12056.
- [5] E. Davies, M. Jenssen, W. Perkins, and B. Roberts. Extremes of the internal energy of the potts model on cubic graphs. *Random Structures & Algorithms*, 53(1):59–75, 2018. doi:10.1002/rsa.20767.
- [6] E. Davies, M. Jenssen, W. Perkins, and B. Roberts. On the average size of independent sets in triangle-free graphs. *Proceedings of the American Mathematical Society*, 146(1):111–124, 2018. doi:10.1090/proc/13728.
- [7] E. Davies, M. Jenssen, W. Perkins, and B. Roberts. Tight bounds on the coefficients of partition functions via stability. *Journal of Combinatorial Theory, Series A*, 160:1–30, 2018. doi:10.1016/j.jcta.2018.06.005.
- [8] I. A. Faradžev. Generation of nonisomorphic graphs with a given distribution of the degrees of vertices. In *Algorithmic studies in combinatorics (Russian)*, pages 11–19, 185. “Nauka”, Moscow, 1978.
- [9] D. Galvin. Maximizing H-colorings of a regular graph. *Journal of Graph Theory*, 73(1):66–84, 2013. doi:10.1002/jgt.21658.

- [10] D. Galvin and P. Tetali. On weighted graph homomorphisms. *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, 63:97–104, 2004. doi:[10.1090/dimacs/063](https://doi.org/10.1090/dimacs/063).
- [11] D. J. Galvin. Bounding the partition function of spin-systems. *Electronic Journal of Combinatorics*, 13(3):R72, 2006.
- [12] M. R. Garey, D. S. Johnson, and L. Stockmeyer. Some Simplified NP-complete Problems. In *Proceedings of the Sixth Annual ACM Symposium on Theory of Computing*, STOC '74, pages 47–63, New York, NY, USA, 1974. ACM. doi:[10.1145/800119.803884](https://doi.org/10.1145/800119.803884).
- [13] E. Ising. Beitrag zur theorie des ferromagnetismus. *Zeitschrift für Physik A Hadrons and Nuclei*, 31(1):253–258, 1925. doi:[10.1007/BF02980577](https://doi.org/10.1007/BF02980577).
- [14] N. Linial. Legal coloring of graphs. In *Foundations of Computer Science, 1983., 24th Annual Symposium on*, pages 470–472. IEEE, 1983. doi:[10.1109/SFCS.1983.28](https://doi.org/10.1109/SFCS.1983.28).
- [15] P.-S. Loh, O. Pikhurko, and B. Sudakov. Maximizing the number of q -colorings. *Proceedings of the London Mathematical Society*, 101(3):655–696, 2010. doi:[10.1112/plms/pdp041](https://doi.org/10.1112/plms/pdp041).
- [16] L. Lovász. *Large networks and graph limits*, volume 60. American Mathematical Society Providence, 2012. doi:[10.1090/coll/060](https://doi.org/10.1090/coll/060).
- [17] B. D. McKay. Practical graph isomorphism. *Congressus Numerantium*, 30:45–87, 1981.
- [18] B. D. McKay. Isomorph-free exhaustive generation. *Journal of Algorithms*, 26(2):306–324, 1998. doi:[10.1006/jagm.1997.0898](https://doi.org/10.1006/jagm.1997.0898).
- [19] M. Mezard and A. Montanari. *Information, physics, and computation*. Oxford University Press, 2009.
- [20] G. Perarnau and W. Perkins. Counting independent sets in cubic graphs of given girth. 2016, [arXiv:1611.01474](https://arxiv.org/abs/1611.01474).
- [21] R. B. Potts. Some generalized order-disorder transformations. In *Mathematical proceedings of the Cambridge philosophical society*, volume 48, pages 106–109. Cambridge University Press, 1952. doi:[10.1017/S0305004100027419](https://doi.org/10.1017/S0305004100027419).
- [22] R. C. Read. Every one a winner or how to avoid isomorphism search when cataloguing combinatorial configurations. *Annals of Discrete Mathematics*, 2:107–120, 1978. doi:[10.1016/S0167-5060\(08\)70325-X](https://doi.org/10.1016/S0167-5060(08)70325-X). Algorithmic aspects of combinatorics (Conf., Vancouver Island, B.C., 1976).
- [23] E. Vigoda. Improved bounds for sampling colorings. *Journal of Mathematical Physics*, 41(3):1555–1569, 2000. doi:[10.1063/1.533196](https://doi.org/10.1063/1.533196).
- [24] F.-Y. Wu. The Potts model. *Reviews of modern physics*, 54(1):235, 1982. doi:[10.1103/RevModPhys.54.235](https://doi.org/10.1103/RevModPhys.54.235).
- [25] Y. Zhao. Extremal regular graphs: independent sets and graph homomorphisms. *American Mathematical Monthly*, 124(9):827–843, 2017. doi:[10.4169/amer.math.monthly.124.9.827](https://doi.org/10.4169/amer.math.monthly.124.9.827).