



HAL
open science

Uniformly Random Colourings of Sparse Graphs

Eoin Hurley, François Pirot

► **To cite this version:**

Eoin Hurley, François Pirot. Uniformly Random Colourings of Sparse Graphs. STOC '23: 55th Annual ACM Symposium on Theory of Computing, Jun 2023, Orlando FL USA, France. pp.1357-1370, 10.1145/3564246.3585242 . hal-04245446

HAL Id: hal-04245446

<https://hal.science/hal-04245446v1>

Submitted on 12 Jan 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

UNIFORMLY RANDOM COLOURINGS OF SPARSE GRAPHS

EOIN HURLEY AND FRANÇOIS PIROT

ABSTRACT. We analyse uniformly random proper k -colourings of sparse graphs with maximum degree Δ in the regime $\Delta < k \ln k$. This regime corresponds to the lower side of the shattering threshold for random graph colouring, a paradigmatic example of the shattering threshold for random Constraint Satisfaction Problems. We prove a variety of results about the solution space geometry of colourings of fixed graphs, generalising work of Achlioptas, Coja-Oghlan [ACO08], and Molloy [Mol12] on random graphs, and justifying the performance of stochastic local search algorithms in this regime. Our central proof relies only on elementary techniques, namely the first-moment method and a quantitative induction, yet it strengthens list-colouring results due to Vu [Vu02], and more recently Davies, Kang, P., and Sereni [DKPS20], and generalises state-of-the-art bounds from Ramsey theory in the context of sparse graphs. It further yields an approximately tight lower bound on the number of colourings, also known as the partition function of the Potts model, with implications for efficient approximate counting.

1. INTRODUCTION

Theoretical Computer Science (TCS), Statistical Physics, and Combinatorics are nearest neighbours in the network of the sciences. Their stated aims typically differ, but occasionally they converge, leading to a fruitful exchange of ideas. A central example of this overlap is graph colouring; for TCS this is a constraint satisfaction problem (CSP), for Statistical Physics this is the anti-ferromagnetic Potts model, and for Combinatorics this is a 150-year-old well of interesting problems. In particular, colourings of sparse graphs and random graphs of average degree Δ with $k \ln k = \Omega(\Delta)$ colours have proved to be a rich source of phase transition type behaviour. Our focus is on uniformly random k -colourings of graphs of maximum degree Δ , satisfying some sparsity conditions (e.g. triangle-freeness) and under the assumption that $\Delta < k \ln k$, in other words right below the phase transition.

The bulk of the paper is centred around the TCS perspective. We build on the rich literature surrounding phase transitions in random CSPs and the so-called shattering threshold by proving that, above the shattering threshold, properties of the solution space geometry of random graphs — namely those investigated in [ACO08, ART06, Mol12] — in fact hold for any fixed sparse graph. For Statistical Physics, we prove an approximately tight lower bound on the number of colourings of sparse graphs; this corresponds to a lower bound on the partition function of the k -state anti-ferromagnetic Potts model at temperature 0. These bounds in turn suggest interesting possibilities for limits of sequences of graphs with growing girth, as well as efficient approximate counting of colourings. For Combinatorics, our result also extends to local list-colourings, where list sizes are defined according to local parameters (namely the degree of each vertex), and allow for a certain density of edges within the neighbourhoods. Our results in this direction strengthen all previous results, and do so with a simpler proof. The bounds that can be derived from our work are tight barring a breakthrough on a 50-year-old Ramsey theory problem.

We now state a simplified version of our primary result (Theorem 2.10), from which all other results follow. We write $[k] := \{1, \dots, k\}$. For every graph G we let $\mathcal{C}_k(G)$ be the set of proper k -colourings of G (we omit the subscript when k is clear from the context). For every vertex $v \in V(G)$,

The research leading to these results was partially supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – 428212407 (E. Hurley).

every subgraph $H \subseteq G$, and every k -colouring $\sigma \in \mathcal{C}_k(H)$, we let $L_\sigma(v) := [k] \setminus \sigma(N(v) \cap V(H))$ be the list of colours available for v given σ , and we let $\ell_\sigma(v) := |L_\sigma(v)|$ be its order.

Theorem 1.1. *Let $\varepsilon \in (0, 1)$ be fixed, let Δ be sufficiently large (in terms of ε), and let G be a triangle-free graph of maximum degree at most Δ . Fix k such that $(1 - \varepsilon)k \ln k \geq \Delta$ and set $\ell := \Delta^{\varepsilon/2}$. Then, for every $v \in V(G)$, the uniformly random proper k -colouring σ of $G \setminus v$ satisfies $\mathbb{E}[\ell_\sigma(v)] \geq \ell$.*

We observe that the conclusion of Theorem 1.1 also holds for the uniformly random proper k -colouring τ of G . Indeed, the number of extensions of $\sigma \in \mathcal{C}(G \setminus v)$ to $\mathcal{C}(G)$ is precisely $\ell_\sigma(v)$, and thus $\mathbb{E}[\ell_\tau(v)] = \mathbb{E}[\ell_\sigma(v)^2] \geq \mathbb{E}[\ell_\sigma(v)]$.

The proof of Theorem 1.1 relies on the following Coupon-Collector Lemma which is a randomised version of a result due to Molloy [Mol19].

Lemma 1.2 (Coupon-Collector Lemma). *Suppose we have random non-empty lists $\mathbf{L}_1, \dots, \mathbf{L}_d$, each of which takes values in the finite subsets of \mathbb{N} . Fix some integer $t \geq 1$, and define the random variable $\mathbf{X} := \#\{i \in [d] : |\mathbf{L}_i| \leq t\}$. Now choose an element $\sigma(i)$ of \mathbf{L}_i uniformly at random for each $i \in [d]$ and define the random variable $\mathbf{L} := [k] \setminus \{\sigma(i) : i \in [d]\}$. Then*

$$\mathbb{E}[|\mathbf{L}|] \geq k_0 e^{-(1+\frac{1}{t})\frac{d}{k_0}}, \quad \text{where } k_0 = k - \mathbb{E}[\mathbf{X}].$$

For exposition's sake, let us show how Theorem 1.1 follows from the Coupon-Collector Lemma. The proof of the more general statement of Theorem 2.10 uses the same idea, but in order to get the explicit value of ε and handle the more general set-up we need an extra layer of technicality that obscures its conceptual simplicity. Let us mention that the proof relies on an induction hypothesis that is (a rephrasing of) the standard Rosenfeld counting setup [Ros20, WW20].

Proof of Theorem 1.1. We prove the result by induction on n , the order of G . If G has only 1 vertex we are trivially done. Suppose the statement holds for all H satisfying the conditions of the statement and with order at most $n-1$. Let G be as in the statement with order n and let $v \in V(G)$ be arbitrary. Letting $H := G \setminus v$, we will show that the uniformly random k -colouring $\sigma \in \mathcal{C}(H)$ satisfies $\mathbb{E}[\ell_\sigma(v)] \geq \ell$. Observe that this is equivalent to the statement $|\mathcal{C}(G)| \geq \ell |\mathcal{C}(G \setminus v)|$.

Let us fix some integer t . Given the realisation of σ , we say that a neighbour $u \in N(v)$ of v has a short list if $\ell_\sigma(u) \leq t$, and we let S_σ be the set of vertices $u \in N(v)$ with short lists. First we observe that the expected size of S_σ is small. Indeed, consider the probability that a vertex $u \in N(v)$ belongs to S_σ ; we have

$$(1) \quad \mathbb{P}[\ell_\sigma(u) \leq t] = \frac{\#\{\sigma \in \mathcal{C}(H) : \ell_\sigma(u) \leq t\}}{|\mathcal{C}(H)|} \leq \frac{t |\mathcal{C}(H \setminus u)|}{\ell |\mathcal{C}(H \setminus u)|} = \frac{t}{\ell},$$

where the bottom line of the inequality follows from the induction hypothesis applied on H and u . So $\mathbb{E}[|S_\sigma|] \leq \Delta t / \ell$.

Let us write $N(v) = \{u_1, \dots, u_d\}$. We are now ready to apply Lemma 1.2 to the random lists $L_\sigma(u_1), \dots, L_\sigma(u_d)$, with $\mathbf{X} = |S_\sigma|$ counting the number of short lists. We simultaneously resample $\sigma(u_i)$ uniformly at random from $L_\sigma(u_i)$ for every $u_i \in N(v)$ and note that this is precisely the set-up of Lemma 1.2. We obtain that $\mathbb{E}[L_\sigma(v)] \geq k_0 e^{-(1+1/t)d/k_0}$ where $k_0 = k - \mathbb{E}[|S_\sigma|] \geq k - \Delta t / \ell$. Crucially, the realisation of the lists $L_\sigma(u_1), \dots, L_\sigma(u_d)$ is determined by the restriction of σ to $H \setminus N(v)$, so it is not affected by the resampling. It follows that σ remains uniformly distributed in $\mathcal{C}(H)$ after this resampling. To finish the proof, we may set $t := \lceil 5/\varepsilon \rceil$, and assume that Δ is large enough so that $\Delta t / \ell \leq \varepsilon k / 5$. We have

$$\begin{aligned} \mathbb{E}[|L_{\sigma}(v)|] &\geq \left(k - \frac{\Delta t}{\ell}\right) e^{-\frac{(1+1/t)\Delta}{k-\Delta t/\ell}} \geq (1 - \varepsilon/5) k e^{-\frac{(1+\varepsilon/5)(1-\varepsilon)k \ln k}{(1-\varepsilon/5)k}} \geq (1 - \varepsilon/5) k e^{(-1-\varepsilon/5+4\varepsilon/5) \ln k} \\ &\geq (1 - \varepsilon/5) k^{3\varepsilon/5} \geq \Delta^{\varepsilon/2}, \end{aligned}$$

assuming once again that Δ is large enough. □

The statement of Theorem 1.1 concerns only the first moment of list orders. We will show that, from this statement, we can derive many properties that the random proper k -colouring σ satisfies with high probability. To this, we need to prove concentration of certain random variables; this is usually achieved by expressing said variable as a sum of (approximately) negatively correlated binary random variables. Here, we wish to prove concentration for the number of short lists that appear in the neighbourhood of a vertex in a uniformly random colouring. We are unfortunately in a context where negative correlation does not hold and thus we introduce another property, namely Bernoulli-domination, that suffices to have high concentration inequalities in that context. This is obtained, in Theorem 3.6, by combining Theorem 1.1 with repeated applications of (1).

This work constitutes, to our knowledge, the first rigorous analysis of the properties of a typical proper colouring of an arbitrary deterministic graph in this regime, and we suspect our methods may be adapted to answer many more questions about such typical colourings and CSP solutions.

1.1. Organisation of Paper. In Section 2, we give problem-specific context and state our headline results. In Subsection 2.1, we review the random CSP's literature before introducing our results on the solution space geometry for colourings of sparse graphs. In Subsection 2.2 we discuss the Combinatorics literature and the extensions to list-colourings and graphs of bounded local density. In Subsection 2.3 we state our lower bound on the number of colourings and discuss the relevance to graph limits and efficient approximate counting. In Section 3 we introduce the probabilistic machinery (the Coupon-Collector-type results in Subsection 3.1 and the ones using Bernoulli-domination in Subsection 3.2) that we rely on, and apply it to sketch the proofs of our results from Subsection 2.1 in Subsection 3.3. All remaining proofs lie in the Appendix.

2. CONTEXT AND STATEMENT OF RESULTS

All graphs considered will be simple loopless graphs, denoted by $G = (V(G), E(G))$ or H and typically on n vertices. For a vertex $v \in V(G)$, we denote by $N_G(v)$ the neighbourhood of v in G , and by $N_G[v] := \{v\} \cup N_G(v)$ its closed neighbourhood. We omit the subscript if it is clear from the context. For a subset $U \subseteq V(G)$ let $G[U] \subseteq G$ denote the subgraph of G induced by U . Throughout, when we say that a property holds with high probability (w.h.p.), we mean with probability tending to 1 as n tends to ∞ . As is standard, we denote by $o_y(x)$ a real-valued function $f(x, y)$ such that $f(x, y)/x \rightarrow 0$ as $y \rightarrow \infty$ for all $x > 0$. Further, we denote by $\Omega(x)$ a real-valued function $f(x)$ for which there exists $C > 0$ such that $f(x) \geq Cx$ for all x sufficiently large.

2.1. Solution Space Geometry. For over half a century we have known that many Constraint Satisfaction Problems (CSPs) are NP-hard (in the worst case). It was hoped that average case complexity would paint a brighter picture, and indeed early results suggested this might be so [Tur88, DF86, Kuč89]. These results yielded polynomial-time algorithms for many random CSPs, including for k -colouring almost all graphs of chromatic number at most k . However, a more challenging¹ test case soon emerged; finding a polynomial-time algorithm that, with high probability, finds a solution to a random CSP with a *fixed ratio of constraints to variables*. When the CSP is graph colouring this test case becomes: k -colour the Erdős-Rényi random graph $G_{n, dn/2}$. That is,

¹More challenging and perhaps more natural, as the assumed condition is checkable in polynomial time.

the graph chosen uniformly at random from all graphs with n vertices and $dn/2$ edges, so of average degree d . In [AM97] it was shown that w.h.p. the following algorithm suffices for $d \leq k \ln k$:

- (1) choose a vertex v uniformly at random from those with the least number of available colours;
- (2) colour v uniformly at random with one of its available colours.

Remarkably, while $G_{n,dn/2}$ is known to be w.h.p. k -colourable for $d \sim 2k \ln k$, no algorithm has been proven to succeed for any $d > (1 + o_k(1))k \ln k$. This extremely naive algorithm is the state of the art. As observed by Achlioptas and Coja-Oghlan in [ACO08], this state of affairs appears to be universal in that it holds for “nearly every random CSP of interest” (when d represents the ratio of constraints to variable). That is, extremely naive algorithms (known for decades) work up to some threshold (d_{Sh}), which is well below the SAT-UNSAT threshold (d_{SU}), but no algorithms are known to work for $d > d_{Sh}$. The value of d_{Sh} is not arbitrary; it coincides with the conjectured shattering threshold [KMRT⁺07] from the statistical physics literature (which is defined in terms of expected spatial correlation decay). Drawing inspiration from these conjectures, and building on [ART06], Achlioptas and Coja-Oghlan rigorously demonstrated that, at least in the case of random graph colouring, random k -SAT and random k -uniform hypergraph 2-colouring, this was no coincidence. As stated in earlier work [ART06] a key idea underlying their approach is the following.

Hypothesis 1. The geometry of the solution space of a given CSP instance dictates the performance of algorithms on that instance²

They proved that w.h.p. two phase transitions, one local and one global, occur in the geometry of the solution space as d increases past this shattering threshold ($d_{Sh} k \ln k$ for random graph colouring). *Globally*, the solution space starts as a large well-connected ball before shattering into exponentially many, exponentially small, well-separated clusters. *Locally*, typical solutions start as able to change the assignment of almost any variable by changing $o(n)$ other variables (almost all variables are “loose”) and end up with this not being the case for almost all variables (almost all variables are “rigid”) (see Section 2.1.1). They hypothesised that these phase transitions are *the* barriers to efficient algorithms. That is, these properties of the geometry of the solution space explain both the existence of algorithms for $d < d_{Sh}$ and the difficulty in finding such algorithms for $d_{SU} > d > d_{Sh}$.

In [Mol12], the local phase transition — also known as the freezing threshold — was further analysed in the case of random graph colouring, showing that an even stronger statement, Theorem 2.4, holds. The local freezing threshold is of particular interest because, in [KMRT⁺07, Zde08, ZK07], it (as opposed to the global clustering threshold) was hypothesised to be the primary barrier to efficient algorithms. Indeed there is evidence that supports this; while the thresholds are asymptotically equal (in k), they are not for small k , and in [AM02] a random graph colouring algorithm is provided for $d = 4$, for values of k above the clustering threshold but below the freezing threshold (the latter is always higher).

Thus while random CSPs have proved to be challenging problems, there was a consistent explanation of why algorithms succeed or not, based on the typical geometry of the solution space of a random problem instance. However, a crack appeared in this picture when, in [Mol19], Molloy used entropy compression (inspired by Moser-Tardos [MT10]) to prove that the following clever and simple algorithm yields a k -colouring of *any* triangle-free graph G of maximum degree $d \leq (1 - o_k(1))k \ln k$ (right up to the shattering threshold!):

- (1) start from an arbitrary partial proper k -colouring of G ;
- (2) while there exists a **Bad** vertex v , resample a partial proper colouring of the neighbourhood of v uniformly at random³;

²This has a very similar flavour to the Overlap Gap Property, see [Gam21].

³In fact in this step Molloy gave priority to the **Bad** vertices that were most recently not **Bad**, but that is not relevant to our discussion.

(3) extend the partial colouring to a complete proper colouring.

This is of course a stochastic local search algorithm [Ili19]. Molloy defined **Bad** vertices to be those with too few colours available, or too many uncoloured neighbours that might create a conflict on a given colour. Via entropy compression, he demonstrated that after resampling neighbourhoods polynomially many times (in expectation), no vertices would be **Bad**. Then, a direct application of the algorithmic Lovász Local Lemma [MT10] suffices to complete the colouring in polynomial-time. In fact, Bernshteyn [Ber19] showed that the initial part of the algorithm could also be viewed as an application of the Lovász Local Lemma [AS16].

The reason this constitutes a crack in our picture is that our barriers to the existence of efficient algorithms above the shattering threshold were hypothesised based on the analysis of random graphs. But Molloy’s algorithm works for *all triangle-free graphs*, not just random graphs. Thus, working off of Hypothesis 1, we would expect that the geometry of the space of colourings of all triangle-free graphs of max degree $d < k \ln k$ is similar to that of $G_{n,d/n}$ as described by Achlioptas, Coja-Oghlan [ACO08]. If, for example, there existed a d -regular triangle-free graph for which the typical k -colouring was rigid, or for which the solution space consisted of exponentially many well-separated components, then this would rule out these phenomena (the local freezing threshold and the global clustering threshold respectively) as barriers to efficient algorithms. The central contribution of this work, see Theorems 2.7 and 2.8, is to show that no such graph exists in the former case, supporting the possibility that the local freezing threshold is a barrier to efficient algorithms. Establishing positive results on the connectivity of the solution space is also particularly relevant to understanding the performance of stochastic local search algorithms, such as Molloy’s [Mol19].

2.1.1. Rigid, Loose, Frozen, Thawed. In this section, we formally define the local “freezing” phase transition discussed in the introduction. For a discussion and formal definition of the global clustering phase transition see [ACO08, Section 2.1]. We say “local” because it concerns the perspective from a single colouring, rather than viewing the whole solution space at once. For two colourings σ and τ of a graph G , let $\text{dist}(\tau, \sigma) := \#\{v \in V(G) : \sigma(v) \neq \tau(v)\}$ be the Hamming distance.

Definition 2.1. For $t, k \in \mathbb{N}$ and any graph G , we define the distance- t k -colouring graph $\mathcal{H}(G, k, t)$ as follows. The vertices of \mathcal{H} are the proper k -colourings of G and two distinct vertices τ, σ are adjacent in \mathcal{H} if $\text{dist}(\tau, \sigma) \leq t$. Further, we define the t -clusters to be the connected components of $\mathcal{H}(G, k, t)$ (when $t = 1$ we just say clusters).

Thus τ and σ are in the same t -cluster if and only if there is a sequence of proper k -colourings $\tau = \tau_1, \dots, \tau_s = \sigma$ such that $\text{dist}(\tau_i, \tau_{i+1}) \leq t$ for all $i \in [s - 1]$ (note that s may be arbitrarily large). When k and t are clear from the context we will just write $\mathcal{H}(G)$.

We now state the definitions of [ACO08, Mol12] in these terms.

Definition 2.2. Given $k, t \in \mathbb{N}$, a graph G , a vertex $v \in V(G)$ and $\tau \in \mathcal{C}_k(G)$, we say that v in τ :

- is t -loose if for all $j \in [k]$ there exists a neighbour $\sigma \in \mathcal{H}(G)$ of τ with $\sigma(v) = j$;
- is t -thawed if for all $j \in [k]$ there exists $\sigma \in \mathcal{H}(G)$ in the same t -cluster as τ with $\sigma(v) = j$;
- is t -rigid if $\tau(v) = \sigma(v)$ for all neighbours $\sigma \in \mathcal{H}(G)$ of τ ;
- is t -frozen if $\tau(v) = \sigma(v)$ for all $\sigma \in \mathcal{H}(G)$ in the same t -cluster as τ .

When we omit t in the notations above, it means that $t = 1$ by default. Note that the definitions are ordered from the most connected to the least connected, and that t -loose implies t -thawed and t -frozen implies t -rigid. Another way to view the definition of thawed and frozen is to consider the projections of the cluster containing τ onto v . If this projection contains one/all element(s) of $[k]$ then v is frozen/thawed in τ . Note that loose/rigid are both about neighbours, not clusters; in particular a vertex v can be loose only if it is isolated, and it is rigid if $N[v]$ spans all k colours.

While in this paper we discuss uniformly random colourings σ of a fixed graph G , in [ACO08, Mol12] they discussed uniformly random colourings of uniformly random graphs. We make this precise as follows. First fix $d, k \in \mathbb{N}$, and now consider $n \in \mathbb{N}$ and $m := \frac{dn}{2}$ (we will let n tend to infinity). We now let \mathbf{G} be chosen uniformly at random from all graphs on n vertices with m edges. Finally we let σ be a proper k -colouring of \mathbf{G} chosen uniformly at random. The random pair (\mathbf{G}, σ) is the subject of [ACO08, Mol12] (although they analyse it via the powerful “planted” model). Note that conditioned on $\mathbf{G} = G$ our random variable σ is the same. Given k, n , and d , we will refer to the above described (\mathbf{G}, σ) as the uniformly random instance-solution pair.

We now review what is known about rigid, loose and frozen variables (thawed was introduced in this paper). Let $k, d \in \mathbb{N}$ such that

$$(2) \quad (1 + o_k(1))k \ln k \leq d \leq (2 - o_k(1))k \ln k.$$

Theorem 2.3. [ACO08, Theorem 4] *For k, d as in (2) the uniformly random instance-solution pair (\mathbf{G}, σ) w.h.p. contains at least $(1 - o_k(1))n$ vertices that are $\Omega(n)$ -rigid.*

Theorem 2.4. [Mol12, Theorem 2.4] *For k, d as in (2), there exists $0 < \alpha, \beta < 1$ such that w.h.p. the uniformly random instance-solution pair (\mathbf{G}, σ) contains $(\beta + o_n(1))n$ vertices that are αn -frozen.*

These results make a strong statement about the solution space of a random graph for $k \ln k < d$. In the same papers, the opposite side of the threshold was also described. Suppose

$$(3) \quad d \leq (1 - o_k(1))k \ln k.$$

Theorem 2.5. [ACO08, Theorem 5] *For k, d as in (3), in the uniformly random instance-solution pair (\mathbf{G}, σ) w.h.p. every vertex is $o(n)$ -loose.*

Theorem 2.6. [Mol12, Theorem 2.4] *For k, d as in (3), there exists a constant $C > 0$ such that w.h.p. the uniformly random instance-solution pair (\mathbf{G}, σ) contains at most $o(n)$ vertices that are $(C \ln n)$ -frozen.*

We remark that Theorems 2.3 and 2.5 also apply to 2-colouring k -uniform hypergraphs and k -CNF. We build on Theorems 2.5 and 2.6, showing that in fact that we can convert these into statements about uniformly random colourings of deterministic graphs of large enough girth.

2.1.2. Original Results on Solution Space Geometry. Our first theorem is about triangle-free graphs, but note that it holds with probability 1 as d tends to infinity, rather than as n tends to infinity. This is the best one could expect by only analysing the graph within a constant radius r of a vertex v , because random fluctuations will occur with probability depending only on r and Δ . Analysing the graph at any radius larger than the girth is of course much more challenging.

Theorem 2.7. *For all $\varepsilon > 0$ the following holds for all k sufficiently large and $\Delta \leq (1 - \varepsilon)k \ln k$. Let G be a triangle-free graph of maximum degree at most Δ and let σ be the uniformly random proper k -colouring of G . Then for every vertex $v \in V(G)$ the following holds with probability at least $1 - e^{-\Delta^{\varepsilon/3}}$:*

- (1) *in σ , v is $(\Delta + 1)$ -loose;*
- (2) *in σ , v is thawed.*

In fact what we will show is that, with the above probability, one can change the colour of v to any colour just by changing the colours of v and its neighbours.

As mentioned above, due to the dependency of the probability on d , this does not directly strengthen Theorems 2.5 and 2.6, although it does preclude the possibility of triangle-free graphs of maximum degree Δ with the properties guaranteed by Theorems 2.3 and 2.4. If we strengthen our girth condition, we do obtain strengthenings of Theorems 2.5 and 2.6 (these are not possible for triangle-free graphs, see Section 3.4).

Theorem 2.8. *For all $\varepsilon > 0$ the following holds for all k sufficiently large and $\Delta \leq (1 - \varepsilon)k \ln k$. Suppose G is a graph on n vertices with maximum degree Δ and girth at least $\ln \ln n$. Then for the uniformly random proper k -colouring σ of G , the following holds w.h.p. for every $v \in V(G)$:*

- (1) in σ , v is $O((\ln n)^2)$ -loose⁴;
- (2) in σ , v is thawed.

Once again what we actually show is that we can change the colour of v to whatever we please by only changing the colours of vertices at distance at most $O(\ln \ln n / \ln \Delta)$ from v . Theorem 2.8 implies that the reconfiguration graph of the proper k -colourings of G (where the reconfiguration step is to change the colour of any vertex) contains a connected component that covers all but $o_\Delta(1)$ colourings. To see that these are indeed strengthenings, recall that with high probability (assuming that n is large enough) there are no triangles and all cycles of length at most $o(\log_d n)$ are vertex disjoint in the Erdős-Rényi random graph $G_{n, dn/2}$ (this is shown for the random regular graph in [MWW04, Equation (2.8)]). Thus, in the proof of Theorem 2.8, when we consider the set of vertices at distance at most $o(\log_d n)$ from some vertex v , we can make the induced graph a tree by deleting a matching. This will not affect the analysis. Theorem 2.8 directly strengthens Theorem 2.6, by replacing $(C \ln n)$ -frozen by 1-frozen (in [ACO08] it had already been remarked that $O(\ln n)$ -looseness is possible).

We conjecture that the global connectivity properties proved in [ACO08] should also hold for all triangle-free graphs.

Conjecture 2.9. *For ε, Δ, k and G as in Theorem 1.1, there exists a cluster of $\mathcal{H}(G)$ containing at least a $(1 - o_\Delta(1))$ -proportion of $\mathcal{C}_k(G)$.*

2.2. Local List Colouring and Density. Given a graph G , a list-assignment of G is a map $L: V(G) \rightarrow 2^{\mathbb{N}}$, and a proper L -colouring σ of G is a map $c: V(G) \rightarrow \mathbb{N}$ such that $c(v) \in L(v)$ for every vertex $v \in V(G)$, and $c(u) \neq c(v)$ for every edge $uv \in E(G)$. Thus list-colouring is a CSP which includes colouring as a special case. If there exists an L -colouring of G , we say that G is L -colourable. The list chromatic number of G , denoted $\chi_\ell(G)$, is the minimum k such that G is L -colourable for every list-assignment L with $|L(v)| \geq k$ for every vertex $v \in V(G)$. Note that the classical chromatic number $\chi(G)$ is always at most $\chi_\ell(G)$.

In a seminal result, Johansson [Joh96] showed that the list chromatic number of triangle-free graphs of maximum degree Δ never exceeds $O(\Delta / \ln \Delta)$ as $\Delta \rightarrow \infty$. Two decades later, Molloy [Mol19] showed with the help of entropy compression that this bound can be tightened to $(1 + o(1))\Delta / \ln \Delta$ as $\Delta \rightarrow \infty$; Bernshteyn [Ber19] then showed that it was possible to replace the use of entropy compression with an application of the lopsided Lovász Local Lemma and obtain a similar result. Following that breakthrough, there has been a lot of interest for triangle-free graph colourings and extensions [AIS19, ABD21, BKNP18, DdJdVVKP20, DKPS20].

Although quite different, all the proofs in those works rely on a 2-step colouring procedure. This involves first finding a partial proper colouring with appropriate properties and then extending it in a second step. This second step is usually called the Finishing Blow. In contrast, our proof does not require the use of the Finishing Blow, and is the first of that nature to the best of our knowledge. Our result also includes graphs where the density is bounded within each neighbourhood, a natural extension of triangle-free graphs. The first result in that vein is due to Alon, Krivelevich, and Sudakov [AKS99]. They relied on the result of Johansson in order to prove that if a graph G of maximum degree Δ is such that each neighbourhood spans at most Δ^2/f edges, where $1 \leq f \leq \Delta^2 + 1$, then its chromatic number is at most $O(\Delta / \ln f)$ as $f \rightarrow \infty$. This result was then extended to the list chromatic number by Vu [Vu02]. More recently, it has been proved by the second author together with Davies, Kang, and Sereni [DKPS20, Corrolary 24], that this upper bound can be

⁴With a more careful (and technical) analysis, the authors believe that this could be replaced with $O(\ln n)$ -loose.

tightened to $(2 + \varepsilon)\Delta/\ln f$, provided that $f \geq (\ln \Delta)^{2/\varepsilon}$, for every $\varepsilon > 0$. Our main theorem is a strengthening of these results which is close to being optimal.

We were inspired by a counting argument which was first used in the context of graph colouring by Rosenfeld [Ros20] (see [Woo20, Section 3.5] for an introduction to the method), and later devised more generally for hypergraph colouring by Wanless and Wood [WW20]. We note however that our proof deviates from the standard arguments in order to exploit the relationship between uniformly random colourings of a graph G and uniformly random colourings of induced subgraphs $H \subset G$.

Given a list-assignment $L: V(G) \rightarrow 2^{\mathbb{N}}$ of G , a proper L -colouring σ of a subgraph H of G and a vertex $v \in V(G)$, we denote $L_\sigma(v) := L(v) \setminus \sigma(N_H(v))$ the set of colours available at v given σ , and $\ell_\sigma(v) := |L_\sigma(v)|$ its size.

Theorem 2.10. *Let G be a graph of maximum degree Δ such that every graph induced by a neighbourhood in G has average degree at most $d \leq \frac{\Delta}{6} - 1$. Write $\rho := \Delta/(d + 1)$ and let $\ell \geq (d + 1)(\ln \rho)^3$. Then for every list-assignment $L: V(G) \rightarrow 2^{\mathbb{N}}$ with*

$$|L(v)| \geq \left(1 + \frac{2}{\ln \rho}\right) \frac{\deg(v)}{W\left(\frac{\deg(v)}{\ell}\right)}$$

for every vertex $v \in V(G)$, the following holds. For all $v \in V(G)$, the uniformly random proper L -colouring σ of $G \setminus v$ satisfies $\mathbb{E}[\ell_\sigma(v)] \geq \ell$.

As before, the conclusion also holds for the uniformly random k -colouring of G . In the statement of Theorem 2.10, we use the W -Lambert function $z \mapsto W(z)$ which is defined as the reciprocal of the function $z \mapsto ze^z$. In the proof of Theorem 2.10, we will use the well-known fact that $e^{W(z)} = z/W(z)$. Moreover, we note that $W(z) = \ln z - \ln \ln z + o(1)$ as $z \rightarrow \infty$; hence, by fixing $d := 0$, Theorem 2.10 implies Theorem 1.1. Since we have already given a proof of Theorem 1.1, whose statement suffices for most applications in this paper, we defer the proof of Theorem 2.10 to the Appendix.

In the context of Vu's result [Vu02], Theorem 2.10 implies the following; this is proved in the Appendix.

Corollary 2.11. *Let G be a graph of maximum degree Δ , such that every neighbourhood spans at most Δ^2/f edges, for some $1 \leq f \leq \Delta^2 + 1$. Then $\chi_\ell(G) \leq (1 + o(1))\frac{\Delta}{\ln \min\{\Delta, f\}}$ as $f \rightarrow \infty$.*

It has been observed in [DdJdVVKP20] that there exist graphs satisfying the hypothesis of Corollary 2.11 and with chromatic number $(\frac{1}{2} - o(1))\Delta/\ln \min\{f, \Delta\}$. So the bound in Corollary 2.11 is sharp up to an asymptotic factor 2. Reducing this gap would constitute a breakthrough in Ramsey theory, as this would imply an improvement of the estimate of the off-diagonal Ramsey numbers $R(3, t)$ for large $t \in \mathbb{N}$, a long-standing open problem.

We note that the results of this section also hold in the context of DP-colouring, an interesting extension of list colouring introduced by Dvořák and Postle [DP18], which is once again a CSP. We have decided to state those results in the context of list colouring in order to avoid the verbose formalism of DP-colouring, but the proof of the results in this section readily adapt to that context.

2.3. Partition Functions and Approximation Schemes. Partition functions are important objects in statistical physics. They count the weighted number of states of a system where said weight is proportional to the probability of seeing the system in said state. Thus they also serve as the normalising constant when we write the probability of seeing a particular state. It seems natural to expect that our lower bound on the expected number of colours available at each vertex would yield a lower bound on the number of colourings and thus said partition function. By slightly adjusting the proof of Theorem 2.10 we can show the following.

Theorem 2.12. *Let G be an n -vertex graph of maximum degree Δ such that every graph induced by a neighbourhood in G has average degree at most $d \leq \frac{\Delta}{6} - 1$. Let $f := \Delta/(d+1)$ and suppose $L: V(G) \rightarrow 2^{\mathbb{N}}$ is a list-assignment with $|L(v)| \geq \left(1 + \frac{1}{\ln \rho}\right) q(v)$, where*

$$q(v) \geq \left(1 + \frac{1}{\ln \rho}\right) \frac{\deg(v)}{W\left(\frac{\deg(v)}{(d+1)(\ln \rho)^3}\right)}$$

for every vertex $v \in V(G)$. Then there are at least $(q/\sqrt{D/(d+1)})^n$ proper L -colourings of G , where D is the geometric mean of the degrees in G , and q is the geometric mean of $\{q(v)\}_{v \in V(G)}$.

Shortly before an early version of this manuscript was published in a preprint repository, independent work by Bernshteyn, Brazelton, Cao, and Kang appeared online [BBCK21]. They also rely on the counting argument used by Rosenfeld in [Ros20], and derive a lower-bound on $|\mathcal{C}_k(G)|$ for every triangle-free graph G of maximum degree $\Delta \leq (1-\varepsilon)k \ln k$ large enough in terms of ε , and show that it is asymptotically sharp for Δ -regular triangle-free graphs. In particular, they prove that a random Δ -regular triangle-free graph will almost surely admit no more than $(1 - \frac{1}{k})^{\Delta n/2} \left(1 + \frac{2 \ln n}{n}\right)^n k$ k -colourings as $n \rightarrow \infty$.

Theorem 2.13 (Bernshteyn, Brazelton, Cao, Kang; 2021). *For every $\varepsilon > 0$, there exists Δ_0 such that the following holds. Let G be an n -vertex triangle-free graph of maximum degree $\Delta \geq \Delta_0$ and with m edges. Then, for every $k \geq (1+\varepsilon)\Delta/\ln \Delta$, we have $|\mathcal{C}_k(G)| \geq (1-1/k)^m ((1-\delta)k)^n$, where $\delta = \frac{4}{k}e^{\Delta/k}$.*

Let us pause to consider how many colourings we should expect. Suppose we choose the colour of each vertex independently and uniformly at random. A moment's reflection shows that if there are no cycles in our graph then the events that two distinct edges are monochromatic are independent. Thus if the graph we are colouring has n vertices, m edges and no cycles, then the number of proper k -colourings is exactly $k^n(1 - \frac{1}{k})^m$. Of course, no such graph exists in general, but we can still compute the number of colourings. If we consider the quantity $f(G, k) := \frac{1}{n} \ln |\mathcal{C}_k(G)|$, known as the free energy per variable, then we see that for a Δ -regular ‘‘graph’’ G with no cycles we have $\frac{m}{n} = \frac{\Delta}{2}$ and $f(G, k) = \ln(k(1 - \frac{1}{k})^{\Delta/2})$. Such a graph exists only if we allow an infinite number of vertices, and is unique; it is the infinite Δ -regular tree \mathbb{T}_Δ (if you enjoy this perspective please see [Csi16]). Of course, the number of colourings is also infinite, but we may extend the notion of free energy per variable to infinite graphs, and for \mathbb{T}_Δ this yields $f(\mathbb{T}_\Delta, k) := \ln(k(1 - \frac{1}{k})^{\Delta/2}) \approx \ln k - \frac{\Delta}{2k}$.

Let us shift now to considering $h(G, k) := \frac{f(G, k)}{f(\mathbb{T}_\Delta, k)}$ for Δ -regular graphs G . This measures the free energy per variable of G relative to \mathbb{T}_Δ . If G has many colourings then $h(G, k)$ is larger than 1; if it has few then $h(G, k)$ is smaller than 1. Theorems 2.12 and 2.13 yield the bound⁵ $h(G, k) \geq (1 - o_\Delta(1))$. Further for the random regular graph $G_{n, \Delta}$, [Ber19] shows that almost surely $h(G_{n, \Delta}, k) \leq (1 + o_n(1))$. We conjecture that both of these bounds are asymptotically tight for sequences of graphs with growing girth.

Conjecture 2.14. *For each $\Delta \in \mathbb{N}$ fix a sequences $\{G_i^\Delta\}_{i=1}^\infty$ of Δ -regular graphs with girth tending to ∞ . Then for all functions $k: \mathbb{N} \rightarrow \mathbb{N}$ satisfying $k(\Delta) > (1 + \varepsilon)\frac{\Delta}{\ln \Delta}$ for some $\varepsilon > 0$ it holds that*

$$\sup_{\Delta \in \mathbb{N}} \limsup_{i \rightarrow \infty} h(G_i^\Delta, k(\Delta)) \leq \lim_{\Delta \rightarrow \infty} \liminf_{i \rightarrow \infty} h(G_i^\Delta, k(\Delta)) = 1.$$

We have established the lower bound and we now chart a path toward the upper bound. First note that, writing $i(G)$ for the number of independent sets in G , it was shown in [Csi16] that if each graph in the aforementioned sequence $\{G_i\}_{i=1}^\infty$ is bipartite, then $\lim_{j \rightarrow \infty} i(G_j)$ exists (it corresponds in a meaningful way to the infinite regular tree). We conjecture that the same result

⁵This is conjectured to hold with half the number of colours in [Ber19].

should hold when we replace $i(G)$ by $f(G, k)$ for $k \ln k \geq (1 + o_\Delta(1))\Delta$ and that the limit should be $f(\mathbb{T}_\Delta, k)$ (again corresponding to the infinite regular tree)⁶. Secondly, in [SSSZ20] it was shown that for all k and Δ -regular graphs G on n vertices we have $f(G, k) \leq f(K_{\Delta, \Delta}, k)$, where $K_{\Delta, \Delta}$ is the complete Δ -regular bipartite graph. That is, it was shown that among Δ -regular graphs, complete bipartite graphs maximise the free energy per variable (for k -colourings). We conjecture that for all g , among Δ -regular graphs G with girth at least g , the supremum of $f(G, k)$ is always achieved by a sequence of bipartite graphs⁷. In fact one only needs to show that the supremum over bipartite graphs of girth $\omega(g)$ is at least that of general graphs of girth g for some $\omega(g)$ tending to infinity. If both of these conjectures are true, then the limit for bipartite graphs would be an upper bound for all sequences of graphs with girth tending to ∞ .

We note that this would imply that for all $\varepsilon > 0$ there exists Δ_0 such that for all $\Delta > \Delta_0$ we have an efficient approximation scheme for $f(G, k)$ for Δ -regular graphs of large girth within a factor of $(1 + \varepsilon)$. On the other hand in [GŠV16] it was shown that for fixed Δ there exists no Fully Polynomial-Time Randomised Approximation Scheme for approximating $f(G, k)$ within an additive error of ε unless $\text{RP} = \text{NP}$. It would be very interesting to pin down exactly where algorithmic complexity appears. We briefly remark that work for graphs of high girth and infinite trees was the launchpad for efficient approximate counting of graph k -colourings when k is some constant factor larger than the maximum degree Δ [BG08, GK12, Wei06]. Perhaps, under appropriate girth conditions, related approximation schemes are possible for $k < \Delta$.

3. TECHNICAL RESULTS AND SELECTED PROOFS

In this section we collect our core results about uniformly random colourings and finish by proving Theorem 2.7 and sketching the proof of Theorem 2.8.

3.1. Coupon-Collector-type results. We begin with a proof of the Coupon-Collector Lemma, of which we recall the statement hereafter.

Lemma 3.1 (Coupon-Collector Lemma). *Suppose we have random non-empty lists $\mathbf{L}_1, \dots, \mathbf{L}_d$, each of which takes values in the finite subsets of \mathbb{N} . Fix some integer $t \geq 1$, and define the random variable $\mathbf{X} := \#\{i \in [d] : |\mathbf{L}_i| \leq t\}$. Now choose an element $\sigma(i)$ of \mathbf{L}_i uniformly at random for each $i \in [d]$ and define the random variable $\mathbf{L} := [k] \setminus \{\sigma(i) : i \in [d]\}$. Then*

$$\mathbb{E}[|\mathbf{L}|] \geq k_0 e^{-(1+\frac{1}{t})\frac{d}{k_0}}, \quad \text{where } k_0 = k - \mathbb{E}[\mathbf{X}].$$

Proof. Let us fix a realisation L_1, \dots, L_d of the random lists. We let $S = \{i : |L_i| \leq t\}$ be the set of short lists, and $B = [d] \setminus S$ be the set of big lists. We observe that since $|L_i| \geq t + 1$ for every $i \in B$, we have $\frac{1}{|L_i|-1} \leq (1 + \frac{1}{t})\frac{1}{|L_i|}$. Combining this observation with one due to Molloy [Mol19], we obtain

$$(4) \quad \sum_{x \in [k]} \sum_{\substack{i \in B \\ x \in L_i}} \frac{1}{|L_i|-1} \leq \left(1 + \frac{1}{t}\right) \sum_{x \in [k]} \sum_{\substack{i \in B \\ x \in L_i}} \frac{1}{|L_i|} \leq \left(1 + \frac{1}{t}\right) \sum_{i \in B} \sum_{x \in L_i} \frac{1}{|L_i|} \leq \left(1 + \frac{1}{t}\right) d.$$

The observation of Molloy is that each list L_i appears in the first double sum exactly $|L_i|$ times, and it does so with weight $1/|L_i|$, meaning its contribution is always 1.

Let us fix $\sigma(i) = x_i \in L_i$ for every $i \in S$, and let $L_0 := [k] \setminus \{x_i : i \in S\}$. Note that $|L_0| \geq k - |S|$. We may now pick $\sigma(i)$ uniformly at random from L_i for every $i \in B$, and let $\mathbf{L}_0 := L_0 \setminus \{\sigma(i) : i \in B\}$. Note that \mathbf{L}_0 is precisely \mathbf{L} under the condition that $\sigma(S) = \{x_i : i \in S\}$. We have

⁶Indeed it was already shown in [CL16] that $f(\mathbb{T}_\Delta, k)$ is a lower bound for bipartite graphs, and thus our lower bounds can be viewed as an approximate cousin of [SSSZ20, Theorem 8].

⁷Of course the most elegant solution would be if G_i was just a smallest (in terms of vertices) Δ -regular bipartite graph with girth at least g for all i , but this is stronger than what we need, nevermind that finding such graphs is an extremely challenging open problem [FNV06].

$$\begin{aligned}
 \mathbb{E}[|\mathbf{L}_0|] &= \sum_{x \in L_0} \mathbb{P}[x \notin \sigma(B)] = \sum_{x \in L_0} \prod_{\substack{i \in B \\ x \in L_i}} \left(1 - \frac{1}{|L_i|}\right) \\
 &\geq \sum_{x \in L_0} \exp\left(-\sum_{\substack{i \in B \\ x \in L_i}} \frac{1}{|L_i| - 1}\right) && \text{since } 1 - \frac{1}{z} > e^{-\frac{1}{z-1}} \text{ for every } z > 1; \\
 &\geq |L_0| \exp\left(-\frac{1}{|L_0|} \sum_{x \in L_0} \sum_{\substack{i \in B \\ x \in L_i}} \frac{1}{|L_i| - 1}\right) && \text{by convexity of exp;} \\
 &\geq (k - |S|) \exp\left(-\frac{(1 + \frac{1}{t})d}{k - |S|}\right) && \text{by } |L_0| > k - |S| \text{ and (4).}
 \end{aligned}$$

It remains to average over all possible realisations of $\mathbf{L}_1, \dots, \mathbf{L}_d$. We use Jensen's inequality together with the convexity of the function $z \mapsto ze^{-C/z}$ for every $C > 0$ over the interval $(0, +\infty)$, and obtain that

$$\mathbb{E}[|\mathbf{L}|] \geq \mathbb{E}\left[(k - \mathbf{X}) \exp\left(-\frac{(1 + \frac{1}{t})d}{k - \mathbf{X}}\right)\right] \geq \mathbb{E}[(k - \mathbf{X})] \exp\left(-\frac{(1 + \frac{1}{t})d}{\mathbb{E}[k - \mathbf{X}]}\right) = k_0 e^{-(1 + \frac{1}{t})\frac{d}{k_0}},$$

where $k_0 = k - \mathbb{E}[\mathbf{X}]$. □

A straightforward application of this result yields the following (deterministic) existence of a colouring that induces a big list at v , via the first moment method. This can be obtained by following the exact same computation as that in the proof of Theorem 1.1.

Corollary 3.2. *Let $\varepsilon \in (0, 1)$ be fixed and let Δ be sufficiently large (in terms of ε). Fix k such that $(1 - \varepsilon)k \ln k \geq \Delta$. Let $L_1, \dots, L_\Delta \subset [k]$ be given, and suppose that the number of short lists is $\#\{i : |L_i| \leq 5/\varepsilon\} \leq \varepsilon k/5$. Then there exists $\sigma(i) \in L_i$ for $i \in [\Delta]$ such that $|[k] \setminus \{\sigma(i) : i \in [\Delta]\}| \geq \Delta^{\varepsilon/2}$.*

3.2. Bernoulli-domination. For the rest of the proofs, we will need high-concentration bounds, and to that end we introduce the concept of Bernoulli-domination. Let us first state the following useful Chernoff-Hoeffding type tail bound for binary random variables that are dominated by independent random variables. This is a special case of [PS97, Theorem 3.4].

Theorem 3.3. *Suppose $\mathbf{X}_1, \dots, \mathbf{X}_s$ are binary random variables and $\mathbf{Y}_1, \dots, \mathbf{Y}_s$ are independent binary random variables. Let $\mathbf{Y} := \sum_{i \in [s]} \mathbf{Y}_i$ and $\mathbf{X} := \sum_{i \in [s]} \mathbf{X}_i$. Write $\mu := \mathbb{E}[\mathbf{Y}]$. Then if*

$$\mathbb{E}\left[\prod_{i \in J} \mathbf{X}_i\right] \leq \mathbb{E}\left[\prod_{i \in J} \mathbf{Y}_i\right],$$

for all $J \subset [s]$, it follows that for all $\delta > 0$,

$$\mathbb{P}[\mathbf{X} \geq (1 + \delta)\mu] \leq \left[\frac{e^\delta}{(1 + \delta)^{1 + \delta}}\right]^\mu.$$

In particular, for every $\sigma \geq 6\mu$, we have

$$\mathbb{P}[\mathbf{X} \geq \sigma] \leq e^{-\sigma}.$$

Motivated by the above theorem in the case where each \mathbf{Y}_i follows the law of Bernoulli $\text{Ber}(p)$, we introduce the following definition.

Definition 3.4. We say binary random variables $(\mathbf{X}_1, \dots, \mathbf{X}_s)$ are Ber(p)-dominated if

$$\mathbb{E} \left[\prod_{i \in J} \mathbf{X}_i \right] \leq p^{|J|},$$

for all $J \subseteq [s]$.

One can understand Theorem 3.3 as saying that the upper tail of the random variable is almost as small as that of a binomial random variable. In fact, one can bootstrap Theorem 1.1 to make this analogy even stronger, by showing that large upward deviations on disjoint sets are also Ber(p)-dominated for some p . The random variable \mathbf{R}_i in the next theorem is the indicator function of a large upward deviation on some set $Q_i \subset [s]$.

Corollary 3.5. *Suppose $(\mathbf{X}_1, \dots, \mathbf{X}_s)$ are Ber(p)-dominated. Let disjoint subsets $Q_1, \dots, Q_m \subseteq [s]$ of equal order be given, fix $\delta > 0$ and define the binary random variables*

$$\mathbf{R}_i := \left\{ \sum_{j \in Q_i} \mathbf{X}_j > (1 + \delta)p|Q_i| \right\},$$

for $i \in [m]$. Let q be the upper bound on $\mathbb{P}[\mathbf{R}_i = 1]$ given by Theorem 3.3. Then $(\mathbf{R}_1, \dots, \mathbf{R}_m)$ are Ber(q)-dominated.

Proof. Consider $J \subseteq [m]$; we wish to show that

$$\mathbb{E} \left[\prod_{i \in J} \mathbf{R}_i \right] \leq q^{|J|}.$$

To this end, let $I := \cup_{j \in J} Q_j \subset [s]$ and observe that, by assumption and Theorem 3.3, we have

$$\mathbb{P} \left[\sum_{i \in I} \mathbf{X}_i > (1 + \delta)p|I| \right] \leq \left[\frac{e^\delta}{(1 + \delta)^{1 + \delta}} \right]^{p|I|} = \prod_{j \in J} \left[\frac{e^\delta}{(1 + \delta)^{1 + \delta}} \right]^{p|Q_j|} = q^{|J|}.$$

To conclude we observe that if $\mathbf{R}_j = 1$ for all $j \in J$ then it follows that $\sum_{i \in I} \mathbf{X}_i > (1 + \delta)p|I|$. \square

We now show that the appearance of short lists at vertices in an independent set is Ber(p)-dominated for an appropriate p . This is crucial for both Theorems 2.7 and 2.8, and this could have many applications in other works. Theorem 3.6 can be seen as a strengthening of (1) from Theorem 1.1. In Theorem 1.1, we could derive (1) from one application of the induction hypothesis; going deeper into the induction allows us to derive the following stronger statement.

Theorem 3.6. *Let $\varepsilon, \Delta, k, \ell$ and G be as in Theorem 1.1. Let $0 < p < 1$ and let I be an independent set of G . Let σ be a uniformly random proper k -colouring of G . Then the random variables $\mathbf{X}_v := \{\ell_\sigma(v) \leq p\ell\}$ for $v \in I$ are Ber(p)-dominated.*

Proof. Let $J \subseteq I$, and label the vertices of J by $\{v_i\}_{i \in [s]}$. For every graph $H \subseteq G$, we may apply Theorem 1.1 and obtain that, for all $v \in V(H)$, if σ is drawn uniformly at random from $\mathcal{C}(H \setminus v)$, then $\mathbb{E}[\ell_\sigma(v)] \geq \ell$. In other words, we have $|\mathcal{C}(H)| \geq \ell |\mathcal{C}(H \setminus v)|$. Repeated applications of this identity yield that $|\mathcal{C}(G)| \geq |\mathcal{C}(G \setminus J)| \ell^s$. So we have

$$\mathbb{E} \left[\prod_{v \in J} \mathbf{X}_v \right] = \frac{\#\{\sigma \in \mathcal{C}(G) : \ell_\sigma(v_1), \dots, \ell_\sigma(v_s) \leq p\ell\}}{|\mathcal{C}(G)|} \leq \frac{|\mathcal{C}(G \setminus J)| \cdot (p\ell)^s}{|\mathcal{C}(G \setminus J)| \cdot \ell^s} = p^s.$$

The upper bound for the numerator in the above inequality comes from two facts that rely on J being an independent set. First, given a colouring $\sigma_0 \in \mathcal{C}(G \setminus J)$, the number of extensions of σ_0 to a colouring $\sigma \in \mathcal{C}(G)$ is $\prod_{v \in J} \ell_{\sigma_0}(v)$. Second, given such an extension σ , we have $\ell_\sigma(v) = \ell_{\sigma_0}(v)$ for every $v \in J$. We conclude that $(\mathbf{X}_v)_{v \in I}$ are Ber(p)-dominated. \square

Leveraging Theorem 3.6, and the fact that neighbourhoods induce independent sets, we can prove the following exponential upper bound on the likelihood of short lists. It also requires Lemma 3.7, the proof of which is adapted from that of [Mol19, Lemma 7] and is deferred to the Appendix.

Lemma 3.7. *Let G be a triangle-free graph, let $v \in V(G)$ and let σ_0 be a proper k -colouring of $G \setminus N[v]$ (for some k), with at least one extension to G . Then if σ is the uniformly random extension of σ_0 to G , writing $\ell := \mathbb{E}[\ell_\sigma(v)]$, we have*

$$\mathbb{P}[\ell_\sigma(v) \leq (1 - \delta)\ell] \leq e^{-\frac{\delta^2 \ell}{2}},$$

for all $\delta \in (0, 1)$.

Corollary 3.8. *Let $\varepsilon, \Delta, k, \ell$ and G be as in Theorem 1.1. Then for all $\delta \in (0, 1)$ and $v \in V(G)$*

$$\mathbb{P}[\ell_\sigma(v) < (1 - \delta)\ell] \leq 2e^{-\frac{\delta^2 \ell}{2}},$$

where σ is the uniformly random k -colouring of $G \setminus v$.

Proof. For every $u \in N(v)$, we let B_u be the random event that u has a short list, i.e. $\ell_\sigma(u) \leq t$ for $t := \lceil 5/\varepsilon \rceil$. Since G is triangle-free, $N(v)$ is an independent set. So by Theorem 3.6, the events $(B_u)_{u \in N(v)}$ are $\text{Ber}(t/\ell)$ -dominated. We apply Theorem 3.3, and obtain that the probability that more than $6t\Delta/\ell$ neighbours of v have a short list is at most $e^{-6t\Delta/\ell}$.

Now, let σ_0 be a possible realisation of $\sigma|_{G \setminus N[v]}$ such that no more than $6t\Delta/\ell$ neighbours of v have a short list in σ_0 . When Δ is large enough, this is at most $\varepsilon k/5$. In that case, we can repeat the computation in the proof of Theorem 1.1, and obtain that

$$\mathbb{E}[\ell_\sigma(v) \mid \sigma|_{G \setminus N[v]} = \sigma_0] \geq \ell.$$

By applying Lemma 3.7, we obtain that

$$\mathbb{P}[\ell_\sigma(v) \leq (1 - \delta)\ell \mid \sigma|_{G \setminus N[v]} = \sigma_0] \leq e^{-\frac{\delta^2 \ell}{2}}.$$

Overall, the probability that $\ell_\sigma(v) \leq (1 - \delta)\ell$ is therefore at most $e^{-6t\Delta/\ell} + e^{-\frac{\delta^2 \ell}{2}} \leq 2e^{-\frac{\delta^2 \ell}{2}}$ \square

It is interesting to observe that it is not possible to have a stronger form of Theorem 3.6 where we replace Bernoulli-domination with negative correlation. Indeed, given a graph G , if $u, v \in N(v)$ share the same neighbourhood, then the lists of u and v are perfectly correlated.

We finish this subsection by proving the following bootstrap percolation result which may be of independent interest. The setup is as follows. Consider the rooted Δ -ary tree T of depth f . That is, the tree constructed by starting from the root r , adding Δ children, and then adding Δ children to each of the leaves, and repeating this a total f times, so that the distance from r to each of the Δ^f leaves is exactly f . At step 1, we randomly activate a subset of the leaves. At each step $i \geq 2$, we activate a vertex if at least s of its children have been activated at step $i - 1$. We call this process s -upward percolation. Clearly, this process reaches a stable state in at most f steps. We will be interested in the probability that the root r is activated at the end of the process. Note that once we have chosen which leaves to activate, the process is deterministic. We are interested in the situation where the activation probabilities for the leaves are $\text{Ber}(p)$ -dominated.

We note that it is essential that the activation events are $\text{Ber}(p)$ -dominated. Indeed, an adversary only needs to activate s^f leaves in order to activate the root and if $s < \Delta$, then this is a vanishing proportion of all leaves as $f \rightarrow \infty$. However, in order to do so the adversary's activated leaves must be the leaves of an s -ary tree. This a very low entropy strategy, and we can use the renormalisation properties of $\text{Ber}(p)$ -dominated random variables (Corollary 3.5) to show that this is very unlikely.

Lemma 3.9. *Let $0 < p < 1$ be a real value, and let $\Delta \geq 2$, $s \geq \max\{6p\Delta, 3 \ln \Delta\}$, and $f \geq 1$ be integers. Suppose we perform s -upward percolation on a Δ -ary tree T of depth f rooted in a vertex r , where the events that the leaves are activated are $\text{Ber}(p)$ -dominated. Then we have⁸*

$$\mathbb{P}[r \text{ is activated}] \leq \exp\left(-s^{\lceil f/2 \rceil}\right).$$

Proof. Given a vertex $v \in V(T)$, let $T[v]$ denote the unique Δ -ary subtree of T rooted at v . For every $i \leq f$, we let $V_i(T)$ denote the set of nodes at depth i in T (in particular, we have $V_0(T) = \{r\}$). If r is activated, then there is an s -ary subtree of T of depth f that contains only activated vertices. In particular, at each depth $i \leq f$, there are at least s^i activated vertices. We conclude that, for every fixed $i \leq f$, we have

$$\mathbb{P}[r \text{ is activated}] \leq \mathbb{P}[V_i(T) \text{ contains at least } s^i \text{ activated vertices}].$$

Let T_1, \dots, T_n be a collection of disjoint trees of depth f , of respective roots r_1, \dots, r_n . We perform s -upward percolation on each of them, and assume that the events that the leaves are activated are $\text{Ber}(p)$ -dominated. Denote by A_j the random event that r_j is activated, for every $j \in [n]$. We will show by induction on f that the events (A_j) are $\text{Ber}(q)$ -dominated, for $q = \exp(-s^{\lceil f/2 \rceil})$. For the base case $f = 1$, for every $j \in [n]$ we have

$$\mathbb{P}[r_j \text{ is activated}] \leq \mathbb{P}[V_1(T_j) \text{ contains at least } s \text{ activated vertices}] \leq e^{-s} = q,$$

by Theorem 3.3 applied with $\mu := p\Delta$ and $\sigma := s \geq 6\mu$. We are in the setting of Corollary 3.5, so the events (A_j) are $\text{Ber}(q)$ -dominated.

We now assume that $f \geq 2$, and we fix $i := \lceil f/2 \rceil$. Let us consider any tree T_j for $j \in [n]$. For every $x \in V_i(T_j)$, let A_x be the event that x is activated. We apply induction on the collection of trees $T_j[x]$ of depth $f - i \geq (f - 1)/2$, for $x \in V_i(T_j)$, and obtain that the events $(A_x)_{x \in V_i(T_j)}$ are $\text{Ber}(q')$ -dominated, with $q' = \exp(-s^{\lceil (f-1)/4 \rceil})$. We apply Theorem 3.3 with $\mu := q'\Delta^i$ and $\sigma := s^i$ and obtain that

$$\mathbb{P}[r_j \text{ is activated}] \leq \mathbb{P}[V_i(T_j) \text{ contains at least } s^i \text{ activated vertices}] \leq e^{-s^i},$$

if we can prove that $\sigma \geq 6\mu$. We are again in the setting of Corollary 3.5, so the events $(A_j)_{j \in [n]}$ are $\text{Ber}(q)$ -dominated with $q = \exp(-s^{\lceil f/2 \rceil})$, as desired.

We now prove that we have $\sigma \geq 6\mu$. If $f \leq 5$, we have $q' = e^{-s}$ and $i \leq 3$, so this reduces to $(s/\Delta)^3 \geq 6e^{-s}$. Since $s^3 \geq (3 \ln 2)^3 > 6$, it suffices to prove that $e^s \geq \Delta^3$, which holds by assumption on s . If $f > 5$, it suffices to prove that $s^{\lceil (f-1)/4 \rceil} \geq \lceil f/2 \rceil \ln \Delta$. Since $s \geq 3 \ln \Delta \geq 2$, it suffices to prove that $2^{\lceil (f-5)/4 \rceil} \geq \frac{1}{3} \lceil f/2 \rceil$. It is easy to check that this holds for every integer $f > 5$. \square

3.3. Proofs of Theorem 2.7 and Theorem 2.8. Relying on the tail-bound given by Corollary 3.8 for the probability of having a short list, we can prove Theorem 2.7.

Proof of Theorem 2.7. Fix a vertex $v \in V(G)$, and let σ be the uniformly random k -colouring of $G \setminus v$. We will actually prove that, with high probability (with respect to Δ), we have $\ell_\sigma(u) \geq 2$ for every $u \in N(v)$. Under that condition, for every colour $x \in [k]$, we may sequentially resample $\sigma(u)$ from $L_\sigma(u) \setminus x$ for every $u \in N(v)$ (since $N(v)$ is an independent set, the lists $(L_\sigma(u))_{u \in N(v)}$ are not affected by this resampling). Then we can set $\sigma(v) \leftarrow x$. This proves that v is both $(\Delta + 1)$ -loose and thawed.

For every $u \in N(v)$, we let B_u be the random event that u has a short list, i.e. $\ell_\sigma(u) \leq \ell/2$, where $\ell = \Delta^{\varepsilon/2}$. By Corollary 3.8 we have $\mathbb{P}[B_u] \leq e^{-\ell/8}$. So by a union bound, the probability that no event B_u occurs is at least $1 - \Delta e^{-\ell/8} \geq 1 - e^{-\Delta^{\varepsilon/3}}$, assuming that Δ is large enough. The conclusion follows. \square

⁸With a more careful analysis one can replace $\lceil f/2 \rceil$ by $(1 - o_f(1))f$.

Relying on the percolation result stated in Lemma 3.9, we can sketch a proof of Theorem 2.8. The full-length proof lies in the Appendix.

Proof Sketch for Theorem 2.8. We let $g := (2 + o(1)) \ln \ln n / \ln \Delta$, and assume that the girth is at least $2g + 2$ (this holds under the assumption that it is at least $\ln \ln \Delta$ when Δ is large enough). We start by fixing a vertex v and sampling from σ . We then deterministically recolour the vertices at distance $g, g - 1, g - 2, \dots$ from v layer by layer (each layer is an independent set), so as to make the list sizes in the next layer as large as possible. If at the end of this process, all $u \in N(v)$ have at least 2 colours on their list, then it is straightforward to recolour v as we please by first recolouring its neighbours and so v is thawed and clearly we have changed the colours of at most $\Delta^g = O((\ln n)^2)$ vertices. Thus v is also $O((\ln n)^2)$ -loose. Using Corollary 3.2, we note that after the above process a vertex w with $\text{dist}(v, w) < g$ will only have a short list if at least $s := \Omega(\Delta / \ln \Delta)$ of its children have short lists. This allows us to bound from above the probability that $u \in N(v)$ has a short list by the probability that the s -upward percolation process on the $(\Delta - 1)$ -ary tree rooted at u ends with u activated. Because the vertices w at distance g from v form an independent set, Theorem 3.6 tells us that the distribution of short lists for said w are $\text{Ber}(p)$ -dominated. Thus we can use the percolation bound from Lemma 3.9 to derive the desired result. \square

3.4. Constructions. One could wonder whether the statement of Theorem 2.8 (w.h.p. all vertices are simultaneously thawed) can be extended to the case where G is triangle-free (rather than having girth $\ln \ln n$). We now show that this is not possible, even if the girth of G is an arbitrarily large constant.

We rely on the following construction from [BBP21].

Proposition 3.10 (Bonamy, Bousquet, Perarnau; 2021). *For every integers $d, g \geq 3$ there exists a d -regular graph G of girth at least g that has a proper $(d + 1)$ -colouring where every vertex is frozen.*

Proposition 3.11. *Let $g \geq 3$ be a given integer. Then for every integer $d \geq 3$ and n large enough, there exists an n -vertex d -regular graph G of girth at least g such that, letting σ be a uniformly random proper $(d + 1)$ -colouring of G , there are w.h.p. $\Theta(n)$ frozen vertices in σ as $n \rightarrow \infty$.*

Proof. Let G_0 be the d -regular graph G of girth g given by Proposition 3.10. We let $n_0 := |V(G_0)|$, and let $\sigma_0 \in \mathcal{C}_{d+1}(G_0)$ be such that every vertex $v \in V(G_0)$ is frozen in σ_0 . Let G consist of n/n_0 disjoint copies of G_0 , and let σ be a uniformly random proper colouring of G . For every copy H of G_0 within G , let E_H be the event that $\sigma|_{V(H)} = \sigma_0$. We have $\mathbb{P}[E_H] \geq 1/(d + 1)^{n_0}$, and the random events (E_H) are independent. So the number of copies of G_0 in G that are entirely frozen in σ follows the binomial distribution $\mathcal{B}(n/n_0, (d + 1)^{-n_0})$. By the standard estimates given by Chernoff-Hoeffding bounds (see Lemma A.3), the probability that less than $\frac{n}{2n_0(d+1)^{n_0}}$ copies of G_0 are entirely frozen is at most $\exp(-\frac{n}{8n_0(d+1)^{n_0}}) \xrightarrow{n \rightarrow \infty} 0$. We conclude that w.h.p. the number of frozen vertices in σ is at least $\frac{1}{2}(d + 1)^{-n_0} \cdot n = \Theta(n)$ as $n \rightarrow \infty$. \square

4. ACKNOWLEDGEMENT

A substantial part of this work has been done during the online workshop Entropy Compression and Related Methods which took place in March 2021. We are thankful to the organisers, Ross J. Kang and Jean-Sébastien Sereni, and more generally to the Sparse Graph Coalition for making that work possible. We are grateful to Felix Joos for proofreading the early versions of that paper. We also thank Matthieu Rosenfeld and Mike Molloy for insightful discussions.

REFERENCES

- [ABD21] James Anderson, Anton Bernshteyn, and Abhishek Dhawan. Coloring graphs with forbidden bipartite subgraphs. [arXiv preprint arXiv:2107.05595](https://arxiv.org/abs/2107.05595), 2021.

- [ACO08] Dimitris Achlioptas and Amin Coja-Oghlan. Algorithmic barriers from phase transitions. In 2008 49th Annual IEEE Symposium on Foundations of Computer Science, pages 793–802. IEEE, 2008.
- [AIS19] Dimitris Achlioptas, Fotis Iliopoulos, and Alistair Sinclair. Beyond the Lovász local lemma: Point to set correlations and their algorithmic applications. In 2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS), pages 725–744. IEEE, 2019.
- [AKS99] Noga Alon, Michael Krivelevich, and Benny Sudakov. Coloring graphs with sparse neighborhoods. Journal of Combinatorial Theory, Series B, 77(1):73–82, 1999.
- [AM97] Dimitris Achlioptas and Michael Molloy. The analysis of a list-coloring algorithm on a random graph. In Proceedings 38th Annual Symposium on Foundations of Computer Science, pages 204–212. IEEE, 1997.
- [AM02] Dimitris Achlioptas and Cristopher Moore. Almost all graphs with average degree 4 are 3-colorable. In Proceedings of the thirty-fourth annual ACM symposium on Theory of computing, pages 199–208, 2002.
- [ART06] Dimitris Achlioptas and Federico Ricci-Tersenghi. On the solution-space geometry of random constraint satisfaction problems. In Proceedings of the thirty-eighth annual ACM symposium on Theory of computing, pages 130–139, 2006.
- [AS16] Noga Alon and Joel H Spencer. The probabilistic method. John Wiley & Sons, 2016.
- [BBCK21] Anton Bernshteyn, Tyler Brazelton, Ruijia Cao, and Akum Kang. Counting colorings of triangle-free graphs. arXiv preprint arXiv:2109.13376, 2021.
- [BBP21] Marthe Bonamy, Nicolas Bousquet, and Guillem Perarnau. Frozen $(\delta+1)$ -colourings of bounded degree graphs. Combinatorics, Probability and Computing, 30(3):330–343, 2021.
- [Ber19] Anton Bernshteyn. The Johansson-Molloy theorem for DP-coloring. Random Structures & Algorithms, 54(4):653–664, 2019.
- [BG08] Antar Bandyopadhyay and David Gamarnik. Counting without sampling: Asymptotics of the log-partition function for certain statistical physics models. Random Structures & Algorithms, 33(4):452–479, 2008.
- [BKNP18] Marthe Bonamy, Tom Kelly, Peter Nelson, and Luke Postle. Bounding χ by a fraction of Δ for graphs without large cliques. arXiv preprint arXiv:1803.01051, 2018.
- [CL16] Péter Csikvári and Zhicong Lin. Sidorenko’s conjecture, colorings and independent sets. arXiv preprint arXiv:1603.05888, 2016.
- [Csi16] Péter Csikvári. Extremal regular graphs: the case of the infinite regular tree. arXiv preprint arXiv:1612.01295, 2016.
- [DdJdVVKP20] Ewan Davies, Rémi de Verclos, Ross J Kang, and François Pirot. Coloring triangle-free graphs with local list sizes. Random Structures & Algorithms, 57(3):730–744, 2020.
- [DF86] Martin E Dyer and Alan M Frieze. Fast solution of some random np-hard problems. In 27th Annual Symposium on Foundations of Computer Science (sfcs 1986), pages 331–336. IEEE Computer Society, 1986.
- [DKPS20] Ewan Davies, Ross J Kang, François Pirot, and Jean-Sébastien Sereni. Graph structure via local occupancy. arXiv preprint arXiv:2003.14361, 2020.
- [DP18] Zdeněk Dvořák and Luke Postle. Correspondence coloring and its application to list-coloring planar graphs without cycles of lengths 4 to 8. Journal of Combinatorial Theory, Series B, 129:38–54, 2018.
- [FNV06] Zoltan Füredi, Assaf Naor, and Jacques Verstraëte. On the turán number for the hexagon. Advances in Mathematics, 203(2):476–496, 2006.
- [Gam21] David Gamarnik. The overlap gap property: A topological barrier to optimizing over random structures. Proceedings of the National Academy of Sciences, 118(41):e2108492118, 2021.
- [GK12] David Gamarnik and Dmitriy Katz. Correlation decay and deterministic fptas for counting colorings of a graph. Journal of Discrete Algorithms, 12:29–47, 2012.
- [GŠV16] Andreas Galanis, Daniel Štefankovič, and Eric Vigoda. Inapproximability of the partition function for the antiferromagnetic ising and hard-core models. Combinatorics, Probability and Computing, 25(4):500–559, 2016.
- [Ili19] Fotios Iliopoulos. Stochastic Local Search and the Lovasz Local Lemma. University of California, Berkeley, 2019.
- [Joh96] Anders Johansson. The choice number of sparse graphs. Preliminary version, April, 1996.
- [KMRT⁺07] Florent Krzakala, Andrea Montanari, Federico Ricci-Tersenghi, Guilhem Semerjian, and Lenka Zdeborová. Gibbs states and the set of solutions of random constraint satisfaction problems. Proceedings of the National Academy of Sciences, 104(25):10318–10323, 2007.
- [Kuč89] Luděk Kučera. Graphs with small chromatic numbers are easy to color. Information processing letters, 30(5):233–236, 1989.

- [Mol12] Michael Molloy. The freezing threshold for k -colourings of a random graph. In Proceedings of the forty-fourth annual ACM symposium on Theory of computing, pages 921–930, 2012.
- [Mol19] Michael Molloy. The list chromatic number of graphs with small clique number. Journal of Combinatorial Theory, Series B, 134:264–284, 2019.
- [MT10] Robin A Moser and Gábor Tardos. A constructive proof of the general lovász local lemma. Journal of the ACM (JACM), 57(2):1–15, 2010.
- [MWW04] Brendan D McKay, Nicholas C Wormald, and Beata Wysocka. Short cycles in random regular graphs. the electronic journal of combinatorics, pages R66–R66, 2004.
- [PS97] Alessandro Panconesi and Aravind Srinivasan. Randomized distributed edge coloring via an extension of the chernoff–hoeffding bounds. SIAM Journal on Computing, 26(2):350–368, 1997.
- [PS21] François Pirot and Jean-Sébastien Sereni. Fractional chromatic number, maximum degree, and girth. SIAM Journal on Discrete Mathematics, 35(4):2815–2843, 2021.
- [Ros20] Matthieu Rosenfeld. Another approach to non-repetitive colorings of graphs of bounded degree. The Electronic Journal of Combinatorics, pages P3–43, 2020.
- [SSSZ20] Ashwin Sah, Mehtaab Sawhney, David Stoner, and Yufei Zhao. A reverse sidorenko inequality. Inventiones mathematicae, 221(2):665–711, 2020.
- [Tur88] Jonathan S Turner. Almost all k -colorable graphs are easy to color. Journal of algorithms, 9(1):63–82, 1988.
- [Vu02] Van H Vu. A general upper bound on the list chromatic number of locally sparse graphs. Combinatorics, Probability and Computing, 11(1):103–111, 2002.
- [Wei06] Dror Weitz. Counting independent sets up to the tree threshold. In Proceedings of the thirty-eighth annual ACM symposium on Theory of computing, pages 140–149, 2006.
- [Woo20] David R Wood. Nonrepetitive graph colouring. arXiv preprint arXiv:2009.02001, 2020.
- [WW20] Ian M Wanless and David R Wood. A general framework for hypergraph colouring. arXiv preprint arXiv:2008.00775, 2020.
- [Zde08] Lenka Zdeborová. Statistical physics of hard optimization problems. arXiv preprint arXiv:0806.4112, 2008.
- [ZK07] Lenka Zdeborová and Florent Krzakaa. Phase transitions in the coloring of random graphs. Physical Review E, 76(3):031131, 2007.

APPENDIX A. REMAINING PROOFS AND REQUIRED RESULTS

In this section we collect the remaining proofs and results. They are written in order of dependency, so that all results have either been cited or proved by the time they are used. Let G be a graph and $X \subset V(G)$. We write $G[X]$ for the subgraph of G induced by X . We need the following simple observation.

Lemma A.1. *Let H be a graph, and L a list-assignment of H such that $|L(v)| \geq \deg(v) + 1$ for every vertex $v \in V(H)$. If σ is a uniformly random proper L -colouring of H , then given some colour $x \in \bigcup_{v \in V(H)} L(v)$, the probability that $\sigma(v) \neq x$ for all $v \in V(H)$ is at least*

$$\prod_{v \in V(H)} \left(1 - \frac{1}{|L(v)| - \deg(v)} \right).$$

Proof. Since H is greedily L -colourable, we can sample a uniform proper L -colouring σ of H . Let $v \in V(H)$. For every L -colouring σ' of $H' := H \setminus v$, one has

$$\mathbb{P}[\sigma(v) = x \mid \sigma|_{H'} = \sigma'] \leq \frac{1}{|L(v)| - \deg(v)},$$

since after removing the colours in $\sigma'(N(v))$ from $L(v)$, there remains at least $|L(v)| - \deg(v)$ possible choices for $\sigma(v)$ which are equiprobable. Sampling a uniformly random L -colouring and then resampling the colour of each vertex once gives the result, as the resulting random L -colouring is also uniformly distributed. \square

The proof of Theorem 2.10 will rely on a more advanced version of the Coupon–Collector Lemma that was needed to prove Theorem 1.1. The statement is about the expected number of colours

that a uniformly random list-colouring of a given graph H does not use, when the input is a random list-assignment of H .

Lemma A.2 (Generalised Coupon-Collector Lemma). *Let H be a graph on n vertices, and \mathbf{L} a random list-assignment of H , such that H is deterministically \mathbf{L} -colourable. Fix some $t \geq 1$, and define the random variable $\mathbf{X} := \#\{v \in V(H) : |\mathbf{L}(v)| < (\deg(v) + 1)(t + 1)\}$. Let σ be a uniformly random \mathbf{L} -colouring of H , and define the random variable $L_\sigma := [k] \setminus \sigma(V(H))$. Then*

$$\mathbb{E}[|L_\sigma|] \geq k_0 e^{-(1+\frac{1}{t})\frac{n}{k_0}}, \quad \text{where } k_0 = k - \mathbb{E}[\mathbf{X}].$$

Proof. Let us fix a realisation $(L(v))_{v \in V(H)}$ of \mathbf{L} . We let $S = \{v : |L(v)| < (\deg(v) + 1)(t + 1)\}$ be the set of vertices with a short list, and $B = V(H) \setminus S$ be the set of vertices with a big list. We observe that since $|L(v)| \geq (\deg(v) + 1)(t + 1)$ for every $v \in B$, we have $\frac{1}{|L(v)| - \deg(v) - 1} \leq (1 + \frac{1}{t})\frac{1}{|L(v)|}$. Combining this observation with one due to Molloy [Mol19], we obtain

$$(5) \quad \sum_{x \in [k]} \sum_{\substack{v \in B \\ x \in L(v)}} \frac{1}{|L(v)| - \deg(v) - 1} \leq (1 + \frac{1}{t}) \sum_{x \in [k]} \sum_{\substack{v \in B \\ x \in L(v)}} \frac{1}{|L(v)|} \leq (1 + \frac{1}{t}) \sum_{v \in B} \sum_{x \in L(v)} \frac{1}{|L(v)|} \leq (1 + \frac{1}{t}) n.$$

Let us fix the realisation σ_0 of $\sigma|_S$, and let $X_0 := [k] \setminus \sigma_0(S)$. Note that $|X_0| \geq k - |S|$. We define $\mathbf{X}_0 := X_0 \setminus \sigma(B)$; note that \mathbf{X}_0 is precisely L_σ under the condition that $\sigma|_S = \sigma_0$. We let $L_0(v) := L(v) \setminus \sigma_0(N(v) \cap S)$ for every $v \in B$; so we have $|L_0(v)| \geq |L(v)| - \deg_S(v)$. Conditioned on $\sigma|_S = \sigma_0$, σ induces a uniformly random L_0 -colouring σ_1 of $H[B]$. We apply Lemma A.1 on $H[B]$ with the list-assignment L_0 , and obtain

$$\begin{aligned} \mathbb{E}[|\mathbf{X}_0|] &= \sum_{x \in X_0} \mathbb{P}[x \notin \sigma_1(B)] \\ &\geq \sum_{x \in X_0} \prod_{\substack{v \in B \\ x \in L_0(v)}} \left(1 - \frac{1}{|L_0(v)| - \deg_B(v)}\right) \geq \sum_{x \in X_0} \prod_{\substack{v \in B \\ x \in L(v)}} \left(1 - \frac{1}{|L(v)| - \deg(v)}\right) \\ &\geq \sum_{x \in X_0} \exp\left(-\sum_{\substack{v \in B \\ x \in L(v)}} \frac{1}{|L(v)| - \deg(v) - 1}\right) \quad \text{since } 1 - \frac{1}{z} > e^{-\frac{1}{z-1}} \text{ for every } z > 1; \\ &\geq |X_0| \exp\left(-\frac{1}{|X_0|} \sum_{x \in X_0} \sum_{\substack{v \in B \\ x \in L(v)}} \frac{1}{|L(v)| - \deg(v) - 1}\right) \quad \text{by convexity of exp;} \\ &\geq (k - |S|) \exp\left(-\frac{(1 + \frac{1}{t})n}{k - |S|}\right) \quad \text{by (5).} \end{aligned}$$

There remains to average over all possible realisations of \mathbf{L} . We use Jensen's inequality together with the convexity of the function $z \mapsto ze^{-C/z}$ for every $C > 0$ over the interval $(0, +\infty)$, and obtain that

$$\mathbb{E}[|\mathbf{L}|] \geq \mathbb{E}\left[(k - \mathbf{X}) \exp\left(-\frac{(1 + \frac{1}{t})n}{k - \mathbf{X}}\right)\right] \geq \mathbb{E}[(k - \mathbf{X})] \exp\left(-\frac{(1 + \frac{1}{t})n}{\mathbb{E}[k - \mathbf{X}]}\right) = k_0 e^{-(1+\frac{1}{t})\frac{n}{k_0}},$$

where $k_0 = k - \mathbb{E}[\mathbf{X}]$. □

Before proceeding with the formal proof of Theorem 2.10, we describe the random experiment at its core. Let G be a graph and $v \in V(G)$. We wish to prove that for σ drawn uniformly at random from the set of proper k -colourings of $G' := G \setminus v$, the list of available colours at v given σ is large in expectation. We do so by analysing the following random procedure.

- (i) Sample a proper k -colouring of G' uniformly at random;
- (ii) mark all vertices in $N(v)$ that have short lists due to the colouring on $G' \setminus N(v)$;
- (iii) uncolour all unmarked vertices in $N(v)$;
- (iv) choose a proper re-colouring of the uncoloured vertices in $N(v)$ uniformly at random.

Conveniently, the random proper colouring obtained at the end of this experiment is once again uniformly distributed across all proper colourings of G' . This is because the vertices we mark in step (ii) are selected solely based on the colouring of $G' \setminus N(v)$, which remains fixed after step (i). The proof proceeds by computing a lower bound on the expected size of the list of available colours at v after step (iv), thus proving the desired induction hypothesis. We restate the theorem for the reader's convenience.

Theorem 2.10. *Let G be an n -vertex graph of maximum degree Δ such that every graph induced by a neighbourhood in G has average degree at most $d \leq \frac{\Delta}{6} - 1$. Write $\rho := \Delta/(d+1)$ and let $\ell \geq (d+1)(\ln \rho)^3$. Then for every list-assignment $L: V(G) \rightarrow 2^{\mathbb{N}}$ with*

$$|L(v)| \geq \left(1 + \frac{2}{\ln \rho}\right) \frac{\deg(v)}{W\left(\frac{\deg(v)}{\ell}\right)}$$

for every vertex $v \in V(G)$, the following holds. For all $v \in V(G)$, the uniformly random proper L -colouring σ of $G \setminus v$ satisfies $\mathbb{E}[\ell_{\sigma}(v)] \geq \ell$.

Proof of Theorem 2.10. Fix $\rho := \frac{\Delta}{d+1} \geq 6$, $t := (d+1)(\ln \rho + 1)$, and $\ell \geq (d+1)(\ln \rho)^3$. For every $v \in V(G)$, let $k(v) := \left(1 + \frac{2}{\ln \rho}\right) \frac{\deg(v)}{W\left(\frac{\deg(v)}{\ell}\right)}$. Note that $k(v) \geq \frac{\deg(v)}{W\left(\frac{\deg(v)}{\ell}\right)} = \ell e^{W(\deg(v)/\ell)} \geq \ell$, because $W(x) \geq 0$ for every $x \geq 0$.

Let L be any list-assignment of G such that $|L(v)| \geq k(v)$ for every vertex $v \in V(G)$. Let $\mathcal{C}(H)$ denote the set of proper L -colourings of H . We show by induction that for all induced subgraphs $H \subseteq G$ we have

$$(\star) \quad |\mathcal{C}(H)| \geq \ell |\mathcal{C}(H \setminus v)|,$$

for all $v \in V(H)$. Observe that, given a colouring $\sigma \in \mathcal{C}(H \setminus v)$, the number of extensions of σ to a colouring in $\mathcal{C}(H)$ is precisely $\ell_{\sigma}(v)$. Hence (\star) is equivalent to $\mathbb{E}[\ell_{\sigma}(v)] \geq \ell$, for a uniformly random colouring $\sigma \in \mathcal{C}(H \setminus v)$.

The base case with $H = v$ for some $v \in V(G)$ follows as we have $|\mathcal{C}(\emptyset)| = 1$, and $|\mathcal{C}(H)| = k(v) \geq \ell$. Suppose now that $|V(H)| \geq 2$, and that the induction hypothesis (\star) holds for all induced subgraphs of $H' := H \setminus v$. We let σ be drawn uniformly at random from $\mathcal{C}(H')$, and we let σ_0 be obtained from σ by uncolouring $N_H(v)$; thus σ_0 is a proper L -colouring of $H_0 := H' \setminus N(v)$. For every $u \in N_H(v)$, we denote by d_u the degree of u within $H[N_H(v)]$, and we let $t_u := (d_u + 1)(\ln \rho + 1)$. Hence the average of t_u over all $u \in N_H(v)$ is at most t . Given the realisation of σ , we say that a neighbour $u \in N(v)$ of v has a short list if $\ell_{\sigma_0}(u) \leq t_u$, and we let S_{σ} be the set of vertices $u \in N(v)$ with short lists. First we observe that the expected size of S_{σ} is small. By the induction hypothesis (\star) , and using again the observation that the number of extensions of a colouring σ to an additional vertex u is $\ell_{\sigma}(u)$, we know that

$$(6) \quad \mathbb{P}[\ell_{\sigma}(u) \leq t_u] = \frac{\#\{\sigma' \in \mathcal{C}(H') : \ell_{\sigma'}(u) \leq t_u\}}{|\mathcal{C}(H')|} \leq \frac{t_u \cdot |\mathcal{C}(H' \setminus u)|}{\ell \cdot |\mathcal{C}(H' \setminus u)|} \leq \frac{t_u}{\ell}.$$

Summing (6) over all $u \in N(v)$, and since $\ell_{\sigma}(u) \leq \ell_{\sigma_0}(u)$ for every $u \in N(v)$, we obtain that

$$\begin{aligned} \mathbb{E}[|S_{\sigma}|] &= \sum_{u \in N(v)} \mathbb{P}[\ell_{\sigma_0}(u) \leq t_u] \leq \sum_{u \in N(v)} \mathbb{P}[\ell_{\sigma}(u) \leq t_u] \leq \sum_{u \in N(v)} \frac{t_u}{\ell} \\ &\leq \frac{t \deg(v)}{\ell} \leq \frac{1}{\ln \rho} \frac{\deg(v)}{\ln \rho - 1} \leq \frac{1}{\ln \rho} \frac{\deg(v)}{W\left(\frac{\deg(v)}{\ell}\right)}, \end{aligned}$$

where we use that $\ln \rho - 1 \geq W(\rho/(\ln \rho)^3) \geq W(\deg(v)/\ell)$, since $\rho \geq 6$. To see this, observe that the function $x \mapsto \ln x - 1 - W(x/(\ln x)^3)$ is increasing when $x \in (1, +\infty)$, and has a positive value at $x = 6$.

We are now going to apply Lemma A.2 to the graph $H[N(v)]$ on $\deg(v)$ vertices, with the random list-assignment L_{σ_0} , where the colours have been renamed in such a way that $L(v) = [k(v)]$. We let σ_1 be a uniformly random L_{σ_0} -colouring of $H[N(v)]$; by Lemma A.2 we have

$$(7) \quad \mathbb{E}[\ell_{\sigma_1}(v)] \geq k_0 e^{-\left(1 + \frac{1}{\ln \rho}\right) \frac{\deg(v)}{k_0}},$$

where $k_0 = k(v) - \mathbb{E}[|S_{\sigma}|] \geq \left(1 + \frac{1}{\ln \rho}\right) \frac{\deg(v)}{W\left(\frac{\deg(v)}{\ell}\right)}$. We observe that, by construction, $\sigma_0 \cup \sigma_1$ and σ are identically distributed, hence we have

$$\mathbb{E}[\ell_{\sigma}(v)] = \mathbb{E}[\ell_{\sigma_1}(v)] \geq \left(1 + \frac{1}{\ln \rho}\right) \frac{\deg(v)}{W\left(\frac{\deg(v)}{\ell}\right)} e^{-W\left(\frac{\deg(v)}{\ell}\right)} = \left(1 + \frac{1}{\rho}\right) \ell.$$

This ends the proof of the induction. □

We now prove that Theorem 2.10 directly implies Corollary 2.11.

Corollary 2.11. *Let G be a graph of maximum degree Δ , such that every neighbourhood spans at most Δ^2/f edges, for some $1 \leq f \leq \Delta^2 + 1$. Then $\chi_{\ell}(G) \leq (1 + o(1)) \frac{\Delta}{\ln \min\{\Delta, f\}}$ as $f \rightarrow \infty$.*

Proof of Corollary 2.11. Let G satisfy the hypothesis of Corollary 2.11. It is well-known that there exists a Δ -regular graph H and a mapping $\varphi: V(H) \rightarrow V(G)$ such that G is an induced subgraph of H , and for every vertex $v \in V(H)$ the number of edges in $H[N(v)]$ equals that in $G[N(\varphi(v))]$ (see for instance the construction in [PS21, Lemma 6]). So we may assume that G is regular. The average degree in $G[N(v)]$ is therefore at most $d := 2\Delta/f$, for every $v \in V(G)$. We have

$$\frac{\Delta}{d+1} = \frac{\Delta f}{f+2\Delta} \geq \begin{cases} \frac{f}{3} & \text{if } f \leq \Delta, \\ \frac{\Delta}{3} & \text{if } f \geq \Delta. \end{cases}$$

Setting $\rho := \frac{\Delta}{d+1} \geq \min\{f/3, \Delta/3\}$, we let $L: V(G) \rightarrow 2^{\mathbb{N}}$ be any list-assignment of G with $|L(v)| \geq (1 + 2/\ln \rho) \frac{\Delta}{W(\rho/(\ln \rho)^3)}$. By Theorem 2.10, G is L -colourable. Hence

$$\chi_{\ell}(G) \leq \left(1 + \frac{2}{\ln \rho}\right) \frac{\Delta}{W\left(\frac{\rho}{(\ln \rho)^3}\right)} \leq (1 + o(1)) \frac{\Delta}{\ln \min\{\Delta, f\}},$$

as $f \rightarrow \infty$ (and therefore also $\Delta \rightarrow \infty$). □

Next we prove Theorem 2.12, which yields a lower bound on the number of colourings of a graph in the setting of Theorem 2.10.

Theorem 2.12. *Let G be an n -vertex graph of maximum degree Δ such that every graph induced by a neighbourhood in G has average degree at most $d \leq \frac{\Delta}{6} - 1$. Let $f := \Delta/(d+1)$ and suppose $L: V(G) \rightarrow 2^{\mathbb{N}}$ is a list-assignment with $|L(v)| \geq \left(1 + \frac{1}{\ln \rho}\right) q(v)$, where*

$$q(v) \geq \left(1 + \frac{1}{\ln \rho}\right) \frac{\deg(v)}{W\left(\frac{\deg(v)}{(d+1)(\ln \rho)^3}\right)}$$

for every vertex $v \in V(G)$. Then there are at least $(q/\sqrt{D/(d+1)})^n$ proper L -colourings of G , where D is the geometric mean of the degrees in G , and q is the geometric mean of $\{q(v)\}_{v \in V(G)}$.

Proof of Theorem 2.12. Following the same set-up as in the proof of Theorem 2.10 we have $k_0 = k(v) - \mathbb{E}[|S_{\sigma}|] \geq q(v)$, for every vertex $v \in V(H)$. We also note that in (7), we can replace $\deg(v)$ with $\deg_H(v)$. Then

$$(8) \quad \mathbb{E}[\ell_{\sigma}(v)] \geq k_0 e^{-\left(1 + \frac{1}{\ln \rho}\right) \frac{\deg_H(v)}{k_0}} \geq q(v) e^{-\left(1 + \frac{1}{\ln \rho}\right) \frac{\deg_H(v)}{q(v)}}.$$

We now let v_1, \dots, v_n be an ordering of $V(G)$ such that $(q(v_i))_{i=1}^n$ is non-decreasing. Letting H_1 be the empty graph, and $H_i := G[v_1, \dots, v_{i-1}]$ for every $2 \leq i \leq n$, we apply (8) on the pairs (H_i, v_i) for every $1 \leq i \leq n$ and obtain that the number of L -colourings of G is

$$\begin{aligned} |\mathcal{C}(G)| &\geq \prod_{i=1}^n q(v_i) e^{-\left(1 + \frac{1}{\ln \rho}\right) \deg_{H_i}(v_i)/q(v_i)} \\ &= q^n \exp\left(-\left(1 + \frac{1}{\ln \rho}\right) \sum_{uv \in E(G)} \min\left\{\frac{1}{q(u)}, \frac{1}{q(v)}\right\}\right) \\ &\geq q^n \exp\left(-\left(1 + \frac{1}{\ln \rho}\right) \sum_{uv \in E(G)} \left(\frac{1}{2q(u)} + \frac{1}{2q(v)}\right)\right) \\ &= q^n \exp\left(-\left(1 + \frac{1}{\ln \rho}\right) \sum_{i=1}^n \frac{\deg(v_i)}{2q(v_i)}\right) \\ &\geq q^n \exp\left(-\frac{1}{2} \sum_{i=1}^n \ln \frac{\deg(v_i)}{d+1}\right) = q^n \left(\frac{D}{d+1}\right)^{-n/2}. \quad \square \end{aligned}$$

Let us recall the statement of Lemma 3.7.

Lemma 3.7. *Let G be a triangle-free graph, let $v \in V(G)$ and let σ_0 be a proper k -colouring of $G \setminus N[v]$ (for some k), with at least one extension to G . Then if σ is the uniformly random extension of σ_0 to G , writing $\ell := \mathbb{E}[\ell_{\sigma}(v)]$, we have*

$$\mathbb{P}[\ell_{\sigma}(v) \leq (1 - \delta)\ell] \leq e^{-\frac{\delta^2 \ell}{2}},$$

for all $\delta \in (0, 1)$.

This is proved in the context of proper partial k -colourings in [Mol19, Lemma 7]. For completeness, we repeat the proof which holds almost readily in the context of Lemma 3.7. In particular, we will need the Chernoff-Hoeffding bound stated in [Mol19, Lemma 3(b)]. We say that a set of binary random variables $(\mathbf{Y}_i)_{i \in [s]}$ are negatively correlated if

$$\mathbb{E}\left[\prod_{i \in [s]} \mathbf{Y}_i\right] \leq \prod_{i \in [s]} \mathbb{E}[\mathbf{Y}_i].$$

Lemma A.3. *Suppose $\mathbf{X}_1, \dots, \mathbf{X}_s$ are binary random variables. Set $\mathbf{Y}_i := 1 - \mathbf{X}_i$, and $\mathbf{X} := \sum_{i \in [s]} \mathbf{X}_i$. Write $\mu := \mathbb{E}[\mathbf{X}]$. If $(\mathbf{Y}_i)_{i \in [s]}$ are negatively correlated, then for all $0 < \delta < 1$,*

$$\mathbb{P}[\mathbf{X} \leq (1 - \delta)\mu] \leq e^{-\frac{\delta^2 \mu}{2}}.$$

Proof of Lemma 3.7. For every $x \in L(v)$, we denote E_x the event that $x \notin L_\sigma(v)$. We first argue that the events $(E_x)_{x \in L(v)}$ are negatively correlated, i.e.

$$(9) \quad \mathbb{P} \left[\bigwedge_{x \in X} E_x \right] \leq \prod_{x \in X} \mathbb{P}[E_x],$$

for every $X \subseteq L(v)$. First observe that if $\mathbb{P}[E_x] = 0$ for some $x \in X$, then (9) trivially holds since both terms equal zero. Second, if $\mathbb{P}[E_x] = 1$ for some $x \in X$, then (9) is equivalent to its statement when we remove x from X . So, in order to prove (9), let us show that $\mathbb{P} \left[E_x \mid \bigwedge_{y \in Y} E_y \right] \leq \mathbb{P}[E_x]$, for every $x \in L(v)$ and $Y \subseteq L(v) \setminus x$, where we assume that $0 < \mathbb{P}[E_x] < 1$. This is equivalent to

$$(10) \quad \mathbb{P} \left[\bigwedge_{y \in Y} E_y \mid \overline{E_x} \right] \geq \mathbb{P} \left[\bigwedge_{y \in Y} E_y \right].$$

We show that (10) holds with a coupling argument⁹. For every $u \in N(v)$, we define $\sigma'(u) := \sigma(u)$ if $\sigma(u) \neq x$; otherwise $\sigma'(u)$ is drawn uniformly at random from $L(u) \setminus x$ (this is non-empty, otherwise we would have $\mathbb{P}[E_x] = 1$). Then σ' follows the distribution of σ under the condition $\overline{E_x}$. Moreover, by construction, whenever we have $Y \subseteq \sigma(N(v))$, we deterministically have $Y \subseteq \sigma'(N(v))$. So

$$\mathbb{P} \left[\bigwedge_{y \in Y} E_y \mid \overline{E_x} \right] = \mathbb{P} [Y \subseteq \sigma'(N(v))] \geq \mathbb{P} [Y \subseteq \sigma(N(v))] = \mathbb{P} \left[\bigwedge_{y \in Y} E_y \right],$$

as desired.

To finish the proof, we apply Lemma A.3. We have that $\ell_\sigma(v)$ is the number of colours $x \in L(v)$ such that E_x does not hold, and since we have shown that the events $(E_x)_{x \in L(v)}$ are negatively correlated, we infer that

$$\mathbb{P}[\ell_\sigma(v) \leq (1 - \delta)\ell] \leq e^{-\frac{\delta^2 \ell}{2}},$$

as desired. □

We finish with a fully-detailed proof of Theorem 2.8, which is our central result on the geometry of the solution space for colourings of high girth graphs.

Theorem 2.8. *For all $\varepsilon > 0$ the following holds for all k sufficiently large and $\Delta \leq (1 - \varepsilon)k \ln k$. Suppose G is a graph on n vertices with maximum degree Δ and girth at least $\ln \ln n$. Then for the uniformly random proper k -colouring σ of G , the following holds w.h.p. for every $v \in V(G)$:*

- (1) in σ , v is $O((\ln n)^2)$ -loose¹⁰;
- (2) in σ , v is thawed.

Proof of Theorem 2.8. Let us fix a vertex $v \in V(G)$. We write $s := \frac{\varepsilon k}{5} \geq \frac{\varepsilon}{5} \frac{\Delta}{\ln \Delta}$ and $g := \lceil \frac{2 \ln \ln n^3}{\ln(s-1)} + 1 \rceil$. When Δ is large enough, we have $2g + 2 \leq \ln \ln n$. Let U_i denote the set of vertices at distance exactly i from v . Thus $U_1 = N(v)$, and since the girth of G is at least $2g + 2$, it holds that U_i is an independent set for every $i \leq g$. Let σ_0 be chosen uniformly at random from $\mathcal{C}_k(G)$. We will show that, with probability at least $1 - 1/n^2$, we can find a sequence $\sigma_1, \dots, \sigma_{g-1} \in \mathcal{C}_k(G)$ with the following properties:

⁹This coupling did not appear in earlier proofs of this statement, (10) was assumed.

¹⁰With a more careful (and technical) analysis, the authors believe that this could be replaced with $O(\ln n)$ -loose.

- (1) for all $i \in [g - 1]$, $\sigma_i|_{V(G) \setminus U_{g-i+1}} = \sigma_{i-1}|_{V(G) \setminus U_{g-i+1}}$;
- (2) $\ell_{\sigma_{g-1}}(u) > 5/\varepsilon$ for all $u \in N(v)$.

We construct our sequence (deterministically) as follows. Given σ_{i-1} , for $i \in [g - 1]$, we let $\sigma_i \in \mathcal{C}_k(G)$ be such that $\ell_{\sigma_i}(w)$ is maximum for all $w \in U_{g-i}$, given that $\sigma_i|_{V(G) \setminus U_{g-i+1}} = \sigma_{i-1}|_{V(G) \setminus U_{g-i+1}}$. We note that we can simultaneously maximise $\ell_{\sigma_i}(w)$ for all $w \in U_{g-i}$ because every vertex in U_{g-i+1} is adjacent to exactly one vertex of U_{g-i} .

We first deal with the case where G is Δ -regular and for convenience we write $\ell_i(u)$ for $\ell_{\sigma_i}(u)$. Now suppose $\ell_i(u) \leq 5/\varepsilon$ for some $u \in U_{g-i}$ (we say that u has a short list in σ_i). By our definition of σ_i , there must be no colouring $\tau \in \mathcal{C}_k(G)$, such that $\tau(w) = \sigma_{i-1}(w)$ for all $w \in V(G) \setminus U_{g-i+1}$ and $\ell_\tau(u) > 5/\varepsilon$. When Δ is large enough, we have $\Delta^{\varepsilon/2} > 5/\varepsilon$, so by Corollary 3.2 we must have $\ell_{i-1}(w) \leq 5/\varepsilon$ for at least s neighbours w of u , $s - 1$ of which must lie in U_{g-i+1} .

We now observe the following. Fix some $r \in N(v)$ and let T be the $(\Delta - 1)$ -ary tree of depth $g - 1$ obtained by deleting v and then looking at the induced subgraph of $G - v$ within distance $g - 1$ of r . Let $X := \{w \in U_g \cap V(T) : \ell_0(w) \leq 5/\varepsilon\}$ be the leaves of T that have short lists in σ_0 . Then $\ell_i(u) \leq 5/\varepsilon$ only if u is eventually activated in the $(s - 1)$ -upward percolation on T where the set of leaves that are initially activated is X .

We know that U_g is an independent set, so by Theorem 3.6 the events that $\ell_0(w) < 5/\varepsilon$ for $w \in U_g$ are Ber(p)-dominated for $p = \frac{5}{\varepsilon \ell}$, where $\ell = \Delta^{\varepsilon/2}$. Hence, when Δ is large enough, we have $s - 1 \geq \max\{6p\Delta, 3 \ln \Delta\}$. We apply Lemma 3.9 with $f = g - 1$ and obtain that

$$\begin{aligned} \mathbb{P}[\ell_{g-1}(r) \leq 5/\varepsilon] &\leq \mathbb{P}[r \text{ is activated}] \leq \exp\left(-(s - 1)^{f/2}\right) \\ &\leq \exp\left(-e^{\frac{1}{2}(g-1)\ln(s-1)}\right) \leq \exp\left(-e^{\ln \ln n^3}\right) = \frac{1}{n^3}. \end{aligned}$$

Union bounding over all neighbours r of v we find that

$$\mathbb{P}[\ell_{g-1}(r) > 5/\varepsilon \text{ for every } r \in N(v)] \geq 1 - \frac{\Delta}{n^3} > 1 - \frac{1}{n^2}.$$

We do a union bound over all $v \in V(G)$, and obtain that with probability at least $1 - 1/n$ we find the desired sequence $\sigma_0, \dots, \sigma_{g-1}$ for every $v \in V(G)$.

We now observe that given such a colouring σ_{g-1} , we can change the colour of v to any other by first recolouring each $r \in N(v)$ to avoid the desired colour. Thus the existence of $\sigma_0, \dots, \sigma_{g-1}$ implies that v is thawed in σ_0 and is $O((\ln n)^2)$ -loose in σ_0 , since we only need to recolour vertices at distance at most g from v , and there are at most $\Delta^g = O((\ln n)^2)$ of them. Union bounding over all choices of v yields the result.

Finally we deal with the case where G is not regular. For each vertex w at distance less than $g + 1$ from v , with degree less than Δ we simply add rooted $(\Delta - 1)$ -ary trees of depth g and connect w to the root, until w has degree Δ . We then treat all newly added edges as invisible in that they can be monochromatic. Thus the lists of the newly added vertices are always just $[k]$, and their colours do not affect the lists of any other vertices. Therefore we once again have the crucial property that if some vertex $u \in U_{g-i}$ has $\ell_i(u) \leq 5/\varepsilon$, then at least s of its children w must have $\ell_{i-1}(w) \leq 5/\varepsilon$. Further the events that the leaves of T have short lists are once again Ber(p) dominated for the same value of p . This allows us to repeat the rest of the arguments. \square

UNAFFILIATED

Email address: ecoin.hurley@umail.ucc.ie

UNIVERSITÉ PARIS-SACLAY, FRANCE.

Email address: francois.pirrot@lisen.fr