THEORY OF COMPUTING, Volume 22 (2091), 2026, pp. 1–38
www.theoryofcomputing.org

An Algorithmic Framework for Colouring Locally Sparse Graphs

Ewan Davies* Ross J. Kang† François Pirot
Jean-Sébastien Sereni

Received November 11, 2021; Revised May 24, 2024; Published February 1, 2026

Abstract. We develop an algorithmic framework for graph colouring that reduces the problem to verifying a local probabilistic property of the independent sets. With this we give, for any fixed $k \geq 3$ and $\varepsilon > 0$, a randomised polynomial-time algorithm for colouring graphs of maximum degree Δ in which each vertex is contained in at most t copies of a cycle of length k, where $1 \leq t \leq \Delta^{\frac{2\varepsilon}{1+2\varepsilon}}/(\log\Delta)^2$, with $\lfloor(1+\varepsilon)\Delta/\log(\Delta/\sqrt{t})\rfloor$ colours.

This unifies, generalises and improves upon several notable results including those of Kim (1995) and Alon, Krivelevich and Sudakov (1999), and more recent ones of Molloy (2019) and Achlioptas, Iliopoulos and Sinclair (2019). The bound on the chromatic number is tight up to an asymptotic factor 2 due to the random regular graph, and as such our work coincides with a famous algorithmic barrier to colouring random graphs, greatly expanding the range of known algorithms for graph colouring that meet this barrier.

*The research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP7/2007–2013) / ERC grant agreement № 339109. Part of this work was done while the author was visiting the Simons Institute for the Theory of Computing.

†Supported by a Vidi grant (639.032.614) of the Dutch Research Council (NWO).

ACM Classification: F.2.2, G.2.2

AMS Classification: 68R10, 68R25, 68W20

Key words and phrases: algorithms, graph algorithms, graph coloring, randomized, partition function

© 2026 Ewan Davies, Ross J. Kang, François Pirot, and Jean-Sébastien Sereni
Licensed under a Creative Commons Attribution License (CC-BY)
DOI: 10.4086/toc.2026.v022a2091

1 Introduction

Let $G = (V, E)$ be a graph. An *independent set* in G is a vertex subset that induces no edges. The *independence number* $\alpha(G)$ of G is the cardinality of a largest independent set in G. The *chromatic number* $\chi(G)$ of G is the least number of parts in a partition of V into independent sets. Determining and bounding these structural parameters has been of fundamental importance to algorithms, optimisation, and operations research [32]. Moreover, these parameters have been central in the development of combinatorial mathematics, especially with respect to random graphs and extremal combinatorics [42, 23, 20, 21].

The algorithmic and combinatorial perspectives on these parameters are inextricably linked. As an example, with an interpretation of the binomial random graph $G_{n,1/2}$ as a model of average-case behaviour, Karp [33] asked in 1976 if for some positive ϵ there is a polynomial-time algorithm that outputs an independent set in $G_{n,1/2}$ of size $(1+\epsilon)\log_2 n$ with probability tending to 1 as the number of vertices n tends to infinity, that is, *with high probability (w.h.p.)*. It is a basic fact that $\alpha(G_{n,1/2}) \sim 2\log_2 n$ w.h.p., so Karp's question is about getting within a factor less than 2 of this existential result with an efficient algorithm. Karp's question remains open and has helped to provoke an influential, sustained series of investigations in random graph theory, see, e. g., [31] for a survey from the perspective of graph colouring. As another example, there are notorious gaps between the best-known upper and lower estimates on classical Ramsey numbers, but bounds have nevertheless proven useful towards approximation algorithms, see, e. g., [28, 7].

Our main contribution is a novel framework for graph colouring that can be applied in significant generality by verifying a simple probabilistic property of the independent sets of the graph to be coloured. Such verification is routine and can be done essentially optimally in classes of graphs where we have some local sparsity condition, including the classically important case of triangle-free graphs as well as natural generalisations such as graphs with few edges in any induced neighbourhood subgraph, and graphs with a fixed cycle as an excluded subgraph. In this work we focus on algorithmic aspects of our framework. We design a hybrid algorithm that resamples colours of vertices and backtracks when undesirable conditions are found, showing correctness and efficiency with a sophisticated criterion for such stochastic local search algorithms from [2]. In its simplest form, one can derive upper bounds on the chromatic number with our framework by analysing a suitable random colouring with the lopsided Lovász local lemma (as in the work of Bernshteyn [8]). In a companion paper [18] we explore a broader array of applications with this formulation, omitting any technicalities that one must overcome to obtain efficient algorithms with our approach. Here we reformulate our framework in the considerably more technical language of stochastic local search in order to show that efficient algorithms follow from our framework, and we tackle a technicality related to backtracking that is only relevant to the efficiency and not the implied upper bounds on the chromatic number. The motivation for the present work is essentially Karp's question (with a focus on colouring and the random regular graph), and we give an algorithm for colouring graphs with few copies of a fixed cycle (as a subgraph) that is tight up to the *algorithmic barrier*, i. e., a factor of 2 akin to the one that Karp asks us to beat.

One old and basic starting point for this research is the pursuit of global asymptotic structure in triangle-free graphs. The search for large independent sets in this context corresponds to the classic off-diagonal case of Ramsey numbers [42, 23, 3, 4, 43, 35, 10, 11, 26, 17], a foundational and profoundly difficult problem in combinatorics. The search for good colourings in this context, a related but more delicate task, is also an important challenge of classic origins, see [48, 45]. There is particular interest in graphs of bounded maximum degree, with natural links to approximation algorithms, see, e. g., [28, 7]. For colouring this interest originated in a question of Vizing [46] from 1968: what is the largest chromatic number taken over all triangle-free graphs of maximum degree Δ? (Without the triangle-free condition a straightforward greedy argument yields an upper bound of $\Delta + 1$, which is sharp for odd cycles and cliques.) Simultaneously strengthening a seminal result of Ajtai, Komlós and Szemerédi [4] for the independence number and answering Vizing's question up to the choice of leading asymptotic constant, Johansson [30] devised a sophisticated semi-random colouring procedure to establish an upper bound of $O(\Delta/\log\Delta)$. This semi-random colouring approach was refined by Pettie and Su [41], who simultaneously showed that the leading asymptotic constant above can be taken as $4 + o(1)$, and gave an efficient, distributed algorithm (in the $\mathcal{LOCAL}$ model of computation) for finding such colourings. Recently, in a dramatic advance, Molloy [36] employed *entropy compression* for a simplified proof of and an intriguing improvement over Johansson's result, quantitatively matching an analogous independence number bound of Shearer [43] as well as giving an efficient algorithm. Throughout the paper, log is the natural logarithm.

Theorem 1.1 (Molloy [36])**.** *For all $\epsilon > 0$, there exists Δ_0 such that if $\Delta \geq \Delta_0$, then $\chi(G) \leq (1+\epsilon)\Delta/\log\Delta$ for any given triangle-free graph G of maximum degree Δ. There is a randomised algorithm that in polynomial time w.h.p. constructs a certificate colouring of G.*

The fact that Molloy's result is algorithmic is especially important, as this means that it is tight up to the algorithmic barrier, which we now discuss briefly. There is a factor 2 gap between the *existence* result [27] stating that the random Δ-regular graph has chromatic number $(1/2 + o(1))\Delta/\log\Delta$ w.h.p., and the best-known *algorithmic* result, which w.h.p. finds a colouring with $(1 + o(1))\Delta/\log\Delta$ colours in polynomial time. One possible intuition for this is that the 'solution space' of q-colourings of the random Δ-regular graph undergoes a change in geometry as q goes from $(1+\epsilon)\Delta/\log\Delta$ to $(1-\epsilon)\Delta/\log\Delta$ [1]. The underlying phenomenon is that the collection of independent sets, as considered under a suitable and natural metric, abruptly transitions from a well-connected space into one with exponentially many well-separated pieces. After this transition, any algorithm for finding independent sets of the desired size, let alone colourings whose average part size is at least the desired size, must ably navigate this shattered space.

Theorem 1.1 provides an algorithm for colouring the random regular graph up to the barrier because for constant Δ the random Δ-regular graph is essentially a high girth graph. For any fixed cycle length we expect at most a constant number of cycles of at most this length, so the random regular graph is close to a triangle-free graph of large chromatic number. This statement tells us that one should not hope to easily improve upon Molloy's upper bound by adding a high-girth assumption to G as the algorithmic barrier remains, and also that perhaps

there is nothing special about triangles here. Indeed, by a result of Alon, Krivelevich, and Sudakov [6] we know that when G is a graph of maximum degree Δ that contains no cycles of a fixed length $k \geq 3$ we have $\chi(G) \leq O(\Delta/\log\Delta)$; however, this result has a large leading constant and so loses sight of the algorithmic barrier.[1]

1.1 Our contributions

We explicitly confirm the idea—motivated above via the random regular graph—that excluding a fixed cycle length permits an algorithm that is tight up to the algorithmic barrier. This is the main message of this work: our framework is precise enough to hit the algorithmic barrier under the natural and general condition of a single (fixed) excluded cycle length, and the following application of the framework is the first result of this precision that does not rely on the special structure of triangle-free graphs in some way—indeed, it *cannot* do so, since the condition ostensibly allows arbitrarily many triangles.

Theorem 1.2. *Fix an integer $k \geq 3$. For all $\epsilon > 0$, there exists some Δ_0 such that if $\Delta \geq \Delta_0$, then $\chi(G) \leq (1+\epsilon)\Delta/\log\Delta$ for any given C_k-free graph G of maximum degree Δ. There is a randomised algorithm that in polynomial time w.h.p. constructs a certificate colouring of G.*

As a further example of the strength of our framework, we actually prove a smooth generalisation of this result that permits each vertex to be contained in few copies of C_k where the number of colours required degrades slowly with the parameter quantifying 'few'.

Theorem 1.3. *Fix an integer $k \geq 3$. For all $\epsilon > 0$, there exists Δ_0 such that if $\Delta \geq \Delta_0$ and $1 \leq t \leq \Delta^{\frac{2\epsilon}{1+2\epsilon}}/(\log\Delta)^2$, then $\chi(G) \leq (1+\epsilon)\Delta/\log(\Delta/\sqrt{t})$ for any given graph G of maximum degree Δ where each vertex of G is contained in at most t copies of the cycle C_k. There is a randomised algorithm that in polynomial time w.h.p constructs a certificate colouring of G.*

Theorem 1.2 follows from the statement of Theorem 1.3 for any constant $t > 0$. We disallow the case $t = 0$ in Theorem 1.3 for technical reasons, but for C_k-free graphs the case $t = 1$ applies and gives the required bound. In the proofs it is natural to work with a slightly weaker condition where C_k is replaced by F_k, the graph formed from a path P_{k-1} on $k-1$ vertices by adding a vertex joined to all vertices of the path (see Figure 1). Since F_k contains the cycle C_k (and indeed all smaller cycles), a graph with at most t copies of C_k must have at most t copies of F_k. We thus obtain the above results from the version assuming the weaker condition of a bound on the number of copies of F_k.

The generalisation to few cycles is interesting as it allows our method to be applied directly to arbitrary graphs, and this result hits the algorithmic barrier for the binomial random graph, as shown [2] for the case $k = 3$. These applications of our framework are the most interesting and important from the perspective of algorithms as they are known to meet the algorithmic barrier. We point out that our framework can be applied to many more types of graph, and has implications on the probabilistic structure of independent sets that are interesting in their

[1] An efficient algorithm was not discussed in [6] either, but the algorithmic aspects of this problem were treated in [2], which we discuss in detail later.

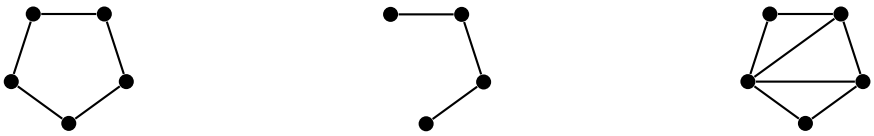


Figure 1: Example cycle, path and fan graphs.

own right and which can also be used to prove in certain cases that the method is tight, see [18] for more details. Indeed, we also show in Section 5 that the graph property our proofs of Theorem 1.1–Theorem 1.3 rely on is quantitatively optimal, and hence to improve upon the number of colours required necessitates a more sophisticated, and perhaps more global, approach. This gives an additional perspective on the algorithmic barrier to colouring random regular graphs (complementing the aforementioned shattering phenomenon) and highlights a central difficulty in studying Karp's question, as it seems hard to exploit global structure efficiently.

An affirmative answer to Karp's question or its analogue for random regular graphs would be considered an unexpected and sensational achievement. In this work we must be satisfied instead with results that greatly extend our understanding of independent sets and graph colouring up to the barrier that Karp asks us to cross. For another hint of how difficult it might be to improve on our results (and the prior works [36, 2]), particularly with respect to the asymptotic leading constant of 1, one can also consider that lowering the leading constant would improve upon the best-known lower bounds on the classical off-diagonal Ramsey numbers due to Shearer [43], as such constituting a dramatic advance in quantitative Ramsey theory. That is, regarding the existence of independent sets, let alone colourings or efficient constructions, any improvement to the leading constant in the number of colours that we require would be a major breakthrough.

1.2 Technical overview

The main advance of this work is a framework which gives sufficient conditions for a sophisticated stochastic local search algorithm to rapidly converge to a *flawless* partial colouring of a graph. We give a randomised algorithm exploring partial colourings of a graph that can resample colours in neighbourhoods and backtrack in the sense of uncolouring vertices when problematic conditions arise. The goal of this algorithm is to find a proper colouring of an induced subgraph of the input graph such that any uncoloured vertex u has (i) many available colours to extend the proper colouring to u, and (ii) there is not too much competition for the colours that are available for u amongst the neighbours of u. In this way our algorithm is a direct descendant of one given by Molloy [36]. In more detail, Molloy gave two algorithms, one for triangle-free graphs and one for the more general condition of K_r-freeness. The algorithm for triangle-free graphs is tight up to the algorithmic barrier, but relies heavily on the fact that a neighbourhood is an independent set. Our algorithm generalises Molloy's K_r-free approach, and one major

advance is to give a precise enough analysis to meet the algorithmic barrier despite the presence of edges in neighbourhoods. Once we have a flawless partial colouring, to obtain a proper colouring of the entire input graph it suffices to randomly assign colours to the uncoloured vertices. This is a standard step that is present in many previous works.

In order to obtain the desired applications, Theorem 1.2 and Theorem 1.3, we must verify that the sufficient conditions hold in these cases. The first step is to solve a simple optimisation problem to show that in a graph satisfying the assumptions of Theorem 1.3, independent sets are suitably large and well distributed to permit our framework to apply, and the second step is to prove a technical result related to sampling independent sets in graphs with no copy of the path on $k-1$ vertices in time polynomial in the number of vertices. The fact that each of these steps requires the proof of a simple statement with no reference to graph colouring demonstrates the power of our framework and that it can be easily applied 'as a black box'. Furthermore, if one is not interested in efficient algorithms then it suffices to perform only the first step.

To obtain a polynomial-time algorithm that operates in the presence of few copies of C_k we incorporate a key idea of Achlioptas, Iliopoulos, and Sinclair [2], who showed a 'few triangles' generalisation of Theorem 1.1 (see Theorem 2.1 in the following section). That proof breaks down if there are too many edges in the neighbourhood of a vertex: it proceeds by ignoring such edges, essentially reducing to the triangle-free case, and showing that when there are few ignored edges their influence on the search for a flawless partial colouring can be managed. In order to obtain a significantly more general algorithm we adapt this idea, but we do not reduce to the triangle-free case as this is too costly. We show that the sophisticated local search result they developed [2] can be applied when, instead of ignoring all the edges in a neighbourhood, we efficiently identify a small set of carefully chosen edges to ignore. That is, the present work contains both (i) a generalisation of Molloy's result (Theorem 1.1) to arbitrary fixed cycle lengths, and (ii) a corresponding generalisation of Achlioptas, Iliopoulos, and Sinclair's technique that permits few copies of the excluded graph. In a sense our technical achievements are to obtain results in the absence of the strong structure provided by graphs with few triangles that (i) are tight up to the algorithmic barrier, and (ii) smoothly degrade if few copies of the excluded graph are introduced.

A crucial feature of the algorithms presented here is that they run in time polynomial in the parameters n, Δ, and q (the number of vertices and maximum degree of the input graph, and the size of the colour lists). Our approach involves sampling independent sets in a *cover graph* that one builds from the input graph and colour lists. This is trivially possible in time $\exp(O(\Delta q))$, and indeed one can deduce algorithms with this behaviour by combining Molloy's approach with the probabilistic analysis in our companion paper [18]. This work, focusing on more efficient algorithms, contains the significant extra sophistication necessary to reduce this dependence to a polynomial.

1.3 Organisation

In Section 2 we discuss the history of the problems we study and the methods involved. We also give the necessary preliminary definitions and state the results that constitute our algorithmic framework for graph colouring. In Section 3 we verify the probabilistic information necessary

to apply our framework to prove Theorem 1.2 and Theorem 1.3, showing how our framework distils the problem of graph colouring into self-contained steps. The technical details of our framework are largely confined to Section 4 where we give the necessary proofs. In Section 5 we discuss one way in which our methods are provably tight and we give concluding remarks in Section 6.

2 Preliminaries and framework statements

Another starting point for our line of research is the modern investigation of stochastic local search algorithms. In broad terms, given a state space equipped with a probability measure that has designated *flawed subsets* (or *flaws*), under what circumstances is there an efficient randomised algorithm, performing *local* moves, to arrive at a flawless state? (One can think of a flawless state as, say, a proper colouring or a satisfying assignment to a CNF formula.) In a remarkable breakthrough, Moser [38] (see [39]) showed that the Lovász local lemma [24]—a fundamental result for proving the existence of combinatorial structures with the probabilistic method—follows from an elementary stochastic search algorithm based on resampling parts of the current state. In his analysis, Moser devised the entropy compression method mentioned earlier, and this has since found wide applicability to various search algorithms that backtrack to avoid problematic regions of the state space, see, e. g., [25]. Achlioptas, Iliopoulos, and Sinclair [2] recently gave a powerful algorithmic form of the local lemma that permits the analysis of *hybrid* algorithms, that can both resample and backtrack. As their main application, they gave the following generalisation of Theorem 1.1, under a smooth relaxation of the triangle-free condition.

Theorem 2.1 (Achlioptas, Iliopoulos, and Sinclair [2])**.** *For all $\epsilon > 0$, there exists Δ_0 such that if $\Delta \geq \Delta_0$ and $1 \leq t \leq \Delta^{\frac{2\epsilon}{1+2\epsilon}}/(\log\Delta)^2$, then $\chi(G) \leq (1+\epsilon)\Delta/\log(\Delta/\sqrt{t})$ for any given graph G of maximum degree Δ where each vertex of G is contained in at most t triangles. There is a randomised algorithm that in polynomial time w.h.p. constructs a certificate colouring of G.*

This result is an example showing the inextricable algorithmic and combinatorial perspectives on graph colouring: as an existence result it was a breakthrough and yet it also entails an efficient construction. The existence result is a sharpening of an aforementioned theorem of Alon, Krivelevich, and Sudakov [6] who used a probabilistic argument to reduce to the triangle-free case and applied Johansson's result (which we recall is a precursor to Theorem 1.1 with a larger leading constant). This brings us to another body of research motivating the present work, namely the results of Shearer [43, 44] and Alon [5] on the independence number of 'locally sparse' graphs. Johansson devised a number of technically impressive graph colouring results (largely unpublished, see [30, 37, 7]) using insights gained from these works on independence number and a careful iterative colouring procedure. The proofs are somewhat technically involved and drawbacks of the method include the need for delicate changes to adapt to different local sparsity conditions, as well as an inevitable loss with respect to the leading constants in the number of colours with this approach.

Molloy's main achievement in [36] was to give a much simpler method (further simplified by Bernshteyn [8]) that works for each of the local sparsity conditions studied by Johansson, but only for triangle-free graphs is the method precise enough to hit the algorithmic barrier. Indeed, for more general definitions of locally sparse Molloy's method relies upon insights gained from [44, 5] and does not surpass the precision of these works. Bonamy, Kelly, Nelson, and Postle [12] further refined the method and embraced a 'unified' approach for different definitions of local sparsity, but again they rely on statements about independent sets from previous works (including also [7]). In this work and our previous one [18], we develop a general framework for global graph structure that simultaneously strengthens Theorem 1.1 and Theorem 2.1 to give Theorem 1.2 and Theorem 1.3, and directly implies sharper versions of the independence number results in [44, 5] and colouring results in [2, 6, 8, 15, 14, 30, 34, 36, 47]. The framework in general reduces the main task to the verification of a probabilistic property of the independent sets that we call *local occupancy*. In several applications this verification is straightforward, resulting in simplified proofs of existing results with matching or improved bounds, see [18]. Moreover, subject to mild extra conditions, we can give polynomial-time constructions, which is our focus here. We overcome the two drawbacks in prior work identified above: our colouring method is generic and does not need tailoring to each application, and with it we expand the situations in which we can colour up to the algorithmic barrier.

Another benefit of our framework is that it implies a lower bound on the expected size of a random independent set from a distribution known as the hard-core model, or equivalently the *occupancy fraction* of the graph. We prove that our analysis of local occupancy behind Theorem 1.3 is asymptotically tight—that is, within an asymptotic factor 2 gap—as $\Delta \to \infty$ because the consequence in terms of occupancy fraction is tight. Unsurprisingly the construction showing tightness is the random regular graph, further strengthening the message that our work effectively hits the algorithmic barrier.

2.1 Local occupancy

We write $\mathcal{I}(G)$ for the set of independent sets in a graph G, and $\mu_{G,\lambda}$ for the *hard-core model on G at fugacity λ*, the probability measure on $\mathcal{I}(G)$ with

$$\mu_{G,\lambda}(I) = \frac{\lambda^{|I|}}{Z_G(\lambda)},$$

where $\lambda > 0$ is the *fugacity parameter*, and $Z_G(\lambda) = \sum_{I \in \mathcal{I}(G)} \lambda^{|I|}$ is the *partition function*. Its *occupancy fraction* is $E[|\mathbf{I}|]/|V(G)|$, where $\mathbf{I} \sim \mu_{G,\lambda}$, i. e., the expected fraction of the vertices in a random sample from $\mu_{G,\lambda}$. We frequently drop subscripts when they are clear from context.

Our framework outputs a range of structural information for a graph in terms of a condition related to the hard-core model which we call *local occupancy*. This systematic approach began in previous works [16, 17, 37, 14, 15] on occupancy fraction and fractional colouring.

Definition 2.2. We say that the hard-core model on a graph G at fugacity λ has *strong local*

(β, γ)-*occupancy* if, for every vertex $u \in V(G)$ and subgraph $F \subset G[N(u)]$ we have

$$\beta \frac{\lambda}{1+\lambda} \frac{1}{Z_F(\lambda)} + \gamma \frac{\lambda Z_F'(\lambda)}{Z_F(\lambda)} \geq 1 \,.$$

One can check that this condition implies $\beta \Pr(u \in \mathbf{I}) + \gamma E[|\mathbf{I} \cap N(u)|] \geq 1$ for any vertex $u \in V(G)$, where $\mathbf{I} \sim \mu_{G,\lambda}$. Since $\Pr(u \in \mathbf{I})$ can be seen as the occupancy probability of u and $E[|\mathbf{I} \cap N(u)|]$ as the expected number of occupied neighbours, this motivates the name of the property.[2]

2.2 List colouring

A key insight that we use to connect local occupancy to colouring is that given a graph G one can construct an auxiliary graph H in which independent sets correspond to (partial) proper colourings of G. This construction also comfortably incorporates the more general concept of *list colouring*. A *q-list assignment* of G is a function L such that $L(u)$, for each vertex $u \in V(G)$, is a list of colours (arbitrary labels that we conventionally take as natural numbers) of size q, and an *L-colouring* of G is a colouring with no monochromatic edges such that the colour of u is a member of $L(u)$ for each vertex u. Then the *list chromatic number* $\chi_\ell(G)$ of G is the least integer q such that every q-list assignment L admits an L-colouring. We can observe that a proper q-colouring of G is an L-colouring where $L(u) = \{1, \dots, q\}$ for each vertex $u \in V(G)$; hence $\chi(G) \leq \chi_\ell(G)$ always. Note that Johansson's and Molloy's bounds mentioned above were also shown for χ_ℓ. We represent list colourings through an auxiliary *cover graph* with notation as in [19].

Definition 2.3. Given a graph G, a *cover* of G is a pair $\mathscr{H} = (L, H)$, consisting of a graph H and a mapping $L\colon V(G) \to 2^{V(H)}$, satisfying the following requirements:

1. the sets $L(u)$ for $u \in V(G)$ form a partition of $V(H)$;
2. for every $u \in V(G)$, the graph $H[L(u)]$ is complete;
3. if $E_H(L(u), L(v)) \neq \emptyset$, then either $u = v$ or $uv \in E(G)$;
4. if $uv \in E(G)$, then $E_H(L(u), L(v))$ is a matching (possibly empty).

A cover $\mathscr{H} = (L, H)$ of G is *q-fold* if $|L(u)| = q$ for all $u \in V(G)$. An *$\mathscr{H}$-colouring* of G is an independent set in H of size $|V(G)|$, that is an independent transversal of the sets $\{L(u) : u \in V(G)\}$.

Covers as defined here capture a more general notion, but most of our results are restricted to list colouring. Given a q-list assignment $\tilde{L}$ of G we create a q-fold cover $\mathscr{H} = (L, H)$ of G such that $\mathscr{H}$-colourings of G correspond to L-colourings of G by making the sets $L(u)$ formally disjoint copies of the lists $\tilde{L}(u)$, and for every edge $uv \in E(G)$ adding an edge between $x \in L(u)$ and $y \in L(v)$ whenever x and y are two copies of the same colour. In an attempt to avoid confusion we will refer to elements of the lists $\tilde{L}(u)$ as labels, and we will refer to vertices of H as colours. We shorten the phrase 'cover that arises from a list assignment' to 'list-cover'.

[2]The reason for including the adjective 'strong' relates to some technical subtleties we discuss later and in [18].

2.3 Statement of the framework

In an abbreviated form, the colouring aspect of our framework is the following theorem.

Theorem 2.4. *Let G be a graph of maximum degree $\Delta \geq 2^6$ such that the hard-core model on G at fugacity λ has strong local (β, γ)-occupancy for some $\lambda, \beta, \gamma > 0$. Suppose there exists $\ell > 7 \log \Delta$ such that for all vertices $u \in V(G)$ and subgraphs $F \subset G[N(u)]$ on at least $\ell/8$ vertices we have $Z_F(\lambda) \geq 8\Delta^4$. Then the list-chromatic number of G is at most q, where*

$$q = r\left(\beta + \gamma \frac{\Delta}{r}\right) \qquad \text{and} \qquad r = \frac{\lambda}{1+\lambda} \frac{\ell}{1 - \sqrt{7(\log \Delta)/\ell}}.$$

This essentially reduces the task of bounding the chromatic number from above to minimising $\beta + \gamma\Delta/r$ subject to local (β, γ)-occupancy, which is often a routine optimisation problem, see [18]. For efficient algorithms we need some more assumptions, as stated in the following result.

Theorem 2.5. *Suppose that the conditions of Theorem 2.4 hold and let $n = |V(G)|$. Suppose also that there is a class of graphs $\mathscr{G}$ and an integer t such that the following hold.*

1. *For each $u \in V(G)$, each induced subgraph $F \subset G[N(u)]$, and any list-cover $\mathscr{H}' = (L', H')$ of F with lists of size at most Δ, we have a procedure* Remove(F, H')*, running in time $T_r \geq \Delta$, for finding a set R of edges in F such that $|R| \leq t$ and the graph $\hat{H}$ obtained from H' by removing every edge between $L'(v)$ and $L'(w)$ for $vw \in R$ satisfies $\hat{H} \in \mathscr{G}$.*
2. *For each $\hat{H} \in \mathscr{G}$ we have a procedure* Sample$(\hat{H}, \lambda)$ *for sampling from the hard-core model on $\hat{H}$ at fugacity λ in time T_s.*
3. *The integer t satisfies $0 \leq t \leq \ell/40$.*

Then, for any q-list-assignment $\tilde{L}$ of G, there is a randomised algorithm, defined in terms of a procedure called AddressB*, that constructs, for any $c \in (0, 1)$, an $\tilde{L}$-colouring of G as guaranteed by Theorem 2.4 in time*

$$O\big((T_r + T_s)\Delta \log(\max(\ell/\lambda, \ell))n + n^{1+c}\big)$$

with probability at least $1 - 2/n^c$.

Each graph F that must be considered above is an induced subgraph of a neighbourhood, so can have up to Δ vertices. This means that there are trivial implementations of Remove and Sample which satisfy our assumptions. When Remove does nothing, Sample$(\hat{H}, \lambda)$ can be a naive exponential-time algorithm that finds every independent set I in $\hat{H}$, its weight $\lambda^{|I|}$, and samples accordingly. This gives $T_r + T_s = \exp(O(q\Delta)) = \exp(O(\Delta^2))$ because in applications we will only be interested in $q \leq \Delta$ as $\chi_\ell(G) \leq \Delta + 1$ by a direct greedy argument. This may be permissible for small constant Δ, but we are not content with exponential dependence on Δ so we work much harder to show that suitable implementations exist in the setting of Theorem 1.3 with $T_r + T_s$ polynomial in Δ.

We give a full proof of Theorem 2.4 and Theorem 2.5 in Section 4.

3 Application to graphs with few cycles

Recall that we work with the more general condition of F_k-freeness instead of C_k-freeness, where F_k is the graph formed from the path P_{k-1} on $k-1$ vertices by adding a new vertex adjacent to every vertex of the path. To prove Theorem 1.3 via the framework results of Theorem 2.4 and Theorem 2.5 we must establish suitable strong local occupancy in graphs with few copies of F_k, and we must give suitable implementations of REMOVE and SAMPLE.

3.1 Local occupancy

It is convenient to give the proof in terms of a *maximum average degree* parameter

$$\operatorname{mad}(G) = \max_{\substack{F \subset G \\ |V(F)| \geq 1}} \left\{ \frac{2|E(F)|}{|V(F)|} \right\},$$

in neighbourhoods. We also write W for the (upper real branch of the) *Lambert W-function* that is the inverse of $x \mapsto xe^x$ defined on $[-1, \infty)$. We use the basic property that as $x \to \infty$ we have $W(x) = (1 - o(1)) \log x$, see, e. g., [13].

Lemma 3.1. *Let $a \geq 0$ and G be a graph such that $\operatorname{mad}(G[N(u)]) \leq a$ for each $u \in V(G)$. Then the following statements hold for any $\lambda > 0$.*

1. *For any $d > 0$, there exist $\beta, \gamma > 0$ such that the hard-core model on G at fugacity λ has strong local (β, γ)-occupancy and*

$$\beta + \gamma d = \frac{1+\lambda}{\lambda} \frac{d(1+\lambda)^a \log(1+\lambda)}{W(d(1+\lambda)^a \log(1+\lambda))}.$$

2. *For any vertex $u \in V(G)$ and any subgraph F of $G[N(u)]$ on y vertices we have*

$$\log Z_F(\lambda) \geq y \log(1+\lambda) \left(1 - \frac{a}{2} \log(1+\lambda)\right).$$

The proof of Lemma 3.1 relies on the following elementary lemma which already appeared in [15], though we give the short proof for completeness.

Lemma 3.2. *For any graph F on y vertices with average degree at most $a > 0$,*

$$\frac{\lambda Z_F'(\lambda)}{Z_F(\lambda)} \geq \frac{\lambda}{1+\lambda} y (1+\lambda)^{-a}, \qquad \text{and} \qquad \log Z_F(\lambda) \geq \frac{y}{a}(1 - (1+\lambda)^{-a}).$$

Proof. Let $\mathbf{S}$ be a random independent set from the hard-core model at fugacity λ on F. First, for any $u \in V(F)$,

$$\Pr(u \in \mathbf{S}) = \frac{\lambda}{1+\lambda} \Pr(\mathbf{S} \cap N(u) = \emptyset) \geq \frac{\lambda}{1+\lambda} (1+\lambda)^{-\deg(u)},$$

because the spatial Markov property gives that $\mathbf{S} \cap N(u)$ is a random independent set drawn from the hard-core model on the subgraph $F[N(u)]$ induced by the externally uncovered neighbours of u. That is, when $\mathbf{U} = N(u) \setminus N(\mathbf{S} \setminus N(u))$ is the set obtained by revealing $\mathbf{S} \setminus N(u)$ and removing from $N(u)$ any vertex with a neighbour in $\mathbf{S} \setminus N(u)$, we write $\mathbf{F}_{N(u)} = F[\mathbf{U}]$, and then $\mathbf{S} \cap N(u)$ is distributed according to the hard-core model on $\mathbf{F}_{N(u)}$ at fugacity λ. The final inequality comes from the fact that any realisation of $\mathbf{F}_{N(u)}$ has $Z_{\mathbf{F}_{N(u)}}(\lambda) \le (1+\lambda)^{\deg(u)}$. The lemma now follows by convexity:

$$E[|\mathbf{S}|] = \sum_{u \in V(F)} \Pr(u \in \mathbf{S}) \ge \frac{\lambda}{1+\lambda} \sum_{u \in V(F)} (1+\lambda)^{-\deg(u)} \ge \frac{\lambda}{1+\lambda} y (1+\lambda)^{-a},$$

and since

$$E[|\mathbf{S}|] = \frac{\lambda Z_F'(\lambda)}{Z_F(\lambda)} = \lambda \frac{\partial}{\partial \lambda} \log Z_F(\lambda),$$

integrating this bound gives the required lower bound on $\log Z_F(\lambda)$. □

Proof of Lemma 3.1. Let u be an arbitrary vertex of G, and suppose that $F \subset G[N(u)]$ has y vertices. By assumption, we know that F has average degree at most a. If $a = 0$ then F contains no edges and $Z_F(\lambda) = (1+\lambda)^y$ so Item 2 holds, and otherwise by Lemma 3.2 we have

$$\log Z_F(\lambda) \ge \frac{y}{a}(1 - (1+\lambda)^{-a}) \ge y \log(1+\lambda) \left(1 - \frac{a}{2}\log(1+\lambda)\right),$$

where the second inequality holds by the fact that for $z \ge 0$ we have $1 - \mathrm{e}^{-z} \ge z(1 - z/2)$ and the substitution $z = a\log(1+\lambda)$. This establishes Item 2 for any $a \ge 0$.

For Item 1 we note that $Z_F(\lambda) \le (1+\lambda)^y$ and hence by Lemma 3.2 we have

$$\beta \frac{\lambda}{1+\lambda} \frac{1}{Z_F(\lambda)} + \gamma \frac{\lambda Z_F'(\lambda)}{Z_F(\lambda)} \ge \frac{\lambda}{1+\lambda} \left(\beta (1+\lambda)^{-y} + \gamma y (1+\lambda)^{-a}\right), \tag{3.1}$$

and we define the right-hand side to be $g(y)$.

The function g is strictly convex with a stationary minimum at

$$y^* = a + \frac{\log\left(\frac{\beta}{\gamma}\log(1+\lambda)\right)}{\log(1+\lambda)},$$

and if we set $g(y^*) = 1$ for strong local (β, γ)-occupancy, and solve for β we obtain

$$\beta = \frac{\gamma(1+\lambda)^{\frac{(1+\lambda)^{1+a}}{\gamma\lambda} - a}}{\mathrm{e}\log(1+\lambda)}. \tag{3.2}$$

To see this, write $\Xi = \log(1+\lambda)$ and observe that

$$g(y^*) = \frac{\gamma\lambda}{1+\lambda}\frac{\mathrm{e}^{-a\Xi}}{\Xi}\left(1 + a\Xi + \log(\Xi\beta/\gamma)\right) = \frac{\gamma\lambda}{1+\lambda}\frac{\mathrm{e}^{-a\Xi}}{\Xi}\left(a\Xi + \log(\mathrm{e}\Xi\beta/\gamma)\right).$$

Then $g(y^*) = 1$ implies

$$\frac{e\beta\Xi}{\gamma} = \exp\left(\frac{1+\lambda}{\gamma\lambda}\frac{\Xi}{e^{-a\Xi}} - a\Xi\right),$$

which gives β as in equation 3.2. The next step is to choose γ to minimize $\beta + d\gamma$, which can be done with intricate but elementary calculus. We compute derivatives of $h(\gamma) = \beta + \gamma d$ as follows.

$$h'(\gamma) = d + \frac{1}{e}(1+\lambda)^{\frac{(1+\lambda)^{1+a}}{\gamma\lambda}}\left(\frac{(1+\lambda)^{-a}}{\log(1+\lambda)} - \frac{1+\lambda}{\gamma\lambda}\right),$$
$$h''(\gamma) = \frac{1}{e}\frac{(1+\lambda)^{2+a}}{\gamma^3\lambda^2}(1+\lambda)^{\frac{(1+\lambda)^{1+a}}{\gamma\lambda}}\log(1+\lambda).$$

Thus, h is strictly convex on the domain $\gamma > 0$ and there is a unique solution to $h'(\gamma) = 0$ with $\gamma > 0$. We claim the unique minimum of h on $\gamma > 0$ occurs at

$$\gamma = \frac{1+\lambda}{\lambda}\frac{(1+\lambda)^a\log(1+\lambda)}{1+W(d(1+\lambda)^a\log(1+\lambda))}.$$

To prove the claim, we set γ to the announced value and write D for the expression $d(1+\lambda)^a\log(1+\lambda)$. This gives

$$\frac{(1+\lambda)^{1+a}}{\gamma\lambda} = \frac{1+W(D)}{\log(1+\lambda)} \tag{3.3}$$

and hence

$$h'(\gamma) = d + \frac{1}{e}e^{1+W(D)}\frac{(1+\lambda)^{-a}}{\log(1+\lambda)}\Big(1-(1+W(D))\Big)$$
$$= d - W(D)e^{W(D)}\frac{(1+\lambda)^{-a}}{\log(1+\lambda)} = 0,$$

because the Lambert W-function satisfies $W(D)e^{W(D)} = D$.

The statement Lemma 3.1 Item 1 follows because the chosen values of β and γ give, via another use of equation 3.3 and the fact that $e^{W(D)} = D/W(D)$,

$$\beta = \frac{\gamma(1+\lambda)^{\frac{(1+\lambda)^{1+a}}{\gamma\lambda}-a}}{e\log(1+\lambda)} = \frac{1+\lambda}{\lambda}\cdot\frac{e^{W(D)}}{1+W(D)} = \frac{1+\lambda}{\lambda}\cdot\frac{D}{W(D)\cdot(1+W(D))}.$$

Then we have

$$\beta + \gamma d = \frac{1+\lambda}{\lambda}\left(\frac{D}{W(D)(1+W(D))} + \frac{D}{1+W(D)}\right) = \frac{1+\lambda}{\lambda}\frac{D}{W(D)},$$

as required for Item 1. □

Though intricate, this proof is an indication of a key benefit of our framework: technical verification of local occupancy properties is entirely separate from the analysis of the coloring algorithm. To apply the above result to G as in Theorem 1.3 we prove a suitable bound on the maximum average degree. With this we have already established conditions sufficient for an application of Theorem 2.4 to give a non-constructive proof of Theorem 1.2. In the following section, we show how to further verify the conditions of Theorem 2.5 for the full algorithmic form.

Lemma 3.3. *For $k \geq 3$ suppose that a vertex $u \in V(G)$ is contained in at most t copies of the fan F_k. Then the average degree of any graph $F \subset G[N(u)]$ is at most $k - 3 + \sqrt{2t}$.*

Proof. Note that since $F \subset G[N(u)]$ at u is contained in at most t copies of the fan F_k it must be the case that F contains at most t copies of P_{k-1}. This is because F_k is formed by adding an additional vertex adjacent to each vertex in the path P_{k-1}.

Let F have y vertices. We assert that the average degree of F is at most

$$\min\left\{y-1,\ k-3+\frac{2t}{y}\right\} \leq k-3+\sqrt{2t}\,.$$

The first bound is straightforward as there are at most $y-1$ possible neighbours for any vertex in F, and the second follows from a theorem of Erdős and Gallai [22, Theorem 2.6] that bounds the average degree of P_{k-1}-free graphs. By removing at most t edges from F we can remove all copies of P_{k-1}, and hence the resulting graph has at most $y(k-3)/2$ edges. This means that F has at most $y(k-3)/2 + t$ edges. The first expression in the assertion follows, and we consider the subcases $y \leq \sqrt{2t}$ and $y > \sqrt{2t}$ to crudely bound from above the minimum. □

3.2 Removing edges and sampling independent sets

For the rest of this section let G be as in Theorem 1.3. Let $\mathcal{F}$ be the class of P_{k-1}-free graphs on at most Δ vertices and let $\mathcal{G}$ be the class of list-covers of graphs in $\mathcal{F}$ with at most Δ colours in each list. For any vertex $u \in V(G)$, subgraph F of $G[N(u)]$, and list-cover $\mathscr{H} = (L', H')$ of F with at most Δ colours in each list we can identify all copies of P_{k-1} in F in time $O(\Delta^k)$ by enumerating all ordered sets of $k-1$ vertices in F. To implement Remove we simply choose an arbitrary edge vw in each P_{k-1} found in this way and remove all edges from H' between $L(v)$ and $L(w)$. For each of the at most $t \leq \Delta^{2\epsilon}$ copies of P_{k-1} found, this removal takes time at most Δ^2 so Remove as in Theorem 2.5 can be done in time $\Delta^{O(k)}$.

Our next result implies that we can implement Sample on $\mathcal{G}$ in time $(k\Delta)^{O(k^3)}$, and although it is elementary it is crucial for our main message on meeting the algorithmic barrier.

Theorem 3.4. *Let $k \geq 3$ and F be a P_{k-1}-free graph on y vertices, and let $\hat{\mathscr{H}} = (\hat{L}, \hat{H})$ be a list-cover of F with at most q colours in each list. Then there is an absolute constant c such that for any $\lambda > 0$ we can sample from the hard-core model on $\hat{H}$ in time $y^{3k^2}(1+q)^{k^3/2}(ck)^{k^3}$.*

We give a short proof in which, rather than focusing on obtaining the best quantitative bounds we can, we merely emphasise that for fixed k the running time is polynomial in q and y.

We rely on the well-known (and straightforward to realise) fact that in a connected graph any two longest paths must have a common vertex, which we state as follows.

Lemma 3.5. *Let $k \geq 3$ and $F = (V, E)$ be a connected P_{k-1}-free graph. If P is the vertex set of a longest path in F, then $F[V \setminus P]$ is P_{k-2}-free.*

We now prove a warm-up to Theorem 3.4 that shows we can calculate the partition function and certain useful probabilities efficiently.

Lemma 3.6. *Let $k \geq 3$ and F be a P_{k-1}-free graph on y vertices, and let $\mathscr{H} = (L, H)$ be a cover of F with at most q colours in each list. Then there is an absolute constant c such that for any positive λ we can evaluate $Z_H(\lambda)$ in time $y^{3k}(1+q)^{k^2/2}(ck)^{k^2}$.*

Further, let $F_1, \dots, F_r$ be the connected components of F. Write $H_i = H[L(V(F_i))]$ for the covers of the components of F given by H, and let $\mathbf{I}$ be a random independent set from the hard-core model at fugacity λ on H. Then in time $y^{3k}(1+q)^{k^2/2}(ck)^{k^2}$ we can compute for every i in $[r]$ a longest path P_i in F_i, and the probabilities $\Pr(\mathbf{I} \cap L(P_i) = J_i)$ for every independent set $J_i \subset L(P_i)$.

Proof. In the proof we write $c_1, c_2, \dots$ for some unspecified absolute constants, and let $f(y, q, k)$ be the upper bound on running time that we wish to calculate.

If $k = 3$ then F is edgeless and hence

$$Z_H(\lambda) = \prod_{v \in V(F)} (1 + |L(v)|\lambda)\,,$$

which is computable in time $c_1 y$. This gives the base case $f(y, q, 3) = c_1 y$.

For $k > 3$, we note that F has at most y connected components and deal with each one separately. The function $Z_H(\lambda)$ is multiplicative over the induced covers of the components of F, so

$$Z_H(\lambda) = \prod_{i=1}^{r} Z_{H_i}(\lambda)\,.$$

Let $F' = (V', E')$ be a component of F, and write $H' = H[L(V')]$ for the induced cover of F'. Let $P = \{v_1, \dots, v_s\}$ (where $s \leq k-2$) be the vertex set of a longest path in F', which we can find in time $c_2(k-3)!2^{k-2}y$ by executing for each j from $k-3$ to 1 a fixed-parameter algorithm that finds a path of length j (if one exists) in time $O(j!2^j y)$, see [9].

Since F' is connected, Lemma 3.5 implies that $F'[V' \setminus P]$ is P_{k-2}-free. We use the fact that the sum over independent sets I in H' that gives $Z_{H'}(\lambda)$ can be split into terms according to $I \cap L(P)$. Let $H^P = H'[L(P)]$ and $H^{V' \setminus P} = H'[L(V' \setminus P)]$. Then

$$Z_{H'}(\lambda) = \sum_{J \in \mathcal{I}(H^P)} \lambda^{|J|} \cdot Z_{H^{V' \setminus P}_J}(\lambda)\,.$$

It is possible to iterate over the required $J \in \mathcal{I}(H^P)$ in time $c_3(k-2)^2(1+q)^{k-2}$ by iterating over all $(1+q)^s \leq (1+q)^{k-2}$ sets J with $|J \cap L(v_i)| \leq 1$ and checking each for independence. The

independence check takes time $O\big((k-2)^2\big)$ as we must verify that each pair of vertices in J is absent from $E(H)$. Then constructing $H_J^{V\setminus P}$ can be done in time $c_4(k-2)y$ because $|J| \le k-2$ and each colour in J can conflict with at most y colours which are removed from $H^{V\setminus P}$ to form $H_J^{V\setminus P}$. Now $Z_{H_J^{V\setminus P}}(\lambda)$ can be computed in time $f(y,q,k-1)$ by induction, so we can compute $Z_{H'}(\lambda)$ in time

$$c_2(k-3)!2^{k-2}y \cdot c_3(k-2)^2(1+q)^{k-2} \cdot c_4(k-2)y \cdot f(y,q,k-1)\,.$$

Since there are at most y components of F, we have the recurrence

$$f(y,q,k) \le y \cdot c_2(k-3)!2^{k-2}y \cdot c_3(k-2)^2(1+q)^{k-2} \cdot c_4(k-2)y \cdot f(y,q,k-1)\,.$$

With the base case of $f(y,q,3) = c_1 y$ we have

$$f(y,q,k) \le c_1 y (c_2c_3c_4y^3)^{k-3}(2(1+q))^{k(k-3)/2}((k-2)!)^3 \cdot \prod_{j=1}^{k-3} j!\,,$$

which for a large enough constant c is at most $y^{3k}(1+q)^{k^2/2}(ck)^{k^2}$.

For the final statement, observe that in the above argument we compute for each $i \in [r]$ an evaluation of $Z_{H_{J_i}[V(F_i)\setminus P_i]}$ for some longest path P_i in F_i, and each independent set $J_i \subset L(P_i)$ along the way to computing $Z_{H_i}(\lambda)$. But

$$\Pr(\mathbf{I} \cap L(P_i) = J_i) = \frac{\lambda^{|J_i|} Z_{H_{J_i}[V(F_i)\setminus P_i]}(\lambda)}{Z_{H_i}(\lambda)}\,,$$

so with some straightforward extra bookkeeping we have the required probabilities. □

With this result we can give the sampling algorithm required for Theorem 3.4.

Proof of Theorem 3.4. We write $\mathbf{I}$ for a random independent set from the hard-core model on $\hat{H}$.

If $k = 3$ then F is edgeless, and it suffices to sample independently for each vertex $v \in V$. With probability $1/(1+|\hat{L}(v)|\lambda)$ take $\mathbf{I} \cap \hat{L}(v) = \emptyset$, otherwise let $\mathbf{I} \cap \hat{L}(v)$ be a uniform random element of $\hat{L}(v)$. This can be done in time $O(qy)$ provided sampling from a biased coin takes time $O(1)$ and sampling uniformly from a list of length q takes time $O(q)$.

If $k > 3$ then let $F_1, \dots, F_r$ be the components of F. By Lemma 3.6 we can compute in time $y^{3k}(1+q)^{k^2/2}(ck)^{k^2}$ a longest path P_i in F_i and the probabilities $\Pr(\mathbf{I} \cap \hat{L}(P_i) = J_i)$ for all independent sets $J_i \subset \hat{L}(P_i)$ and for all i. Hence, we can sample $\mathbf{I} \cap \hat{L}(P_i)$ for all i in this time. Then we can construct $H_i' = \hat{H}_{\mathbf{I}\cap\hat{L}(P_i)}[\hat{L}(V(F_i)\setminus P_i)]$ for each i in time $O(rky)$, and use the fact that $\mathbf{I} \cap \hat{L}(V(F_i)\setminus P_i)$ is distributed according to the hard-core model at fugacity λ on H_i'.

With this scheme it is straightforward to show by induction that the time taken is at most

$$y^{3k^2}(1+q)^{k^3/2}(ck)^{k^3}\,.$$ □

We can now finish the deduction of Theorem 1.3 from Theorem 2.4 and Theorem 2.5.

Proof of Theorem 1.3. Fix an arbitrary vertex $u \in V(G)$ and an arbitrary subgraph $F \subset G[N(u)]$. By Lemma 3.3 we have $\mathrm{mad}(F) \le a = k - 3 + \sqrt{2t}$. It is convenient to exclude the case $a = 0$ in the argument, which is one place the assumption $t \ge 1$ comes in useful, giving $a > 0$.

We want to apply Theorem 2.5 to conclude the proof. Above we defined a class $\mathscr{G}$, and the procedures REMOVE and SAMPLE such that Item 1 and Item 2 hold with $T_r = \Delta^{O(k)}$ and $T_s = (k\Delta)^{O(k^3)}$. We now define the parameters so that the remaining requirements of Theorem 2.5 are satisfied: those are $t \le \ell/40$ (Item 3) and the hypotheses of Theorem 2.4. Lemma 3.1 and Lemma 3.3 provide the requisite strong local occupancy and the condition on $Z_F(\lambda)$. Indeed, with $a = k - 3 + \sqrt{2t}$, given any $\lambda > 0$ and $\ell > 7\log\Delta$, and with r and q as in Theorem 2.4, Item 1 of Lemma 3.1 gives us β and γ such that the hard-core model on G at fugacity λ has strong local (β, γ)-occupancy with

$$q = r\left(\beta + \gamma\frac{\Delta}{r}\right) = \frac{1+\lambda}{\lambda}\frac{\Delta(1+\lambda)^a\log(1+\lambda)}{W(\Delta(1+\lambda)^a\log(1+\lambda)/r)}\,.$$

We set

$$\log(1+\lambda) = \frac{1}{a\log(\Delta/\sqrt{t})}\,, \qquad \text{and} \qquad \ell = \frac{40a}{\log(\Delta/\sqrt{t})}\left(\frac{\Delta}{\sqrt{t}}\right)^{\frac{\epsilon}{1+\epsilon}}.$$

First, recall that $t \le \Delta^{\frac{2\epsilon}{1+2\epsilon}}/(\log\Delta)^2$, and hence $40\cdot t/\ell = O\left((\log\Delta)^{-\varepsilon/(1+\varepsilon)}\right)$ tends to 0 as Δ goes to infinity, so t is indeed at most $\ell/40$ if Δ_0 is large enough.

Second, we must show that

$$\frac{\ell}{8} \ge \frac{\log(8\Delta^4)}{\log(1+\lambda)\left(1-\frac{a}{2}\log(1+\lambda)\right)} = \frac{a\log(\Delta/\sqrt{t})\log(8\Delta^4)}{1-\frac{1}{2\log(\Delta/\sqrt{t})}}\,, \tag{3.4}$$

so that Item 2 of Lemma 3.1 ensures that $Z_G(\lambda) \ge 8\Delta^4$ for any subgraph F of any $G[N(u)]$ on at least $\ell/8$ vertices. The right-hand side of equation 3.4 is $O\left(a\log(\Delta/\sqrt{t})\log\Delta\right)$ as Δ and hence $\Delta/\sqrt{t}$ tend to infinity. This is less than ℓ for large enough Δ_0 in terms of ϵ because the bound $t \le \Delta^{\frac{2\epsilon}{1+2\epsilon}}/(\log\Delta)^2$ gives

$$\log(\Delta/\sqrt{t}) \ge \frac{1+\epsilon}{1+2\epsilon}\log\Delta - \log\log\Delta = \Omega(\log\Delta)\,,$$

so that $a\log(\Delta/\sqrt{t})\log\Delta = O(a(\log(\Delta/\sqrt{t}))^2) = o(\ell)$.

As $\Delta \to \infty$ we now have $\lambda = o(1)$, $a\lambda = o(1)$, $7\log\Delta = o(\ell)$, $\ell = o(\Delta)$, $a = O(\sqrt{t})$, and

$$q \sim \frac{\Delta}{\log(\Delta/\ell)} \sim \frac{\Delta}{\log\left((\Delta/\sqrt{t})^{1/(1+\epsilon)}\log(\Delta/\sqrt{t})\right)}\,.$$

To obtain Theorem 1.3, note that for large enough Δ_0 and $\Delta \geq \Delta_0$,

$$(1+\epsilon)\frac{\Delta}{\log(\Delta/\sqrt{t})} \geq q\,,$$

so the result follows from the application of Theorem 2.5. □

4 Proof of the framework

4.1 An algorithm for graph colouring

We start by giving an overview of the two-phase method. We define (precisely later) a *flaw* for each vertex u to capture the problem of having too few available colours or too much competition for available colours. The first phase is that subject to the conditions of Theorem 2.5, a *flawless partial colouring of G* can be found efficiently.

Lemma 4.1. *Suppose that the conditions of Theorem 2.5 hold, and let $\mathscr{H} = (L, H)$ be a q-fold list-cover of G. Then there is an algorithm that constructs a flawless partial $\mathscr{H}$-colouring of G in time $O\big((T_r + T_s)\Delta \log(\max(\ell/\lambda, \ell)) \cdot n\big)$ with probability at least $1 - 2^{-n}$.*

The second phase is that a flawless partial colouring can be efficiently completed to a list colouring of G. This 'efficient finishing blow' is standard and was established in earlier work [36, 2].

Lemma 4.2. *Suppose that the conditions of Theorem 2.5 hold so that,* inter alia*, G is an n-vertex graph. Suppose also that $\mathscr{H} = (L, H)$ is a q-fold list-cover of G, and let σ be a flawless partial $\mathscr{H}$-colouring of G. Then there is an algorithm that finds, for each $c \in (0, 1)$, an $\mathscr{H}$-colouring of G in time $O\left(n^{1+c}\right)$ with probability at least $1 - 1/n^c$.*

Note that Lemma 4.1 followed by Lemma 4.2 directly implies Theorem 2.5.

The algorithm for Lemma 4.2 selects a colour for the remaining vertices uniformly at random, resampling if there are any conflicts. For completeness, we give a sketch proof in Section 4.3.

The algorithm for Lemma 4.1 explores the space of partial colourings of G, starting with a trivial colouring in which every vertex is *coloured blank*. We define an order on flaws, and while the current partial colouring σ is flawed we execute a resampling action to address the least flaw present in σ and move to a new partial colouring. We give the proof of Lemma 4.1 in the rest of this section with some details deferred to Section 4.2.

4.1.1 Notation

Given the setup of Theorem 2.4, we work with a q-fold list-cover $\mathscr{H} = (L, H)$ of G. For a set $S \subset V(G)$ we write $L(S) = \bigcup_{u \in S} L(u)$. We refer to the vertices of H as colours, and write H^* for the graph obtained from H by removing all edges inside the sets $L(u)$ for all $u \in V(G)$. Then for $u \in V(G)$ and $x \in L(u)$ we write $\deg^*_{\mathscr{H}}(x)$ for the degree in H^* of a colour x, which is the number of colours on lists of neighbours of u in G that conflict with x.

Writing $\mathfrak{B}$ for a special blank colour, and borrowing from [2], a *partial $\mathscr{H}$-colouring* σ of G is a function from $V(G)$ to $\{\mathfrak{B}\} \cup V(H) \cup E(G)$ such that the following hold.

1. For all $u \in V(G)$, either $\sigma(u) = \mathfrak{B}$, $\sigma(u) \in L(u)$, or $\sigma(u) = e \in E(G)$ with $u \in e$.
2. Restricting the image of σ to $V(H)$ gives an independent set: $\sigma(V(G)) \cap V(H) \in \mathcal{I}(H)$.

We write Ω for the set of such partial $\mathscr{H}$-colourings σ, and omit the prefix $\mathscr{H}$ when it is clear.

Given $\sigma \in \Omega$ we have *blank vertices* $\mathrm{bla}(\sigma) = \{u \in V(G) : \sigma(u) = \mathfrak{B}\}$ for which $\sigma(u) = \mathfrak{B}$, *coloured vertices* $\mathrm{col}(\sigma) = \{u \in V(G) : \sigma(u) \in L(u)\}$ for which $\sigma(u) \in L(u)$, and *uncoloured vertices* $\mathrm{unc}(\sigma) = \{u \in V(G) : \sigma(u) \in E(G)\}$ for which $\sigma(u)$ is an edge of G containing u. We also write $\mathrm{ind}(\sigma) = \sigma(V(G)) \cap V(H) = \sigma(\mathrm{col}(\sigma))$ for the independent set in H signified by σ.

We also require some notation for the cover that remains on the blank vertices. We write $G_\sigma = G[\mathrm{bla}(\sigma)]$ for the subgraph of G induced by $\mathrm{bla}(\sigma)$, and then write $\mathscr{H}_\sigma = (L_\sigma, H_\sigma)$ for the cover of G_σ obtained by setting $L_\sigma(u) = L(u) \setminus N_H[\mathrm{ind}(\sigma)]$ for $u \in \mathrm{bla}(\sigma)$ and $H_\sigma = H[L_\sigma(\mathrm{bla}(\sigma))]$. Note that $\mathscr{H}_\sigma$ is a list-cover of G_σ. When there are no uncoloured vertices, these definitions agree with those of [8, 18] if σ is identified with $\mathrm{ind}(\sigma)$.

To interpret this notation, note that the list $L_\sigma(u)$ contains the colours in $L(u)$ that do not conflict with colours of the vertices in $\mathrm{col}(\sigma)$. This means that if $I \in \mathcal{I}(H_\sigma)$, then $\mathrm{ind}(\sigma) \cup I \in \mathcal{I}(H)$. In particular, if σ has no uncoloured vertices, and if we can find $I \in \mathcal{I}(H_\sigma)$ of size $|V(G_\sigma)|$ then $\mathrm{ind}(\sigma) \cup I$ is an $\mathscr{H}$-colouring of G. This is exactly how the two-phase method proceeds.

4.1.2 The flaws

We define a flaw for each vertex u, writing

$$B_u = \{\sigma \in \Omega : u \notin \mathrm{col}(\sigma) \text{ and either } |L_\sigma(u)| < \ell, \text{ or } \exists x \in L_\sigma(u) \text{ with } \deg^*_{\mathscr{H}_\sigma}(x) > \ell/8\}\,.$$

We also define a flaw for each pair (u, e) where $u \in e \in E(G)$, which represents the fact that u is an uncoloured vertex with $\sigma(u) = e$, writing $U^e_u = \{\sigma \in \Omega : \sigma(u) = e\}$. We write $F_B = \{B_u : u \in V(G)\}$ and $F_U = \{U^e_u : u \in e \in E(G)\}$ so that $F = F_B \cup F_U$ is the set of all flaws. Note that $|F_B| = |V(G)| = n$, and $|F_U| = 2|E(G)| \le \Delta n$.

It is important that we address the flaws in a sensible order, and any fixed order that puts every flaw of the form B_u before any flaw of the form U^e_v suffices. To be explicit, consider an arbitrary ordering of the vertices, and the induced lexicographic ordering on edges where uv is ordered according to the pair (u, v) with $u < v$. We address the flaws consistent with the order that puts flaws of the form B_u first, ordered according to u, and then puts the U^e_u ordered according to u and then e.

4.1.3 The actions

To address the flaw B_u at state σ we execute the action AddressB(u, σ) defined in Algorithm 1 in terms of the procedures Remove and Sample guaranteed by the assumptions of Theorem 2.5.

Algorithm 1 has three distinct parts. The first is some setup in which we define a partial colouring σ' by reassigning coloured vertices in $N_G(u)$ to $\mathfrak{B}$, which gives us an induced

Algorithm 1

procedure AddressB(u, σ)
  let σ' be obtained from σ by setting $\sigma'(v) = \mathfrak{B}$ for all $v \in N_G(u) \setminus \mathrm{unc}(\sigma)$, and $\sigma'(v) = \sigma(v)$ otherwise
  let $F = G_{\sigma'}[N_G(u)]$, $H' = H_{\sigma'}[L(N_G(u))]$, and let $\hat{H} = \textsc{Remove}(F, H')$
  let $\mathbf{J}_0 = \textsc{Sample}(\hat{H}, \lambda)$, and **let** $\tau_0(v) = \sigma'(v)$ unless $v \in N_G(u)$ and $\mathbf{J}_0 \cap L(v) = \{y\}$, in which case $\tau_0(v) = y$
  let $i = 0$
  while $\mathbf{J}_i \notin \mathcal{I}(H')$ **do**
    let vw be the lowest-indexed edge in $G[N(u)]$ for which $\mathbf{J}_i$ spans an edge of H' going from $L(v)$ to $L(w)$, and suppose that $v < w$
    let τ_{i+1} be obtained from τ_i by setting $\tau_{i+1}(v) = vw$ (uncolouring v) and setting τ_{i+1} to agree with τ_i elsewhere
    let $\mathbf{J}_{i+1} = \mathbf{J}_i \setminus L(v)$
    increment i
  let $\tau = \tau_i$ and **let** $\mathbf{J} = \mathbf{J}_i$
  return τ

subgraph F of $G[N_G(u)]$ and a list-cover $H' = H_{\sigma'}[L(N_G(u))]$ of F. We then use the procedure Remove to remove some edges from F and any corresponding edges in H', which results in a cover $\hat{H}$ of F in the class $\mathscr{G}$. The second is sampling an independent set $\mathbf{J}_0$ in $\hat{H}$ and a partial colouring-like object τ_0 corresponding to $\mathbf{J}_0$. Here we say 'partial colouring-like' because although $\mathbf{J}_0$ is independent in $\hat{H}$, it is not necessarily independent in $H_{\sigma'}$ so τ_0 is not necessarily a valid partial $\mathscr{H}$-colouring of G. In the third part we iterate over a loop variable i starting at 0 and uncolour vertices in $\mathbf{J}_i$ that participate in edges of H, making a sequence of corresponding τ_i as we go. When the loop exits $\mathbf{J}_i$ is independent in $H_{\sigma'}$, and so the final τ_i is a valid partial colouring.

To address the flaw U_u^e at state σ we simply resample $\sigma(u)$ from the hard-core model as follows. Let σ' be obtained from σ by letting $\sigma'(v) = \sigma(v)$ for $v \neq u$, and letting $\sigma'(u) = \mathfrak{B}$ with probability $1/(1 + |L_\sigma(u)|\lambda)$, and otherwise setting $\sigma'(u)$ to be a uniform colour from $L_\sigma(u)$.

4.1.4 Proving termination

In the analysis of the algorithm we discuss the transition probabilities induced by these actions, writing $\rho_f(\sigma, \tau)$ for the probability that the final state is τ when addressing the flaw f at state σ. Let $\tilde{\mu}$ be the probability measure on Ω given by

$$\tilde{\mu}(\sigma) = \frac{\lambda^{|\mathrm{col}(\sigma)|+|\mathrm{unc}(\sigma)|}}{\sum_{\tau\in\Omega}\lambda^{|\mathrm{col}(\tau)|+|\mathrm{unc}(\tau)|}},$$

and write $\tilde{Z}_H(\lambda)$ for the denominator. We note that $\tilde{\mu}$ is inspired by the hard-core model on H; its definition is motivated by the fact that creating an uncoloured vertex costs weight λ when sampling $\mathbf{J}_0$ in the procedure AddressB(u, σ).

Given subsets of flaws $S, S' \subset F$, we say that S' *covers* S if[3]

$$S' \cap F_U = S \cap F_U \text{ and } S' \cap F_B \supset S \cap F_B .$$

For any flaw f and subset $S \subset F$ of flaws, we define[4]

$$\begin{aligned}\mathrm{In}_f^S(\tau) = \{\sigma \in f : \rho_f(\sigma,\tau) > 0 \text{ and the set of flaws} \\ \text{introduced by the transition } \sigma \to \tau \text{ covers } S\} ,\end{aligned}$$

and we define the *charge* $c^S(f)$ to be

$$c^S(f) = \max_{\tau\in\Omega} \left\{ \sum_{\sigma\in \mathrm{In}_f^S(\tau)} \frac{\tilde{\mu}(\sigma)}{\tilde{\mu}(\tau)} \rho_f(\sigma,\tau) \right\}, \tag{4.1}$$

which represents a kind of compatibility between the measure $\tilde{\mu}$ and the transitions ρ_f induced by the actions for flaws $f \in F$.

We can now state the main theorem of Achlioptas et al. [2, Theorem 2.4], specialised to our setting, to show that our stochastic local search algorithm terminates quickly with high probability. To compare with the original, more general statement, we point out that we use their Remark 2.4 and that we start our algorithm in the all-$\mathfrak{B}$ partial colouring which has measure $1/\tilde{Z}_H(\lambda)$ and such that the only flaws present in the initial state are of the form B_u (since there are no uncoloured vertices).

Theorem 4.3 (Achlioptas, Iliopoulos, and Sinclair [2])**.** *If there exist positive numbers* $(\psi_f)_{f\in F}$ *such that for every* $f \in F$ *we have*

$$\zeta_f = \frac{1}{\psi_f} \sum_{S\subset F} c^S(f) \prod_{g\in S} \psi_g < 1 ,$$

then for $s \geq 0$ *Algorithm 1 reaches a flawless state in* $(T_0 + s)/\delta$ *steps with probability at least* $1 - 2^{-s}$*, where* $\delta = 1 - \max_{f\in F}\{\zeta_f\}$ *and*

$$T_0 = \log_2 \tilde{Z}_H(\lambda) + \sum_{u\in V(G)} \log_2(1 + \psi_{B_u}) + \log_2 \left(\max_{S\subset F} \frac{1}{\prod_{f\in S} \psi_f} \right) .$$

[3]For readers familiar with the definitions of [2], this is because flaws in F_U are *primary* while those in F_B are not. We avoid making precise what primary means here, see [2] for details.

[4]Note that this definition differs from that given in [2] as we have additionally restricted to mappings σ such that $\rho_f(\sigma,\tau) > 0$. The other states contribute zero to the charges we need to analyse, so this change simply means that we avoid explicitly having to exclude such states in our analysis of charges.

4.1.5 Bounding charges

By design the compatibility between $\tilde{\mu}$ and our algorithm is good enough for the following result to control the charges $c^S(B_u)$.

Lemma 4.4. *Suppose that the conditions of Theorem 2.5 hold, and let $\mathscr{H} = (L, H)$ be a q-fold list-cover of G. Then for any partial $\mathscr{H}$-colouring σ of G and any $u \in V(G)$ such that $u \notin \mathrm{col}(\sigma)$, the following holds. If τ is the random partial colouring of G that results from the procedure* AddressB(B_u, σ)*, then* $\Pr(\tau \in B_u) \le 1/(4\Delta^3)$.

A similar version of this is key to the lopsided local lemma formulation of our methods [18, Lemma 16]. To obtain an algorithm we essentially take advantage of some subtle extra strength arising from the fact that the lemma holds for any σ rather than only when σ has no uncoloured vertices and $\mathrm{ind}(\sigma)$ is sampled from the hard-core model on H. The proof of Lemma 4.4 is given in Section 4.2.

The following Lemma 4.5 and Lemma 4.6 comprise the bounds on charges to which we need to apply Theorem 4.3. Let $S(U_u^e) = \{B_v : v \in N^2[u]\}$ and $S(B_u) = \{B_v : v \in N^3[u]\} \cup \{U_v^e : v \in e \in E(G[N(u)])\}$. We will see that these are the only flaws that addressing U_u^e and B_u, respectively, can introduce.

Lemma 4.5. *For every vertex $u \in V(G)$, the following hold.*

1. *If $S \not\subset S(B_u)$ then $c^S(B_u) = 0$.*
2. *If S contains more than t flaws of the form U_v^{vw} with $v, w \in N(u)$ then $c^S(B_u) = 0$.*
3. $\max_{S \subset F}\{c^S(B_u)\} \le 1/(4\Delta^3)$.

Proof. Addressing B_u by executing AddressB(u, σ) only modifies σ on $N(u)$ and hence can only introduce flaws B_v for $v \in N^3[u]$, or U_v^{vw} for $v, w \in N(u)$ and $vw \in E(G[N(u)])$.

In the procedure AddressB we can only introduce at most t flaws of the form U_v^{vw} since the procedure Remove removes at most t edges from F that can be monochromatic, and each monochromatic edge leads to at most one uncolouring.

If addressing B_u results in τ with positive probability then the previous state σ must agree with τ outside $N(u)$, and inside $N(u)$ we must have $\mathrm{unc}(\tau) \supset \mathrm{unc}(\sigma)$ as no uncoloured vertices are coloured by the procedure AddressB. In fact, by the definition of covering for sets of flaws, when bounding $c^S(B_u)$ we can restrict our attention to triples (S, σ, τ) such that τ and σ agree outside $N(u)$ and

$$\mathrm{unc}(\tau) = \mathrm{unc}(\sigma) \cup \{v \in N(u) : U_v^e \in S \text{ for some } e \in E(G)\}\,.$$

That is, we can restrict attention to triples (S, σ, τ) where S carries the information necessary to deduce $\mathrm{unc}(\sigma)$ from $\mathrm{unc}(\tau)$. To this end, write

$$\nu(S) = \{v \in N(u) : U_v^e \in S \text{ for some } e \in E(G)\}\,,$$

so that $\nu(S)$ is the set of uncoloured vertices present in τ but not in σ for any $\sigma \in \mathrm{In}^S_{B_u}(\tau)$. We now assert that for such triples (S, σ, τ) we have

$$\rho_{B_u}(\sigma, \tau) = \frac{\lambda^{|\operatorname{col}(\tau)\cap N(u)|+|\nu(S)|}}{Z_{\hat{H}}(\lambda)},$$

where $\hat{H}$ is the graph $\hat{H}$ appearing in AddressB(u, σ); repeated here for convenience, $\hat{H}$ is constructed as follows. Let σ' be obtained from σ by setting $\sigma'(v) = \mathfrak{B}$ for all $v \in N(u) \setminus \operatorname{unc}(\sigma)$, and $\sigma'(v) = \sigma(v)$ otherwise, and then let $F = G_{\sigma'}[N_G(u)]$ and $H' = H_{\sigma'}[L(N_G(u))]$. Then $\hat{H} = \textsc{Remove}(F, H')$. Note that $\hat{H}$ does not depend on σ in the sense that when (S, σ, τ) are as above we can construct $\hat{H}$ from S and τ alone. Here we crucially exploit the definition of covering for sets of flaws. The key point is that σ' as above can be constructed from S and τ because the uncoloured vertices of σ are given by $\operatorname{unc}(\sigma) = \operatorname{unc}(\tau) \setminus \nu(S)$.

The assertion now follows from the definition of AddressB(u, σ) because from (S, τ) we can recover the size of the independent set $\mathbf{J}$ sampled in the procedure when $\tau =$ AddressB(u, σ). More accurately, we can determine $\operatorname{col}(\mathbf{J})$, which must be the disjoint union of $\operatorname{col}(\tau) \cap N(u)$ and the set $\nu(S)$ of vertices uncoloured during the execution of AddressB(u, σ). Given $\operatorname{col}(\mathbf{J})$ the uncolouring steps are deterministic, so we have the required expression for $\rho_{B_u}(\sigma, \tau)$ by the definition of the hard-core model. We also see that

$$\frac{\tilde{\mu}(\sigma)}{\tilde{\mu}(\tau)} = \frac{\lambda^{|\operatorname{col}(\sigma)\cap N(u)|+|\operatorname{unc}(\sigma)\cap N(u)|}}{\lambda^{|\operatorname{col}(\tau)\cap N(u)|+|\operatorname{unc}(\tau)\cap N(u)|}} = \frac{\lambda^{|\operatorname{col}(\sigma)\cap N(u)|}}{\lambda^{|\operatorname{col}(\tau)\cap N(u)|+|\nu(S)|}},$$

which holds by the definition of ν and because σ and τ only differ in $N(u)$. Then we have shown that given S and τ, for any $\sigma \in \mathrm{In}^S_{B_u}(\tau)$ we have

$$\frac{\tilde{\mu}(\sigma)}{\tilde{\mu}(\tau)}\rho_{B_u}(\sigma, \tau) = \frac{\lambda^{|\operatorname{col}(\sigma)\cap N(u)|}}{Z_{\hat{H}}(\lambda)}, \tag{4.2}$$

where $\hat{H}$ can be obtained from S and τ alone.

Turning to the charge $c^S(B_u)$, we have

$$c^S(B_u) = \max_{\tau\in\Omega}\left\{\sum_{\sigma\in\mathrm{In}^S_{B_u}(\tau)}\frac{\tilde{\mu}(\sigma)}{\tilde{\mu}(\tau)}\rho_{B_u}(\sigma, \tau)\right\} = \max_{\substack{\tau\in\Omega:\\ u\notin\operatorname{col}(\tau)}}\left\{\sum_{\sigma\in\mathrm{In}^S_{B_u}(\tau)}\frac{\lambda^{|\operatorname{col}(\sigma)\cap N(u)|}}{Z_{\hat{H}}(\lambda)}\right\}, \tag{4.3}$$

where $\hat{H}$ can be computed from S and τ as above. The restriction to τ such that $u \notin \operatorname{col}(\tau)$ is valid because for any $\tau \in \Omega$ with $u \in \operatorname{col}(\tau)$ we have $\mathrm{In}^S_{B_u}(\tau) = \emptyset$ since $\sigma \in B_u$ means that $u \notin \operatorname{col}(\sigma)$, but then $\rho_{B_u}(\sigma, \tau) = 0$ because the procedure AddressB(u, σ) does not alter $\sigma(u)$.

Nearing conclusion, we now argue that Lemma 4.4 implies $c^S(B_u) \le 1/(4\Delta^3)$. This holds because given S and τ we can construct the graph $\hat{H}$ occurring in equation 4.3, and note that it is a bona fide cover of some induced subgraph $F \subset G[N(u)]$. That is, there is a cover $\hat{\mathscr{H}} = (L, \hat{H})$ of G with the following two properties.

1. One obtains $\hat{H}$ from H by removing edges in H', and hence $\mathscr{H}$ satisfies the hypotheses of Theorem 2.5.

2. There is an independent set $\hat{J} \in \mathcal{I}(\hat{H})$ such that $\hat{H} = H_{\hat{J}}[L(N_G(u))]$.

Then the sum in equation 4.3 over the states $\sigma \in \mathrm{In}^S_{B_u}(\tau) \subset B_u$ can be interpreted as the probability that when $\mathrm{col}(\sigma) \cap N(v)$ is a random independent set from the hard-core model on $\hat{H}$ at fugacity λ, we have $\sigma \in \mathrm{In}^S_{B_u}(\tau)$. Since $\mathrm{In}^S_{B_u}(\tau) \subset B_u$, we can bound this from above by the probability that σ belongs to B_u given this random experiment; and Lemma 4.4 shows this probability to be at most $1/(4\Delta^3)$. □

Lemma 4.6. *For every u and e such that $u \in e \in E(G)$, the following hold.*

1. *If $S \not\subset S(U^e_u)$ then $c^S(U^e_u) = 0$.*

2. $\max_{S \subset F}\{c^S(U^e_u)\} \le \lambda/(1+\ell\lambda)$.

Proof. Addressing U^e_u at state σ can only introduce flaws of the form B_v for v in $N^2[u]$ because it causes no uncolouring, and only affects $\sigma(u)$.

If addressing U^e_u results in τ with positive probability then the previous state σ must be obtained from τ by setting $\tau(u) = e$. That is, for $S \subset F$, the set $\mathrm{In}^S_{U^e_u}(\tau)$ is either empty, or contains exactly one state σ. In the former case the charge is zero; and in the latter case, if $u \in \mathrm{bla}(\tau)$ then $\tilde{\mu}(\sigma)/\tilde{\mu}(\tau) = \lambda$ and $\rho_{U^e_u}(\sigma,\tau) = 1/(1+|L_\sigma(u)|\lambda)$, and if $u \in \mathrm{col}(\tau)$ then $\tilde{\mu}(\sigma)/\tilde{\mu}(\tau) = 1$ and $\rho_{U^e_u}(\sigma,\tau) = \lambda/(1+|L_\sigma(u)|\lambda)$, hence both possibilities yield that

$$\frac{\tilde{\mu}(\sigma)}{\tilde{\mu}(\tau)}\rho_{U^e_u}(\sigma,\tau) = \frac{\lambda}{1+|L_\sigma(u)|\lambda}\,.$$

Then $c^S(U^e_u) \le \lambda/(1+\ell\lambda)$ because the ordering on flaws ensures that $\sigma \notin B_u$ whenever we are addressing U^e_u, and hence u has at least ℓ available colours. □

4.1.6 Finishing the proof

We can now choose parameters ψ_f for $f \in F$ such that the desired result follows from Theorem 4.3. For this we take a positive real ψ to be determined later and set $\psi_f = \psi/(4\Delta^3)$ for all $f \in F_B$ and $\psi_f = \psi\lambda/(1+\ell\lambda)$ for all $f \in F_U$.

By Lemma 4.5 we know that $c^S(B_u) = 0$ unless all B_v flaws in S correspond to vertices v in $N^3[u]$, and there is a set $R \subset E(G[N(u)])$ of most t edges such that if S contains a flaw of the form U^{vw}_v, then $vw \in R$. Since $|N^3[u]| \le \Delta^3$, we deduce from Lemma 4.5 that for each $u \in V(G)$, setting

$$\psi = \frac{4(1+\ell\lambda)}{1+\ell\lambda+4t\lambda}\,,$$

we have

$$\frac{1}{\psi_{B_u}}\sum_{S\subset F} c^S(B_u)\prod_{g\in S}\psi_g \le \frac{1}{\psi}\prod_{v\in N^3[u]}\left(1+\frac{\psi}{4\Delta^3}\right)\prod_{vw\in R}\left(1+\frac{\psi\lambda}{1+\ell\lambda}\right)$$
$$\le \frac{1}{\psi}\exp\left(\frac{\psi}{4}\right)\exp\left(\frac{t\psi\lambda}{1+\ell\lambda}\right) = \frac{\mathrm{e}}{4}\left(1+\frac{4t\lambda}{1+\ell\lambda}\right) \le \frac{3}{4},$$

because by assumption $t \le \ell/40$, and hence $4t\lambda \le (1+\ell\lambda)(3/\mathrm{e}-1)$.

Similarly, by Lemma 4.6 and the facts that $\Delta \ge 2$ and $|N^2[u]| \le 1+\Delta^2$, we have for each pair (u,e) with $u \in e \in E(G)$,

$$\frac{1}{\psi_{U_u^e}}\sum_{S\subset F} c^S(U_u^e)\prod_{g\in S}\psi_g \le \frac{1}{\psi}\prod_{v\in N^2[u]}\left(1+\frac{\psi}{4\Delta^3}\right) \le \frac{1}{\psi}\exp\left(\psi\frac{1+\Delta^2}{4\Delta^3}\right)$$
$$= \frac{1}{4}\left(1+\frac{4t\lambda}{1+\ell\lambda}\right)\exp\left(\frac{1+\Delta^2}{\Delta^3}\left(1-\frac{4t\lambda}{1+\ell\lambda+4t\lambda}\right)\right)$$
$$\le \frac{\mathrm{e}^{(1+\Delta^2)/\Delta^3}}{4}\left(1+\frac{4t\lambda}{1+\ell\lambda}\right) \le \frac{3}{4}.$$

Hence, we can apply Theorem 4.3 with parameters ψ_{B_u} and $\psi_{U_u^e}$ such that

$$\frac{\mathrm{e}}{3\Delta^3} \le \psi_{B_u} = \frac{1}{\Delta^3}\cdot\frac{1+\ell\lambda}{1+\ell\lambda+4t\lambda} \le \frac{1}{\Delta^3},$$
$$\frac{4\mathrm{e}\lambda}{3(1+\ell\lambda)} \le \psi_{U_u^e} = \frac{4\lambda}{1+\ell\lambda+4t\lambda} \le \frac{4\lambda}{1+\ell\lambda},$$

giving $\delta = 1/4$ and

$$T_0 \le \log_2 \tilde{Z}_H(\lambda) + n\log_2\left(1+\frac{1}{\Delta^3}\right) + n\log_2\left(\frac{3\Delta^3}{\mathrm{e}}\right) + n\Delta\log_2\left(\frac{3(1+\ell\lambda)}{4\mathrm{e}\lambda}\right).$$

We have

$$\tilde{Z}_H(\lambda) = \sum_{\tau\in\Omega}\lambda^{|\operatorname{col}(\tau)|+|\operatorname{unc}(\tau)|} \le (1+2\Delta\lambda)^n,$$

because for each $u \in V(G)$ we can have either $\tau(u) = \mathfrak{B}$, which does not contribute to the exponent of λ, or $\tau(u) \in L(u) \cup \{e \in E(G) : u \in e\}$, which contributes 1. There are at most 2Δ choices in the latter case. Then

$$T_0 = O(n\log\Delta + n\log\lambda + n\Delta\log(1/\lambda) + n\Delta\log\ell).$$

Therefore, $T_0 = O(n\Delta\log(\ell))$ if $\lambda \ge 1$ while $T_0 = O(n\Delta\log(\ell/\lambda))$ if $\lambda < 1$. Since $\delta = 1/4$ and $T_0 \ge n$, setting $s = n$ yields that the probability that the algorithm finds a flawless partial colouring in at most $2T_0$ steps is at least $1-2^{-n}$. Each step takes time $O(T_r+T_s)$ because if we are addressing a flaw $f \in F_B$ then we execute action AddressB(u,σ) which executes Remove and Sample once each in series, while the action to address a flaw in F_U is simply sampling from a distribution supported on at most $\Delta+1$ outcomes with probabilities of the form $1/(1+y\lambda)$ and $\lambda/(1+y\lambda)$ where $\ell \le y \le \Delta$. This completes the proof of Lemma 4.1, showing that the first, main phase of the algorithm works as desired.

4.2 Proof of Lemma 4.4

We require the following standard concentration inequality. Given a probability space, the $\{0,1\}$-valued random variables $\mathbf{X}_1, \dots, \mathbf{X}_n$ are *negatively correlated* if for each subset S of the set $\{1, \dots, n\}$,

$$\Pr\left(\mathbf{X}_i = 1, \forall i \in S\right) \leq \prod_{i \in S} \Pr(\mathbf{X}_i = 1)\,.$$

Lemma 4.7 (Panconesi and Srinivasan [40])**.** *Given a probability space, let $\mathbf{X}_1, \dots, \mathbf{X}_n$ be $\{0,1\}$-valued random variables. Set $\mathbf{X} = \sum_{i=1}^{n} \mathbf{X}_i$ and $\mathbf{Y}_i = 1 - \mathbf{X}_i$ for each $i \in \{1, \dots, n\}$. If the variables $\mathbf{Y}_1, \dots, \mathbf{Y}_n$ are negatively correlated, then for any $\eta \in (0,1)$,*

$$\Pr\left(\mathbf{X} \leq (1-\eta)E[\mathbf{X}]\right) \leq \mathrm{e}^{-\eta^2 E[\mathbf{X}]/2}\,.$$

Recall that for Lemma 4.4 we assume the hypotheses of Theorem 2.5 (and hence also Theorem 2.4), and we consider the transition from an arbitrary partial $\mathscr{H}$-colouring σ of G induced by the action AddressB(B_u, σ) for some $u \notin \mathrm{col}(\sigma)$. In what follows, we write τ for the random state arising from this transition, and Pr and E represent probabilities and expectations over the randomness in the action AddressB(B_u, σ), respectively.

We first argue that $|L_\tau(u)|$ is large and concentrated around its expectation, and second that for each $x \in L(u)$ the probability that x is in $L_\tau(u)$ and has large degree in H^*_τ is small. For this we require some additional notation. We reuse the notation of the procedure AddressB, writing σ' for the partial $\mathscr{H}$-colouring of G obtained from σ by setting $\sigma'(v) = \mathfrak{B}$ for all vertices $v \in N_G(u) \setminus \mathrm{unc}(\sigma)$, and $\sigma'(v) = \sigma(v)$ otherwise. We also have the definitions $H' = H_{\sigma'}[L(N_G(u))]$ and $\hat{H} = \textsc{Remove}(G_{\sigma'}[N_G(u)], H')$, and $\mathbf{J}_0$ is an independent set in $\hat{H}$ sampled from the hard-core model at fugacity λ. If we write $L'(v) = L(v) \cap V(H')$ then $\mathscr{H}$ and $\mathscr{H}' = (L', H')$ are list-covers of G and $G_{\sigma'}[N_G(u)]$, respectively. This provides the additional structure that each colour $x \in V(H)$ is considered a copy of some label c in the list-assignment, and xy is an edge of H^* if and only if $x \in L(u)$ and $y \in L(v)$ are copies of the same label and $uv \in E(G)$.

For each $x \in L(u)$, let Λ_x be the *layer* of x, given by $\Lambda_x = N_{H^*_{\sigma'}}(x)$. This consists of the colours in $L_{\sigma'}(N_G(u))$ that conflict with x, and hence Λ_x consists of every colour $y \in V(H')$ that is a copy of a fixed label, written c. So for distinct $x, y \in L(u)$ the layers Λ_x and Λ_y are necessarily disjoint. The fact that $\mathscr{H}'$ is a list-cover of $G[N_G(u)]$ means that every edge leaving Λ_x in H' joins two colours belonging to some set $L(v)$ with $v \in N_G(u)$, which means there are no edges in $H^*_{\sigma'}$ between one layer Λ_x and another Λ_y and facilitates the analysis of each layer separately. The set $\mathbf{J}_0$ sampled in step 4 of the procedure AddressB is not necessarily independent in H' and the uncolouring steps 5–11 yield the subset $\mathbf{J} \subset \mathbf{J}_0$ which is independent in H'. For any $x \in L(u)$ we note that $\mathbf{J} \setminus \Lambda_x$ depends only on $\mathbf{J}_0 \setminus \Lambda_x$ and not on $\mathbf{J}_0 \cap \Lambda_x$ because of the above property of edges leaving Λ_x. To see this, observe that every uncolouring step is due to some tuple (v_1, v_2, y_1, y_2) such that $\{v_1, v_2\} \subset N_G(u)$, $v_1v_2 \in E(G)$, $y_1 \in L(v_1) \cap \mathbf{J}_0$, $y_2 \in L(v_2) \cap \mathbf{J}_0$, and $y_1y_2 \in E(H')$. But then y_1 and y_2 must be copies of the same label c' and hence either both are in Λ_x or neither one is in Λ_x. This means the uncolouring steps due to edges of H' inside Λ_x are independent of the other uncolouring steps, and hence $\mathbf{J} \setminus \Lambda_x$ depends only on $\mathbf{J}_0 \setminus \Lambda_x$.

We now show some key properties of how $\mathbf{J}_0$ is distributed on the sets Λ_x, and how this affects $L_\tau(u)$ and the flaw B_u. Let us write $\mathbf{U}_0(x)$ for the set of vertices obtained by revealing $\mathbf{J}_0 \setminus \Lambda_x$ and taking those vertices in Λ_x that in the graph $\hat{H}$ are not adjacent to any vertex of $\mathbf{J}_0 \setminus \Lambda_x$. That is, $\mathbf{U}_0(x) = \Lambda_x \setminus N_{\hat{H}}(\mathbf{J}_0 \setminus \Lambda_x)$. Then write $\mathbf{F}_0(x) = \hat{H}[\mathbf{U}_0(x)]$. By the spatial Markov property of the hard-core model, $\mathbf{J}_0 \cap \Lambda_x$ is distributed according to the hard-core model on the graph $\mathbf{F}_0(x)$ at fugacity λ. It is important to observe that $\hat{H}[\Lambda_x]$ is isomorphic to a subgraph of $G[N(u)]$, as is $\mathbf{F}_0(x)$, so the assumptions of Theorem 2.4 give that $\mathbf{F}_0(x)$ has strong local (β, γ)-occupancy.

The above definitions deal with the cover $\hat{H}$ in which we sample, but we must also deal with the original cover H'. To this end, we analogously write $\mathbf{U}(x)$ for the set of vertices obtained by revealing $\mathbf{J} \setminus \Lambda_x$ and taking those vertices in Λ_x that in the graph H' are not adjacent to any vertex of $\mathbf{J} \setminus \Lambda_x$. That is, $\mathbf{U}(x) = \Lambda_x \setminus N_{H'}(\mathbf{J} \setminus \Lambda_x)$. We also write $\mathbf{F}(x) = H'[\mathbf{U}(x)]$. We now note the following facts that hold for all $x \in L(u)$.

Fact 1 $x \in L_\tau(u)$ if and only if $\mathbf{J}_0 \cap \Lambda_x = \emptyset$.

Fact 2 $\mathbf{J} \subset \mathbf{J}_0$ and hence $\mathbf{U}(x) \supset \mathbf{U}_0(x)$.

Fact 3 If $x \in L_\tau(u)$ then $\deg^*_{\mathscr{H}_\tau}(x) = |\mathbf{U}_0(x)|$.

To see Fact 1, observe that $x \in L_\tau(u)$ if and only if for every $v \in \mathrm{col}(\tau) \cap N_G(u)$ we have $x\tau(v) \notin H$. That is, if x is a copy of the label c then $x \in L_\tau(u)$ if and only if no neighbour of u is coloured with a copy of c under τ. This clearly holds if $\mathbf{J}_0 \cap \Lambda_x = \emptyset$, and also if $\mathbf{J}_0 \cap \Lambda_x \neq \emptyset$ then at least one neighbour of u is coloured with a copy of c. But in each uncolouring step we only uncolour one end of a monochromatic edge, so that if $\mathbf{J}_0 \cap \Lambda_x \neq \emptyset$ we must also have $\mathbf{J} \cap \Lambda_x \neq \emptyset$.

For Fact 2 note that in each uncolouring step we remove a vertex from $\mathbf{J}_i$ to form $\mathbf{J}_{i+1}$ and so $\mathbf{J}$ (which is the set after all uncolouring steps) is a subset of $\mathbf{J}_0$.

Finally, for Fact 3 suppose that x is a copy of the label c and note that $\deg^*_{\mathscr{H}_\tau}(x)$ counts the number of colours in $L_\tau(N_G(u))$ that conflict with x, or equivalently the number of neighbours $v \in N_G(u)$ that are coloured blank by τ and that have a copy of c present in their list $L_\tau(v)$. A colour in $L_\tau(N_G(u))$ that conflicts with x must be present in $\mathbf{U}(x)$, and in order that $x \in L_\tau(u)$ we must have $\mathbf{J} \cap \mathbf{U}(x) = \emptyset$ by Fact 1 and Fact 2. This means that every colour $y \in \mathbf{U}(x)$ such that $y \in L(v)$ for some $v \in \mathrm{bla}(\tau)$ contributes to $\deg^*_{\mathscr{H}_\tau}(x)$. Then it suffices to show that for $y \in L(v)$ we have the following two properties

$$y \in \mathbf{U}_0(x) \Rightarrow v \in \mathrm{bla}(\tau)\,, \tag{4.4}$$

$$y \in \mathbf{U}(x) \setminus \mathbf{U}_0(x) \Rightarrow v \in \mathrm{unc}(\tau)\,. \tag{4.5}$$

This crucially exploits the fact that $\mathscr{H}'$ is a list-cover. To prove equation 4.4, observe that if $y \in \mathbf{U}_0(x)$ then no neighbour of y in $\hat{H}$ belongs to $\mathbf{J}_0 \setminus \Lambda_x$, and since $x \in L_\tau(u)$ we have $\mathbf{J}_0 \cap \Lambda_x = \emptyset$ by Fact 1. Together these facts mean that v is neither coloured nor uncoloured, and hence belongs to $\mathrm{bla}(\tau)$ as required. To prove equation 4.5, suppose that $y \in \mathbf{U}(x) \setminus \mathbf{U}_0(x)$. Then there must be some $z \in (\mathbf{J}_0 \setminus \Lambda_x) \setminus (\mathbf{J} \setminus \Lambda_x)$ such that yz is an edge of H'. As reported earlier, since $\mathscr{H}'$ is a list-cover we know that every edge leaving Λ_x is inside some set $L(w)$, and as $y \in L(v)$ we

deduce that $z \in L(v)$. But then the reason we have $z \in (\mathbf{J}_0 \setminus \Lambda_x) \setminus (\mathbf{J} \setminus \Lambda_x)$ is that z belongs to $\mathbf{J}_0$ and was then removed from some $\mathbf{J}_i$ in the creation of the next $\mathbf{J}_{i+1}$, and so necessarily $v \in \mathrm{unc}(\tau)$ as required.

We are now ready to show that $|L_\tau(u)|$ is likely to be at least ℓ.

Lemma 4.8. *Writing*

$$m = \frac{1+\lambda}{\beta\lambda}(q - \gamma\Delta) \qquad \textit{and} \qquad \eta = \sqrt{7(\log\Delta)/\ell}\,,$$

we have $\Pr\left(|L_\tau(u)| \le (1-\eta)m\right) \le \mathrm{e}^{-\eta^2 m/2} \le 1/(8\Delta^3)$.

Proof. We first note that $\mathrm{e}^{-\eta^2 m/2} \le 1/(8\Delta^3)$ holds because the parameter choices in Theorem 2.4 give

$$m = \frac{\ell}{1-\eta} \ge \ell = \frac{7\log\Delta}{\eta^2} \ge \frac{6\log(2\Delta)}{\eta^2}\,.$$

Now by Fact 1, we have $x \in L_\tau(u)$ if and only if $\mathbf{J}_0 \cap \Lambda_x = \emptyset$. As reported earlier, $\mathbf{J}_0 \cap \Lambda_x$ is distributed according to the hard-core model on a graph $\mathbf{F}_0(x)$ that is isomorphic to a subgraph of $G[N_G(u)]$. Then $\mathbf{F}_0(x)$ has strong local (β,γ)-occupancy and so

$$\beta\frac{\lambda}{1+\lambda}\Pr(x \in L_\tau(u)) + \gamma E[|\mathbf{J}_0 \cap \Lambda_x|] \ge 1\,, \tag{4.6}$$

because elementary calculations with the hard-core model (see, e. g., [17]) now give

$$\Pr(\mathbf{J}_0 \cap \Lambda_x = \emptyset) = \frac{1}{Z_{\mathbf{F}_0(x)}(\lambda)} \qquad \text{and} \qquad E[|\mathbf{J}_0 \cap \Lambda_x|] = \frac{\lambda Z'_{\mathbf{F}_0(x)}(\lambda)}{Z_{\mathbf{F}_0(x)}(\lambda)}\,.$$

We sum equation 4.6 over all q colours $x \in L(u)$ to obtain

$$q \le \beta\frac{\lambda}{1+\lambda}E[|L_\tau(u)|] + \gamma\sum_{x\in L(u)} E[|\mathbf{J}_0 \cap \Lambda_x|] \le \beta\frac{\lambda}{1+\lambda}E[|L_\tau(u)|] + \gamma\Delta\,,$$

where the last inequality holds because $\sum_{x\in L(u)} E[|\mathbf{J}_0 \cap \Lambda_x|] = E\left[\sum_{x\in L(u)} |\mathbf{J}_0 \cap \Lambda_x|\right]$ and every neighbour of u contributes at most 1 to the sum as $|\mathbf{J}_0 \cap L(v)| \le 1$ for every vertex v. Rearranging immediately yields $E[|L_\tau(u)|] \ge m$, and the result will follow from an application of Lemma 4.7.

For this application, note that $E[|L_\tau(u)|]$ is a sum over $x \in L(u)$ of the indicator variables $\mathbf{X}_x$ for the events $\{\mathbf{J}_0 \cap \Lambda_x = \emptyset\}$. We can apply Lemma 4.7 if we show that the random variables $\mathbf{Y}_x = 1 - \mathbf{X}_x$ are negatively correlated. This correlation was shown formally by Bernshteyn [8] in the triangle-free case, and is somewhat intuitive here. Consider the random set $\mathbf{J}_0$. Given $x \in L(u)$, if $\mathbf{J}_0 \cap \Lambda_x = \emptyset$ then no colours conflicting with x are chosen for vertices in $N(u)$. This makes other colours more likely to be chosen, such as those which conflict with $x' \in L(u) \setminus \{x\}$. We repeat Bernshteyn's argument for completeness.

It is enough to show that for all $x \in L(u)$ and $Y \subset L(u) \setminus \{x\}$ we have

$$\Pr\left(x \notin L_\tau(u) \,\middle|\, Y \cap L_\tau(u) = \emptyset\right) \leq \Pr(x \notin L_\tau(u))\,,$$

which is equivalent to

$$\Pr\left(Y \cap L_\tau(u) = \emptyset \,\middle|\, x \in L_\tau(u)\right) \geq \Pr(Y \cap L_\tau(u) = \emptyset)\,,$$

which we can write (using Fact 1) as

$$\Pr\left(\mathbf{J}_0 \cap \Lambda_y \neq \emptyset \text{ for all } y \in Y \,\middle|\, \mathbf{J}_0 \cap \Lambda_x = \emptyset\right) \geq \Pr\left(\mathbf{J}_0 \cap \Lambda_y \neq \emptyset \text{ for all } y \in Y\right).$$

This holds because the layers Λ_z for $z \in L(u)$ are pairwise disjoint. □

We now prove a result designed to handle the degree condition in the flaw B_u.

Lemma 4.9. *For any $x \in L(u)$, writing*

$$D_x = \left\{x \in L_\tau(u) \text{ and } \deg^*_{\mathscr{H}_\tau}(x) > \ell/8\right\},$$

we have $\Pr(D_x) \leq 1/(8q \cdot \Delta^3)$.

Proof. By Fact 3 we know that if $x \in L_\tau(u)$ then $\deg^*_{\mathscr{H}_\tau}(x) = |\mathbf{U}_0(x)|$, so it suffices to show whenever $|\mathbf{U}_0(x)| > \ell/8$ that

$$\Pr(x \in L_\tau(u)) \leq \frac{1}{8q\Delta^3}\,.$$

As already reported, by Fact 1 (and an elementary property of the hard-core model) we have $\Pr(x \in L_\tau(u)) = 1/Z_{\mathbf{F}_0(x)}(\lambda)$. We recall that $\mathbf{F}_0(\lambda)$ is isomorphic to a subgraph of $G[N(u)]$, so that the upper bound on $\Pr(x \in L_\tau(u))$ follows directly from the assumptions on ℓ and $Z_F(\lambda)$ for $F \subset G[N(u)]$ stated in Theorem 2.4. There we assume that for all such F on at least $\ell/8$ vertices we have $Z_F(\lambda) \geq 8\Delta^4$. Consequently, noticing that without loss of generality we have $q \leq \Delta$ (for otherwise a greedy argument finds any q-list colouring of G), we infer that

$$\Pr(x \in L_\tau(u) \,|\, |\mathbf{U}_0(x)| > \ell/8) \leq \frac{1}{8\Delta^4} \leq \frac{1}{8q\Delta^3}\,.$$

The result follows. □

The combination of Lemma 4.8 and Lemma 4.9 completes the proof of Lemma 4.4.

4.3 Proof of Lemma 4.2

We can prove Lemma 4.2 with the local lemma of Moser and Tardos [39]. Let σ be a flawless partial colouring of G, and if necessary remove colours from H_σ such that $|L_\sigma(u)| = \ell$ for all $u \in G_\sigma$. We let $I' \subset V(H_\sigma)$ be obtained by choosing for each remaining blank vertex $u \in \mathrm{bla}(\sigma)$ a uniform random colour from $L_\sigma(u)$. Since σ is flawless there are no uncoloured vertices, so the conclusion

follows if I' is independent. Then for this algorithmic local lemma application we have a flaw A_{xy} for each edge xy of H^*_σ, and $I' \in A_{xy}$ if and only if $\{x, y\} \subset I'$. To address A_{xy} we resample $I' \cap L(u)$ and $I' \cap L(v)$ independently, uniformly at random.

Suppose that $x \in L(u)$, $y \in L(v)$, and $xy \in E(H^*_\sigma)$. Then A_{xy} is independent of any $A_{x'y'}$ such that x' and y' are not in $L_\sigma(u) \cup L_\sigma(v)$. In this simple setting the full charge machinery of Theorem 4.3 is not necessary, and it suffices to use the result of Moser and Tardos [39, Thm. 1.2]. We use the slightly stronger version found in Iliopoulos' doctoral thesis [29, Thm. 3.9, p. 17]: it implies the original formulation by setting $x_A = \frac{\psi_A}{1+\psi_A} \in (0, 1)$.

Theorem 4.10 (Iliopoulos [29])**.** *Let $\mathcal{P}$ be a finite set of mutually independent random variables, and let $\mathcal{A}$ be a finite set of events determined by these variables. Let $\Gamma(A) = \{B \in \mathcal{A} : A \text{ and } B \text{ are dependent}\}$, and let μ be the probability measure that results from sampling the variables $\mathcal{P}$. If for each $A \in \mathcal{A}$ there exists a positive real ψ_A such that*

$$\frac{\mu(A)}{\psi_A} \sum_{S \subset \Gamma(A)} \prod_{B \in S} \psi_B \leq 1\,,$$

then there is a randomised algorithm that finds an assignment to $\mathcal{P}$ violating none of the events in $\mathcal{A}$ with expected number of resamplings $\sum_{a \in \mathcal{A}} \psi_A$.

Note that we have $A \in \Gamma(A)$ since without loss of generality $\mu(A) \in (0, 1)$ and hence A depends on itself.

In our setting, the variables in $\mathcal{P}$ are $I' \cap L(u)$ for $u \in V(G_\sigma)$ and the set of events is $\mathcal{A} = \{A_{xy} : xy \in H^*_\sigma\}$. Supposing that $x \in L(u)$, $y \in L(v)$, and $xy \in E(H^*_\sigma)$, the fact that σ is flawless implies that

$$\mu(A_{xy}) = \frac{1}{|L(u)||L(v)|} = \ell^{-2}\,,$$

and also that

$$|\Gamma(A_{xy})| \leq \sum_{x' \in L(u)} \deg^*_{H_\sigma}(x') + \sum_{y' \in L(v)} \deg^*_{H_\sigma}(y') \leq \ell^2/4\,.$$

Hence, it suffices to take $\psi = 4\ell^{-2}$ to have, for each $A \in \mathcal{A}$,

$$\frac{\mu(A)}{\psi_A} \sum_{S \subset \Gamma(A)} \prod_{B \in S} \psi_B \leq \frac{1}{4} \cdot (1 + 4/\ell^2)^{\ell^2/4} \leq \frac{\mathrm{e}}{4} < 1\,.$$

The expected number of resamplings is then at most $n/4$, because H^*_σ, having no more than $n \cdot \ell$ vertices and maximum degree no more than $\ell/8$, contains at most $n\ell^2/16$ edges. Consequently, Markov's inequality implies that, for any real $c \in (0, 1)$, the algorithm succeeds in $O\left(n^{1+c}\right)$ resamplings with probability at least $1 - 1/n^c$. Since a resampling is done in constant time, this constitutes a proof of Lemma 4.2.

5 Tightness

In this section we state a proposition indicating that for a large range of λ, the local occupancy of Lemma 3.1 is asymptotically best possible for F_k-free graphs. First, we note an extra (but foundational) component of our framework essentially originating in [17], see [18].

Theorem 5.1. *Let G be a graph of maximum degree Δ such that the hard-core model on G at fugacity λ has (strong) local (β, γ)-occupancy for some $\lambda, \beta, \gamma > 0$. Then the occupancy fraction of G at fugacity λ is at least $1/(\beta + \gamma\Delta)$.*

Suppose that k is a fixed integer greater than 2 and that we choose λ and t as functions of Δ such that $\lambda = o(1)$ and $\lambda\sqrt{t} = o(1)$ as $\Delta \to \infty$. Let G be a graph of maximum degree Δ in which each vertex is the centre of at most t copies of F_k. Let $a = k - 3 + \sqrt{2t}$. We know from Lemma 3.1 applied with this a and $d = \Delta$ (as permitted by Lemma 3.3), that there exist $\beta, \gamma > 0$ such that G has strong local (β, γ)-occupancy and

$$\beta + \gamma\Delta = \frac{1+\lambda}{\lambda} \frac{\Delta(1+\lambda)^a \log(1+\lambda)}{W(\Delta(1+\lambda)^a \log(1+\lambda))}.$$

As $\Delta \to \infty$, using that $k = O(1)$, $\lambda = o(1)$, and $\lambda\sqrt{t} = o(1)$ we have $a\log(1+\lambda) = o(1)$ and hence by Theorem 5.1 the occupancy fraction of G at fugacity λ is at least

$$(1 - o(1))\frac{W(\Delta\lambda)}{\Delta}.$$

But the occupancy fraction is monotone increasing in λ (see [17]) so this lower bounds holds for all larger values of λ too. This is asymptotically tight at least for $\lambda \le \Delta^{1+o(1)}$ but in fact no improvement to this lower bound is known for larger values of λ, even for the case of triangle-free graphs (that corresponds to $k = 3$). In addition, when $t = \Delta^{o(1)}$ we can take $\lambda = 1/\log\Delta$ and obtain a lower bound on the occupancy fraction which is $(1 - o(1))\log\Delta/\Delta$, and hence an improvement to the leading order for any λ larger than $1/\log\Delta$ would immediately lead to an improvement to Shearer's result [43] that every n-vertex triangle-free graph of maximum degree Δ contains an independent set of size at least $(1 - o(1))n\log\Delta/\Delta$.

The following proposition shows that the above lower bound on the occupancy fraction is asymptotically tight, so that any leading-order improvement to Shearer's bound cannot be obtained via quantitatively improving local occupancy parameters in triangle-free graphs.

Proposition 5.2 ([17])**.** *Given $\epsilon > 0$ there is Δ_0 such that for all fixed $\Delta > \Delta_0$ and $\lambda \le \Delta^{1+o(1)}$, there is a Δ-regular F_k-free graph G with occupancy fraction at most $(1+\epsilon)W(\Delta\lambda)/\Delta$.*

6 Concluding remarks

In this section we compare the algorithmic framework given here with the more combinatorial treatment in the companion paper [18]. While in the introduction we deliberately omitted

all mention of stronger graph colouring concepts, here we describe some subtleties related to several strengthenings handled within our framework, including local and correspondence colouring.

First, our methods immediately generalise to a 'local' formulation where each vertex is given a list of size depending primarily on $\deg(v)$ instead of Δ. See [14, 18] for details, and in particular an interesting minimum list size phenomenon that arises. Secondly, when applying Theorem 2.5 with a trivial implementation of REMOVE that does nothing, our proof gives two extra properties. When REMOVE is trivial, Theorem 2.5 works for correspondence colouring; one can dispense with the stated list-cover assumption in this case. Moreover, if REMOVE is trivial and we do have a list-cover then Theorem 2.5 only requires weak local occupancy: a variant of local occupancy that applies only to *induced subgraphs* F of $G[N(u)]$ instead of to arbitrary subgraphs.

This means that our methods give an algorithmic version of Bernshteyn's strengthening of Theorem 1.1 to correspondence colouring [8], as well as the generalisation to F_k-free graphs for any $k \geq 3$ (recall that F_3 is a triangle and $F_3, C_k \subset F_k$). In cases where we have local occupancy (e. g., one of the many settings covered in [18]), the primary bottleneck for an efficient algorithm is SAMPLE, and one of our contributions here is a polynomial-time implementation of SAMPLE that suffices for Theorem 1.2. Without seeking an efficient algorithm, raising t presents no serious challenge as our quantitative local occupancy guarantee degrades smoothly and slowly as t increases. This is one of the key breakthroughs of our framework, see [18, Sec. 5]. By contrast, the efficient implementation of SAMPLE is extremely fragile and fails completely at $t = 1$. With the power of REMOVE we can handle t from 1 up to the stated bound in Theorem 1.3, but the cost of this is twofold. First, we require a list-assignment as the analysis of the uncolouring steps depends crucially upon this. Second, we require strong instead of weak local occupancy because removing edges can create arbitrary subgraphs of neighbourhoods $G[N(u)]$ at the relevant sampling stage, even with a list-assignment.

Here summarises some key differences among these strengthened colourings as treated both here and in [18]. Consider the problem of colouring a graph G of maximum degree Δ in which each vertex is contained in at most t copies of F_k with $(1+\epsilon)\Delta/\log(\Delta/\sqrt{t})$ colours, or lists of this size. Let Δ_0 be large enough in terms of ϵ and $\Delta \geq \Delta_0$, and let $c = c(\epsilon)$ be a large enough constant. For usual graph colouring our *existence* methods work up to $t \leq \Delta^2/c$, but for the largest t this requires a reduction that does not apply to list colouring. For list colouring we require $t \leq \Delta^2/(\log\Delta)^{2/\epsilon}$, and our *algorithmic* methods work for $t \leq \Delta^{\frac{2\epsilon}{1+2\epsilon}}/(\log\Delta)^2$. For correspondence colouring our existence methods behave the same as for list colouring, but for an efficient algorithm we need $t < 1$. It would be very interesting to learn if these differences are essential, or whether refined techniques can unify these results.

7 Acknowledgements

We thank Alistair Sinclair, Fotis Iliopoulos, and Charlie Carlson for insightful discussions and we thank the anonymous referees for their comments.

References

[1] DIMITRIS ACHLIOPTAS AND AMIN COJA-OGHLAN: Algorithmic barriers from phase transitions. In *Proc. 49th FOCS*, pp. 793–802. IEEE Comp. Soc., 2008. [doi:10.1109/FOCS.2008.11] 3

[2] DIMITRIS ACHLIOPTAS, FOTIS ILIOPOULOS, AND ALISTAIR SINCLAIR: Beyond the Lovász Local Lemma: Point to set correlations and their algorithmic applications. In *Proc. 60th FOCS*, pp. 725–744. IEEE Comp. Soc., 2019. [doi:10.1109/FOCS.2019.00049] 2, 4, 5, 6, 7, 8, 18, 19, 21

[3] MIKLÓS AJTAI, JÁNOS KOMLÓS, AND ENDRE SZEMERÉDI: A note on Ramsey numbers. *J. Combin. Theory–A*, 29(3):354–360, 1980. [doi:10.1016/0097-3165(80)90030-8] 3

[4] MIKLÓS AJTAI, JÁNOS KOMLÓS, AND ENDRE SZEMERÉDI: A dense infinite Sidon sequence. *Europ. J. Combinat.*, 2(1):1–11, 1981. [doi:10.1016/S0195-6698(81)80014-5] 3

[5] NOGA ALON: Independence numbers of locally sparse graphs and a Ramsey type problem. *Random Struct. Algor.*, 9(3):271–278, 1996. [doi:10.1002/(SICI)1098-2418(199610)9:3<271::AID-RSA1>3.0.CO;2-U] 7, 8

[6] NOGA ALON, MICHAEL KRIVELEVICH, AND BENNY SUDAKOV: Coloring graphs with sparse neighborhoods. *J. Combin. Theory–B*, 77(1):73–82, 1999. [doi:10.1006/jctb.1999.1910] 4, 7, 8

[7] NIKHIL BANSAL, ANUPAM GUPTA, AND GURU GURUGANESH: On the Lovász theta function for independent sets in sparse graphs. *SIAM J. Comput.*, 47(3):1039–1055, 2018. Preliminary version in STOC'15. [doi:10.1137/15M1051002] 2, 3, 7, 8

[8] ANTON BERNSHTEYN: The Johansson–Molloy theorem for DP-coloring. *Random Struct. Algor.*, 54(4):653–664, 2019. [doi:10.1002/rsa.20811] 2, 8, 19, 28, 32

[9] HANS L. BODLAENDER: On linear time minor tests with depth-first search. *J. Algorithms*, 14(1):1–23, 1993. [doi:10.1006/jagm.1993.1001] 15

[10] TOM BOHMAN: The triangle-free process. *Adv. Math.*, 221(5):1653–1677, 2009. [doi:10.1016/j.aim.2009.02.018] 3

[11] TOM BOHMAN AND PETER KEEVASH: Dynamic concentration of the triangle-free process. *Random Struct. Algor.*, 58(2):221–293, 2021. [doi:10.1002/rsa.20973] 3

[12] MARTHE BONAMY, TOM KELLY, PETER NELSON, AND LUKE POSTLE: Bounding χ by a fraction of Δ for graphs without large cliques. *J. Combin. Theory–B*, 157:263–282, 2022. [doi:10.1016/j.jctb.2022.06.002] 8

[13] ROB M. CORLESS, GASTON H. GONNET, DAVID E. G. HARE, DAVID J. JEFFREY, AND DONALD E. KNUTH: On the Lambert *W* function. *Adv. Comput. Math.*, 5(1):329–359, 1996. [doi:10.1007/BF02124750] 11

[14] EWAN DAVIES, RÉMI DE JOANNIS DE VERCLOS, ROSS J. KANG, AND FRANÇOIS PIROT: Coloring triangle-free graphs with local list sizes. *Random Struct. Algor.*, 57(3):730–744, 2020. [doi:10.1002/rsa.20945] 8, 32

[15] EWAN DAVIES, RÉMI DE JOANNIS DE VERCLOS, ROSS J. KANG, AND FRANÇOIS PIROT: Occupancy fraction, fractional colouring, and triangle fraction. *J. Graph Theory*, 97(4):557–568, 2021. [doi:10.1002/jgt.22671] 8, 11

[16] EWAN DAVIES, MATTHEW JENSSEN, WILL PERKINS, AND BARNABY ROBERTS: Independent sets, matchings, and occupancy fractions. *J. London Math. Soc.*, 96(1):47–66, 2017. [doi:10.1112/jlms.12056] 8

[17] EWAN DAVIES, MATTHEW JENSSEN, WILL PERKINS, AND BARNABY ROBERTS: On the average size of independent sets in triangle-free graphs. *Proc. AMS*, 146(1):111–124, 2018. [doi:10.1090/proc/13728] 3, 8, 28, 31

[18] EWAN DAVIES, ROSS J. KANG, FRANÇOIS PIROT, AND JEAN-SÉBASTIEN SERENI: Graph structure via local occupancy. 2020. [arXiv:2003.14361] 2, 5, 6, 8, 9, 10, 19, 22, 31, 32

[19] ZDENĚK DVOŘÁK AND LUKE POSTLE: Correspondence coloring and its application to list-coloring planar graphs without cycles of lengths 4 to 8. *J. Combin. Theory–B*, 129:38–54, 2018. [doi:10.1016/j.jctb.2017.09.001] 9

[20] PAUL ERDŐS: Some remarks on the theory of graphs. *Bull. AMS*, 53(4):292–294, 1947. [doi:10.1090/S0002-9904-1947-08785-1] 2

[21] PAUL ERDŐS: Graph theory and probability. *Canad. J. Math.*, 11:34–38, 1959. [doi:10.4153/CJM-1959-003-9] 2

[22] PAUL ERDŐS AND TIBOR GALLAI: On maximal paths and circuits of graphs. *Acta Math. Hungar.*, 10(3–4):337–356, 1959. [doi:10.1007/BF02024498] 14

[23] PAUL ERDŐS AND GEORGE SZEKERES: A combinatorial problem in geometry. *Compositio Math.*, 2:463–470, 1935. Available at EuDML. 2, 3

[24] PAUL ERDŐS AND LÁSZLÓ LOVÁSZ: Problems and results on 3-chromatic hypergraphs and some related questions. In ANDRÁS HAJNAL, RICHARD RADO, AND VERA T. SÓS, editors, *Infinite and finite sets: To Paul Erdős on his 60th birthday*, volume II of *Colloq. Math. Soc. János Bolyai, 10*, pp. 609–627. János Bolyai Math. Soc. and North-Holland, 1975. 7

[25] LOUIS ESPERET AND ALINE PARREAU: Acyclic edge-coloring using entropy compression. *Europ. J. Combinat.*, 34(6):1019–1027, 2013. [doi:10.1016/j.ejc.2013.02.007] 7

[26] GONZALO FIZ PONTIVEROS, SIMON GRIFFITHS, AND ROBERT MORRIS: The triangle-free process and the Ramsey number $R(3,k)$. *Mem. AMS*, 263(1274), 2020. [doi:10.1090/memo/1274] 3

[27] ALAN FRIEZE AND TOMASZ ŁUCZAK: On the independence and chromatic numbers of random regular graphs. *J. Combin. Theory–B*, 54(1):123–132, 1992. [doi:10.1016/0095-8956(92)90070-E] 3

[28] MAGNÚS M. HALLDÓRSSON AND JAIKUMAR RADHAKRISHNAN: Improved approximations of independent sets in bounded-degree graphs via subgraph removal. *Nord. J. Comput.*, 1(4):475–492, 1994. Preliminary version in SWAT'94. 2, 3

[29] FOTIS ILIOPOULOS: *Stochastic Local Search and the Lovász Local Lemma*. Ph. D. thesis, EECS Department, University of California, Berkeley, 2019. Available at UC Berkeley. 30

[30] ANDERS JOHANSSON: Asymptotic choice number for triangle-free graphs. Technical Report 91-5, DIMACS, 1996. 3, 7, 8

[31] ROSS J. KANG AND COLIN MCDIARMID: Colouring random graphs. In LOWELL W. BEINEKE AND ROBIN J. WILSON, editors, *Topics in chromatic graph theory*, volume 156 of *Encyclopedia Math. Appl.*, pp. 199–229. Cambridge Univ. Press, 2015. [doi:10.1017/CBO9781139519793.013] 2

[32] RICHARD M. KARP: Reducibility among combinatorial problems. In RAYMOND E. MILLER, JAMES W. THATCHER, AND JEAN D. BOHLINGER, editors, *Complexity of Computer Computations*, pp. 85–103. Springer, 1972. [doi:10.1007/978-1-4684-2001-2_9] 2

[33] RICHARD M. KARP: The probabilistic analysis of some combinatorial search algorithms. In *Algorithms and complexity (Proc. Sympos., Carnegie-Mellon Univ., Pittsburgh, Pa., 1976)*, pp. 1–19, 1976. Available at UC Berkeley. 2

[34] JEONG HAN KIM: On Brooks' theorem for sparse graphs. *Combin. Probab. Comput.*, 4(2):97–132, 1995. [doi:10.1017/S0963548300001528] 8

[35] JEONG HAN KIM: The Ramsey number $R(3,t)$ has order of magnitude $t^2/\log t$. *Random Struct. Algor.*, 7(3):173–207, 1995. [doi:10.1002/rsa.3240070302] 3

[36] MICHAEL MOLLOY: The list chromatic number of graphs with small clique number. *J. Combin. Theory–B*, 134:264–284, 2019. [doi:10.1016/j.jctb.2018.06.007] 3, 5, 8, 18

[37] MICHAEL MOLLOY AND BRUCE REED: *Graph Colouring and the Probabilistic Method*. Volume 23 of *Algorithms and Combinatorics*. Springer, 2002. [doi:10.1007/978-3-642-04016-0] 7, 8

[38] ROBIN A. MOSER: A constructive proof of the Lovász Local Lemma. In *Proc. 41st STOC*, pp. 343–350. ACM Press, 2009. [doi:10.1145/1536414.1536462] 7

[39] ROBIN A. MOSER AND GÁBOR TARDOS: A constructive proof of the general Lovász Local Lemma. *J. ACM*, 57(2):11:1–15, 2010. [doi:10.1145/1667053.1667060] 7, 29, 30

[40] ALESSANDRO PANCONESI AND ARAVIND SRINIVASAN: Randomized distributed edge coloring via an extension of the Chernoff–Hoeffding bounds. *SIAM J. Comput.*, 26(2):350–368, 1997. [doi:10.1137/S0097539793250767] 26

[41] SETH PETTIE AND HSIN-HAO SU: Distributed coloring algorithms for triangle-free graphs. *Inform. Comput.*, 243:263–280, 2015. Preliminary version in ICALP'13. [doi:10.1016/J.IC.2014.12.018] 3

[42] FRANK P. RAMSEY: On a problem of formal logic. *Proc. London Math. Soc.*, 30(4):264–286, 1929. [doi:10.1112/plms/s2-30.1.264] 2, 3

[43] JAMES B. SHEARER: A note on the independence number of triangle-free graphs. *Discr. Math.*, 46(1):83–87, 1983. [doi:10.1016/0012-365X(83)90273-X] 3, 5, 7, 31

[44] JAMES B. SHEARER: On the independence number of sparse graphs. *Random Struct. Algor.*, 7(3):269–271, 1995. [doi:10.1002/rsa.3240070305] 7, 8

[45] PETER UNGAR AND BLANCHE DESCARTES: Solution for 4526: k-chromatic graphs without triangles. *Amer. Math. Monthly*, 61(5):352–353, 1954. [doi:10.2307/2307489] 3

[46] VADIM G. VIZING: Some unsolved problems in graph theory. *Russ. Math. Surv.*, 23(6):125–141, 1968. [doi:10.1070/RM1968v023n06ABEH001252] 3

[47] VAN H. VU: A general upper bound on the list chromatic number of locally sparse graphs. *Combin. Probab. Comput.*, 11(1):103–111, 2002. [doi:10.1017/S0963548301004898] 8

[48] ALEXANDER A. ZYKOV: On some properties of linear complexes. *Sbornik Math.*, 66(2):163–188, 1949. Available at Math-Net.Ru. 3

AUTHORS

Ewan Davies
Assistant professor
Department of Computer Science
Colorado State University
Fort Collins, Colorado, USA
ewan.davies@colostate.edu
https://ewandavies.org/

Ross J. Kang
Associate professor
Korteweg–de Vries Institute for Mathematics
University of Amsterdam
Amsterdam, the Netherlands
ross.kang@gmail.com
https://staff.fnwi.uva.nl/j.r.kang/

François Pirot
Associate professor (*Maître de conférences*)
Laboratoire interdisciplinaire des sciences du numérique (LISN)
Universié Paris-Saclay
Gif-sur-Yvette, France
francois.pirot@lisn.fr
https://www.lri.fr/~fpirot/

Jean-Sébastien Sereni
Directeur de recherche at C.N.R.S.
Centre de Recherche en Biomédecine de Strasbourg
ICube/UMR7357, CSTB
Strasbourg, France
jean-sebastien.sereni@cnrs.fr
https://lbgi.fr/~sereni/

ABOUT THE AUTHORS

EWAN DAVIES received his Ph.D. in Discrete Mathematics from the London School of Economics and Political Science in 2017, advised by Peter Allen and Jozef Skokan. During postdoctoral positions at the University of Amsterdam, the Simons Institute for the Theory of Computing, and the University of Colorado Boulder he became increasingly interested in algorithms and the theory of computer science. His research interests now include approximate counting and sampling, probabilistic combinatorics, and extremal graph theory. He enjoys spending time in the mountains.

ROSS J. KANG is a Canadian mathematician based in Amsterdam. After obtaining a DPhil at the University of Oxford in 2008, he held postdoctoral positions in Canada and the UK before settling in the Netherlands, based at the University of Amsterdam since 2022. He has a broad base of interests across combinatorics, probability and algorithms, and he returns time and again to random graphs, probabilistic combinatorics, and graph colouring. In recent years he has nurtured an active interest in fostering diversity and openness in the (mathematical) sciences.

FRANÇOIS PIROT obtained the degrees of doctor in Mathematics from Radboud University and in Computer Sciences from the Université de Lorraine in 2019. After postdoctoral positions in Université Libre de Bruxelles (Belgique), G-SCOP (Grenoble, France), and Inria Sophia Antipolis (France), he became associate professor in Université Paris-Saclay (France) in 2021. His main research activities focus on the colouring problem in various contexts, including graph powers, locally sparse graphs, distributed colouring, and randomised algorithms. He has also worked on problems from Bioinformatics related to circular codes, using efficient graph theoretical tools.

JEAN-SÉBASTIEN SERENI obtained the degree of doctor at the University of Nice-Sophia Antipolis (France) in 2006, advised by Frédéric Havet. He is a full-time researcher at Centre National de la Recherche Scientifique, currently based at the University of Strasbourg.