# An algorithmic framework for colouring locally sparse graphs 

Ewan Davies, Ross Kang, François Pirot, Jean-Sébastien Sereni

## To cite this version:

Ewan Davies, Ross Kang, François Pirot, Jean-Sébastien Sereni. An algorithmic framework for colouring locally sparse graphs. 2020. hal-02548191

## HAL Id: hal-02548191 <br> https://hal.science/hal-02548191

Preprint submitted on 20 Apr 2020

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

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

# AN ALGORITHMIC FRAMEWORK FOR COLOURING LOCALLY SPARSE GRAPHS 

EWAN DAVIES, ROSS J. KANG, FRANÇOIS PIROT, AND JEAN-SÉBASTIEN SERENI


#### 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 $\Delta$ in which each vertex is contained in at most $t$ copies of a cycle of length $k$, where $1 / 2 \leq t \leq \Delta^{\frac{2 \varepsilon}{1+2 \varepsilon}} /(\log \Delta)^{2}$, with $\lfloor(1+\varepsilon) \Delta / \log (\Delta / \sqrt{t})\rfloor$ colours.

This 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). This bound on the chromatic number is tight up to an asymptotic factor 2 and it coincides with a famous algorithmic barrier to colouring random graphs.


## 1. Introduction

Let $G=(V, E)$ be a graph. An independent set of $G$ is a vertex subset that induces an edgeless subgraph of $G$. The independence number $\alpha(G)$ of $G$ is the cardinality of a largest independent set of $G$. The chromatic number $\chi(G)$ of $G$ is the least number of parts in a partition of $V$ into independent sets of $G$. Determining or bounding these structural parameters have been of fundamental importance to algorithms, optimisation, and operations research [30]. Moreover, they have been central in the development of combinatorial mathematics, especially with respect to random graphs and extremal combinatorics [39, 21, 19, 20].

The algorithmic and combinatorial perspectives are inextricably linked. As an example, with an interpretation of the random graph $G_{n, 1 / 2}$ as a model of average-case behaviour, Karp asked in 1976 [31] if for some positive $\varepsilon$ there is a polynomial-time algorithm that outputs an independent set in $G_{n, 1 / 2}$ of size $(1+\varepsilon) \log _{2} n$ 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 existentially we have $\alpha\left(G_{n, 1 / 2}\right) \sim 2 \log _{2} n$ w.h.p.) Karp's question remains open and has helped to provoke an influential, sustained series of investigations in random graph theory, cf. e.g. [29] 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, cf. e.g. [26, 6.

Our main contribution is a novel framework for the asymptotic global structure - in terms of independent sets or colourings-of graphs that satisfy some local sparsity condition, having e.g. few edges in any induced neighbourhood subgraph. This framework is built around the establishment of elementary local properties of the so-called hard-core model on a graph, a probabilistic approach having its roots in statistical physics. Our work lies near the interface between the above-mentioned parallel perspectives, and in fact is closely related to the two examples above. In this extended abstract we focus on algorithmic aspects of our framework, through one specific (and important) application, and show a comfortable incorporation of modern stochastic local search machinery to improve on the state of the art. In a companion paper [17] we explore a broader but also more combinatorial array of applications, prioritising not-necessarily-algorithmic existential results.

[^0]One old and basic starting point for this research is the pursuit of global asymptotic structure in triangle-free graphs, that is, in graphs having no edges whatsoever in any induced neighbourhood subgraph. The search for large independent sets in this context corresponds to the classic off-diagonal case of Ramsey numbers [39, 21, 3, 4, 40, 33, 9, 10, 25, 16], 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, cf. [45, 41].

There is particular interest in graphs of bounded maximum degree, with natural links to approximation algorithms, cf. e.g. [26, 6]. For colouring this interest originated in a question of Vizing from 1968 [42]: what is the largest chromatic number taken over all triangle-free graphs of maximum degree $\Delta$ ? (Even without the triangle-free condition a trivial 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 [28] devised a sophisticated semirandom colouring procedure to establish an upper bound of $O(\Delta / \log \Delta)$ as $\Delta \rightarrow \infty$. Recently, in a dramatic advance, Molloy [34] employed entropy compression for a simplified proof and an intriguing improvement over Johansson's result, quantitatively matching an analogous independence number bound of Shearer [40].
Theorem 1 (Molloy [34]). For all $\varepsilon>0$, there exists $\Delta_{0}$ such that if $\Delta \geq \Delta_{0}$, then $\chi(G) \leq$ $(1+\varepsilon) \Delta / \log \Delta$ for any given triangle-free graph $G$ of maximum degree $\Delta$. There is a randomised algorithm that in polynomial time w.h.p. constructs a certificate colouring of $G$.

For a hint of how difficult it might be to improve on this result, particularly with respect to the asymptotic leading constant of 1 , one can take two issues into consideration. First, lowering the constant appreciably would by the same token improve upon the best to date lower bounds on the classical off-diagonal Ramsey numbers (which are due to Shearer [40] as alluded to above), and it would constitute a breakthrough in quantitative Ramsey theory. Second, a lowering of the constant and with a polynomial-time algorithm would essentially imply a positive answer to the direct analogue of Karp's question above, for the random $\Delta$-regular rather than binomial random graph. Indeed, the following result is well known in random graph theory.
Proposition 2. For all $\varepsilon>0$, there exists some $\Delta_{0}$ such that for all fixed $\Delta \geq \Delta_{0}$, we have the following for all $n$ sufficiently large. With probability at least $1-\varepsilon$, the random $\Delta$-regular graph $G_{n, \Delta}$ on $n$ vertices is triangle-free and satisfies $\alpha\left(G_{n, \Delta}\right) \in(2 \pm \varepsilon)(n \log \Delta) / \Delta$.
Since $\alpha(G) \geq|V| / \chi(G)$ for all $G=(V, E)$, this shows the asymptotic leading term in Molloy's result (and the corresponding result of Shearer) to be correct up to a factor 2.

We offer a more general principle behind Theorem 1, through locally-defined probabilistic properties of the independent sets. Through this, one may witness that certain methods behind Theorem 1 are sharp and cannot be improved asymptotically; we discuss this in Subsection 1.2 . Important too is that the principle is flexible enough for a host of applications, which we partially present through this extended abstract (with more treated in the companion paper [17]). To give a first flavour of the extra breadth in our approach, here is a prototypical version of our main result in this extended abstract. For $k \geq 3$, let us define the fan $F_{k}$ of order $k$ as the graph formed from a path on $k-1$ vertices by adding a vertex joined to all vertices of the path. We call a graph $F_{k}$-free if it does not contain the fan $F_{k}$ as a subgraph.
Theorem 3. Fix an integer $k \geq 3$. For all $\varepsilon>0$, there exists some $\Delta_{0}$ such that if $\Delta \geq \Delta_{0}$, then $\chi(G) \leq(1+\varepsilon) \Delta / \log \Delta$ for any given $F_{k}$-free graph $G$ of maximum degree $\Delta$. There is a randomised algorithm that in polynomial time w.h.p. constructs a certificate colouring of $G$.

Note $F_{k}$ contains a cycle of each length between 3 and $k$, and so this strengthens Theorem 1 in a natural way. For $k>3$, earlier work in this direction [5, 43, 2] was not enough to obtain a leading asymptotic constant of 1 (even without demanding a polynomial-time algorithm). Keeping in mind Proposition 2, this constant is at most twice the optimal value, just as for Theorem 1 .

Another basic but more modern starting point for this research is the 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 satisfying assignment or a colouring.) In a remarkable breakthrough, Moser [36] (cf. [37]), showed that the Lovász local lemma [23]-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, cf. e.g. [24]. 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 , under a smooth relaxation of the triangle-free condition.
Theorem 4 (Achlioptas, Iliopoulos, and Sinclair [2]). For all $\varepsilon>0$, there exists $\Delta_{0}$ such that if $\Delta \geq \Delta_{0}$ and $1 / 2 \leq t \leq \Delta^{\frac{2 \varepsilon}{1+2 \varepsilon}} /(\log \Delta)^{2}$, then $\chi(G) \leq(1+\varepsilon) \Delta / \log (\Delta / \sqrt{t})$ for any given graph $G$ of maximum degree $\Delta$ where each vertex of $G$ is contained in at most triangles. There is a randomised algorithm that in polynomial time w.h.p. constructs a certificate colouring of $G$.
1.1. Our contributions. As mentioned above, our main achievement is the development of a general framework for global graph structure that significantly strengthens the results stated above (Theorems 1, 3, and 44). It encompasses or improves upon a long line of earlier work in this area [2, 5, ,7, 13, 14, 28, 32, 34, 43]. 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 for existing results with matching or improved bounds, cf. [17]. Moreover, subject to mild extra conditions, we can give polynomial-time constructions, which is our focus here. Our main application is a common generalisation of Theorems 3 and 4 (and it implies the announced result since $C_{k} \subset F_{k}$ ).
Theorem 5. Fix an integer $k \geq 3$. For all $\varepsilon>0$, there exists $\Delta_{0}$ such that if $\Delta \geq \Delta_{0}$ and $1 / 2 \leq t \leq \Delta^{\frac{2 \varepsilon}{1+2 \varepsilon}} /(\log \Delta)^{2}$, then $\chi(G) \leq(1+\varepsilon) \Delta / \log (\Delta / \sqrt{t})$ for any given graph $G$ of maximum degree $\Delta$ where each vertex of $G$ is contained in at most $t$ copies of the fan $F_{k}$. There is a randomised algorithm that in polynomial time w.h.p. constructs a certificate colouring of $G$.

Our method builds upon Molloy's proof of Theorem 1. starting with a 'blank' partial colouring and resampling the colours of neighbourhoods until a flawless partial colouring is found. We have distilled the graph structure necessary for this resampling to eventually succeed, namely local occupancy, and this strategy alone suffices for the existence of colourings as guaranteed by Theorem 5. Resampling is performed according to the hard-core model and in general this is not known to be possible in polynomial time. A crucial innovation we develop to raise $k$ from 3 (Theorem 4) to an arbitrary integer is an efficient resampling if the neighbourhood contains no long path (Theorem 18). These ideas already suffice for Theorem 3 (which corresponds to the case $t<1$ ), but for Theorem 5 we incorporate an adaptation of the backtracking steps used in [2] for Theorem 4 with the stated upper bound on $t$. Thus we handle few copies of $F_{k}$ by 'removing' an edge in each copy and show that the removed edges are unlikely to stall the algorithm. That is, we can successfully backtrack away from colourings that make the removed edges monochromatic and still show the algorithm will terminate. We remark that the upper bound condition on $t$ in Theorems 4 and 5 stems from the demand for a polynomial-time construction. For existence alone, one can not only substantially relax the condition on $t$ but also allow $k$ to increase as a modest function of $\Delta$, as we show in the companion paper [17].
1.2. Optimality. As intimated earlier, two important and related facts support the idea that our method is optimal. The first concerns Karp's longstanding question mentioned earlier. The difficulty in this is now recognised as deriving from shattering (or dynamic replica symmetry breaking, as it is referred to in statistical physics) in the collection of independent sets of a given size [44, 1, 11]. The rough intuition (stated in terms of Karp's question) is that as the desired set size increases from $(1-\varepsilon) \log _{2} n$ to $(1+\varepsilon) \log _{2} n$, the collection of independent sets
of $G_{n, 1 / 2}$, as considered under a suitable and natural metric, abruptly transitions from a wellconnected 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 set size) must ably navigate this shattered space. Thus $\log _{2} n$ is considered an intuitive algorithmic barrier for the independent set problem in $G_{n, 1 / 2}$, and analogously $n / \log _{2} n$ is an algorithmic barrier for colouring $G_{n, 1 / 2}$. A similar intuition should hold for binomial random graphs throughout the range of choices for the edge probability $p=p(n)$ satisfying $n p=\Omega(1)$ and $p=o(1)$ and analogously also for random $\Delta$-regular graphs with $\Delta$ fixed, with thresholds at around $\frac{1}{p} \log (n p)($ or $n p / \log (n p))$ and $\frac{n}{\Delta} \log \Delta($ or $\Delta / \log \Delta)$, respectively. An affirmative answer to Karp's question, or its analogue for random regular graphs, would be considered an unexpected and sensational achievement. Theorem 5 (just as does Theorem 1 or 4) precisely matches this algorithmic barrier. In particular, by Proposition 2 the random $\Delta$-regular graph provides examples of triangle-free, and so $F_{k}$-free, graphs $G$ such that $\chi(G) \geq(1 / 2-o(1)) \Delta / \log \Delta$ as $\Delta \rightarrow \infty$. By comparison, Theorem 5 with a choice of $\varepsilon=\varepsilon(\Delta)=o(1)$ (and so $t=\Delta^{o(1)}$ ) as $\Delta \rightarrow \infty$ (covering a much more general class of graphs) efficiently certifies $\chi(G) \leq(1+o(1)) \Delta / \log \Delta$.

The second and more concrete fact is that our framework (see Section 2) incorporates the quantitative probabilistic property of local occupancy. We find asymptotically tight parameters for local occupancy in graphs of maximum degree $\Delta$ where each vertex is contained in $\Delta^{o(1)}$ copies of $F_{k}$ (for any fixed $k \geq 3$ ) as $\Delta \rightarrow \infty$, see Appendix A. Then to improve the leading order of the upper bound in Theorem 5 for $t=\Delta^{o(1)}$ must require novel techniques that use more sophisticated knowledge of the underlying structure of the graph. The methods used to prove Theorems 1, 4 and 5 all hit the same obstruction, and our reduction of the problem to local occupancy suggests that surpassing these bounds will require a more global approach. Showing even only the existence of a colouring (so without requiring a polynomial-time construction) that betters these bounds by some constant factor, e.g. via some global data that bypasses the local occupancy bottleneck, would be a breakthrough in classical Ramsey theory and graph colouring.

An alluring feature of our work is the suggestion that local occupancy and the algorithmic barrier might be two sides of the same coin. Could it be the case that for locally sparse $n$-vertex graphs $G$, given $q$ such that there are enough independent sets of size at least $n / q$ in $G$ so that $q$-colouring can be performed efficiently, the collection of independent sets in the graph will be sufficiently rich and 'well connected' to permit local occupancy with parameters that enable our method to show $\chi(G) \leq q$ ?
1.3. Organisation. In Section 2 we introduce the main concepts in our framework and state the key results that establish its efficacy. In Section 2.1 we give a general algorithm for graph colouring and analyse it with our framework. In Section 3 we verify the probabilistic information necessary to apply our framework in the case of graphs with few fans, completing the proof of Theorem 5, and in Appendix A we discuss barriers to improving our results. We defer several technical proofs to Appendices $B$ and $C$.

## 2. The framework

We introduce some extra notation for two concepts central to our framework. First we do so for an important strengthened form of colouring, through which we prove all of our results. 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 (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. By taking $L(u)=\{1, \ldots, q\}$ for each $u$ we see that $\chi(G) \leq \chi_{\ell}(G)$ always. Note that Johansson's and Molloy's bounds mentioned above were also shown in terms of $\chi_{\ell}$.

Second 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 $\lambda$, 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 $\mathbb{E}|\mathbf{I}| /|V(G)|$, 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 with the verification of a condition in terms of the hard-core model, which we call local occupancy. This systematic approach began with a number of previous works [15, 16, [35, 13, 14] on occupancy fraction and fractional colouring. Here we focus on our framework's consequences for efficient (list) colouring, and refer the reader to the companion paper [17] for more general structural implications.

Definition 6. We say that the hard-core model on a graph $G$ at fugacity $\lambda$ has strong local $(\beta, \gamma)$-occupancy if, for every vertex $u \in V(G)$ and subgraph $F \subset G[N(u)]$ we have

$$
\beta \frac{\lambda}{1+\lambda} \frac{1}{Z_{F}(\lambda)}+\gamma \frac{\lambda Z_{F}^{\prime}(\lambda)}{Z_{F}(\lambda)} \geq 1
$$

One easily checks that this condition implies $\beta \operatorname{Pr}(u \in \mathbf{I})+\gamma \mathbb{E}|\mathbf{I} \cap N(u)| \geq 1$ for any $u$, motivating the label. The reason for including the adjective 'strong' relates to some technical subtleties we discuss later as well as in [17. In an abbreviated form, our framework is as follows.

Theorem 7. Let $G$ be a graph of maximum degree $\Delta \geq 2^{6}$ such that the hard-core model on $G$ at fugacity $\lambda$ has strong local $(\beta, \gamma)$-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

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

In this form, the framework essentially reduces the task of bounding the chromatic number from above to minimising $\beta+\gamma \Delta / r$ subject to local $(\beta, \gamma)$-occupancy, an optimisation which can be routinely performed to yield several other applications, cf. [17]. We have given a conceptually elegant proof of Theorem 7 (in the style of Bernshteyn [7]) in [17]. For polynomial-time constructions, here we need a more involved proof that requires some additional assumptions.

For an organic assimilation of the hard-core model in our arguments, it will be helpful to represent list colourings through an auxiliary cover graph as in the work of Dvořák and Postle [18], on a stronger variant of list colouring called correspondence colouring. (This is a key insight.)
Definition 8. Given a graph $G$, a cover of $G$ is a pair $\mathscr{H}=(L, H)$, consisting of a graph $H$ and a mapping $L: V(G) \rightarrow 2^{V(H)}$, satisfying the following requirements:
(i) the sets $\{L(u): u \in V(G)\}$ form a partition of $V(H)$;
(ii) for every $u \in V(G)$, the graph $H[L(u)]$ is complete;
(iii) if $E_{H}(L(u), L(v)) \neq \emptyset$, then either $u=v$ or $u v \in E(G)$;
(iv) if $u v \in E(G)$, 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 $|V(G)|$.
Although covers as defined here capture a more general notion, most of our results here will remain restricted to list colouring. (We discuss the subtleties at the end.) 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 $u v \in E(G)$ 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 natural numbers, and we will refer to vertices of $H$ as colours. We shorten the phrase 'cover that arises from a list assignment' to 'list-cover'.

We now state the extra assumptions needed for our framework to yield an efficient algorithm.
Theorem 9. Suppose that the conditions of Theorem 7 hold, and let $n:=|V(G)|$. Suppose also that there is a class of graphs $\mathcal{C}$ and an integer $t$ such that the following hold.
(i) For each $u \in V(G)$, each induced subgraph $F \subset G[N(u)]$, and any list-cover $\mathscr{H}^{\prime}=\left(L^{\prime}, H^{\prime}\right)$ of $F$ with at most $\Delta$ colours in each list, we have a procedure $\operatorname{Remove}\left(F, H^{\prime}\right)$, 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 $H^{\prime}$ by removing any edge between $L^{\prime}(v)$ and $L^{\prime}(w)$ for $v w \in R$ satisfies $\hat{H} \in \mathcal{C}$.
(ii) For each $\hat{H} \in \mathcal{C}$ we have a procedure $\operatorname{Sample}(\hat{H}, \lambda)$ for sampling from the hard-core model on $\hat{H}$ at fugacity $\lambda$ in time $T_{s}$.
(iii) 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 that constructs, for any $c \in(0,1)$, an $\tilde{L}$-colouring of $G$ as guaranteed by Theorem 7 in time

$$
O\left(\left(T_{r}+T_{s}\right) \Delta \log (\max (\ell / \lambda, \ell)) n+n^{1+c}\right)
$$

with probability at least $1-2 / n^{c}$.
In general the graphs $F$ can have $\Delta$ vertices so the maximum of $T_{s}$ and $T_{r}$ could be exponential in $\Delta$. When $G$ has few copies of $F_{k}$ we show that these conditions hold with $\mathcal{C}$ the class of covers of $F_{k}$-free graphs with at most $\Delta$ colours in each list, and with $T_{r}$ and $T_{s}$ polynomial in $\Delta$, leading to the polynomial running time bound announced in Theorem 5 .
2.1. An algorithm for graph colouring. First here is an overview of the two-phase method. We define (precisely later) a flaw for $u$ to capture the problem of having too few available colours or too much competition for available colours. The first and foremost phase of the proof is that subject to the conditions of Theorem 9, a flawless partial colouring of $G$ can be found efficiently.
Lemma 10. Suppose that the conditions of Theorem 9 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 $O\left(\left(T_{r}+T_{s}\right) \Delta \log (\max (\ell / \lambda, \ell)) \cdot n\right)$ 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 [34, 2].
Lemma 11. Suppose that the conditions of Theorem 9 hold, $\mathscr{H}=(L, H)$ is a q-fold list-cover of $G$, and let $\sigma$ 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 10 followed by Lemma 11 directly implies Theorem 9 .
The algorithm for Lemma 11 selects a colour for the remaining vertices uniformly at random, resampling if there are any conflicts. For completeness we give a sketch proof in Appendix B.4.

The algorithm for Lemma 10 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 $\sigma$ is flawed we execute a resampling action to address the least flaw present in $\sigma$ and move to a new partial colouring. We give the proof of Lemma 10 in the rest of this section with some details deferred to Appendix B.
2.1.1. Notation. Given the setup of Theorem 7 , 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 $\operatorname{deg}_{\mathscr{H}}^{*}(x)$ for the degree in $H^{*}$ of a colour $x$, which is the number of colours on lists of neighbours of $u$ that conflict with $x$.

Writing $\mathfrak{B}$ for a special blank colour, and borrowing from [2], a partial $\mathscr{H}$-colouring $\sigma$ of $G$ is a function from $V(G)$ to $\{\mathfrak{B}\} \cup V(H) \cup E(G)$ such that the following hold.
(i) For all $u \in V(G)$, either $\sigma(u)=\mathfrak{B}, \sigma(u) \in L(u)$, or $\sigma(u)=e \in E(G)$ with $u \in e$.
(ii) Restricting the image of $\sigma$ to $V(H)$ gives an independent set; $\sigma(V(G)) \cap V(H) \in \mathcal{I}(H)$.

We write $\Omega$ for the set of such partial $\mathscr{H}$-colourings $\sigma$, and omit the prefix $\mathscr{H}$ when it is clear.
Given $\sigma \in \Omega$ we have blank vertices bla $(\sigma):=\{u \in V(G): \sigma(u)=\mathfrak{B}\}$ for which $\sigma(u)=\mathfrak{B}$, coloured vertices $\operatorname{col}(\sigma):=\{u \in V(G): \sigma(u) \in L(u)\}$ for which $\sigma(u) \in L(u)$, and uncoloured vertices $\operatorname{unc}(\sigma):=\{u \in V(G): \sigma(u) \in E(G)\}$ for which $\sigma(u)$ is an edge of $G$ containing $u$. We also write $\operatorname{ind}(\sigma):=\sigma(V(G)) \cap V(H)=\sigma(\operatorname{col}(\sigma))$ for the independent set in $H$ signified by $\sigma$.

We also require some notation for the cover that remains on the blank vertices. We write $G_{\sigma}:=G[\operatorname{bla}(\sigma)]$ for the subgraph of $G$ induced by bla $(\sigma)$, and then write $\mathscr{H}_{\sigma}=\left(L_{\sigma}, H_{\sigma}\right)$ for the cover of $G_{\sigma}$ obtained by setting $L_{\sigma}(u):=L(u) \backslash N_{H}[\operatorname{ind}(\sigma)]$ for $u \in \operatorname{bla}(\sigma)$ and $H_{\sigma}:=$ $H\left[L_{\sigma}(\operatorname{bla}(\sigma))\right]$. Note that $\mathscr{H}_{\sigma}$ is a list-cover of $G_{\sigma}$. When there are no uncoloured vertices, these definitions agree with those of [7, 17] if $\sigma$ is identified with $\operatorname{ind}(\sigma)$.

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 $\operatorname{col}(\sigma)$. This means that if $I \in \mathcal{I}\left(H_{\sigma}\right)$, then $\operatorname{ind}(\sigma) \cup I \in$ $\mathcal{I}(H)$. In particular, if $\sigma$ has no uncoloured vertices, and if we can find $I \in \mathcal{I}\left(H_{\sigma}\right)$ of size $\left|V\left(G_{\sigma}\right)\right|$ then $\operatorname{ind}(\sigma) \cup I$ is an $\mathscr{H}$-colouring of $G$. This is exactly how the two-phase method proceeds.

### 2.1.2. The flaws. We define a flaw for each vertex $u$, writing

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

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 $U_{u}^{e}:=\{\sigma \in \Omega: \sigma(u)=e\}$. We write $F_{B}:=\left\{B_{u}: u \in V(G)\right\}$ and $F_{U}:=\left\{U_{u}^{e}: u \in e \in E(G)\right\}$ so that $F:=F_{B} \cup F_{U}$ is the set of all flaws. Note that $\left|F_{B}\right|=|V(G)|=n$, and $\left|F_{U}\right|=2|E(G)| \leq \Delta n$.

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 $U_{v}^{e}$ suffices. To be explicit, consider an arbitrary ordering of the vertices, and the induced lexicographic ordering on edges where $u v$ 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 $U_{u}^{e}$ ordered according to $u$ and then $e$.
2.1.3. The actions. To address the flaw $B_{u}$ at state $\sigma$ we execute the action $\operatorname{AddressB}(u, \sigma)$ defined in Algorithm 1 in terms of the procedures Remove and Sample guaranteed by the assumptions of Theorem 9 .

```
Algorithm 1
    procedure \(\operatorname{AddressB}(u, \sigma)\)
        let \(\sigma^{\prime}\) be obtained from \(\sigma\) by setting \(\sigma^{\prime}(v):=\mathfrak{B}\) for all \(v \in N_{G}(u) \backslash \operatorname{unc}(\sigma)\), and \(\sigma^{\prime}(v):=\sigma(v)\) otherwise
        let \(F:=G_{\sigma^{\prime}}\left[N_{G}(u)\right], H^{\prime}:=H_{\sigma^{\prime}}\left[L\left(N_{G}(u)\right)\right]\), and let \(\hat{H}:=\operatorname{Remove}\left(F, H^{\prime}\right)\)
        let \(\mathbf{J}_{0}:=\operatorname{SAmple}(\hat{H}, \lambda)\), and let \(\tau_{0}(v):=\sigma^{\prime}(v)\) 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 \(\mathbf{J}_{i} \notin \mathcal{I}\left(H^{\prime}\right)\) do
            let \(v w\) be the lowest-indexed edge in \(G[N(u)]\) for which \(\mathbf{J}_{i}\) spans an edge of \(H^{\prime}\) going from \(L(v)\) to \(L(w)\),
            and suppose that \(v<w\)
            let \(\tau_{i+1}\) be obtained from \(\tau_{i}\) by setting \(\tau_{i+1}(v):=v w\) (uncolouring \(v\) ) and setting \(\tau_{i+1}\) to agree with \(\tau_{i}\)
            elsewhere
            let \(\mathbf{J}_{i+1}:=\mathbf{J}_{i} \backslash L(v)\)
            increment \(i\)
        let \(\tau:=\tau_{i}\) and let \(\mathbf{J}:=\mathbf{J}_{i}\)
        return \(\tau\)
```

Algorithm 1 has three distinct parts. The first is some setup in which we define a partial colouring $\sigma^{\prime}$ by reassigning coloured vertices in $N_{G}(u)$ to $\mathfrak{B}$, which gives us an induced subgraph $F$ of $G\left[N_{G}(u)\right]$ and a list-cover $H^{\prime}:=H_{\sigma^{\prime}}\left[L\left(N_{G}(u)\right)\right]$ of $F$. We then use the procedure Remove to remove some edges from $F$ and any corresponding edges in $H^{\prime}$, which results in a cover $\hat{H}$ of $F$ in
the class $\mathcal{C}$. The second is sampling an independent set $\mathbf{J}_{0}$ in $\hat{H}$ and a partial colouring-like object $\tau_{0}$ 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 $H_{\sigma^{\prime}}$ so $\tau_{0}$ 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 $\tau_{i}$ as we go. When the loop exits $\mathbf{J}_{i}$ is independent in $H_{\sigma^{\prime}}$, and so the final $\tau_{i}$ is a valid partial colouring.

To address the flaw $U_{u}^{e}$ at state $\sigma$ we simply resample $\sigma(u)$ from the hard-core model as follows. Let $\sigma^{\prime}$ be obtained from $\sigma$ by letting $\sigma^{\prime}(v):=\sigma(v)$ for $v \neq u$, and letting $\sigma^{\prime}(u):=\mathfrak{B}$ with probability $1 /\left(1+\left|L_{\sigma}(u)\right| \lambda\right)$, and otherwise setting $\sigma^{\prime}(u)$ to be a uniform colour from $L_{\sigma}(u)$.
2.1.4. 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 $\tau$ when addressing the flaw $f$ at state $\sigma$. Let $\tilde{\mu}$ be the probability measure on $\Omega$ given by

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

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 $\lambda$ when sampling $\mathbf{J}_{0}$ in the procedure AddressB $(u, \sigma)$.

Given subsets of flaws $S, S^{\prime} \subset F$, we say that $S^{\prime}$ covers $S$ if ${ }^{1}$

$$
S^{\prime} \cap F_{U}=S \cap F_{U} \text { and } S^{\prime} \cap F_{B} \supset S \cap F_{B} .
$$

For any flaw $f$ and subset $S \subset F$ of flaws, we defin ${ }^{2}$

$$
\begin{aligned}
\operatorname{In}_{f}^{S}(\tau):=\{\sigma \in f: & \rho_{f}(\sigma, \tau)>0 \text { and the set of flaws } \\
& \text { introduced by the transition } \sigma \rightarrow \tau \text { covers } S\},
\end{aligned}
$$

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

$$
\begin{equation*}
c^{S}(f):=\max _{\tau \in \Omega}\left\{\sum_{\sigma \in \operatorname{In}_{f}^{S}(\tau)} \frac{\tilde{\mu}(\sigma)}{\tilde{\mu}(\tau)} \rho_{f}(\sigma, \tau)\right\}, \tag{1}
\end{equation*}
$$

which represents a kind of compatibility between the measure $\tilde{\mu}$ and the transitions $\rho_{f}$ 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 Algorithm 1 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 12 (Achlioptas, Iliopoulos, and Sinclair [2]). If there exist positive numbers $\left(\psi_{f}\right)_{f \in F}$ 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 $\left(T_{0}+s\right) / \delta$ steps with probability at least $1-2^{-s}$, where $\delta:=1-\max _{f \in F}\left\{\zeta_{f}\right\}$ and

$$
T_{0}:=\log _{2} \tilde{Z}_{H}(\lambda)+\sum_{u \in V(G)} \log _{2}\left(1+\psi_{B_{u}}\right)+\log _{2}\left(\max _{S \subset F} \frac{1}{\prod_{f \in S} \psi_{f}}\right)
$$

[^1]2.1.5. Bounding charges. By design the compatibility between $\tilde{\mu}$ and our algorithm is good enough for the following result to control the charges $c^{S}\left(B_{u}\right)$.

Lemma 13. Suppose that the conditions of Theorem 9 hold, and let $\mathscr{H}=(L, H)$ be a q-fold list-cover of $G$. Then for any partial $\mathscr{H}$-colouring $\sigma$ of $G$ and any $u \in V(G)$ such that $u \notin \operatorname{col}(\sigma)$, the following holds. If $\tau$ is the random partial colouring of $G$ that results from the procedure $\operatorname{AdDressB}\left(B_{u}, \sigma\right)$, then $\operatorname{Pr}\left(\tau \in B_{u}\right) \leq 1 /\left(4 \Delta^{3}\right)$.

A similar version of this is key to the lopsided local lemma formulation of our methods [17]. To obtain an algorithm we essentially take advantage of some subtle extra strength from the fact that the lemma holds for any $\sigma$ rather than only when $\sigma$ has no uncoloured vertices and ind $(\sigma)$ is sampled from the hard-core model on $H$. The proof of Lemma 13 is given in Appendix B.1.

The following results comprise the bounds on charges we need to apply Theorem 12, and the proofs are in Appendices B.2 and B.3. Let $S\left(U_{u}^{e}\right):=\left\{B_{v}: v \in N^{2}[u]\right\}$ and $S\left(B_{u}\right):=\left\{B_{v}: v \in\right.$ $\left.N^{3}[u]\right\} \cup\left\{U_{v}^{e}: v \in e \in E(\bar{G}[N(u)])\right\}$. We will see that these are the only flaws that addressing $U_{u}^{e}$ and $B_{u}$ can introduce, respectively.

Lemma 14. For every vertex $u \in V(G)$, the following hold.
(i) If $S \not \subset S\left(B_{u}\right)$ then $c^{S}\left(B_{u}\right)=0$.
(ii) If $S$ contains more than $t$ flaws of the form $U_{v}^{v w}$ with $v, w \in N(u)$ then $c^{S}\left(B_{u}\right)=0$.
(iii) $\max _{S \subset F}\left\{c^{S}\left(B_{u}\right)\right\} \leq 1 /\left(4 \Delta^{3}\right)$.

Lemma 15. For every $u$ and $e$ such that $u \in e \in E(G)$, the following hold.
(i) If $S \not \subset S\left(U_{u}^{e}\right)$ then $c^{S}\left(U_{u}^{e}\right)=0$.
(ii) $\max _{S \subset F}\left\{c^{S}\left(U_{u}^{e}\right)\right\} \leq \lambda /(1+\ell \lambda)$.
2.1.6. Finishing the proof. We can now choose parameters $\psi_{f}$ for $f \in F$ such that the desired result follows from Theorem 12, For this we take a positive real $\psi$ to be determined later and set $\psi_{f}:=\psi /\left(4 \Delta^{3}\right)$ for all $f \in F_{B}$ and $\psi_{f}:=\psi \lambda /(1+\ell \lambda)$ for all $f \in F_{U}$.

By Lemma 14 we know that $c^{S}\left(B_{u}\right)=0$ 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 $U_{v}^{v w}$, then $v w \in R$. Since $\left|N^{3}[u]\right| \leq \Delta^{3}$, we deduce from Lemma 14 that for each $u \in V(G)$, setting

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

we have

$$
\begin{aligned}
\frac{1}{\psi_{B_{u}}} \sum_{S \subset F} c^{S}\left(B_{u}\right) \prod_{g \in S} \psi_{g} & \leq \frac{1}{\psi} \prod_{v \in N^{3}[u]}\left(1+\frac{\psi}{4 \Delta^{3}}\right) \prod_{v w \in R}\left(1+\frac{\psi \lambda}{1+\ell \lambda}\right) \\
& \leq \frac{1}{\psi} \exp \left(\frac{\psi}{4}\right) \exp \left(\frac{t \psi \lambda}{1+\ell \lambda}\right)=\frac{e}{4}\left(1+\frac{4 t \lambda}{1+\ell \lambda}\right) \leq \frac{3}{4}
\end{aligned}
$$

because by assumption $t \leq \ell / 40$, and hence $4 t \lambda \leq(1+\ell \lambda)(3 / e-1)$.
Similarly, by Lemma 15 and the facts that $\Delta \geq 2$ and $\left|N^{2}[u]\right| \leq 1+\Delta^{2}$, we have for each pair ( $u, e$ ) with $u \in e \in E(G)$,

$$
\begin{aligned}
\frac{1}{\psi_{U_{u}^{e}}} \sum_{S \subset F} c^{S}\left(U_{u}^{e}\right) \prod_{g \in S} \psi_{g} & \leq \frac{1}{\psi} \prod_{v \in N^{2}[u]}\left(1+\frac{\psi}{4 \Delta^{3}}\right) \leq \frac{1}{\psi} \exp \left(\psi \frac{1+\Delta^{2}}{4 \Delta^{3}}\right) \\
& =\frac{1}{4}\left(1+\frac{4 t \lambda}{1+\ell \lambda}\right) \exp \left(\frac{1+\Delta^{2}}{\Delta^{3}}\left(1-\frac{4 t \lambda}{1+\ell \lambda+4 t \lambda}\right)\right) \\
& \leq \frac{e^{\left(1+\Delta^{2}\right) / \Delta^{3}}}{4}\left(1+\frac{4 t \lambda}{1+\ell \lambda}\right) \leq \frac{3}{4}
\end{aligned}
$$

Hence we can apply Theorem 12 with parameters $\psi_{B_{u}}$ and $\psi_{U_{u}^{e}}$ such that

$$
\begin{aligned}
& \frac{e}{3 \Delta^{3}} \leq \psi_{B_{u}}=\frac{1}{\Delta^{3}} \cdot \frac{1+\ell \lambda}{1+\ell \lambda+4 t \lambda} \leq \frac{1}{\Delta^{3}}, \\
& \frac{4 e \lambda}{3(1+\ell \lambda)} \leq \psi_{U_{u}^{e}}=\frac{4 \lambda}{1+\ell \lambda+4 t \lambda} \leq \frac{4 \lambda}{1+\ell \lambda},
\end{aligned}
$$

giving $\delta=1 / 4$ and

$$
T_{0} \leq \log _{2} \tilde{Z}_{H}(\lambda)+n \log _{2}\left(1+\frac{1}{\Delta^{3}}\right)+n \log _{2}\left(\frac{3 \Delta^{3}}{e}\right)+n \Delta \log _{2}\left(\frac{3(1+\ell \lambda)}{4 e \lambda}\right)
$$

We have

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

because for each $u \in V(G)$ we can have either $\tau(u)=\mathfrak{B}$, which does not contribute to the exponent of $\lambda$, or $\tau(u) \in L(u) \cup\{e \in E(G): u \in e\}$, which contributes 1 . There are at most $2 \Delta$ 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 $\lambda \geq 1$ while $T_{0}=O(n \Delta \log (\ell / \lambda))$ if $\lambda<1$. Since $\delta=1 / 4$ and $T_{0} \geq n$, setting $s:=n$ yields that the probability that the algorithm finds a flawless partial colouring in at most $2 T_{0}$ steps is at least $1-2^{-n}$. Each step takes time $O\left(T_{r}+T_{s}\right)$ because if we are addressing a flaw $f \in F_{B}$ then we execute action $\operatorname{AddressB}(u, \sigma)$ 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 $\ell \leq y \leq \Delta$. This completes the proof of Lemma 10, showing that the first, main phase of the algorithm works as desired.

## 3. Application to graphs with few fans

To prove Theorem 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. We start with the strong local occupancy, which relies on 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 $x \mapsto x e^{x}$ defined on $[-1, \infty)$. We use the basic property that as $x \rightarrow \infty$ we have $W(x)=(1-o(1)) \log x$, see e.g. [12]. The following result is proved in Appendix C] (see also [14, 17]).
Lemma 16. 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 $\lambda>0$.
(i) For any $d>0$, there exist $\beta, \gamma>0$ such that the hard-core model on $G$ at fugacity $\lambda$ has strong local $(\beta, \gamma)$-occupancy and

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

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

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

To apply the above result to $G$ as in Theorem 5 we prove a suitable mad bound.
Lemma 17. Let $u \in V(G)$ be 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 $k-3+\sqrt{2 t}$.

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

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

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, which means $F$ has at most $y(k-3) / 2+t$ edges. The first expression in the assertion follows, and we consider the subcases $y \leq \sqrt{2 t}$ and $y>\sqrt{2 t}$ to crudely bound from above the minimum.

For the rest of this section let $G$ be as in Theorem 5. Let $\mathcal{F}$ be the class of $P_{k-1}$-free graphs on at most $\Delta$ vertices and let $\mathcal{C}$ be the class of list-covers of graphs in $\mathcal{F}$ with at most $\Delta$ colours in each list. For any vertex $u \in V(G)$, subgraph $F$ of $G[N(u)]$, and list-cover $\mathscr{H}=\left(L^{\prime}, H^{\prime}\right)$ of $F$ with at most $\Delta$ colours in each list we can identify all copies of $P_{k-1}$ in $F$ in time $O\left(\Delta^{k}\right)$ by enumerating all ordered sets of $k-1$ vertices in $F$. To implement Remove we simply choose an arbitrary edge $v w$ of each $P_{k-1}$ found in this way and remove all edges from $H^{\prime}$ between $L(v)$ and $L(w)$. For each of the at most $t \leq \Delta^{2 \varepsilon}$ copies of $P_{k-1}$ found, this removal takes time at most $\Delta^{2}$ so Remove as in Theorem 9 can be done in time $\Delta^{O(k)}$.

The following result which we prove in Appendix C implies that we can implement Sample on $\mathcal{C}$ in time $(k \Delta)^{O\left(k^{3}\right)}$.
Theorem 18. 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 $y^{3 k^{2}}(1+q)^{k^{3} / 2}(c k)^{k^{3}}$.

We can now finish the proof of Theorem 5
Proof of Theorem 5. Fix an arbitrary vertex $u \in V(G)$ and an arbitrary subgraph $F \subset G[N(u)]$. By Lemma 17 we have $\operatorname{mad}(F) \leq a=: k-3+\sqrt{2 t}$. It is convenient to exclude the case $a=0$ in the argument, which is one place the assumption $t \geq 1 / 2$ comes in useful, giving $a \geq 1$.

We want to apply Theorem 9 to conclude the proof. Above we defined a class $\mathcal{C}$, Remove, and Sample such that Items (i) and (ii) hold with $T_{r}=\Delta^{O(k)}$ and $T_{s}=(k \Delta)^{O\left(k^{3}\right)}$. We now define the parameters so that the remaining requirements of Theorem 9 are satisfied: those are $t \leq \ell / 40$ (Item (iii)) and the hypothesis of Theorem 7 . In particular, we need to show that there is strong local $(\beta, \gamma)$-occupancy for the hard-core model on $G$ at fugacity $\lambda$ for some positive reals $\beta, \gamma$ and $\lambda$. To this end we use Lemma 16. which will also provide the requirement on $Z_{F}(\lambda)$. Indeed, with $a=k-3+\sqrt{2 t}$, given any $\lambda>0$ and $\ell>7 \log \Delta$, and with $r$ and $q$ as in Theorem 7 . Item (i) of Lemma 16 gives us $\beta$ and $\gamma$ such that the hard-core model on $G$ at fugacity $\lambda$ has strong local $(\beta, \gamma)$-occupancy with

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

We set

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

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

Second, we must show that

$$
\begin{equation*}
\frac{\ell}{8} \geq \frac{\log \left(8 \Delta^{4}\right)}{\log (1+\lambda)\left(1-\frac{a}{2} \log (1+\lambda)\right)}=\frac{a \log (\Delta / \sqrt{t}) \log \left(8 \Delta^{4}\right)}{1-\frac{1}{2 \log (\Delta / \sqrt{ } t)}}, \tag{2}
\end{equation*}
$$

so that Item (ii) of Lemma 16 ensures that $Z_{G}(\lambda) \geq 8 \Delta^{4}$ for any subgraph $F$ of any $G[N(u)]$ on at least $\ell / 8$ vertices. The right-hand side of (22) is $O(a \log (\Delta / \sqrt{t}) \log \Delta)$ as $\Delta$ and hence $\Delta / \sqrt{t}$
tend to infinity. This is less than $\ell$ for large enough $\Delta_{0}$ in terms of $\varepsilon$ because the bound $t \leq \Delta^{\frac{2 \varepsilon}{1+2 \varepsilon}} /(\log \Delta)^{2}$ gives

$$
\log (\Delta / \sqrt{t}) \geq \frac{1+\varepsilon}{1+2 \varepsilon} \log \Delta-\log \log \Delta=\Omega(\log \Delta)
$$

so that $a \log (\Delta / \sqrt{t}) \log \Delta=O\left(a(\log (\Delta / \sqrt{t}))^{2}\right)=o(\ell)$.
As $\Delta \rightarrow \infty$ we now have $\lambda=o(1), a \lambda=o(1), 7 \log \Delta=o(\ell), a=O(\sqrt{t})$, and

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

To obtain Theorem 5, note that for large enough $\Delta_{0}$ and $\Delta \geq \Delta_{0}$,

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

so the result follows from the application of Theorem 9.

## 4. Concluding remarks

In this section we compare the algorithmic framework given here with the more combinatorial treatment in the companion paper [17]. 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 $\operatorname{deg}(v)$ instead of $\Delta$. See [13, 17] for details, and in particular an interesting minimum list size phenomenon that arises. Secondly, when applying Theorem 9 with a trivial implementation of Remove that does nothing, our proof gives two extra properties. When Remove is trivial Theorem 9 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 9 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 to correspondence colouring [7], 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 [17]), 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 3, or equivalently the case $t<1$ in Theorem 5. 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 [17, 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 55, 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 [17. Consider the problem of colouring a graph $G$ of maximum degree $\Delta$ in which each vertex is contained in at most $t F_{k}$ 's with $(1+\varepsilon) \Delta / \log (\Delta / \sqrt{t})$ colours, or lists of this size. Let $\Delta_{0}$ be large enough in terms of $\varepsilon$ and $\Delta \geq \Delta_{0}$, and let $c=c(\varepsilon)$ 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 $t \leq \Delta^{2} /(\log \Delta)^{2 / \varepsilon}$, and our algorithmic methods work for $t \leq \Delta^{\frac{2 \varepsilon}{1+2 \varepsilon}} /(\log \Delta)^{2}$. 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.

## 5. Acknowledgements

We thank Alistair Sinclair, Fotis Iliopoulos, and Charlie Carlson for insightful discussions.

## Appendix A. Tightness

In this section we state a proposition indicating that for a large range of $\lambda$, the local occupancy of Lemma 16 is asymptotically best possible for $F_{k}$-free graphs. First we note an extra (but foundational) component of our framework essentially originating in [16], cf. [17].
Theorem 19. Let $G$ be a graph of maximum degree $\Delta$ such that the hard-core model on $G$ at fugacity $\lambda$ has (strong) local ( $\beta, \gamma$ )-occupancy for some $\lambda, \beta, \gamma>0$. Then the occupancy fraction of $G$ at fugacity $\lambda$ is at least $1 /(\beta+\gamma \Delta)$.
Supposing that $k$ is a fixed integer greater than 2 and $\lambda=o(1)$ such that $\lambda \sqrt{t}=o(1)$ as $\Delta \rightarrow \infty$ we know from Lemma 16 and Theorem 19 that for any graph $G$ with maximum degree $\Delta$ in which each vertex is the centre of at most $t$ copies of $F_{k}$, the occupancy fraction of $G$ at fugacity $\lambda$ is at least

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

But the occupancy fraction is monotone increasing in $\lambda$ (see [16]) so this lower bounds holds for all larger values of $\lambda$ too. This is asymptotically tight at least for $\lambda \leq \Delta^{1+o(1)}$, but in fact no improvement to this lower bound is known for larger values of $\lambda$ even for the case of triangle-free graphs, that is when $k=3$. 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 $\lambda$ larger than $1 / \log \Delta$ would immediately lead to an improvement to Shearer's result that every $n$-vertex triangle-free graph of maximum degree $\Delta$ contains an independent set of size at least $(1-o(1)) n \log \Delta / \Delta$.
Proposition 20 ([16]). Given $\varepsilon>0$ there is $\Delta_{0}$ such that for all fixed $\Delta>\Delta_{0}$ and $\lambda \leq \Delta^{1+o(1)}$, there is a $\Delta$-regular $F_{k}$-free graph $G$ with occupancy fraction at most $(1+\varepsilon) W(\Delta \lambda) / \Delta$.

## Appendix B. Proofs for the main framework

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

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

Lemma 21 (Panconesi and Srinivasan [38]). Given a probability space, let $\mathbf{X}_{1}, \ldots, \mathbf{X}_{n}$ 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 $i \in\{1, \ldots, n\}$. If the variables $\mathbf{Y}_{1}, \ldots, \mathbf{Y}_{n}$ are negatively correlated, then for any $\eta \in(0,1)$,

$$
\operatorname{Pr}(\mathbf{X} \leq(1-\eta) \mathbb{E} \mathbf{X}) \leq e^{-\eta^{2} \mathbb{E} \mathbf{X} / 2}
$$

B.1. Proof of Lemma 13, Recall that for Lemma 13 we assume the hypotheses of Theorem 9 (and hence also Theorem 7 ), and we consider the transition from an arbitrary partial $\mathscr{H}$-colouring $\sigma$ of $G$ induced by the action $\operatorname{ADDressB}\left(B_{u}, \sigma\right)$ for some $u \notin \operatorname{col}(\sigma)$. In what follows, we write $\tau$ for the random state arising from this transition, and $\operatorname{Pr}$ and $\mathbb{E}$ represent probabilities and expectations over the randomness in the action $\operatorname{ADDRESSB}\left(B_{u}, \sigma\right)$, respectively.

We first argue that $\left|L_{\tau}(u)\right|$ 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 $H_{\tau}^{*}$ is small. For this we require some additional notation. We reuse the notation of the procedure ADDRESSB, writing $\sigma^{\prime}$ for the partial $\mathscr{H}$-colouring of $G$ obtained from $\sigma$ by setting $\sigma^{\prime}(v):=\mathfrak{B}$ for all vertices $v \in$ $N_{G}(u)$, and $\sigma^{\prime}(v):=\sigma(v)$ otherwise. We also have the definitions $H^{\prime}:=H_{\sigma^{\prime}}\left[L\left(N_{G}(u)\right)\right]$ and $\hat{H}:=\operatorname{Remove}\left(G_{\sigma^{\prime}}\left[N_{G}(u)\right], H^{\prime}\right)$, and $\mathbf{J}_{0}$ is an independent set in $\hat{H}$ sampled from the hard-core model at fugacity $\lambda$. If we write $L^{\prime}(v):=L(v) \cap V\left(H^{\prime}\right)$ then $\mathscr{H}$ and $\mathscr{H}^{\prime}:=\left(L^{\prime}, H^{\prime}\right)$ are
list-covers of $G$ and $G_{\sigma^{\prime}}\left[N_{G}(u)\right]$, respectively. This provides the additional structure that each colour $x \in V(H)$ is considered a copy of some natural number $c$ in the list-assignment, and $x y$ is an edge of $H^{*}$ if and only if $x \in L(u)$ and $y \in L(v)$ are copies of the same natural number and $u v \in E(G)$.

For each $x \in L(u)$, let $\Lambda_{x}$ be the layer of $x$, given by $\Lambda_{x}:=N_{H_{\sigma^{\prime}}^{*}}(x)$. This consists of the colours in $L_{\sigma^{\prime}}\left(N_{G}(u)\right)$ that conflict with $x$, and hence $\Lambda_{x}$ consists of every colour $y \in V\left(H^{\prime}\right)$ that is a copy of a fixed natural number, written $c$. So for distinct $x, y \in L(u)$ the layers $\Lambda_{x}$ and $\Lambda_{y}$ are necessarily disjoint. The fact that $\mathscr{H}^{\prime}$ is a list-cover of $G\left[N_{G}(u)\right]$ means that every edge leaving $\Lambda_{x}$ in $H^{\prime}$ joins two colours belonging to some set $L(v)$ with $v \in N_{G}(u)$, which means there are no edges between one layer $\Lambda_{x}$ and another $\Lambda_{y}$ 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 $H^{\prime}$ and the uncolouring steps $5 \backslash 11$ yield the subset $\mathbf{J} \subset \mathbf{J}_{0}$ which is independent in $H^{\prime}$. For any $x \in L(u)$ we note that $\mathbf{J} \backslash \Lambda_{x}$ depends only on $\mathbf{J}_{0} \backslash \Lambda_{x}$ and not on $\mathbf{J}_{0} \cap \Lambda_{x}$ because of the above property of edges leaving $\Lambda_{x}$. To see this, observe that every uncolouring step is due to some tuple ( $v_{1}, v_{2}, y_{1}, y_{2}$ ) such that $\left\{v_{1}, v_{2}\right\} \subset N_{G}(u), v_{1} v_{2} \in E(G), y_{1} \in L\left(v_{1}\right) \cap \mathbf{J}_{0}$, $y_{2} \in L\left(v_{2}\right) \cap \mathbf{J}_{0}$, and $y_{1} y_{2} \in E\left(H^{\prime}\right)$. But then $y_{1}$ and $y_{2}$ must be a copy of the same natural number $c^{\prime}$ and hence either both are in $\Lambda_{x}$ or neither one is in $\Lambda_{x}$. This means the uncolouring steps due to edges of $H^{\prime}$ inside $\Lambda_{x}$ are independent of the other uncolouring steps, and hence $\mathbf{J} \backslash \Lambda_{x}$ depends only on $\mathbf{J}_{0} \backslash \Lambda_{x}$.

We now show some key properties of how $\mathbf{J}_{0}$ is distributed on the sets $\Lambda_{x}$, 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 $\mathbf{J}_{0} \backslash \Lambda_{x}$ and taking those vertices in $\Lambda_{x}$ that in the graph $\hat{H}$ are not adjacent to any vertex of $\mathbf{J}_{0} \backslash \Lambda_{x}$. That is, $\mathbf{U}_{0}(x):=\Lambda_{x} \backslash N_{\hat{H}}\left(\mathbf{J}_{0} \backslash \Lambda_{x}\right)$. Then write $\mathbf{F}_{0}(x):=\hat{H}\left[\mathbf{U}_{0}(x)\right]$. 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 $\lambda$. It is important to observe that $\hat{H}\left[\Lambda_{x}\right]$ is isomorphic to a subgraph of $G[N(u)]$, as is $\mathbf{F}_{0}(x)$, so the assumptions of Theorem 7 give that $\mathbf{F}_{0}(x)$ has strong local ( $\beta, \gamma$ )-occupancy.

The above definitions deal with the cover $\hat{H}$ in which we sample, but we must also deal with the original cover $H^{\prime}$. To this end, we analogously write $\mathbf{U}(x)$ for the set of vertices obtained by revealing $\mathbf{J} \backslash \Lambda_{x}$ and taking those vertices in $\Lambda_{x}$ that in the graph $H^{\prime}$ are not adjacent to any vertex of $\mathbf{J} \backslash \Lambda_{x}$. That is, $\mathbf{U}(x):=\Lambda_{x} \backslash N_{H^{\prime}}\left(\mathbf{J} \backslash \Lambda_{x}\right)$. We also write $\mathbf{F}(x):=H^{\prime}[\mathbf{U}(x)]$. 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 $\mathbf{J}_{0} \cap \Lambda_{x}=\varnothing$.
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 $\operatorname{deg}_{\mathscr{H}_{\tau}}^{*}(x)=\left|\mathbf{U}_{0}(x)\right|$.
To see Fact 1, observe that $x \in L_{\tau}(u)$ if and only if for every $v \in \operatorname{col}(\tau) \cap N_{G}(u)$ we have $x \tau(v) \notin H$. That is, if $x$ is a copy of the natural number $c$ then $x \in L_{\tau}(u)$ if and only if no neighbour of $u$ is coloured with a copy of $c$ under $\tau$. This clearly holds if $\mathbf{J}_{0} \cap \Lambda_{x}=\varnothing$, and also if $\mathbf{J}_{0} \cap \Lambda_{x} \neq \varnothing$ 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 $\mathbf{J}_{0} \cap \Lambda_{x} \neq \varnothing$ we must also have $\mathbf{J} \cap \Lambda_{x} \neq \varnothing$.

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

$$
\begin{align*}
& y \in \mathbf{U}_{0}(x) \Rightarrow v \in \operatorname{bla}(\tau)  \tag{3}\\
& y \in \mathbf{U}(x) \backslash \mathbf{U}_{0}(x) \Rightarrow v \in \operatorname{unc}(\tau) . \tag{4}
\end{align*}
$$

This crucially exploits the fact that $\mathscr{H}^{\prime}$ is a list-cover. To prove (3), observe that if $y \in \mathbf{U}_{0}(x)$ then no neighbour of $y$ in $\hat{H}$ belongs to $\mathbf{J}_{0} \backslash \Lambda_{x}$, and since $x \in L_{\tau}(u)$ we have $\mathbf{J}_{0} \cap \Lambda_{x}=\varnothing$ by Fact 1. Together these facts mean that $v$ is neither coloured nor uncoloured, and hence belongs to bla $(\tau)$ as required. To prove (4), suppose that $y \in \mathbf{U}(x) \backslash \mathbf{U}_{0}(x)$. Then there must be some $z \in\left(\mathbf{J}_{0} \backslash \Lambda_{x}\right) \backslash\left(\mathbf{J} \backslash \Lambda_{x}\right)$ such that $y z$ is an edge of $H^{\prime}$. As reported earlier, since $\mathscr{H}^{\prime}$ is a list-cover we know that every edge leaving $\Lambda_{x}$ 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 $z \in\left(\mathbf{J}_{0} \backslash \Lambda_{x}\right) \backslash\left(\mathbf{J} \backslash \Lambda_{x}\right)$ is because $z$ had been coloured and was then removed from some $\mathbf{J}_{i}$ in the creation of the next $\mathbf{J}_{i+1}$, and so necessarily $v \in \operatorname{unc}(\tau)$ as required.

We are now ready to show that $\left|L_{\tau}(u)\right|$ is likely to be at least $\ell$.
Lemma 22. Writing

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

we have $\operatorname{Pr}\left(\left|L_{\tau}(u)\right| \leq(1-\eta) m\right) \leq e^{-\eta^{2} m / 2} \leq 1 /\left(8 \Delta^{3}\right)$.
Proof. We first note that $e^{-\eta^{2} m / 2} \leq 1 /\left(8 \Delta^{3}\right)$ holds because the parameter choices in Theorem 7 give

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

Now by Fact 1, we have $x \in L_{\tau}(u)$ if and only if $\mathbf{J}_{0} \cap \Lambda_{x}=\varnothing$. 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 $G\left[N_{G}(u)\right]$. Then $\mathbf{F}_{0}(x)$ has strong local $(\beta, \gamma)$-occupancy and so

$$
\begin{equation*}
\beta \frac{\lambda}{1+\lambda} \operatorname{Pr}\left(x \in L_{\tau}(u)\right)+\gamma \mathbb{E}\left|\mathbf{J}_{0} \cap \Lambda_{x}\right| \geq 1, \tag{5}
\end{equation*}
$$

because elementary calculations with the hard-core model (see e.g. [16]) now give

$$
\operatorname{Pr}\left(\mathbf{J}_{0} \cap \Lambda_{x}=\varnothing\right)=\frac{1}{Z_{\mathbf{F}_{0}(x)}(\lambda)} \quad \text { and } \quad \mathbb{E}\left|\mathbf{J}_{0} \cap \Lambda_{x}\right|=\frac{\lambda Z_{\mathbf{F}_{0}(x)}^{\prime}(\lambda)}{Z_{\mathbf{F}_{0}(x)}(\lambda)} .
$$

We sum (5) over all $q$ colours $x \in L(u)$ to obtain

$$
q \leq \beta \frac{\lambda}{1+\lambda} \mathbb{E}\left|L_{\tau}(u)\right|+\gamma \sum_{x \in L(u)} \mathbb{E}\left|\mathbf{J}_{0} \cap \Lambda_{x}\right| \leq \beta \frac{\lambda}{1+\lambda} \mathbb{E}\left|L_{\tau}(u)\right|+\gamma \Delta,
$$

where the last inequality holds because $\sum_{x \in L(u)} \mathbb{E}\left|\mathbf{J}_{0} \cap \Lambda_{x}\right|=\mathbb{E} \sum_{x \in L(u)}\left|\mathbf{J}_{0} \cap \Lambda_{x}\right|$ and every neighbour of $u$ contributes at most 1 to the sum as $\left|\mathbf{J}_{0} \cap L(v)\right| \leq 1$ for every vertex $v$. Rearranging immediately yields $\mathbb{E}\left|L_{\tau}(u)\right| \geq m$, and the result will follow from an application of Lemma 21 .

For this application, note that $\mathbb{E}\left|L_{\tau}(u)\right|$ is a sum over $x \in L(u)$ of the indicator variables $\mathbf{X}_{x}$ for the events $\left\{\mathbf{J}_{0} \cap \Lambda_{x}=\varnothing\right\}$. We can apply Lemma 21 if we show that the random variables $\mathbf{Y}_{x}:=$ $1-\mathbf{X}_{x}$ are negatively correlated. This correlation was shown formally by Bernshteyn [7] in the triangle-free case, and is somewhat intuitive here. Consider the random set $\mathbf{J}_{0}$. Given $x \in L(u)$, if $\mathbf{J}_{0} \cap \Lambda_{x}=\varnothing$ 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 $x^{\prime} \in L(u) \backslash\{x\}$. We repeat Bernshteyn's argument for completeness.

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

$$
\operatorname{Pr}\left(x \notin L_{\tau}(u) \mid Y \cap L_{\tau}(u)=\varnothing\right) \leq \operatorname{Pr}\left(x \notin L_{\tau}(u)\right),
$$

which is equivalent to

$$
\operatorname{Pr}\left(Y \cap L_{\tau}(u)=\varnothing \mid x \in L_{\tau}(u)\right) \geq \operatorname{Pr}\left(Y \cap L_{\tau}(u)=\varnothing\right),
$$

which we can write (using Fact 1) as

$$
\begin{aligned}
& \operatorname{Pr}\left(\mathbf{J}_{0} \cap \Lambda_{y} \neq \varnothing \text { for all } y \in Y \mid \mathbf{J}_{0} \cap \Lambda_{x}=\varnothing\right) \\
& \geq \operatorname{Pr}\left(\mathbf{J}_{0} \cap \Lambda_{y} \neq \varnothing \text { for all } y \in Y\right) .
\end{aligned}
$$

This holds because the layers $\Lambda_{z}$ 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 23. For any $x \in L(u)$, writing

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

we have $\operatorname{Pr}\left(D_{x}\right) \leq 1 /\left(8 q \cdot \Delta^{3}\right)$.
Proof. By Fact 3 we know that if $x \in L_{\tau}(u)$ then $\operatorname{deg}_{\mathscr{\mathscr { H } _ { \tau }}}^{*}(x)=\left|\mathbf{U}_{0}(x)\right|$, so it suffices to show whenever $\mid \overline{\mathbf{U}_{0}(x) \mid}>\ell / 8$ that

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

As already reported, by Fact 1 (and an elementary property of the hard-core model) we have $\operatorname{Pr}\left(x \in L_{\tau}(u)\right)=1 / Z_{\mathbf{F}_{0}(x)}(\lambda)$. We recall that $\mathbf{F}_{0}(\lambda)$ is isomorphic to a subgraph of $G[N(u)]$, so that the upper bound on $\operatorname{Pr}\left(x \in L_{\tau}(u)\right)$ follows directly from the assumptions on $\ell$ and $Z_{F}(\lambda)$ for $F \subset G[N(u)]$ stated in Theorem 7. 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

$$
\operatorname{Pr}\left(x \in L_{\tau}(u)| | \mathbf{U}_{0}(x) \mid>\ell\right) \leq \frac{1}{8 \Delta^{4}} \leq \frac{1}{8 q \Delta^{3}} .
$$

The result follows.
The combination of Lemmas 22 and 23 completes the proof of Lemma 13 .

## B.2. Proof of Lemma 14 .

Proof of Lemma 14. Addressing $B_{u}$ by executing $\operatorname{AddressB}(u, \sigma)$ only modifies $\sigma$ on $N(u)$ and hence can only introduce flaws $B_{v}$ for $v \in N^{3}[u]$, or $U_{v}^{v w}$ for $v, w \in N(u)$ and $v w \in E(G[N(u)])$.

In the procedure AddressB we can only introduce at most $t$ flaws of the form $U_{v}^{v w}$ 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 $\tau$ with positive probability then the previous state $\sigma$ must agree with $\tau$ outside of $N(u)$, and inside $N(u)$ we must have unc $(\tau) \supset \operatorname{unc}(\sigma)$ as no uncoloured vertices are coloured by the procedure AddressB. In fact, by the definition of covering for sets of flaws, when bounding $c^{S}\left(B_{u}\right)$ we can restrict our attention to triples $(S, \sigma, \tau)$ such that $\tau$ and $\sigma$ agree outside $N(u)$ and

$$
\operatorname{unc}(\tau)=\operatorname{unc}(\sigma) \cup\left\{v \in N(u): U_{v}^{e} \in S \text { for some } e \in E(G)\right\} .
$$

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

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

so that $\nu(S)$ is the set of uncoloured vertices present in $\tau$ but not $\sigma$ for any $\sigma \in \operatorname{In}_{B_{u}}^{S}(\tau)$. We now assert that for such triples $(S, \sigma, \tau)$ 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 $\operatorname{AddressB}(u, \sigma)$; repeated here for convenience, $\hat{H}$ is constructed as follows. Let $\sigma^{\prime}$ be obtained from $\sigma$ by setting $\sigma^{\prime}(v):=\mathfrak{B}$ for all $v \in N(u) \backslash \operatorname{unc}(\sigma)$, and $\sigma^{\prime}(v):=\sigma(v)$ otherwise, and then let $F:=G_{\sigma^{\prime}}\left[N_{G}(u)\right]$ and $H^{\prime}:=H_{\sigma^{\prime}}\left[L\left(N_{G}(u)\right)\right]$. Then $\hat{H}:=\operatorname{Remove}\left(F, H^{\prime}\right)$. Note that $\hat{H}$ does not depend on $\sigma$ in the sense that when $(S, \sigma, \tau)$ are as above we can construct $\hat{H}$ from $S$ and $\tau$ alone. Here we crucially exploit the definition of covering for sets of flaws. The key point is that $\sigma^{\prime}$ as above can be constructed from $S$ and $\tau$ because the uncoloured vertices of $\sigma$ are given by unc $(\sigma)=\operatorname{unc}(\tau) \backslash \nu(S)$.

The assertion now follows from the definition of $\operatorname{AddressB}(u, \sigma)$ because from $(S, \tau)$ we can recover the size of the independent set $\mathbf{J}$ sampled in the procedure when $\tau=\operatorname{AddressB}(u, \sigma)$. 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 $\operatorname{AddressB}(u, \sigma)$. 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 $\nu$ and because $\sigma$ and $\tau$ only differ in $N(u)$. Then we have shown that given $S$ and $\tau$, for any $\sigma \in \operatorname{In}_{B_{u}}^{S}(\tau)$ we have

$$
\begin{equation*}
\frac{\tilde{\mu}(\sigma)}{\tilde{\mu}(\tau)} \rho_{B_{u}}(\sigma, \tau)=\frac{\lambda^{|\operatorname{col}(\sigma) \cap N(u)|}}{Z_{\hat{H}}(\lambda)} \tag{6}
\end{equation*}
$$

where $\hat{H}$ can be obtained from $S$ and $\tau$ alone.
Turning to the charge $c^{S}\left(B_{u}\right)$, we have

$$
\begin{equation*}
c^{S}\left(B_{u}\right)=\max _{\tau \in \Omega}\left\{\sum_{\sigma \in \operatorname{In}_{B_{u}}^{S}(\tau)} \frac{\tilde{\mu}(\sigma)}{\tilde{\mu}(\tau)} \rho_{B_{u}}(\sigma, \tau)\right\}=\max _{\substack{\tau \in \Omega: \\ u \notin \operatorname{col}(\tau)}}\left\{\sum_{\sigma \in \operatorname{In}_{B_{u}}^{S}(\tau)} \frac{\lambda^{|\operatorname{col}(\sigma) \cap N(u)|}}{Z_{\hat{H}}(\lambda)}\right\}, \tag{7}
\end{equation*}
$$

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

Nearing conclusion, we now argue that Lemma 13 implies $c^{S}\left(B_{u}\right) \leq 1 /\left(4 \Delta^{3}\right)$. This holds because given $S$ and $\tau$ we can construct the $\hat{H}$ occurring in (7), 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.
(i) One obtains $\hat{H}$ from $H$ by removing edges in $H^{*}$, and hence $\hat{\mathscr{H}}$ satisfies the hypotheses of Theorem 7 .
(ii) There is a partial colouring independent set $\hat{J} \in \mathcal{I}(\hat{H})$ such that $\hat{H}=\hat{H}_{\hat{J}}\left[L\left(N_{G}(u)\right)\right]$. Then the sum in (7) over states $\sigma \in \operatorname{In}_{B_{u}}^{S}(\tau) \subset B_{u}$ can be interpreted as the probability that when $\operatorname{col}(\sigma) \cap N(v)$ is a random independent set from the hard-core model on $\hat{H}$ at fugacity $\lambda$, we have $\sigma \in \operatorname{In}_{B_{u}}^{S}(\tau)$. Since $\operatorname{In}_{B_{u}}^{S}(\tau) \subset B_{u}$, we can bound this from above by the probability that $\sigma$ belongs to $B_{u}$ given this random experiment; and Lemma 13 shows this probability to be at most $1 /\left(4 \Delta^{3}\right)$.

## B.3. Proof of Lemma 15 ,

Proof of Lemma 15. Addressing $U_{u}^{e}$ at state $\sigma$ 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 $U_{u}^{e}$ results in $\tau$ with positive probability then the previous state $\sigma$ must be obtained from $\tau$ by setting $\tau(u)=e$. That is, for $S \subset F$, the set $\operatorname{In}_{U_{u}^{e}}^{S}(\tau)$ is either empty, or contains exactly one state $\sigma$. In the former case the charge is zero; and in the latter case, if $u \in \operatorname{bla}(\tau)$ then $\tilde{\mu}(\sigma) / \tilde{\mu}(\tau)=\lambda$ and $\rho_{U_{u}^{e}}(\sigma, \tau)=1 /\left(1+\left|L_{\sigma}(u)\right| \lambda\right)$, and if $u \in \operatorname{col}(\tau)$ then $\tilde{\mu}(\sigma) / \tilde{\mu}(\tau)=1$ and $\rho_{U_{u}^{e}}(\sigma, \tau)=\lambda /\left(1+\left|L_{\sigma}(u)\right| \lambda\right)$, hence both possibilities yield that

$$
\frac{\tilde{\mu}(\sigma)}{\tilde{\mu}(\tau)} \rho_{U_{u}^{e}}(\sigma, \tau)=\frac{\lambda}{1+\left|L_{\sigma}(u)\right| \lambda} .
$$

Then $c^{S}\left(U_{u}^{e}\right) \leq \lambda /(1+\ell \lambda)$ because the ordering on flaws ensures that $\sigma \notin B_{u}$ whenever we are addressing $U_{u}^{e}$, and hence $u$ has at least $\ell$ available colours.
B.4. Proof of Lemma 11, We can prove Lemma 11 with the local lemma of Moser and Tardos [37]. Let $\sigma$ be a flawless partial colouring of $G$, and if necessary remove colours from $H_{\sigma}$ such that $\left|L_{\sigma}(u)\right|=\ell$ for all $u \in G_{\sigma}$. We let $I^{\prime} \subset V\left(H_{\sigma}\right)$ be obtained by choosing for each remaining blank vertex $u \in \operatorname{bla}(\sigma)$ a uniform random colour from $L_{\sigma}(u)$. Since $\sigma$ is flawless there are no uncoloured vertices, so the conclusion follows if $I^{\prime}$ is independent. Then for this algorithmic local lemma application we have a flaw $A_{x y}$ for each edge $x y$ of $H_{\sigma}^{*}$, and $I^{\prime} \in A_{x y}$
if and only if $\{x, y\} \subset I^{\prime}$. To address $A_{x y}$ we resample $I^{\prime} \cap L(u)$ and $I^{\prime} \cap L(v)$ independently, uniformly at random.

Suppose that $x \in L(u), y \in L(v)$, and $x y \in E\left(H_{\sigma}^{*}\right)$. Then $A_{x y}$ is independent of any $A_{x^{\prime} y^{\prime}}$ such that $x^{\prime}$ and $y^{\prime}$ are not in $L_{\sigma}(u) \cup L_{\sigma}(v)$. In this simple setting the full charge machinery of Theorem 12 is not necessary, and it suffices to use the result of Moser and Tardos [37, Thm. 1.2]. We use the slightly stronger version found in Iliopoulos' doctoral thesis [27, Thm. 3.9, p. 17]: it implies the original formulation by setting $x_{A}:=\frac{\psi_{A}}{1+\psi_{A}} \in(0,1)$.
Theorem 24 (Iliopoulos [27). 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 $\Gamma(A):=\{B \in \mathcal{A}$ : $A$ and $B$ are dependent\}, and let $\mu$ be the probability measure that results from sampling the variables $\mathcal{P}$. If for each $A \in \mathcal{A}$ there exists a positive real $\psi_{A}$ 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 $I^{\prime} \cap L(u)$ for $u \in V\left(G_{\sigma}\right)$ and the set of events is $\mathcal{A}:=\left\{A_{x y}: x y \in H_{\sigma}^{*}\right\}$. Supposing that $x \in L(u), y \in L(v)$, and $x y \in E\left(H_{\sigma}^{*}\right)$, the fact that $\sigma$ is flawless implies that

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

and also that

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

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

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

The expected number of resamplings is then at most $n / 4$, because $H_{\sigma}^{*}$, 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 11.

## Appendix C. Proofs for the application

C.1. Proof of Lemma 16. The proof of Lemma 16 relies on the following elementary lemma, which already appeared [14], though we give the short proof here for completeness.

Lemma 25. For any graph $F$ on $y$ vertices with positive average degree at most $a$,

$$
\frac{\lambda Z_{F}^{\prime}(\lambda)}{Z_{F}(\lambda)} \geq \frac{\lambda}{1+\lambda} y(1+\lambda)^{-a}, \quad \text { and } \quad \log Z_{F}(\lambda) \geq \frac{y}{a}\left(1-(1+\lambda)^{-a}\right)
$$

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

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

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 $\mathbf{U}:=N(u) \backslash N(\mathbf{S} \backslash N(u))$ is the set obtained by revealing $\mathbf{S} \backslash N(u)$ and removing from $N(u)$ any vertex with a neighbour in $\mathbf{S} \backslash N(u)$, 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 $\lambda$. The final inequality comes from the fact that any realisation of $\mathbf{F}_{N(u)}$ has $Z_{\mathbf{F}_{N(u)}}(\lambda) \leq(1+\lambda)^{\operatorname{deg}(u)}$. The lemma now follows by convexity:

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

and since

$$
\mathbb{E}|\mathbf{S}|=\frac{\lambda Z_{F}^{\prime}(\lambda)}{Z_{F}(\lambda)}=\lambda \frac{\partial}{\partial \lambda} \log Z_{F}(\lambda)
$$

integrating this bound gives the required lower bound on $\log Z_{F}(\lambda)$.
Proof of Lemma 16. 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 $Z_{F}(\lambda)=(1+\lambda)^{y}$, otherwise by Lemma 25 we have

$$
\log Z_{F}(\lambda) \geq \frac{y}{a}\left(1-(1+\lambda)^{-a}\right) \geq y \log (1+\lambda)\left(1-\frac{a}{2} \log (1+\lambda)\right),
$$

which establishes Item (ii)] For Item (i)] we note that $Z_{F}(\lambda) \leq(1+\lambda)^{y}$ and hence by Lemma 25 we have

$$
\beta \frac{\lambda}{1+\lambda} \frac{1}{Z_{F}(\lambda)}+\gamma \frac{\lambda Z_{F}^{\prime}(\lambda)}{Z_{F}(\lambda)} \geq \frac{\lambda}{1+\lambda}\left(\beta(1+\lambda)^{-y}+\gamma y(1+\lambda)^{-a}\right),
$$

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 $g\left(y^{*}\right)=1$ for strong local $(\beta, \gamma)$-occupancy, and solve for $\beta$ we obtain

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

Then the function $\beta+\gamma d$ is strictly convex in $\gamma$, and the unique minimiser is attained when

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

One checks that, indeed, setting $\beta$ and $\gamma$ to the announced values, and writing $D$ for the expression $d(1+\lambda)^{a} \log (1+\lambda)$, we have

$$
\beta=\frac{1+\lambda}{\lambda} \cdot \frac{e^{W(D)}}{1+W(D)}=\frac{1+\lambda}{\lambda} \cdot \frac{D}{W(D) \cdot(1+W(D))},
$$

noticing that $(1+\lambda)^{1+a} /(\gamma \lambda)=(1+W(D)) / \log (1+\lambda)$, and hence

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

as announced. Furthermore, $y^{*}=W(D) / \log (1+\lambda)$, and hence indeed

$$
\begin{aligned}
g\left(y^{*}\right) & =\frac{D \cdot(1+\lambda)^{-W(D) / \log (1+\lambda)}}{W(D)(1+W(D))}+\frac{W(D)}{1+W(D)} \\
& =\frac{1}{1+W(D)}+\frac{W(D)}{1+W(D)}=1
\end{aligned}
$$

C.2. Proof of Theorem 18, As is well known (and straightforward to realise), in a connected graph any two longest paths must have a common vertex, which we state as follows.

Lemma 26. 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 $F[V \backslash P]$ is $P_{k-2}$-free.
We now prove a warm-up to Theorem 18 that shows we can calculate the partition function and certain useful probabilities efficiently.
Lemma 27. 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 $\lambda$ we can evaluate $Z_{H}(\lambda)$ in time $y^{3 k}(1+q)^{k^{2} / 2}(c k)^{k^{2}}$.

Further, let $F_{1}, \ldots, F_{r}$ be the connected components of $F$. Write $H_{i}=H\left[L\left(V\left(F_{i}\right)\right)\right]$ 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 $\lambda$ on $H$. Then in time $y^{3 k}(1+q)^{k^{2} / 2}(c k)^{k^{2}}$ we can compute for every $i$ in $[r]$ a longest path $P_{i}$ in $F_{i}$, and the probabilities $\operatorname{Pr}\left(\mathbf{I} \cap L\left(P_{i}\right)=J_{i}\right)$ for every independent set $J_{i} \subset L\left(P_{i}\right)$.

Proof. In the proof we write $c_{1}, c_{2}, \ldots$ 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 $F^{\prime}=\left(V^{\prime}, E^{\prime}\right)$ be a component of $F$, and write $H^{\prime}=H\left[L\left(V^{\prime}\right)\right]$ for the induced cover of $F^{\prime}$. Let $P=\left\{v_{1}, \ldots, v_{k-2}\right\}$ be the vertex set of a longest path $F^{\prime}$, which we can find in time $c_{2}(k-2)!2^{k-2} y$ by executing for each $j$ from $k-1$ to 2 a fixed-parameter algorithm that finds a path of length $j$ (if one exists) in time $O\left(j!2^{j} y\right)$, see [8].

Since $F^{\prime}$ is connected, Lemma 26 implies that $F^{\prime}\left[V^{\prime} \backslash P\right]$ is $P_{k-2}$-free. We use the fact that the sum over independent sets $I$ in $H^{\prime}$ that gives $Z_{H^{\prime}}(\lambda)$ can be split into terms according to $I \cap L(P)$. Let $H^{P}=H^{\prime}[L(P)]$ and $H^{V^{\prime} \backslash P}=H^{\prime}\left[L\left(V^{\prime} \backslash P\right)\right]$. Then

$$
Z_{H^{\prime}}(\lambda)=\sum_{J \in \mathcal{I}\left(H^{P}\right)} \lambda^{|J|} \cdot Z_{H_{J}^{V^{\prime} \backslash P}}(\lambda) .
$$

It is possible to iterate over the required $J \in \mathcal{I}\left(H^{P}\right)$ in time $c_{3}(k-2)^{2}(1+q)^{k-2}$ by iterating over all $(1+q)^{k-2}$ sets $J$ with $\left|J \cap L\left(v_{i}\right)\right| \leq 1$ and checking each for independence. The independence check takes time $O\left((k-2)^{2}\right)$ as we must verify that each pair of vertices in $J$ is absent from $E(H)$. Then constructing $H_{J}^{V \backslash P}$ can be done in time $c_{4}(k-2) y$ because $|J| \leq k-2$ and each colour in $J$ can conflict with at most $y$ colours which are removed from $H^{V \backslash P}$ to form $H_{J}^{V \backslash P}$. Now $Z_{H_{J}^{V \backslash P}}(\lambda)$ can be computed in time $f(y, q, k-1)$ by induction, so we can compute $Z_{H^{\prime}}(\lambda)$ in time

$$
c_{2}(k-2)!2^{k-2} y \cdot c_{3}(k-2)^{2}(1+q)^{k-2} \cdot c_{4}(k-2) y \cdot f(y, q, k-1) .
$$

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

$$
f(y, q, k)=y \cdot c_{2}(k-2)!2^{k-2} y \cdot c_{3}(k-2)^{2}(1+q)^{k-2} \cdot c_{4} k y \cdot f(y, q, k-1) .
$$

With the base case of $f(y, q, 3)=c_{1} y$ we have

$$
f(y, q, k) \leq c_{1} y\left(c_{2} c_{3} c_{4} y^{3}\right)^{k-3}(2(1+q))^{k(k-3) / 2}((k-2)!)^{4} \cdot \prod_{j=1}^{k-3} j!,
$$

which for a large enough constant $c$ is at most $y^{3 k}(1+q)^{k^{2} / 2}(c k)^{k^{2}}$.
For the final statement, observe that in the above argument we compute for each $i \in[r]$ an evaluation of $Z_{H_{J_{i}}\left[V\left(F_{i}\right) \backslash P_{i}\right]}$ for some longest path $P_{i}$ in $F_{i}$, and each independent set $J_{i} \subset L\left(P_{i}\right)$ along the way to computing $Z_{H_{i}}(\lambda)$. But

$$
\operatorname{Pr}\left(\mathbf{I} \cap L\left(P_{i}\right)=J_{i}\right)=\frac{\lambda^{\left|J_{i}\right|} Z_{H_{J_{i}}\left[V\left(F_{i}\right) \backslash P_{i}\right]}(\lambda)}{Z_{H_{i}}(\lambda)},
$$

so with some straightforward extra bookkeeping we have the required probabilities.
With this result we can give the required sampling algorithm, which is restated here.
Theorem 18. 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 $y^{3 k^{2}}(1+q)^{k^{3} / 2}(c k)^{k^{3}}$.
Proof. We write 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 $\mathbf{I} \cap \hat{L}(v)=\varnothing$, otherwise let $\mathbf{I} \cap \hat{L}(v)$ be a uniform random element of $\hat{L}(v)$. This can be done in time $O(q y)$ 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 $F_{1}, \ldots, F_{r}$ be the components of $F$. By Lemma 27 we can compute in time $y^{3 k}(1+q)^{k^{2} / 2}(c k)^{k^{2}}$ a longest path $P_{i}$ in $F_{i}$ and the probabilities $\operatorname{Pr}\left(\mathbf{I} \cap \hat{L}\left(P_{i}\right)=J_{i}\right)$ for all independent sets $J_{i} \subset \hat{L}\left(P_{i}\right)$ and for all $i$. Hence we can sample $\mathbf{I} \cap \hat{L}\left(P_{i}\right)$ for all $i$ in this time. Then we can construct $H_{i}^{\prime}=\hat{H}_{\mathbf{I} \cap \hat{L}\left(P_{i}\right)}\left\{\hat{L}\left(V\left(F_{i}\right) \backslash P_{i}\right)\right]$ for each $i$ in time $O(r k y)$, and use the fact that $\mathbf{I} \cap \hat{L}\left(V\left(F_{i}\right) \backslash P_{i}\right)$ is distributed according to the hard-core model at fugacity $\lambda$ on $H_{i}^{\prime}$.

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

$$
y^{3 k^{2}}(1+q)^{k^{3} / 2}(c k)^{k^{3}}
$$

## References

[1] D. Achlioptas and A. Coja-Oghlan. Algorithmic barriers from phase transitions. 200849 th Annual IEEE Symposium on Foundations of Computer Science, pages 793-802, Oct. 2008. DOI:10.1109/FOCS.2008.11
[2] D. Achlioptas, F. Iliopoulos, and A. Sinclair. Beyond the Lovász Local Lemma: Point to Set Correlations and Their Algorithmic Applications. In 2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS), pages 725-744, Baltimore, MD, USA, Nov. 2019. IEEE. DOI:10.1109/FOCS.2019.00049
[3] M. Ajtai, J. Komlós, and E. Szemerédi. A note on Ramsey numbers. J. Combin. Theory Ser. A, 29(3):354-360, 1980. DOI:10.1016/0097-3165(80)90030-8
[4] M. Ajtai, J. Komlós, and E. Szemerédi. A dense infinite Sidon sequence. European J. Combin., 2(1):1-11, 1981. DOI:10.1016/S0195-6698(81)80014-5.
[5] N. Alon, M. Krivelevich, and B. Sudakov. Coloring Graphs with Sparse Neighborhoods. Journal of Combinatorial Theory, Series B, 77(1):73-82, Sept. 1999. DOI:10.1006/jctb.1999.1910
[6] N. Bansal, A. Gupta, and G. Guruganesh. On the Lovász Theta function for Independent Sets in Sparse Graphs. Apr. 2015, arXiv:1504.04767
[7] A. Bernshteyn. The Johansson-Molloy theorem for DP-coloring. Random Structures \& Algorithms, 54(4):653664, July 2019. DOI:10.1002/rsa. 20811
[8] H. Bodlaender. On Linear Time Minor Tests with Depth-First Search. Journal of Algorithms, 14(1):1-23, Jan. 1993. DOI:10.1006/jagm.1993.1001.
[9] T. Bohman. The triangle-free process. Adv. Math., 221(5):1653-1677, 2009. DOI:10.1016/j.aim.2009.02.018.
[10] T. Bohman and P. Keevash. Dynamic concentration of the triangle-free process. Sept. 2019, arXiv:1302.5963
[11] A. Coja-Oghlan and C. Efthymiou. On independent sets in random graphs. Random Structures Algorithms, 47(3):436-486, 2015. DOI:10.1002/rsa. 20550
[12] R. M. Corless, G. H. Gonnet, D. E. G. Hare, D. J. Jeffrey, and D. E. Knuth. On the Lambert $W$ function. Advances in Computational Mathematics, 5(1):329-359, Dec. 1996. DOI:10.1007/BF02124750.
[13] E. Davies, R. de Joannis de Verclos, R. J. Kang, and F. Pirot. Colouring triangle-free graphs with local list sizes. Dec. 2018, arXiv:1812.01534.
[14] E. Davies, R. de Joannis de Verclos, R. J. Kang, and F. Pirot. Occupancy fraction, fractional colouring, and triangle fraction. Dec. 2018, arXiv:1812.11152
[15] E. Davies, M. Jenssen, W. Perkins, and B. Roberts. Independent sets, matchings, and occupancy fractions. Journal of the London Mathematical Society, 96(1):47-66, Aug. 2017. DOI:10.1112/jlms.12056.
[16] E. Davies, M. Jenssen, W. Perkins, and B. Roberts. On the average size of independent sets in triangle-free graphs. Proc. Amer. Math. Soc., 146(1):111-124, 2018. DOI:10.1090/proc/13728
[17] E. Davies, R. J. Kang, F. Pirot, and J.-S. Sereni. Graph structure via local occupancy. Mar. 2020, arXiv:2003.14361.
[18] Z. Dvořák and L. Postle. Correspondence coloring and its application to list-coloring planar graphs without cycles of lengths 4 to 8. Journal of Combinatorial Theory, Series B, 129:38-54, Mar. 2018. DOI:10.1016/j.jctb.2017.09.001.
[19] P. Erdős. Some remarks on the theory of graphs. Bull. Amer. Math. Soc., 53:292-294, 1947. DOI:10.1090/S0002-9904-1947-08785-1.
[20] P. Erdős. Graph theory and probability. Canadian J. Math., 11:34-38, 1959. DOI:10.4153/CJM-1959-003-9.
[21] P. Erdős and G. Szekeres. A combinatorial problem in geometry. Compositio Math., 2:463-470, 1935. URL http://www.numdam.org/item?id=CM_1935__2__463_0.
[22] P. Erdős and T. Gallai. On maximal paths and circuits of graphs. Acta Mathematica Academiae Scientiarum Hungaricae, 10(3-4):337-356, Sept. 1959. DOI:10.1007/BF02024498.
[23] P. Erdős and L. Lovász. Problems and results on 3-chromatic hypergraphs and some related questions. In Infinite and Finite Sets (Colloq., Keszthely, 1973; Dedicated to P. Erdős on His 60th Birthday), Vol. II, pages 609-627. Colloq. Math. Soc. János Bolyai, Vol. 10. North-Holland, Amsterdam, 1975.
[24] L. Esperet and A. Parreau. Acyclic edge-coloring using entropy compression. European J. Combin., 34(6):10191027, 2013. DOI:10.1016/j.ejc.2013.02.007.
[25] G. Fiz Pontiveros, S. Griffiths, and R. Morris. The Triangle-Free Process and the Ramsey Number $R(3, k)$. Memoirs of the American Mathematical Society, 263(1274), Jan. 2020. DOI:10.1090/memo/1274.
[26] M. M. Halldórsson and J. Radhakrishnan. Improved approximations of independent sets in bounded-degree graphs. In G. Goos, J. Hartmanis, E. M. Schmidt, and S. Skyum, editors, Algorithm Theory - SWAT '94, volume 824, pages 195-206. Springer Berlin Heidelberg, 1994. DOI:10.1007/3-540-58218-5_18.
[27] F. Iliopoulos. Stochastic Local Search and the Lovasz Local Lemma. PhD thesis, EECS Department, University of California, Berkeley, 2019.
[28] A. Johansson. Asymptotic choice number for triangle-free graphs. Technical Report 91-5, DIMACS, 1996.
[29] R. J. Kang and C. McDiarmid. Colouring random graphs. In Topics in chromatic graph theory, volume 156 of Encyclopedia Math. Appl., pages 199-229. Cambridge Univ. Press, Cambridge, 2015.
[30] R. M. Karp. Reducibility among Combinatorial Problems. In R. E. Miller, J. W. Thatcher, and J. D. Bohlinger, editors, Complexity of Computer Computations, pages 85-103. Springer US, Boston, MA, 1972. DOI:10.1007/978-1-4684-2001-2_9.
[31] R. M. Karp. The probabilistic analysis of some combinatorial search algorithms. In Algorithms and complexity (Proc. Sympos., Carnegie-Mellon Univ., Pittsburgh, Pa., 1976), pages 1-19, 1976.
[32] J. H. Kim. On Brooks' Theorem for Sparse Graphs. Combinatorics, Probability and Computing, 4(2):97-132, June 1995. DOI:10.1017/S0963548300001528
[33] J. H. Kim. The Ramsey number $R(3, t)$ has order of magnitude $t^{2} / \log t$. Random Structures Algorithms, $7(3): 173-207,1995$. DOI:10.1002/rsa.3240070302.
[34] M. Molloy. The list chromatic number of graphs with small clique number. Journal of Combinatorial Theory, Series B, 134:264-284, Jan. 2019. DOI:10.1016/j.jctb.2018.06.007.
[35] M. Molloy and B. Reed. Graph Colouring and the Probabilistic Method. Algorithms and Combinatorics. Springer-Verlag, Berlin Heidelberg, 2002. DOI:10.1007/978-3-642-04016-0
[36] R. A. Moser. A constructive proof of the Lovász local lemma. In STOC'09—Proceedings of the 2009 ACM International Symposium on Theory of Computing, pages 343-350. ACM, New York, 2009.
[37] R. A. Moser and G. Tardos. A constructive proof of the general Lovász local lemma. Journal of the ACM, 57(2):1-15, Jan. 2010. DOI:10.1145/1667053.1667060
[38] A. Panconesi and A. Srinivasan. Randomized Distributed Edge Coloring via an Extension of the ChernoffHoeffding Bounds. SIAM Journal on Computing, 26(2):350-368, Mar. 1997. DOI:10.1137/S0097539793250767
[39] F. P. Ramsey. On a Problem of Formal Logic. Proc. London Math. Soc. (2), 30(4):264-286, 1929. DOI:10.1112/plms/s2-30.1.264.
[40] J. B. Shearer. A note on the independence number of triangle-free graphs. Discrete Mathematics, 46(1):83-87, 1983. DOI:10.1016/0012-365X(83)90273-X
[41] P. Ungar and B. Descartes. Advanced Problems and Solutions: Solutions: 4526. The American Mathematical Monthly, 61(5):352, May 1954. DOI:10.2307/2307489.
[42] V. G. Vizing. Some unsolved problems in graph theory. Uspehi Mat. Nauk, 23(6 (144)):117-134, 1968.
[43] V. H. Vu. A General Upper Bound on the List Chromatic Number of Locally Sparse Graphs. Combinatorics, Probability and Computing, 11(1):103-111, Jan. 2002. DOI:10.1017/S0963548301004898
[44] L. Zdeborová and F. Krząkała. Phase transitions in the coloring of random graphs. Physical Review E, 76(3):031131, Sept. 2007. DOI:10.1103/PhysRevE.76.031131
[45] A. A. Zykov. On some properties of linear complexes. Mat. Sbornik N.S., 24(66):163-188, 1949.

Department of Computer Science, University of Colorado Boulder, USA
Email address: maths@ewandavies.org
Department of Mathematics, Radboud University Nijmegen, Netherlands.
Email address: ross.kang@gmail.com
G-SCOP, CNRS, Univ. Grenoble Alpes, Grenoble, France
Email address: francois.pirot@grenoble-inp.fr
Service Public Français de la Recherche, Centre National de la Recherche Scientifique, CStB (ICube), Strasbourg, France

Email address: sereni@kam.mff.cuni.cz


[^0]:    (E. Davies) 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.
    (R. J. Kang) Supported by a Vidi grant (639.032.614) of the Netherlands Organisation for Scientific Research (NWO).

[^1]:    ${ }^{1}$ 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.
    ${ }^{2}$ Note that this definition differs from that given in [2] as we have additionally restricted to mappings $\sigma$ 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.

