Asymptotically optimal frugal colouring

Michael Molloy* Bruce Reed†

June 18, 2009

Abstract

We prove that every graph with maximum degree Δ can be properly $(\Delta+1)$-coloured so that no colour appears more than $O(\log\Delta/\log\log\Delta)$ times in the neighbourhood of any vertex. This is best possible up to the constant multiple in the $O(-)$ term.

1 Introduction

In [9], Hind, Molloy and Reed defined a proper vertex coloring to be β-frugal if no vertex has more than β members of any colour class in its neighbourhood. It is very easy to $(\Delta+1)$-colour any graph with maximum degree Δ. The main result of that paper was to show that every graph with maximum degree Δ in fact has a β-frugal $(\Delta+1)$-colouring with $\beta = O(\log^8\Delta)$. Pemmaraju and Srinivasan[23] recently improved this to $\beta = O(\log^2\Delta/\log\log\Delta)$.

Alon (see [9]) provided a class of examples that do not have a $(\log\Delta/\log\log\Delta)$-frugal $(\Delta+1)$-colouring. In fact, for every $t>0$, and Δ sufficiently large, there is a graph which does not have a $(\log\Delta/\log\log\Delta)$-frugal $t\Delta$-colouring. In this paper, we close that asymptotic gap by proving:

Theorem 1.1 *There exists a constant Δ_0 such that every graph G with maximum degree $\Delta\geq\Delta_0$ has a $(50\log\Delta/\log\log\Delta)$-frugal $(\Delta+1)$-colouring.*

*Dept of Computer Science, University of Toronto, molloy@cs.toronto.edu. Research supported by NSERC.

†Canada Research Chair in the Combinatorics of Complex Networks, School of Computer Science, McGill University, breed@cs.mcgill.ca. Research supported by NSERC.

*A preliminary version of this paper appeared in the proceedings of SODA 2009.

We do not specify Δ_0; we just assume that it is large enough to satisfy several inequalities scattered throughout the paper. We made no attempt to optimize the constant "50", and it is chosen mainly for its "roundness". In fact, it is very easy to lower it. However, we don't see a way to get it close to 1.

For a graph with maximum degree $\Delta < \Delta_0$, any $(\Delta+1)$-colouring is Δ_0-frugal. Therefore, Theorem 1.1 implies that every graph has a $(T \log \Delta / \log\log \Delta)$-frugal $(\Delta+1)$-colouring with $T \leq \max\{\Delta_0, 50\}$. (In our proof, Δ_0 is much bigger than 50.)

The main motivation of the initial study of frugal colouring in [9] was an application to total colouring, where one colours the vertices and edges of a graph so that the same colour does not appear on any two adjacent vertices, incident edges, or an edge and its endpoint. In [10] we proved that every graph with maximum degree Δ has a $\Delta + O(\log^8 \Delta)$ total colouring, by beginning with a $O(\log^8 \Delta)$-frugal vertex colouring, and then carefully colouring the edges. This result was later improved to $\Delta + O(1)$ in [16]. A well-known conjecture is that it can be improved to $\Delta + 2$ (see eg. [13]).

Amini, Esperet and van den Heuvel[2] study frugal colourings of planar graphs, as a generalization of the problem of bounding the chromatic number of the square of a planar graph. In [29], Yuster introduced linear colorings, which are proper colourings that are both acyclic (the union of any two colour classes induces a forest) and 2-frugal; this is equivalent to saying that the union of any two colour classes is a forest of paths. In their aforementioned paper[23], Pemmaraju and Srinivasan show that every triangle-free graph has an $O(\log^2 \Delta)$-frugal $O(\Delta/\log \Delta)$-colouring, and that every d-degenerate graph has a β-frugal $(d+1)$-colouring for $\beta \approx O(\frac{\Delta}{d} \log^2 \Delta)$.

Our proof is probabilistic. We use a randomized procedure to $(\Delta+1)$-colour the graph, and we show that, with positive probability, the colouring produced will be β-frugal with $\beta = 50 \log \Delta / \log\log \Delta$. Suppose that the neighbours of a vertex v all received independently chosen uniformly random colours from $\{1, ..., \Delta+1\}$. A simple calculation shows that the expected number of colours chosen more than roughly β times is $o(1)$ and so, with high probability, no colour is chosen more than β times. By applying the Lovasz Local Lemma (see Section 2.3), one can often move from "with high probability the neighbourhood of one vertex is fine" to "with postitive probability the neighbourhood of every vertex is fine".

Of course, we can't always ensure that the colours appearing on $N(v)$ are chosen independently. For example, this is impossible if there are many edges in $N(v)$. But on an intuitive level, having many edges in $N(v)$ should be to our advantage since they make it even less likely that a colour would appear more than β times in $N(v)$.

So our aim is to choose a randomized procedure in which, very roughly speaking, the colours appearing on each $N(v)$ are chosen in a manner that is similar to uniform and independent. Similar enough to allow us to show that, with positive probability, no vertex appears β times in any neighbourhood. Our procedure is a version of what is often called the "Rodl Nibble", the "semi-random method" or the "Naive Colouring Procedure" (see [26, 19]).

In the case that the graph has girth at least 5, the resulting colouring on each neighbourhood is in many senses very close to being a set of independent uniform colours. In the case that 3- and 4-cycles are present, the procedure still works quite well so long as no $N(v)$ is close to being a Δ-clique (eg. it is sufficient for every $N(v)$ to contain at most $(1-\epsilon)\binom{\Delta}{2}$ edges for some constant $\epsilon > 0$). But in the presence of vertices with very dense neighbourhoods, we need to modify the procedure further. To do so, we make use of the dense decomposition introduced by Reed in [24], in which vertices with very dense neighbourhoods are isolated so that they can be coloured more carefully. This decomposition has been applied in [16, 18, 20, 24, 25]. (See [19] for a thorough presentation of this and related techniques.

Our procedure has three phases. The first two are along the same lines as similar procedures from other applications of this dense decomposition. Most of the new ideas in this paper appear in the third phase. There, we colour vertices one-at-a-time in a manner that creates far too much dependency for us to be able to apply the most common form of the Local Lemma; instead, we apply the Lopsided Local Lemma. Doing so requires an analysis of the probability of a set of vertices X all receiving the same colour, conditioned on the assignments of colours to vertices not in X. Conditioning on the assignments to vertices coloured before those in X is straightforward, but conditioning on assignments made after (or between) the vertices of X is the sort of thing that is often very problematic (see for example, Kahn's discussion in the epilogue of [14]). Fortunately, in this particular setting, we were able to handle the conditioning adequately.

In the next section, we present the decomposition and our key probabilistic tools. We close that section by giving an overview of our colouring procedure. In the following three sections, we prove Theorem 1.1 by describing and analyzing our random colouring procedure in three phases.

2 Preliminaries

2.1 A Dense Decomposition

We begin by describing the graph decomposition introduced in [24].

Consider any graph G with maximum degree Δ. It will be convenient to assume that G is Δ-regular, which we can do since every graph with maximum degree Δ is easily seen to be a subgraph of a Δ-regular graph (see eg. [24]).

We begin by decomposing G into *dense sets* $D_1, ..., D_\ell$ and a collection S of *sparse vertices* in the same way that we did in [16]. We set $\epsilon = 10^{-6}$. Section 2 of [16], in particular Lemmas 2.1(b,d), and 2.2 imply:

Lemma 2.1 *For each D_i:*

(a) $\Delta - 5\epsilon\Delta < |D_i| < \Delta + 2\epsilon\Delta$*;*

(b) there are at most $4\epsilon\Delta^2$ *edges from* D_i *to* $G - D_i$*;*

(c) every vertex $v \in S$ *has at least* $\epsilon\binom{\Delta}{2}$ *pairs of non-adjacent vertices in its neighbourhood;*

(d) each vertex is in D_i *iff it has at least* $\frac{3}{4}\Delta$ *neighbours in* D_i*.*

We define $\mathcal{D} = \cup_{i=1}^{\ell} D_i$. Note that $S = V(G) - \mathcal{D}$.

We wish to $\Delta + 1$ colour each D_i. We do so by partitioning it into a set of colour classes $\mathcal{C}_i$ each of size 1 or 2 so that either (i) the number of classes of size 2 is exactly $\lfloor 10\epsilon\Delta \rfloor$ or (ii) the number of classes of size 2 is less than $\lfloor 10\epsilon\Delta \rfloor$ and the vertices in the classes of size 1 form a clique. Lemma 2.4 of [16] and the Fact preceding it say:

Lemma 2.2 *For each* D_i*:*

(a) $\Delta - 15\epsilon\Delta \leq |\mathcal{C}_i| \leq \Delta + 1$*;*

(b) each colour class in $\mathcal{C}_i$ *has at most* $(\frac{1}{4} + 4\sqrt{\epsilon})\Delta < \frac{1}{3}\Delta$ *external neighbours.*

For each $v \in S$, it will be convenient to consider the set $\{v\}$ to be a colour class so that every vertex in G belongs to one colour class.

2.2 Ornery dense sets, kernels, Big_i and $\mathrm{Notbig}(i, x)$

For each vertex v in any D_i, we define:

- Out_v is the set of neighbours of v that are not in D_i;
- each member of Out_v is an *external neighbour* of v.

We say that D_i is *ornery* if $|\mathcal{C}_i| > \Delta - \log^4 \Delta$. For each ornery D_i, we define:

- K_i, its *kernel*, is the set of vertices in D_i with at most $\log^6 \Delta$ external neighbours.
- Big_i is the set of vertices outside D_i which have at least $\Delta^{7/8}$ neighbours in D_i.
- $\mathrm{Notbig}(i, x)$ is the set of vertices in D_i which do not have any external neighbours in $G - \mathrm{Big}_i$ with colour x. Note that $\mathrm{Notbig}(i, x)$ can change during the course of our colouring procedure.

We say that u, v are *big-neighbours* if they are both in Big_i for some i. In the first two phases of our colouring procedure, we will require that no two big-neighbours receive the same colour. Lemma 2.5 of [16] implies:

Lemma 2.3 *For each ornery D_i:*

(a) $|D_i - K_i| < 3\log^5 \Delta$.

(b) $|E(D_i, G - D_i)| < \Delta \log^7 \Delta$.

(c) $\mathcal{C}_i$ *has at most* $2\log^5 \Delta$ *colour classes of size 2.*

Proof Parts (a,b) are Lemma 2.5(b,c) of [16]. Lemma 2.5(a) of [16] says that $|D_i| < \Delta + \log^5 \Delta$. Since $|\mathcal{C}_i| > \Delta - \log^4 \Delta$, this implies that $\mathcal{C}_i$ has at most $\log^5 \Delta + \log^4 \Delta < 2\log^5 \Delta$ classes of size 2. □

Corollary 2.4 *Every vertex in G has at most $\Delta^{1/4} \log^7 \Delta$ big-neighbours.*

Proof Each vertex lies in Big_i for at most $\Delta/\Delta^{7/8} = \Delta^{1/8}$ dense sets D_i. By Lemma 2.3(b), each Big_i contains at most $|E(D_i, G - D_i)|/\Delta^{7/8} \leq \Delta^{1/8} \log^7 \Delta$ vertices. □

We say that $u, v \in G$ are *strongly non-adjacent* if they do not both lie in one dense set and if no member of the colour class containing v is a neighbour or big-neighbour of any member of the colour class containing u.

Lemma 2.5 *(a) Every $v \in S$ has at least $\frac{\epsilon}{80}\Delta^2$ pairs of strongly non-adjacent vertices in $N(s)$.*

(b) Every $v \in D_i$ has at least $\frac{\Delta}{10}|\mathrm{Out}_v|$ pairs of strongly non-adjacent vertices u, w where $u \in \mathrm{Out}_v$ and $w \in N(v)$ is a colour class of size one in $\mathcal{C}_i$.

Proof (a) follows from Lemma 2.9 of [16]. For (b): consider some $v \in D_i$. If $|\mathrm{Out}_v| \leq \frac{\Delta}{10}$ then by our assumption that G is regular, v has at least $\frac{9}{10}\Delta$ neighbours in D_i. There are at most $\lfloor 10\epsilon\Delta \rfloor$ colour classes of size two in $\mathcal{C}_i$, so at least $\frac{9}{10}\Delta - 20\epsilon\Delta$ neighbours of v are colour classes of size one in $\mathcal{C}_i$. Consider any $u \in \mathrm{Out}_v$. If $u \in S$ then by Lemma 2.1(d) and Corollary 2.4, u has fewer than $\frac{3}{4}\Delta + \Delta^{1/4}\log^7 \Delta$ neighbours or big-neighbours in D_i. If $u \in \mathcal{D}$ then by Lemmas 2.2(b) and Corollary 2.4, the colour class containing u has at most $\frac{\Delta}{3} + 2\Delta^{1/4}\log^7 \Delta$ neighbours or big-neighbours in D_i. So u is strongly non-adjacent to at least $\frac{9}{10}\Delta - 20\epsilon\Delta - (\frac{3}{4}\Delta + \Delta^{1/4}\log^7 \Delta) > \frac{\Delta}{10}$ neighbours of v that are colour classes of size one in $\mathcal{C}_i$.

If $|\mathrm{Out}_v| > \frac{\Delta}{10}$ then, by Lemma 2.1(d) and an argument like that above, there are at least $|\mathrm{Out}_v|(\frac{3}{4}\Delta - 20\epsilon\Delta)$ pairs of vertices u, w where $u \in \mathrm{Out}_v$ and $w \in N(v)$ is a colour class of size one in $\mathcal{C}_i$. By Lemma 2.1(b) there are at most $4\epsilon\Delta^2$ edges from D_i to Out_v, and each such edge can cause at most two of these pairs to be strongly non-adjacent. (The worse case is when the edge is incident to one of two vertices forming a colour class in $\mathcal{C}_j$, $j \neq i$.) By Corollary 2.4, the colour class containing any member of Out_v has at most $2\Delta^{1/4}\log^7\Delta$ big-neighbours in D_i. It follows that the number of pairs of strongly non-adjacent vertices u, w where $u \in \mathrm{Out}_v$ and $w \in N(v)$ is a colour class of size one in $\mathcal{C}_i$ is at least $|\mathrm{Out}_v|(\frac{3}{4}\Delta - 20\epsilon\Delta - 4\Delta^{1/4}\log^7\Delta) - 8\epsilon\Delta^2 > \frac{\Delta}{10}|\mathrm{Out}_v|$ (where the last inequality uses $|\mathrm{Out}_v| > \frac{\Delta}{10}$). $\square$

2.3 Probabilistic Preliminaries

In this section, we present a few probabilistic tools that we will use in this paper. First, we often use the following straightforward bound:

$$\binom{a}{b} \leq \left(\frac{ea}{b}\right)^b.$$

The following tool is crucial in this paper, as it is in many applications of the probabilistic method:

The Lovasz Local Lemma[6] *Let $\mathcal{A} = \{A_1, ..., A_n\}$ be a set of random events so that for each $1 \leq i \leq n$:*

(i) $\mathbf{Pr}(A_i) \leq p$*; and*

(ii) A_i *is mutually independent of all but at most d other events.*

If $pd \leq \frac{1}{4}$ then $\mathbf{Pr}(\overline{A_1} \cap ... \cap \overline{A_n}) > 0$.

In our final application of the Local Lemma, we will not have strict independence. Fortunately, we can get away with something weaker, using the following version which follows from the usual proof of the Local Lemma and was first used in [8]:

The Lopsided Local Lemma *Let $\mathcal{A} = A_1, ..., A_n$ be a set of random events. Suppose that for each A_i, we have a subset $B_i \subseteq \mathcal{A}$ such that:*

(i) for any subset $B \subset \mathcal{A} - B_i$,

$$\mathbf{Pr}(A_i | \cap_{A_j \in B} \overline{A_j}) \leq p;$$

(ii) $|B_i| \leq d$.

If $pd \leq \frac{1}{4}$ then $\mathbf{Pr}(\overline{A_1} \cap ... \cap \overline{A_n}) > 0$.

The binomial random variable $BIN(n, p)$ is the sum of n independent $0-1$ random variables where each is equal to 1 with probability p. The following is a derivation from Chernoff's original bound[5]. It follows from, eg Corollary A.1.10 and Theorem A.1.13 from Appendix A of [1].

The Chernoff Bound [5] *For any* $0 < t \leq np$*:*

$$\mathbf{Pr}(|BIN(n,p) - np| > t) < 2e^{-t^2/3np}.$$

Our next concentration bound is the Hoeffding-Azuma Inequality. Rather than using the original statements from [3, 11], we will use the following common corollary (see eg. Corollary 2.27 of [12]):

The Hoeffding-Azuma Inequality: *Let X be a non-negative random variable determined by the independent trials $T_1, ..., T_n$. Suppose that for every set of possible outcomes of the trials, we have:*

(i) changing the outcome of T_i can affect X by at most c_i.

Then for any $t \geq 0$, we have

$$\mathbf{Pr}(|X - \mathbf{Exp}(X)| > t) \leq 2e^{-\frac{t^2}{2\sum_{i=1}^{n} c_i^2}}.$$

Talagrand's Inequality requires another condition, but often provides a stronger bound when $\mathbf{Exp}(X)$ is much smaller than n. Rather than using Talagrand's original statement from [28], we will use the following useful reworking, which is proved in [20]:

Talagrand's Inequality *Let X be a non-negative random variable determined by the independent trials $T_1, ..., T_n$. Suppose that for every set of possible outcomes of the trials, we have:*

(i) changing the outcome of any one trial can affect X by at most c; and

(ii) for each $s > 0$, if $X \geq s$ then there is a set of at most rs trials whose outcomes certify that $X \geq s$.

Then for any $t \geq 50c\sqrt{r\mathbf{Exp}(X))} + 256c^2r$ we have

$$\mathbf{Pr}(|X - \mathbf{Exp}(X)| > t) \leq 4e^{-\frac{t^2}{32c^2r(\mathbf{Exp}(X)+t)}}.$$

McDiarmid extended Talagrand's Inequality to the setting where X depends on independent trials and permutations, a setting that arises often in this paper. Again, we present a useful reworking rather than the original inequality; this reworking is also proved in [20]. Talagrand[28] derived this for the case where there is exactly one permutation.

In the context of this inequality, a *choice* means either (a) the outcome of a random trial or (b) the position that a particular element gets mapped to in a permutation.

McDiarmid's Inequality[15] *Let X be a non-negative random variable determined by independent trials $T_1, ..., T_n$ and independent permutations $\Pi_1, ..., \Pi_m$. Suppose that for every set of possible outcomes of the trials and permutations, we have:*

(i) changing the outcome of any one trial can affect X by at most c;

(ii) interchanging two elements in any one permutation can affect X by at most c; and

(iii) for each $s > 0$, if $X \geq s$ then there is a set of at most rs choices whose outcomes certify that $X \geq s$.

Then for any $t \geq 50c\sqrt{r\mathbf{Exp}(X)}) + 256c^2r$ we have

$$\mathbf{Pr}(|X - \mathbf{Exp}(X)| > t) \leq 4e^{-\frac{t^2}{128c^2r(\mathbf{Exp}(X)+t)}}.$$

2.4 An overview

We begin by taking a dense decomposition of the graph. Our colouring procedure then proceeds in three phases.

In the first phase, we assign each vertex in S an independently and uniformly chosen colour from $\{1, ..., \Delta + 1\}$. For each D_i, we assign an independently and uniformly chosen permutation of colours to the colour classes $\mathcal{C}_i$. We then correct pairs of neighbours having the same colour by uncolouring some vertices in S and labelling some vertices in $\mathcal{D}$ as being only *temporarily coloured*. The same simple analysis described in the opening of this paper allows us to show that with positive probability, the resulting partial colouring is $20 \log \Delta / \log \log \Delta$-frugal. We also show that it has several other useful properties that bound the size of every $\text{Notbig}(i, x)$ and the number of temporarily coloured vertices in each D_i, and imply that every vertex not lying in a kernel will always have many available colours in future phases.

In the second phase, we recolour all the temporarily coloured vertices in the kernels of the ornery dense sets. We do so by swapping their colours with other vertices in the same dense set. It is very useful to know that a vertex with colour x has many choices for a vertex with which to swap. This follows from the bound on $\text{Notbig}(i, x)$, which implies that very few vertices have an external neighbour with colour x. Again, a fairly simple analysis shows that, with positive probability, no colour is assigned to $20 \log \Delta / \log \log \Delta$ vertices in any neighbourhood during this phase.

In the final phase, we colour the remaining vertices one-at-a-time. None of these vertices lie in kernels, and so at its turn, each vertex has a large list of available colours to choose from. We choose one at random. There is an annoying subtlety here: the manner in which we colour the vertices introduces too much dependence for us to apply the straightforward version of the Lovasz Local Lemma. So instead we use the Lopsided Local Lemma to show that, with positive probability, no colour is assigned to $4 \log \Delta / \log \log \Delta$ vertices in any neighbourhood during this phase.

This produces a colouring where no colour is assigned to $(20 + 20 + 4) \log \Delta / \log \log \Delta < 50 \log \Delta / \log \log \Delta$ vertices in any one neighbourhood, as required. (The astute reader will already see one way to reduce the constant "50".)

3 Phase I: An initial colouring

In this phase, we obtain an initial partial colouring using the following random procedure. All random choices are made independently.

1. We assign a uniformly random colour from $\{1, ..., \Delta + 1\}$ to each vertex $v \in S$.
2. For each D_i, we use $|\mathcal{C}_i|$ colours uniformly from $\{1, ..., \Delta + 1\}$ and then assign a random permutation of those colours to $\mathcal{C}_i$.
3. Let $\{(x_1, y_1), ..., (x_\ell, y_\ell)\}$ be the set of all pairs of neighbours or big-neighbours that are assigned the same colour. For each pair in that set, we choose one member, uniformly at random, to *correct*. To correct $v \in S$, we uncolour v. To correct $v \in \mathcal{D}$, we label v as being only *temporarily coloured*.

To clarify: if x, y are both in the same colour class, then they will both receive the same colour in Step 2, but possibly only one of them will be labelled as temporarily coloured in Step 3. We define:

- $U \subseteq S$ is the vertices of S that are uncoloured in Step 3;
- Temp_i is the set of vertices of D_i that are labelled as *temporarily coloured* in Step 3;
- Temp_i^* is the set of vertices $u \in D_i$ such that $\{u, v\} \in \mathcal{C}_i$ for some $v \in \mathrm{Temp}_i$;
- $\mathrm{Temp}_i^+ \subseteq \mathrm{Temp}_i^*$ is the set of vertices $u \in D_i$ such that $\{u, v\} \in \mathcal{C}_i$ for some $v \in \mathrm{Temp}_i$ with $|\mathrm{Out}_v| < |\mathrm{Out}_u|$;
- for each $0 \leq a \leq \Delta$, $\mathrm{Temp}_i(a)$ is the set $v \in \mathrm{Temp}_i$ with $|\mathrm{Out}_v| \leq a$;

- $\mathrm{Temp} = \cup_i \mathrm{Temp}_i$; $\mathrm{Temp}^* = \cup_i \mathrm{Temp}_i^*$; $\mathrm{Temp}^+ = \cup_i \mathrm{Temp}_i^+$.

Note that possibly $\mathrm{Temp} \cap \mathrm{Temp}^+ \neq \emptyset$. All vertices in Temp will be recoloured during Phases II and III. Vertices of Temp^* might also be recoloured during Phase II. Vertices of $\mathrm{Temp}^+ - \mathrm{Temp}$ might be moved to Temp during Phase II; those that are will be recoloured during Phase III.

For each ornery D_i, we will recolour the vertices in $\mathrm{Temp}_i \cap K_i$ during Phase II by swapping their colours with other vertices in D_i. To facilitate this, we carry out one more step:

4. For each ornery D_i, we select uniformly at random a set F_i of $\frac{9}{10}\Delta$ of the vertices of K_i that are colour classes of size one in $\mathcal{C}_i$.

The vertices of F_i will be eligible to swap their colours with the temporarily coloured vertices in K_i. We use F to denote the union over all ornery D_i of F_i.

Lemma 3.1 *With positive probability:*

(a) every $v \in S$ has at least $\frac{\epsilon}{10^9}\Delta$ colours that appear twice in $N(v) - (U \cup \mathrm{Temp} \cup \mathrm{Temp}^* \cup F)$;

(b) every $v \in \mathcal{D}$ with $|\mathrm{Out}_v| \geq \log^3 \Delta$ has at least $\frac{\epsilon}{10^9}|\mathrm{Out}(v)|$ colours that appear twice in $N(v) - (U \cup \mathrm{Temp} \cup \mathrm{Temp}^* \cup F)$;

(c) for each D_i and integer $a \in \{\lceil \log^3 \Delta \rceil, ..., \Delta\}$, we have $|\mathrm{Temp}_i(a)| \leq 2a$;

(d) for each vertex $v \in G$, no colour is assigned in Steps 1 and 2 to more than $20 \log \Delta / \log \log \Delta$ vertices in $N(v)$;

(e) for each colour x and ornery D_i, $|\mathrm{Notbig}(i, x)| \leq \Delta^{19/20}$;

(f) for each vertex $v \in G$,

$$\sum_{u \in N(v) \cap (\mathrm{Temp} \cup \mathrm{Temp}^+)} \frac{1}{\max(|\mathrm{Out}_u|, \log^3 \Delta)} \leq 299999.$$

Lemma 3.1 proves the existence of a partial colouring satisfying properties (a) to (f). For Phase I, we take such a colouring.

Proof We will use the Local Lemma. For each $v \in S$, $A_1(v)$ is the event that v violates part (a). For each $v \in \mathcal{D}$ with $|\mathrm{Out}_v| \geq \log^3 \Delta$, $A_2(v)$ is the event that v violates part (b). For each D_i and $a \geq \log^3 \Delta$, $A_3(i, a)$ is the event that D_i, a violates part (c). For each $v \in G$, $A_4(v)$ is the event that v violates part (d). For each colour x and ornery D_i,

$A_5(i,x)$ is the event that D_i, x violate part (e). For each $v \in G$, $A_6(v)$ is the event that v violates part (f).

We will prove that each event has probability at most Δ^{-8}. The events $A_1(v), A_2(v), A_4(v), A_6(v)$ are determined by the colour assignment and choice of F for dense sets and vertices that have neighbours or big-neighbours in $N(v)$. $A_3(i,a)$ and $A_5(i,x)$ are determined by the colour assignment and choice of F for dense sets and vertices that have neighbours or big-neighbours in D_i. It follows that each of these events is mutually independent of all events $A_1(u), A_2(u), A_3(j,b), A_4(u), A_5(j,y), A_6(u)$ where u and D_j are at distance at least 6 from v or D_i in the graph formed from G by adding edges between every pair of big-neighbours. (Note that, eg., $A_1(v)$ and $A_1(u)$ could be dependent for some u, v of distance 5 in that graph if there is a dense set that is adjacent to a neighbour of v and to a neighbour of u.) Noting that graph has maximum degree less than 2Δ by Corollary 2.4, and multiplying by the 6 types of events, and the at most $\Delta + 1$ choices of colours for $A_5(i,x)$ or values of a for $A_3(i,a)$, we find that each event is mutually independent of all but at most $6(2\Delta)^6 \times (\Delta+1)$ other events. This will imply our lemma since $6(2\Delta)^6(\Delta+1) \times \Delta^{-8} < 1/4$.

The random experiment: We will use McDiarmid's Inequality repeatedly. To do so, we view our random choices as a collection of independent trials and independent permutations as follows: In Step 1, we have a trial for each vertex $v \in S$. In Step 2, we carry out a permutation for each D_i: from the colours $\{1, ..., \Delta+1\}$ to the colour classes of $\mathcal{C}_i$ plus $\Delta + 1 - |\mathcal{C}_i|$ artificial classes. In Step 3, we carry out a trial for each pair of neighbours or big-neighbours which have the same colour. To avoid dependency, we will actually carry out this trial for *every* pair of neighbours or big-neighbours and ignore the outcome for those pairs which do not have the same colour; thus the set of trials is not dependent on the outcomes of Steps 1 and 2. In Step 4, we carry out a permutation for each ornery D_i: from the colour classes of size one in K_i to the integers $\{1, ...\}$; F_i is the set of vertices mapped to $\{1, ..., \frac{1}{10}\Delta\}$. Note that these steps are equivalent to the manner in which the procedure was described above.

In the context of McDiarmid's Inequality, a **"choice"** is either: the colour assigned to a vertex or colour class in Step 1 or Step 2, the vertex chosen to be corrected in Step 3, or the choice as to whether a particular vertex is placed into F_i in Step 4. **"Changing the outcome of a trial"** is changing the colour assigned to a vertex in S during Step 1 or changing which of a pair of vertices is corrected in Step 3. **"Interchanging two elements of a permutation"** is swapping the colours assigned to two colour classes in some $\mathcal{C}_i$ during Step 2 or removing a vertex from F_i and replacing it with another.

$A_1(v)$**:** Let X be the number of colours α satisfying: (a) α is assigned to at least two vertices in $N(v)$ and (b) no vertex in $N(v)$ that is assigned α is in $U \cup \text{Temp} \cup \text{Temp}^* \cup F$. Clearly X is a lower bound on the number of colours that appear at least twice in $N(v) - (U \cup \text{Temp} \cup \text{Temp}^* \cup F)$.

By Lemma 2.5(a), v has at least $\frac{\epsilon}{80}\Delta^2$ pairs of strongly non-adjacent neighbours. X is at least the number of such pairs u, w such that: (i) u, w are assigned the same colour; (ii) that colour is not assigned to any neighbours of v outside of the colour classes containing u, w; (iii) no vertices in the colour classes containing u, w are corrected in Step 3; and (iv) u, w are not selected to be in F.

The probability that some non-adjacent pair $u, w \in N(v)$ satisfies (i) is $\frac{1}{\Delta+1}$. The probability that (ii) is satisfied is easily seen to be at least $\frac{1}{4} \times \frac{3}{4} = \frac{3}{16}$; the extreme case is when v has nearly $\frac{3}{4}\Delta$ neighbours in one dense set and nearly $\frac{1}{4}\Delta$ neighbours in another. If D_i is ornery then by Lemma 2.3(a), the definition of "ornery" and the way the colour classes were formed, there are at least $\Delta - \log^4 \Delta - 3\log^5 \Delta - 2\lceil 10\epsilon\Delta \rceil > \frac{99}{100}\Delta$ vertices eligible to be chosen for F_i. So u, w satisfy (iv) with probability at least $\left(1 - \frac{9\Delta/10}{99\Delta/100}\right)^2 > \frac{1}{150}$.

Let Ψ be the set of neighbours and big-neighbours of the colour classes containing u, w, not including vertices that lie in the same dense set as one of those classes. By Lemma 2.2(b) and Corollary 2.4, $|\Psi| \leq 2 \times (\Delta + \Delta^{1/4} \log^7 \Delta)$ (the extreme case is when $u, w \in S$). So at most two dense sets contain more than $\frac{3}{4}\Delta$ vertices of Ψ. The colour of u, w can be assigned to at most two vertices of Ψ in each of these dense sets. If it is assigned to four such vertices the probability that this does not cause u or w to be corrected is at least $(\frac{1}{2})^8$ as these 4 vertices yield at most 4×2 conflicting pairs involving u or w. That colour is not assigned to any other member of Ψ with probability at least $\frac{1}{2} \times (\frac{1}{4})^2 - o(1)$; the extreme case is when two dense sets each contain $\frac{3}{4}\Delta$ members of Ψ and another contains the remaining $\frac{1}{2}\Delta + o(\Delta)$ members. Therefore, the probability that (iii) is satisfied is at least $(\frac{1}{2})^8 \times (\frac{1}{2})^5 - o(1) > \frac{1}{10^4}$.

Therefore,

$$\mathbf{Exp}(X) \geq \frac{\epsilon}{80}\Delta^2 \times \frac{1}{\Delta+1} \times \frac{3}{16} \times \frac{1}{10^4} \times \frac{1}{150} > \frac{\epsilon}{8 \times 10^8}\Delta.$$

So if $A_1(v)$ holds then X must differ from its mean by at least $\frac{\epsilon}{4 \times 10^9}\Delta$.

We apply McDiarmid's Inequality to show that X is highly concentrated. To do so, we consider two related variables: X_1 is the number of colours assigned to at least 2 vertices in $N(v)$; X_2 is the number of colours assigned to at least 2 vertices in $N(v)$ where at least one of the vertices is in $U \cup \mathrm{Temp} \cup F$. Note that $X = X_1 - X_2$. Thus, if $A_1(v)$ holds then either X_1 or X_2 must differ from its mean by at least $\frac{\epsilon}{8 \times 10^9}\Delta$.

Trivially, $\mathbf{Exp}(X_1) \leq \Delta$. For any s, if $X_1 \geq s$ then there is a set of $2s$ choices whose outcomes certify that $X_1 \geq s$, namely the colour assignments to s pairs of vertices. Changing the colour of a vertex or colour class can only affect X_1 by at most 2 since at worse it affects whether X_1 counts the old colour and the new colour. Similarly, changing the outcome of any other trial or interchanging two elements in any permutation can affect X_1 by at most 2. Therefore McDiarmid's Inequality with $c = r = 2$ yields:

$$\mathbf{Pr}(|X_1 - \mathbf{Exp}(X_1)| > \frac{\epsilon}{8 \times 10^9}\Delta) \leq 4\mathrm{e}^{-\left(\frac{\epsilon}{8 \times 10^9}\right)^2 \Delta^2/(128 \times 8 \times (2\Delta))} < \frac{1}{2}\Delta^{-8}.$$

Similarly: $\mathbf{Exp}(X_2) \leq \Delta$. If $X_2 \geq s$ then there is a set of at most $4s$ choices whose outcomes certify that $X_2 \geq s$: the colour assignment for s pairs of vertices and for each pair: either the choice to put one of the vertices in F or the assignment of the same colour to a neighbour or big-neighbour of one vertex and the choice to correct the vertex. Also, changing the outcome of a trial or interchanging two elements of a permutation can affect X_2 by at most 2 by the same reason as for X_1. Therefore, McDiarmid's Inequality with $c = 2, r = 4$ implies that

$$\mathbf{Pr}(|X_2 - \mathbf{Exp}(X_2)| > \frac{\epsilon}{8 \times 10^9}\Delta) < \frac{1}{2}\Delta^{-8}.$$

Therefore, $\mathbf{Pr}(A_1(v)) \leq \Delta^{-8}$ as required.

$A_2(v)$**:** The same argument as for $A_1(v)$, this time using Lemma 2.5(b) and the fact that $\frac{1}{10}\Delta|\text{Out}_v| > \frac{\epsilon}{80}\Delta|\text{Out}_v|$ yields that $\mathbf{Pr}(A_2(v)) \leq \Delta^{-8}$. The main difference is that we are bounding the probabilities that X_1, X_2 differ from their means by at least $\frac{\epsilon}{8\times 10^9}|\text{Out}_v| \geq \frac{\epsilon}{8\times 10^9}\log^3\Delta$ rather than $\frac{\epsilon}{8\times 10^9}\Delta$. Nevertheless, the bounds obtained are easily seen to be less than $\frac{1}{2}\Delta^{-8}$.

$A_3(i, a)$**:** We first bound $\mathbf{Exp}(|\text{Temp}_i(a)|)$. For each vertex $v \in D_i$ with outdegree at most a, and each external neighbour u of v, the probability that u gets the same colour as v is $1/(\Delta + 1)$. So the probability that at least one external neighbour of v gets the same colour is at most $a/(\Delta + 1)$. So $\mathbf{Exp}(|\text{Temp}_i(a)|) \leq |D_i| \times a/(\Delta + 1) \leq (1 + 2\epsilon)a$ by Lemma 2.1(a).

Now we apply McDiarmid's Inequality to show that $|\text{Temp}_i(a)|$ is highly concentrated. For any s, if $|\text{Temp}_i(a)| \geq s$ then the colours assigned to s members of $\text{Temp}_i(a)$ and to a neighbour of each of them, plus the choices of which of these vertices to correct, will be a set of at most $3s$ trials that certify this fact. Changing the colour of a vertex in S can affect $|\text{Temp}_i(a)|$ by at most 2, since each colour is assigned to at most two vertices in D_i. Similarly, changing the choice of which of a pair of vertices is corrected can affect $|\text{Temp}_i(a)|$ by at most one, and swapping the colours on two members of some D_j can affect $|\text{Temp}_i(a)|$ by at most 4. So McDiarmid's Inequality with $c = 4, r = 3$ yields:

$$\mathbf{Pr}(A_3(i, a)) \leq \mathbf{Pr}(|\text{Temp}_i(a)| - \mathbf{Exp}(|\text{Temp}_i(a)|) > \frac{1}{2}a) \leq 4\mathrm{e}^{-\frac{1}{4}a^2/(128\times 48\times(2a))} < \Delta^{-8},$$

since $a \geq \log^3\Delta$.

$A_4(v)$**:** Consider any particular colour x, and set $t = 20\log\Delta/\log\log\Delta$. If $A_4(v)$ holds, then x is assigned to at least $\frac{1}{2}t$ colour classes that intersect $N(v)$. For any set of $\frac{1}{2}t$ colour classes, the probability that x is assigned to every member of X is at most $(\frac{1}{\Delta+1})^{\frac{1}{2}t}$. Multiplying this by the number of choices for x and the colour classes, yields:

$$\mathbf{Pr}(A_4(v)) \leq (\Delta{+}1)\binom{\Delta}{\frac{1}{2}t}(\Delta{+}1)^{-\frac{1}{2}t} \leq \Delta\left(\frac{2e}{t}\right)^{t/2} < \Delta(\log\Delta/\log\log\Delta)^{-10\log\Delta/\log\log\Delta} < \Delta^{-8}.$$

$A_5(i,x)$**:** We first bound $\mathbf{Exp}(|\text{Notbig}(i,x)|)$. For any $v \in D_i$, the probability that $v \in \text{Notbig}(i,x)$ is at most $|\text{Out}_v|/(\Delta+1)$. By Lemma 2.3(b), the sum of the external degrees of the vertices of D_i is at most $\Delta \log^7 \Delta$. So $\mathbf{Exp}(|\text{Notbig}(i,x)|) \leq \Delta \log^7 \Delta/(\Delta+1) < \log^7 \Delta$ and if $A_5(i,x)$ holds then $|\text{Notbig}(i,x)| - \mathbf{Exp}(|\text{Notbig}(i,x)|) > \frac{1}{2}\Delta^{19/20}$.

Now we apply the Hoeffding-Azuma Inequality. For this bound, it is convenient to regard $\text{Notbig}(i,x)$ as being determined by the following independent choices: for each $v \in S$, the choice of whether to assign x to v and for each $D_j, j \neq i$, the choice of which (if any) colour class in $\mathcal{C}_j$ is assigned x. For each vertex $u \notin D_i \cup \text{Big}_i$, we let $d_i(u)$ be the number of neighbours that u has in D_i. For each colour class α, we let $d_i(\alpha)$ be the sum over $u \in \alpha \backslash \text{Big}_i$ of $d_i(u)$. For each $u \in S$, changing the colour assignment for u can only affect whether the neighbours of u are in $\text{Notbig}(i,x)$ and so will affect $|\text{Notbig}(i,x)|$ by at most $d_i(u)$. Similarly, for each $j \neq i$, changing which colour class of $\mathcal{C}_j$ gets colour x can affect $|\text{Notbig}(i,x)|$ by at most $\max_{\alpha \in \mathcal{C}_j} d_i(\alpha)$. The sum of the squares of these maximum effects is at most $\sum_\alpha d_i(\alpha)^2$. By Lemma 2.3(b), $\sum_\alpha d_i(\alpha) \leq \Delta \log^7 \Delta$, and we have $d_i(\alpha) \leq 2\Delta^{7/8}$ since $|\alpha| \leq 2$ and vertices of Big_i don't contribute to $\delta_i(\alpha)$. Therefore, $\sum_u d_i(u)^2 \leq \frac{1}{2}\Delta^{1/8}\log^7 \Delta \times (2\Delta^{7/8})^2 = 2\Delta^{15/8}\log^7 \Delta$. So the Hoeffding-Azuma Inequality with $t = \frac{1}{2}\Delta^{19/20}$ yields:

$$\mathbf{Pr}(A_5(i,x)) \leq 2\mathrm{e}^{-\frac{(\frac{1}{2}\Delta^{19/20})^2}{2\Delta^{15/8}\log^7 \Delta}} < \Delta^{-8}.$$

$A_6(i,x)$**:** Set

$$\begin{aligned} Y &= \sum_{u \in N(v) \cap \text{Temp}} \frac{1}{\max(|\text{Out}_u|, \log^3 \Delta)} \\ Y^+ &= \sum_{u \in N(v) \cap \text{Temp}^+} \frac{1}{\max(|\text{Out}_u|, \log^3 \Delta)} \end{aligned}$$

If $u \in N(v)$ is in Temp then at least one member of Out_u is assigned the same colour as u. So $\mathbf{Pr}(u \in \text{Temp}) \leq |\text{Out}_u|/(\Delta+1)$. Therefore $\mathbf{Exp}(Y) \leq \sum_{u \in N(v)} \frac{|\text{Out}_u|}{\Delta+1} \times \frac{1}{|\text{Out}_u|} < 1$. If $u \in \text{Temp}^+$ then at least one member of $\text{Out}_{u'}$ is assigned the same colour as u' for the at most one u' with $\{u,u'\} \in \mathcal{C}_i$. We also must have $|\text{Out}_{u'}| < |\text{Out}_u|$ and so the same calculation shows $\mathbf{Exp}(Y^+) < 1$.

We begin by showing Y is concentrated. Again, we apply McDiarmid's Inequality. Doing so is a bit tricky, because single colour assignments can possibly have a very large affect on Y; there are two issues of this sort: The first is that it is possible for all neighbours of v to receive the same colour and then the assignment to a single vertex which has the same neighbourhood as v would cause all neighbours of v to enter Temp. To eliminate this unlikely situation, we consider a set $\text{Temp}' \subseteq \text{Temp}$ defined to contain every vertex $u \in \mathcal{D}$ such that (1) u has a neighbour or big-neighbour w which receives the same colour as u; (2) at most $\log \Delta$ neighbours or big-neighbours of w receive that colour; and (3) u is the vertex from

$\{u,w\}$ chosen to be corrected. We then define

$$Y' = \sum_{u \in N(v) \cap \mathrm{Temp}'} \frac{1}{\max(|\mathrm{Out}_u|, \log^3 \Delta)}.$$

Note that condition (2) in the definition of Temp$'$ implies that the colour assignment to any one vertex can cause at most $\log \Delta$ neighbours or big-neighbours of v to enter Temp$'$. We will show that Y' is highly concentrated. This will be sufficent because $Y' \le Y$ so $\mathbf{Exp}(Y') \le \mathbf{Exp}(Y)$ and because:

Claim: $\mathbf{Pr}(Y \ne Y') < \Delta^{-9}$.

Proof: By Corollary 2.4, at most $\Delta \times (\Delta + \Delta^{1/4} \log^7 \Delta) < 2\Delta^2$ vertices are neighbours or big-neighbours of a neighbour of v. If $Y \ne Y'$ then at least one of those vertices has at least $\log \Delta$ neighbours that receive the same colour. The probability that some set of t colour classes all receive the same colour is at most $(\Delta+1)^{-(t-1)}$ (it is zero if two of them lie in the same dense set). If $Y \ne Y'$ then this occurs for $t \ge \frac{1}{2}\log \Delta$ colour classes that intersect the neighbourhood of that vertex, so

$$\mathbf{Pr}(Y \ne Y') \le 2\Delta^2 \binom{\Delta}{\frac{1}{2}\log \Delta} (\Delta+1)^{-(\frac{1}{2}\log \Delta - 1)} \le 2\Delta^3 \left(\frac{2e}{\log \Delta}\right)^{\frac{1}{2}\log \Delta} < \Delta^{-9}.$$

The other issue we need to deal with is the fact that different vertices can contribute very different amounts to Y', since the sizes of their external neighbourhoods can vary greatly. This creates difficulties when applying McDiarmid's Inequality directly. So instead, we break Y' up into several sums taken over neighbours of v that have external neighbourhoods of similar size.

We define $I_0 = \{0, ..., 2\log^3 \Delta - 1\}$ and we define $I_i = \{2^i \log^3 \Delta, ..., (2^{i+1}) \log^3 \Delta - 1\}$, for each of the roughly $\log \Delta$ values of $i \ge 2$ for which this interval contains some values up to $\frac{1}{3}\Delta$, the maximum possible external degree of a vertex (by Lemma 2.2(b)). For each i, we define $N_i \subseteq N(v) - S$ to be the neighbours of v which have external degree in I_i, and we define

$$Y_i' = \sum_{u \in N_i \cap \mathrm{Temp}'} \frac{1}{\max(|\mathrm{Out}_u|, \log^3 \Delta)}.$$

We will now apply McDiarmid's Inequality to each Y_i'. If $Y_i' \ge s$, then there is a set of at most $3 \times 2^{i+1} \log^3 \Delta \times s$ trials that certify this fact: the assignments to at most $(2^{i+1} \log^3 \Delta)s$ members of $N_i \cap \mathrm{Temp}'$ (as each contributes more than $1/(2^{i+1} \log^3 \Delta)$ to Y_i'), the assignment to a neighbour or big-neighbour of each of those vertices and the choice to correct each of those vertices. Changing the colour assignment to any one colour class or changing the choice of whether to correct a vertex will change Y_i' by at most $\log \Delta / (2^i \log^3 \Delta) = 2^{-i} \log^{-2} \Delta$. (The greatest change is if the colour assignment causes $\log \Delta$ members of N_i to enter Temp_i', each of whom have external degree $2^i \log^3 \Delta$.) Similarly, switching the colours of two colour classes in some $\mathcal{C}_j$ can affect Y_i' by at most $2 \times 2^{-i} \log^{-2} \Delta$. We set $t_i = 30{,}000 \times 2^{-i/2}$, and note that $1 + t_i \leq 30,001$. Note that $\mathbf{Exp}(Y_i') \leq \mathbf{Exp}(Y) \leq 1$. McDiarmid's Inequality with $c = 2 \times 2^{-i} \log^{-2} \Delta$ and $r = 3 \times 2^{i+1} \log^3 \Delta$ yields:

$$\mathbf{Pr}(Y_i' - \mathbf{Exp}(Y_i') > t_i) \leq 4e^{-t_i^2/(128 \times (2 \times 2^{-i} \log^{-2} \Delta)^2 \times 3 \times 2^{i+1} \log^3 \Delta (1+t_i))} < 4e^{-\frac{3 \times 10^8}{8 \times 128(1+t_i)} \log \Delta} < \Delta^{-9.5}.$$

Since there are fewer than $\log \Delta$ intervals, the probability that $Y_i' - \mathbf{Exp}(Y_i) \leq t_i$ for all i is at least $1 - \log \Delta \times \Delta^{-9.5} > 1 - \Delta^{-9}$. If this happens, then

$$Y' - \mathbf{Exp}(Y') < \sum_{i \geq 0} t_i = 30,000 \times \frac{1}{1 - 2^{-1/2}} < 149,999.$$

This implies that $\mathbf{Pr}(Y \geq 149,999) < \frac{1}{2}\Delta^{-9}$. A nearly identical argument shows that $\mathbf{Pr}(Y^+ \geq 149,999) < \frac{1}{2}\Delta^{-9}$. Therefore:

$$\mathbf{Pr}(A_6(i,x)) \leq \Delta^{-9}.$$

$\square$

4 Phase II: The kernels of the ornery dense sets

In this phase, we colour all vertices in kernels that have the same colour as a neighbour. I.e., we will colour the vertices of $\text{Temp}_i(\log^6 \Delta)$ for each ornery D_i. Note that, by Lemma 3.1(c), there are at most $2\log^6 \Delta$ such vertices in any ornery dense set.

We use $\gamma(w)$ to denote the colour of a vertex w at the beginning of this phase; i.e. the colour, if any, that it had at the end of Phase I.

Consider any ornery D_i and any $v \in \text{Temp}_i(\log^6 \Delta)$. We will recolour v by swapping its colour with a vertex in F_i. If $\{v, v'\}$ is a colour class of size two in $\mathcal{C}_i$, and if $v' \in K_i$ then we swap its colour with the same vertex that v swaps with; we refer to v' and v as *swapping partners*. Note that, in this case, $v' \in \text{Temp}^*$. But if $v' \notin K_i$, then we leave the colour of v' as $\gamma(v)$, and we place v' into Temp_i since v' might now conflict with the vertex that v swaps with. Note that, in this case, $v' \in \text{Temp}^+$.

For each $v \in \text{Temp}_i(\log^6 \Delta)$, we define Swappable_v to be the set of vertices $u \in F_i$ such that:

(i) u is a colour class of size one in $\mathcal{C}_i$;

(ii) $u \notin \text{Temp}_i \cup \text{Temp}_i^*$;

(iii) $\gamma(v)$ does not appear on $\text{Out}_u - \text{Temp}$;

(iv) $\gamma(u)$ does not appear on $\text{Out}_v - \text{Temp}$;

(v) $\gamma(u)$ does not appear on $\mathrm{Out}_{v'} - \mathrm{Temp}$ for a swapping partner v' of v.

Lemma 4.1 *For each ornery D_i and each $v \in \mathrm{Temp}_i(\log^6 \Delta)$, $|\mathrm{Swappable}_v| \geq \frac{1}{10}\Delta$.*

Proof The size of F_i is $\frac{9}{10}\Delta$. By Lemma 2.3(c), at most $2\log^5 \Delta$ members of F_i violate condition (i). Since $F_i \subseteq K_i$, every member of F_i has at most $\log^6 \Delta$ external neighbours. Therefore by Lemma 3.1(c), at most $2\log^6 \Delta$ members of F_i are in Temp_i, and condition (i) implies none of them are in Temp_i^*. Thus at most $2\log^6 \Delta$ members of F_i violate condition (ii). Since $|\mathrm{Out}_v| \leq \log^6 \Delta$, at most $\log^6 \Delta$ members of F_i violate condition (iv). Similarly, at most $\log^6 \Delta$ members of F_i violate condition (v).

If two big-neighbours received the same colour in Phase I, then one was corrected; i.e. either uncoloured or placed into Temp. Therefore $\gamma(v)$ appears on at most one member of $\mathrm{Big}_i - \mathrm{Temp}$. This fact, along with Lemmas 3.1(e) and 2.1(d), implies that at most $\frac{3}{4}\Delta + \Delta^{19/20}$ members of F_i violate condition (iii). Since $2\log^5 \Delta + 2\log^6 \Delta + 2\log^6 \Delta + \frac{3}{4}\Delta + \Delta^{19/20} < \frac{8}{10}\Delta$, the lemma follows. □

Each $v \in \mathrm{Temp}_i(\log^6 \Delta)$ will select 20 members of $\mathrm{Swappable}_v$ uniformly at random. These 20 vertices will be denoted the *candidates* for v. (Swapping partners select the same set of candidates.) Each such v will swap its colour with one of its candidates. $\mathrm{Swappable}_v$ was defined to ensure that making a single swap will not create a conflict. But we need to be careful to ensure that conflicts are not created by making multiple swaps. We define a candidate u of v to be *bad* if:

(i) u is a candidate of another vertex;

(ii) v or a swapping partner of v has an external neighbour w that has a candidate w' with $\gamma(w') = \gamma(u)$;

(iii) v or a swapping partner of v has an external neighbour w that is a candidate for exactly one vertex w' where $\gamma(w') = \gamma(u)$;

(iv) u has an external neighbour w that has a candidate w' with $\gamma(w') = \gamma(v)$; or

(v) u has an external neighbour w that is a candidate for exactly one vertex w' where $\gamma(w') = \gamma(v)$.

A *candidate* u of v is *good* if it is not bad.

Lemma 4.2 *With positive probability:*

(a) For each ornery D_i, every vertex in $\mathrm{Temp}_i(\log^6 \Delta)$ has a good candidate.

(b) For each vertex $v \in G$ and each colour x, at most $20\log\Delta/\log\log\Delta$ neighbours of v have a candidate with colour x or are a candidate of a vertex with colour x.

We present the proof below. First we note how this lemma enables us to complete Phase II:

Lemma 4.2 proves the existence of a collection of candidates satisfying properties (a) and (b); we take such a collection. For each $v \in \text{Temp}_i$, we swap the colour of v and any swapping partner of v with that of one of v's good candidates. Property (a) ensures that we can do so. If $\{v, v'\} \in \mathcal{C}_i$ and $v' \notin K_i$ then we place v' in Temp_i.

We remove from Temp all vertices that were successfully coloured; i.e. all those vertices of Temp that are in kernels of ornery sets. Note that no vertices in kernels are added to Temp during this phase. Note also that every vertex whose colour changed during this phase was in $\text{Temp} \cup \text{Temp}^* \cup F$ and that every vertex placed into Temp was in $\text{Temp}^+ \subseteq \text{Temp}^*$.

Our definition of *good* ensures that we have a proper partial colouring on all the vertices outside of $U \cup \text{Temp}$. Property (b) ensures that for every vertex $w \in G$ and each colour x, at most $20 \log \Delta / \log \log \Delta$ neighbours of w are given the colour x during Phase II.

Proof of Lemma 4.2 We will apply the Lovasz Local Lemma. For every ornery D_i and every vertex $v \in \text{Temp}_i(\log^6 \Delta)$, we define $A_1(v)$ to be the event that v does not have a good candidate. For every vertex $v \in G$ and every colour x, we define $A_2(v, x)$ to be the event that v, x violate (b). We will prove that each of these events has probability less than Δ^{-8}.

Each event is determined only by candidate choices in dense sets that are adjacent to v or to another vertex in the same dense set as v. It follows that each event is mutually independent of all but $2\Delta^5 \times (\Delta + 1)$ other events. (The extra $\Delta + 1$ term is for the number of choices of x.) Since $2\Delta^5(\Delta + 1) \times \Delta^{-8} < 1/4$, the Local Lemma shows that with positive probability, none of these events occur.

$A_1(v)$**:** Consider some $v \in \text{Temp}_i(\log^6 \Delta)$ where D_i is ornery; we will bound $\mathbf{Pr}(A_1(v))$. We first choose the candidates for all vertices other than v. Then we let Bad denote the set of vertices in Swappable_v which would be bad candidates for v. We will show that with high probability, $|\text{Bad}| < \log^{13} \Delta$.

By Lemma 3.1(c), $|\text{Temp}_i(\log^6 \Delta)| < 2 \log^6 \Delta$. Therefore, at most $20 \times 2 \log^6 \Delta$ members of Swappable_v meet condition (i) of the definition of *bad*. v has at most $\log^6 \Delta$ external neighbours, and if v has a swapping partner then it also has at most $\log^6 \Delta$ external neighbours. Each of those external neighbours has at most 20 candidates and so at most $40 \log^6 \Delta$ members of Swappable_v meet condition (ii) or (iii) of the definition of *bad*.

Let W be the set of vertices in kernels of other ornery dense sets that have external neighbours in Swappable_v; if two swapping partners are both in W then we remove one of them. Since every member of Swappable_v is in K_i, $|W| \leq |\text{Swappable}_v| \times \log^6 \Delta < \Delta \log^6 \Delta$. Each member of W selects a candidate of colour $\gamma(v)$ with probability at most $20/(\Delta/10) = 200/\Delta$, and these choices are independent. So the probability that more than $600 \log^6 \Delta$ members of W do so is at most

$$\binom{\Delta \log^6 \Delta}{600 \log^6 \Delta} \left(\frac{200}{\Delta}\right)^{600 \log^6 \Delta} < \left(\frac{e\Delta \log^6 \Delta}{600 \log^6 \Delta} \times \frac{200}{\Delta}\right)^{600 \log^6 \Delta} < \Delta^{-9}.$$

If at most $600\log^6\Delta$ members of W select a candidate of colour $\gamma(v)$ then, along with possibly one swapping partner for each of them, at most $1200\log^6\Delta$ vertices in kernels that have external neighbours in Swappable_v do so. Since each of these vertices is in a kernel, it has at most $\log^6\Delta$ neighbours in Swappable_v. Therefore, the above calculations show the probability that more than $1200\log^6\Delta\times\log^6\Delta = 1200\log^{12}\Delta$ members of Swappable_v meet condition (iv) of the definition of *bad* is at most Δ^{-9}.

The proof of the analogous fact for condition (v) is nearly identical: each member of W is a candidate of the at most one colour class in its dense set of colour $\gamma(v)$ with probability at most $200/\Delta$; however, this time the events are not independent. Fortunately, the dependency goes in the right direction and so for every $w_1,...,w_t \in W$, the probability that $w_1,...,w_t$ are all candidates of vertices with colour $\gamma(v)$ is at most $(200/\Delta)^t$. Thus the probability that there are at least $t = 600\log^6\Delta$ such vertices in W is at most $\binom{|W|}{t}\times(200/\Delta)^t < \Delta^{-9}$. So the probability that more than $600\log^{12}\Delta$ members of Swappable_v meet condition (v) of the definition of *bad* is at most Δ^{-9}.

Therefore with probability at least $1-2\Delta^{-9}$ we have $|\text{Bad}| < 80\log^6\Delta + 1800\log^{12}\Delta < \log^{13}\Delta$. If that bound on $|\text{Bad}|$ holds, then the probability that v has no good candidates, i.e. that it only selects candidates from Bad, is at most

$$\left(\frac{|\text{Bad}|}{|\text{Swappable}_v|}\right)^{20} < \left(\frac{\log^{13}\Delta}{\Delta/10}\right)^{20} < \Delta^{-9}.$$

Therefore, $\mathbf{Pr}(A_1(v)) < 3\Delta^{-9} < \Delta^{-8}$.

$A_2(v,x)$**:** Consider any $v \in G$ and any colour x. Set $t = 20\log\Delta/\log\log\Delta$. Similar analysis to that used above for conditions (iv) and (v) show that a particular member of $N(v)$ has a candidate of colour x or is a candidate of a vertex with colour x with probability at most $2\times(20/\frac{\Delta}{10}) = \frac{400}{\Delta}$. For $A_2(v,x)$ to hold, then this must occur for at least $\frac{1}{2}t$ vertices or pairs of swapping partners in $N(v)$. Again, any dependency goes in the right direction and so:

$$\mathbf{Pr}(A_2(v,x) \le \binom{\Delta}{t/2}\times\left(\frac{400}{\Delta}\right)^{t/2} < \left(\frac{e\Delta}{t/2}\times\frac{400}{\Delta}\right)^{t/2} = \left(\frac{40e\log\log\Delta}{\log\Delta}\right)^{10\log\Delta/\log\log\Delta} < \Delta^{-8}.$$

□

5 Phase III: Completing the colouring

In this phase, we complete the colouring of G by assigning colours to every vertex in $U\cup\text{Temp}$. We use the following simple random procedure. At any point, we use $L(u)$ to denote the set of colours that do not appear on any neighbours of u. We are no longer concerned about avoiding colours that appear on big-neighbours of u.

1. Uncolour every vertex in Temp.
2. Let $v_1, ..., v_\ell$ be an ordering of the uncoloured vertices (i.e. $U \cup \text{Temp}$) such that the vertices of Temp appear in non-decreasing order of $|\text{Out}_{v_i}|$.
3. For $i = 1$ to ℓ, assign to v_i a colour chosen uniformly at random from $L(v_i)$.

Of course, we need to know that there will always be at least one colour available for each v_i; in fact, there will be many. We define:

- for each $v \in U$, $Q(v) = \frac{\epsilon}{10^9}\Delta$;
- for each $v \in \text{Temp}$, $Q(v) = \frac{\epsilon}{10^9}\max\{|\text{Out}_v|, \log^3 \Delta\}$.

Lemma 5.1 *(a) when we colour $v \in U \cup \text{Temp}$, we have $|L(v)| \geq Q(v)$;*

(b) for each vertex $v \in G$,

$$\sum_{u \in N(v) \cap (U \cup \text{Temp})} \frac{1}{Q(u)} \leq \frac{3 \times 10^{14}}{\epsilon}.$$

Proof Lemma 3.1(a) and the fact that only vertices in $\text{Temp} \cup \text{Temp}^* \cup F$ are recoloured during Phases II and III imply: At the end of Phase I, each $v \in S$ had at least $\frac{\epsilon}{10^9}\Delta$ colours that appeared twice in its neighbourhood on vertices whose colours do not change in subsequent phases. Since the total number of colours is greater than the degree of v, this implies part (a) for the case $v \in U$.

The case $v \in \text{Temp}$ follows for vertices with $|\text{Out}_v| \geq \log^3 \Delta$ from applying Lemma 3.1(b) in the same manner. If a vertex $v \in \text{Temp}_i$ has $|\text{Out}_v| < \log^3 \Delta$, then D_i is not ornery as otherwise v would have been coloured in Phase II. Therefore, by the definition of ornery, $|\mathcal{C}_i| < \Delta - \log^4 \Delta$. After uncolouring all vertices in Temp, each colour class in $\mathcal{C}_i$ contains at most one colour. Every vertex w coloured before v has $|\text{Out}_w| \leq |\text{Out}_v| < \log^3 \Delta$. Lemma 3.1(c) and the fact that all vertices added to Temp in Phase II have more than $\log^3 \Delta$ external neighbours imply that there are at most $2\log^3 \Delta$ such vertices w. Therefore, when we colour v, there are at most $|\mathcal{C}_i| + |\text{Out}_v| + 2\log^3 \Delta < \Delta - \log^3 \Delta$ colours appearing in $N(v)$ and so $|L(v)| > Q(v) = \frac{\epsilon}{10^9}\log^3 \Delta$.

For each $u \in U$: since $Q(u) = \frac{\epsilon}{10^9}\Delta$ $\sum_{u \in N(v) \cap U} \frac{1}{Q(u)} \leq \frac{10^9}{\epsilon}$. Lemma 3.1(f) and the fact that only vertices of Temp^+ entered Temp during Phase II imply that at the beginning of Phase III: $\sum_{u \in N(v) \cap \text{Temp}} \frac{1}{Q(u)} \leq \frac{10^9}{\epsilon} \times 299999$. This yields part (b). □

Lemma 5.1(a) guarantees that each vertex will always have an available colour, and so our procedure will succeed in producing a proper colouring. Furthermore, it proves that the probability of v receiving a particular colour x is at most $10^9/(\epsilon\Delta)$ for $v \in U$ and at most

$10^9/(\epsilon|\mathrm{Out}_v|)$ for $v \in \mathrm{Temp}$. Part (b) implies that for each $v \in G$, the expected number of neighbours of v to receive x is at most a constant. This will allow us to prove that with high probability, the number of neighbours to receive x is sufficiently low.

Lemma 5.2 *With positive probability: for each vertex $v \in G$ and each colour x, at most $4\log\Delta/\log\log\Delta$ neighbours of v are assigned x during Phase III.*

Lemma 5.2 proves the existence of a colouring of the remaining vertices in which no colour is assigned to more than $4\log\Delta/\log\log\Delta$ vertices in any neighbourhood; For Phase III, we take such a colouring. This, along with property (a) from Lemma 3.1 and Property (b) from Lemma 4.2, ensures that no colour appears more than $50\log\Delta/\log\log\Delta$ times in the neighbourhood of a vertex in our overall colouring of G. This proves Theorem 1.1.

Proof of Lemma 5.2: We will apply the Lopsided Local Lemma. For each vertex $v \in G$, we define $A(v)$ to be the event that there is a colour x which is assigned to at least $t = 4\log\Delta/\log\log\Delta$ neighbours of v. We define $N^2(v)$ to be the set of vertices of distance at most 2 from v and we define $B(v) = \{A(u) : u \in N^2(v)\}$. Note that $|B(v)| \leq \Delta^2$. We will prove that for any collection of events outside of $B(v)$, conditioning on none of them occurring will result in the conditional probability of $A(v)$ being at most $\frac{1}{4}\Delta^2$. The Lopsided Local Lemma then proves Lemma 5.2.

To prove the desired bound on the conditional probabilities, we actually prove something stronger, but conceptually a bit simpler. First, for each $v \in U \cup \mathrm{Temp}$, we define:

- $L_0(v)$ is the set of colours not appearing on any neighbours of v at the beginning of Phase III.

Claim 1: For every $u \in (U \cup \mathrm{Temp}) - N(v)$, choose any colour $c(u) \in L_0(v)$ such that for every adjacent u_1, u_2 we have $c(u_1) \neq c(u_2)$. Conditioning on the event that each such u is assigned $c(u)$ during Phase III, the conditional probability of $A(v)$ is at most $\frac{1}{4}\Delta^2$.

By observing that all events outside of $B(v)$ are completely determined by the colours assigned to $(U \cup \mathrm{Temp}) - N(v)$, it is straightforward to show that Claim 1 will imply the condition required for our application of the Lopsided Local Lemma. Indeed: let B be any collection of events outside of $B(v)$, and define the event $E(B) = \cap_{A\in B}\overline{A}$. For every possible colour assignment σ to the vertices of $U \cup \mathrm{Temp} - N(v)$, let $E(\sigma)$ be the event that σ is selected during Phase III. Since σ determines whether $E(B)$ holds, we have $\mathbf{Pr}(A|E(\sigma) \cap E(B)) = \mathbf{Pr}(A|E(\sigma))$, which Claim 1 implies to be at most $\frac{1}{4}\Delta^2$. Therefore:

$$\mathbf{Pr}(A|E(B)) = \sum_\sigma \mathbf{Pr}(E(\sigma)|E(B)) \times \mathbf{Pr}(A|E(\sigma) \cap E(B)) \leq \sum_\sigma \mathbf{Pr}(E(\sigma)|E(B))\frac{1}{4}\Delta^2 = \frac{1}{4}\Delta^2. \tag{1}$$

To prove Claim 1, we start by proving:

Claim 2: Consider any set of vertices $w_1, ..., w_t$ and any colour x. For every $u \in U \cup \text{Temp} - \{w_1, ..., w_t\}$, choose any colour $c(u) \in L_0(u)$ such that for every adjacent u_1, u_2 we have $c(u_1) \neq c(u_2)$. Conditioning on the event that each such u is assigned $c(u)$ during Phase III, the conditional probability that $w_1, ..., w_t$ are all assigned x is at most $e^{6\times 10^{14}t/\epsilon} \times \prod_{i=1}^{t} \frac{1}{Q(w_i)}$.

At first glance, Claim 2 may appear trivial as Lemma 5.1(a) implies that regardless of what colours are assigned to the vertices preceding w_i in Phase III, the probability that w_i receives x is at most $\frac{1}{Q(w_i)}$. So this should imply Claim 2, without the extra $e^{6\times 10^{14}t/\epsilon}$ term. However, this argument only considers the way that the distribution of the colour assigned to w_i is affected by conditioning on the colours assigned to earlier vertices. We also need to deal with the effect of conditioning on the colours assigned to future vertices. This latter effect is more insidious.

To prove Claim 2, consider any such choice of colours $\mathcal{C} = (c(u) : u \in U)$. Let $\Omega = \Omega(\mathcal{C})$ be the set of all colour assignments $\alpha = (\alpha_1, ..., \alpha_t)$ to $w_1, ..., w_t$ such that α and $\mathcal{C}$ yield a proper colouring of G. The same simple arguments used in the proof of Lemma 5.1 imply that

$$|\Omega(\mathcal{C})| \geq \prod_{i=1}^{t} Q(w_i).$$

We refer to the assignment of α and $\mathcal{C}$ to $U \cup \text{Temp}$ as Θ_α and we use $\rho(\alpha)$ to denote the unconditional probability that Phase III actually produces Θ_α.

For each vertex $v \in U \cup \text{Temp}$ we use $\lambda(v) \leq |L_0(v)|$ to denote the number of colours still available for v when we reach it during Phase III, if each vertex z preceding v was assigned $\Theta_\alpha(z)$. Thus, $\rho(\alpha) = \prod_{v\in U\cup \text{Temp}} 1/\lambda(v)$.

Suppose that we were to carry out Phase III, but skipped the vertices $\{w_1, ..., w_t\}$; i.e. when we reached w_i we did not assign a colour to it. For each vertex $v \in U \cup \text{Temp}$ we use $\lambda'(v) \geq \lambda(v)$ to denote the number of colours still available for v when we reach it, if each vertex $z \notin \{w_1, ..., w_t\}$ preceding v was assigned $\Theta_\alpha(z)$.

Note that, for every choice of α, $\rho(\alpha) \geq \prod_{v\in U\cup \text{Temp}} 1/\lambda'(v)$. We are most interested in the case $\alpha = \alpha^* = (x, x, ..., x)$; i.e. the case where each w_i is assigned the colour x. (We can assume $\alpha^* \in \Omega$, as otherwise the conditional probability of α^* is zero.) Note that in the assignment Θ_{α^*}, we have $\lambda(v) \geq \lambda'(v) - 1$ for all $v \in U \cup \text{Temp}$. Furthermore, using Y to denote the set of vertices with a neighbour in $\{w_1, ..., w_t\}$, we have $\lambda(v) = \lambda'(v)$ for every $v \notin Y$. Therefore, $\rho(\alpha^*) \leq \prod_{v\in Y} 1/(\lambda'(v) - 1) \times \prod_{v\in U\cup \text{Temp}-Y} 1/\lambda'(v)$.

The probability that Phase III assigns α^* to $\{w_1, ..., w_t\}$, conditional on $\mathcal{C}$ being assigned to $U \cup \text{Temp} - \{w_1, ..., w_t\}$ is:

$$\begin{aligned} \frac{\rho(\alpha^*)}{\sum_{\alpha\in\Omega(\mathcal{C})} \rho(\alpha)} &\leq \frac{\prod_{v\in Y} 1/(\lambda'(v) - 1) \times \prod_{v\in U\cup \text{Temp}-Y} 1/\lambda'(v)}{\sum_{\alpha\in\Omega(\mathcal{C})} \prod_{v\in U\cup \text{Temp}} 1/\lambda'(v)} \\ &\leq \frac{1}{|\Omega(\mathcal{C})|} \prod_{v\in Y} \frac{\lambda'(v)}{\lambda'(v) - 1} \end{aligned}$$

$$\leq \frac{1}{\prod_{i=1}^{t} Q(w_i)} \times \prod_{v\in Y}(1+\frac{1}{\lambda'(v)-1})$$
$$< \exp(\sum_{v\in Y}\frac{1}{\lambda'(v)-1}) \times \prod_{i=1}^{t}\frac{1}{Q(w_i)}.$$

By Lemma 5.1(a), $\lambda'(v) - 1 \geq Q(v) - 1 > \frac{1}{2}Q(v)$. That, along with Lemma 5.1(b) yields that this probability is at most:

$$\exp(\sum_{v\in Y}\frac{2}{Q(v)}) \times \prod_{i=1}^{t}\frac{1}{Q(w_i)} < e^{6\times 10^{14}t/\epsilon} \times \prod_{i=1}^{t}\frac{1}{Q(w_i)}.$$

□

We complete our proof by showing how Claim 2 implies Claim 1. There are $\Delta+1$ choices of colour x, and $\binom{\Delta}{t}$ choices of t neighbours of v to which x might be assigned. We consider any choice $\{w_1, ..., w_t\}$ of those neighbours, and use Claim 2 to bound the probability that those neighbours are all assigned x. Claim 2 does not quite apply directly, since the event it conditions on is different than the one that Claim 1 conditions on; specifically, Claim 2 conditions on colour assignments to all vertices outside of $\{w_1, ..., w_t\}$, not just to those vertices outside of $N(v)$. Nevertheless, the same reasoning that was used to derive (1) applies here to show that under the conditioning of Claim 1, the probability that they are all assigned x is at most $e^{6\times 10^{14}t/\epsilon}\prod_{i=1}^{t}\frac{1}{Q(w_i)}$. Therefore, under the conditioning of Claim 1, the probability of $A(v)$ is at most:

$$(\Delta+1) \times \sum_{\{w_1,...,w_t\}\subset N(v)} e^{6\times 10^{14}t/\epsilon}\prod_{i=1}^{t}\frac{1}{Q(w_i)} = (\Delta+1)e^{6\times 10^{14}t/\epsilon} \times \sum_{\{w_1,...,w_t\}\subset N(v)}\prod_{i=1}^{t}\frac{1}{Q(w_i)}.$$

We will bound $S = \sum_{\{w_1,...,w_t\}\subseteq N(v)}\prod_{i=1}^{t}\frac{1}{Q(w_i)}$ subject to:

(i) $\frac{1}{Q(w)} \geq 0$ for all w;

(ii) $\sum_{w\in N(v)}\frac{1}{Q(w)} \leq \frac{3\times 10^{14}}{\epsilon}$ (from Lemma 5.1(b)).

It is straightforward to prove that, subject to these constraints, S is maximized when for all $w \in N(v)$, $Q(w) = \frac{\epsilon|N(v)|}{3\times 10^{14}} \leq \frac{\epsilon\Delta}{3\times 10^{14}}$. To see this, set $q(w) = \frac{1}{Q(w)}$ for each w, and verify that replacing $q(w), q(w')$ both by $\frac{q(w)+q(w')}{2}$ does not decrease S. Therefore, the conditional probability of $A(v)$ is at most:

$$(\Delta+1)e^{6\times 10^{14}t/\epsilon}\binom{\Delta}{t}\left(\frac{3\times 10^{14}}{\epsilon\Delta}\right)^t < (\Delta+1)\left(\frac{e^{6\times 10^{14}/\epsilon}\times 3e\times 10^{14}}{\epsilon t}\right)^t$$
$$< \frac{1}{4}\left(\frac{1}{\log^{1/2}\Delta}\right)^t$$
$$= \frac{1}{4}\Delta^{-2}.$$

□

6 Algorithms

Our proof of Theorem 1 is an existence proof; in this section, we will discuss how to modify the proof to yield an efficient algorithm which produces a frugal colouring. The basic technique we use dates back to Beck's seminal paper[4] in which he showed how to convert some applications of the Local Lemma into efficient algorithms. We remark that the recent work by Moser[21] and Moser and Tardos[22] would also apply to the first two phases of our procedure (and, in fact, would yield much simpler algorithms with no loss in the constants) but it does not seem to apply to the third stage.

Theorem 6.1 *There is a constant $T > 0$ such that there is a randomized polynomial expected-time algorithm which takes as input any graph G on n vertices and outputs a $(T \log \Delta(G)/ \log\log \Delta(G))$-frugal $(\Delta + 1)$-colouring of G. For any constant D, there is a polynomial time deterministic algorithm to produce such a colouring on graphs for which $\Delta(G) \leq D$.*

There might, in fact, be a deterministic polytime algorithm for general graphs, i.e. without bounded maximum degree. The key step required to produce such an algorithm is to devise an efficient way to compute the conditional probabilities that the "bad events" from our applications of the Local Lemma hold, conditioned on the outcomes of a subset of the random trials.

As is usual in this sort of setting, we need to sacrifice a bit in our constants. Our algorithm will find a $(250 \log \Delta/ \log\log \Delta)$-frugal $(\Delta + 1)$-colouring in any graph of maximum degree Δ_0 for a particular constant Δ_0 (which will be larger than the Δ_0 required for Theorem 1.1).

The randomized algorithm to produce the partial colourings for Phases I and II nearly follows from Theorem 2.1 of [27]; the deterministic algorithm nearly follows from Theorem 3.1 of [17] (see also Chapter 25 of [19]). We say "nearly" because both of those theorems apply to settings where the random experiment is a series of independent random choices. But in Step 2 of Phase I, when we assign a random permutation of $|D_i|$ colours to the vertices of D_i, the colour assignments to the vertices are not independent. However, it is straightforward to check that the proofs of those two theorems also carry through for this setting. The important thing to note is that we can choose the random permutation by processing the vertices of D_i in an arbitrary and non-predetermined order; each time we come to a vertex, we give it a uniformly random colour from amongst those not yet assigned to D_i.

To apply each of these theorems to our setting, the main work is to show that we can strengthen the requirement "$pd \leq \frac{1}{4}$" from the Local Lemma to $pd^4 \leq \frac{1}{20}$ and $pd^9 \leq \frac{1}{512}$

respectively. The algorithms that these theorems guarantee are, at heart, much like Beck's algorithm from [4]. They are very similar to the algorithm that we describe below for Phase III.

In Phase I, d is roughly $6\Delta^7$ and so it will suffice if each of our bad events has probability at most Δ^{-64} as $\Delta^{-64} \times (6\Delta^7)^9 < \frac{1}{512}$ for large Δ. The probabilities of events A_1, A_2, A_3, A_5 are all asymptotically lower than the inverse of any poynomial in Δ and hence are less than Δ^{-64} for sufficiently large Δ. Increasing t to $100 \log \Delta / \log\log \Delta$ in the analysis of A_4 will decrease $\mathbf{Pr}(A_4)$ below Δ^{-64}; this requires us to increase "20" to "100" in Lemma 3.1(d). Increasing t_i to $300,000 \times 2^{-i/2}$ in the analysis of A_6 will decrease $\mathbf{Pr}(A_6)$ below Δ^{-64}; this requires us to increase the bound in Lemma 3.1(f) to 2,999,999. This yields the desired algorithms to complete Phase I so that no colour appears more than $100 \log \Delta / \log\log \Delta$ times in any neighbourhood.

In Phase II, d is roughly $2\Delta^6$ and so it will suffice if each of our bad events has probability at most Δ^{-55}. By taking 100 candidates for each $v \in \mathrm{Temp}_i(\log^6 \Delta)$ rather than 20 candidates, we decrease $\mathbf{Pr}(B_1)$ below Δ^{-55}. (This requires a straightforward readjustment of some of the other constants in the analysis.) Increasing t to $100 \log \Delta / \log\log \Delta$ in the analysis of B_2 will decrease $\mathbf{Pr}(B_2)$ below Δ^{-55}; this requires us to increase "20" to "100" in Lemma 4.2(b). This yields the desired algorithms to complete Phase II so that no colour is assigned more than $100 \log \Delta / \log\log \Delta$ times to any neighbourhood.

For Phase III, we apply the Lopsided Local Lemma, and so the theorems from [17] and [27] do not apply directly. Fortunately, Beck's technique works very well for this particular application of the Lopsided Local Lemma. Again, we sacrifice a bit in the constants: our goal is that no colour will be assigned to a neighbourhood more than $24 \log \Delta / \log\log \Delta$ times during this phase; i.e. we increase the constant "4" to "24".

We define a hypergraph H as follows: The vertices of H are the vertices of G which are still uncoloured at the end of Phase II. For every $v \in G$, the vertices of $N(v)$ that are in H form a hyperedge of H. So our goal for the remaining stages is to complete the colouring of G so that no colour appears too many times in any hyperedge of H.

The algorithm runs in 3 stages as follows:

Stage 1: We colour the vertices of G one-at-a-time. When a colour is assigned, during this stage, to $8 \log \Delta / \log\log \Delta$ vertices in a hyperedge of H, then we *freeze* all remaining uncoloured vertices in that hyperedge. When we come to v, if it is frozen then we do not assign it a colour; if it is unfrozen we assign it a uniformly random colour from $L(v)$ and we remove that colour from $L(u)$ for every u that is adjacent to v in G.

H' is the hypergraph formed by removing from H all vertices that are assigned a colour during Stage 1. When some, but not all, of a hyperedge's vertices are removed, the hyperedge itself is not removed - it is merely reduced in size. The main outcome of Stage 1 is that every component of H' is small:

Lemma 6.2 *With probability at least $\frac{1}{2}$, every component of H' has at most $\Delta^2 \log n$ vertices.*

This is a very standard lemma - it's proof is nearly identical to, eg., that of Lemma 25.2 in [19] - so we omit the details. The key fact needed for this proof is: Consider any collection of t disjoint hyperedges. The probability that they all become frozen during Stage 1 is at most $(\frac{1}{4\Delta^4})^t$. This follows from the same analysis as in the proof of Lemma 5.2: Increasing the number of times a colour can be assigned to a neighbourhood from $4\log\Delta/\log\log\Delta$ to $8\log\Delta/\log\log\Delta$ decreases the probability of the bad event from Lemma 5.2 from $\frac{1}{4\Delta^2}$ to $\frac{1}{4\Delta^4}$; i.e. the probability that one particular such hyperedge becomes frozen, even after conditioning on the event that some of the others become frozen, is at most $\frac{1}{4\Delta^4}$. Thus the probability that all become frozen is at most $(\frac{1}{4\Delta^4})^t$.

After running Stage 1, if any components of H' have more than $\Delta^2\log n$ vertices, then we run Stage 1 over again. Lemma 6.2 implies that we probably won't have to restart very often; in fact the expected number of runs is at most 2.

Stage 2: We process the components of H' one-at-a-time. We repeat the procedure of Stage 1 on each component.

H'' is the hypergraph formed by removing from H' all vertices that are assigned a colour during Stage 2. Noting that each edge of H' is no bigger than the corresponding edge of H, the same argument as for Lemma 6.2 yields:

Lemma 6.3 *With probability at least $\frac{1}{2}$, every component of H'' has at most $\Delta^2\log(\Delta^2\log n) < 2\Delta^2\log\log n$ vertices.*

Again, if any components of H'' have more than $2\Delta^2\log\log n$ vertices then we run Stage 2 over again. We probably won't have to restart very often.

Stage 3: In this stage, we colour all the vertices of H''; i.e. all the remaining uncoloured vertices of G so that no colour is assigned to any hyperedge (i.e. any neighbourhood in G) more than $8\log\Delta/\log\log\Delta$ times.

Note that we can process the components of H'' independently of each other. If we colour the vertices of each such component so that no colour appears too many times in any hyperedge of that component, then the overall colouring of G will be as desired. The components of H'' are small enough that it will straightforward to find a suitable colouring for each of them.

Case 1: $\Delta \le \log\log n$.

The analysis from Section 5 implies that the desired colouring exists. For each component of H'', we will find that colouring by exhaustive search. The number of possible colourings is at most $(\Delta+1)^{2\Delta^2\log\log n} < \mathrm{e}^{(\log\log n)^4} = n^{o(1)}$. Each such colouring can be generated and checked in $O(m\Delta)$ time.

Case 2: $\Delta > \log\log n$.

We process the uncoloured vertices of a component of H'' one-at-a-time, as in Section 5. Each vertex appears in at most Δ neighbourhoods (i.e. hyperedges) and so the number of hyperedges of H'' is at most $\Delta \times 2\Delta^2 \log\log n < 2\Delta^4$ The probability that a colour appears too many times in a particular hyperedge is at most $\frac{1}{4\Delta^4}$, and so the expected number of bad hyperedges for which this happens is at most $\frac{1}{2}$. This allows us to find a suitable colouring using the technique of Erdős and Selfridge [7], as follows.

Consider a partial colouring of a component Φ of H'', and consider a hyperedge e and colour c. Let α denote the number of vertices in e with colour c, and let Ψ denote the set of uncoloured vertices in e; we will assume that $\alpha \leq 8\log\Delta/\log\log\Delta$. We define $w(e,c)$ to be the sum over all subsets $\Psi' \subseteq \Psi$ of size $8\log\Delta/\log\log\Delta - \alpha$ of $\prod_{u\in\Psi'}\frac{1}{Q(u)}$. We set $W = \sum_{e,c} w(e,c)$. Suppose that we were to continue colouring the vertices of H' one-at-a-time, as in Phase III, each time choosing for a vertex a uniformly random member of $L(v)$. By Lemma 5.1(a), the expected number of edges of H'' which contain the same colour more than $8\log\Delta/\log\log\Delta$ times is at most W.

Initially, when Φ colours no vertices, the calculations above yield $W \leq \frac{1}{2}$. (Note that the calculations in Section 5 actually bounded W.) So we colour all vertices one-at-a-time. Each time we come to a vertex v, there is at least one colour in $L(v)$ that we can assign to v without increasing W. We can easily find this colour in polytime by simply checking what W would change to under each of the $|L(v)|$ possible assignments to v. When all vertices have been coloured, we will still have $W \leq \frac{1}{2} < 1$ and so no edge will contain the same colour more than $8\log\Delta/\log\log\Delta$ times.

Each stage runs in polytime. The resulting colouring is β-frugal for $\beta = (100 + 100 + 24)\log\Delta/\log\log\Delta < 250\log\Delta/\log\log\Delta$.

Acknowledgement We thank an anonymous referee for a careful reading and some helpful corrections.

References

[1] N. Alon and J. Spencer, The Probabilistic Method. Wiley, New York (1992).

[2] O. Amini, L. Esperet and J. van den Heuvel, *Frugal Colouring of Graphs.* CDAM Research Report CDAM-LSE-2007-11 (2007).

[3] K. Azuma, *Weighted sums of certain dependent random variables.* Tokuku Math. J. **19** (1967), 357 - 367.

[4] J. Beck, *An algorithmic approach to the Lovasz Local Lemma.* Rand. Str. & Alg. **2**, 343 - 365 (1991).

[5] H.Chernoff, *A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations.* Ann. Math. Statist. **23** 493 - 509 (1952).

[6] P. Erdős and L. Lovász, *Problems and Results on 3-Chromatic Hypergraphs and Some Related Questions.* In: “Infinite and Finite Sets” (A. Hajnal et. al. Eds), Colloq. Math. Soc. J. Bolyai 11, North Holland, Amsterdam, 609–627 (1975).

[7] P. Erdős and J. Selfridge, *On a combinatorial game.* J. Comb. Th. (A) **14** 298 - 301 (1973).

[8] P. Erdős and J. Spencer, *Lopsided Lovász Local Lemma and latin transversals.* Disc. Appl. Math. **30** 151 - 154 (1991).

[9] H. Hind, M. Molloy and B. Reed, *Colouring a graph frugally.* Combinatorica **17** 469 - 482 (1997).

[10] H. Hind, M. Molloy and B. Reed, *Total colouring with* $\Delta+poly(\log\Delta)$ *colours.* SIAM J. Computing **28** 816 - 821 (1998).

[11] W. Hoeffding, *Probability inequalities for sums of bounded random variables.* J. Amer. Statist. Assoc. **58**, 713 - 721 (1963).

[12] S. Janson, T. Łuczak and A. Ruciński. Random Graphs. Wiley, New York (2000).

[13] T. Jensen and B. Toft, *Graph Colouring Problems.* Wiley, New York (1995).

[14] J. Kahn, *Asymptotics of the list chromatic index for hypergraphs.* Rand. Struc. & Alg. **17** 117 - 155 (2000).

[15] C. McDiarmid, *Concentration for Independent Permutations.* Combinatorics, Probability and Computing **11**, 163 - 178 (2002).

[16] M. Molloy and B. Reed, *A bound on the total chromatic number.* Combinatorica **18** 241 - 280 (1998).

[17] M. Molloy and B. Reed, *Further algorithmic aspects of the Local Lemma.* Proceedings of the 30th ACM Symposium on Theory of Computing, 524–529 (1998).

[18] M. Molloy and B. Reed, *Colouring graphs whose chromatic number is near their maximum degree.* Proceedings of Latin American Theoretical Informatics, 216 - 225 (1998).

[19] M. Molloy and B. Reed, Graph Colouring and the Probabilistic Method. Springer (2001).

[20] M. Molloy and B. Reed, *Colouring graphs when the number of colours is almost the maximum degree.* Submitted. Preliminary version in: Proceedings of the 33rd ACM Symposium on Computing (STOC) (2001).

[21] R. Moser. *A constructive proof of the Lovasz Local Lemma.* Proceedings of the 41st ACM Symposium on Theory of Computing (2009).

[22] R. Moser and G. Tardos. *A constructive proof of the general Lovasz Local Lemma.* Manuscript.

[23] S. Pemmaraju and A. Srinivasan, *The randomized coloring procedure with symmetry-breaking.* Proceedings of the 35th International Colloquium on Automata, Languages, and Programming (ICALP), 306 - 319 (2008).

[24] B. Reed, ω, Δ *and* χ. J. Graph Th. **27**, 177-212 (1998).

[25] B. Reed, *A strengthening of Brooks' Theorem.* J.Comb.Th.(B) **76**, 136 - 149 (1999).

[26] V. Rodl, *On a packing and covering problem.* Eur. J. Comb. **5** (1985), 69 - 78.

[27] A. Srinivasan, *Improved algorithmic versions of the Lovsz Local Lemma.* Proc. ACM-SIAM Symposium on Discrete Algorithms (SODA), 611 - 620 (2008).

[28] M.Talagrand, *Concentration of measure and isoperimetric inequalities in product spaces.* Instutut Des Hautes Etudes Scientifiques, Publications Mathematiques **81**, 73 - 205 (1995).

[29] R. Yuster, *Linear coloring of graphs.* Discrete Math **185** 293 - 297 (1998).