Asymptotically optimal frugal colouring

Michael Molloy* Bruce Reed†

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 factor in the $O(-)$ term. We also provide an efficient algorithm to produce such a colouring.

1 Introduction

In [7], 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[17] recently improved this to $\beta = O(\log^2\Delta/\log\log\Delta)$.

Alon (see [7]) 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.*

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 G has a $(T\log\Delta/\log\log\Delta)$-frugal $(\Delta+1)$-colouring where $T=\max\{\Delta_0, 50\}$. (In our proof, Δ_0 is much bigger than 50.)

The main motivation of the initial study of frugal colouring in [7] 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 [8] 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 improved to $\Delta+O(1)$ in [12]. A well-known conjecture is that it can be improved to $\Delta+2$ (see eg. [9]).

Amini, Esperet and van den Heuvel[1] 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 [22], 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[17], 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 $\beta\approx O(\frac{\Delta}{d}\log^2\Delta)$. Frugal colourings have been used for channel-allocation schemes in multi-channel, multi-radio wireless networks (see [17]).

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 2.1), one can often move from "with high probability the neighbourhood of one vertex is fine" to "with positive 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

---

*Dept of Computer Science, University of Toronto, molloy@cs.toronto.edu.

†Canadian Research Chair in the Combinatorics of Complex Networks, School of Computer Science, McGill University and Projet Mascotte, CNRS-INRIA, Sophia-Antipolis, France, breed@cs.mcgill.ca.

$N(v)$. 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.

Vertices v for which $N(v)$ has a very high number of edges have to be handled carefully. To do so, we make use of the dense decomposition introduced by Reed in [18], in which vertices with very dense neighbourhoods are isolated so that they can be coloured more carefully. This decomposition has been applied in [12, 14, 16, 18, 19]. See [15] for a thorough presentation of this and related techniques.

A common approach for this sort of problem is to randomly colour the vertices "all-at-once", correct conflicts by uncolouring some vertices, and then iterate. This has the substantial advantage that the random colours assigned to vertices are independent. The tradeoff is that we require several, typically roughly $O(\log n)$, iterations. The reasoning described above can show that, in any one iteration, no colour is assigned to more than $O(\log\Delta/\log\log\Delta)$ vertices in any one neighbourhood; this is how Pemmaraju and Srinivasan[17] obtained an $O(\log^2\Delta/\log\log\Delta)$-frugal colouring. Our main challenge in proving Theorem 1.1 was to produce our colouring using only $O(1)$ rounds, thus obtaining an $O(\log\Delta/\log\log\Delta)$-frugal colouring.

The main new ideas can be found in our Phase III, where we move from an "all-at-once" procedure to a "one-at-a-time" procedure. By colouring the vertices in sequence, we can ensure that there are no conflicts and so we don't need to iterate. The drawback is that we create lots of dependency between the colours assigned to vertices, even to pairs of vertices that are far apart. This prevents us from applying the most common form of the Local Lemma; instead, we must 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 [10]). Fortunately, in this particular setting, we were able to handle the conditioning adequately.

The first two phases of our procedure are used to obtain a partial colouring upon which Phase III will perform well. The main issue is to make sure that every uncoloured vertex has many more available colours than it has uncoloured neighbours, and so when it is coloured, there will be many colours to choose from. Phase I is very much like a similar step in other papers such as [18, 12, 16]. In this extended abstract we highlight the variation that we require for the present application. Phase II is nearly identical to the similar step in those papers, so we omit the details.

In the next section, we present some probabilistic tools, the decomposition and an overview of our colouring procedure. In the following three sections, we present the three phases of our procedure.

2 Preliminaries

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

The Lovasz Local Lemma[4] *Let $\mathcal{A} = \{A_1, ..., A_n\}$ be a set of random events so that for each $1 \le i \le n$:*

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

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

If $pd \le \frac{1}{4}$ *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 [6]:

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 $B \subset \mathcal{A} - B_i$*,*

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

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

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

Our main concentration tool is a variation on Talagrand's Inequality due to McDiarmid[11]. It applies to a random variable X that is determined by independent trials and permutations. Talagrand's original inequality[21] covered the case where X was determined only by independent trials. Talagrand also derived this for the case where there are no independent trials and exactly one permutation. Rather than using McDiarmid's original statement, we will use the following useful reworking, which is proved in [16].

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

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

2.2 A Dense Decomposition We begin by describing the graph decomposition introduced in [18]. 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. [18]). 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 [12]. We define $\mathcal{D} = \cup_{i=1}^{\ell} D_i$, and so $S = V(G) - \mathcal{D}$. We set $\epsilon = 10^{-6}$. Section 2 of [12], 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 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 [12] 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.*

We say that D_i is *ornery* if $|\mathcal{C}_i| > \Delta - \log^4 \Delta$. For each vertex v in any D_i, we define Out_v to be the set of neighbours of v that are not in D_i. Each member of Out_v is said to be an *external neighbour* of v. 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 of any member of the colour class containing u. The proof can be found in the full paper.

Lemma 2.3. *(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 $\mathcal{C}_i$.

2.3 An overview Our colouring procedure will proceed 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. A more complicated analysis shows that it has several other useful properties that bound the number of temporarily coloured vertices in each D_i, and imply that all but a few particularly problematic vertices will always have many available colours in Phase III.

In the second phase, we recolour those particularly problematic temporarily coloured vertices. We do so by swapping their colours with randomly chosen vertices in the same dense set. Properties of the partial colouring from Phase I ensure that we can do so such that 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. Properties of the partial colouring from Phase I ensure that each vertex, at its turn, 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 $(20 + 20 + 4)\log\Delta/\log\log\Delta < 50\log\Delta/\log\log\Delta$ colours 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.

Remark The procedure described here is slightly simplified. In the full analysis, we have to be careful about so-called *big-neighbours*, in order to facilitate Phase II. See the full paper for more details.

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 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*.

Remark: In similar steps in other papers, we correct both vertices in conflicting pairs. By correcting only one, we manage to avoid some lengthy technical issues (eg. Lemmas 16.7 and 16.8 from [15]).

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 set of 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;
- for each $0 \leq a \leq \Delta$, $\mathrm{Temp}_i(a)$ is the set of $v \in \mathrm{Temp}_i$ with $|\mathrm{Out}_v| \leq a$;
- $\mathrm{Temp} = \cup_i \mathrm{Temp}_i$;

All vertices in Temp will be recoloured during Phases II and III. In Phase II, we will recolour the temporarily coloured vertices in ornery sets with at most $\log^6\Delta$ external neighbours 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 temporarily coloured vertices during Phase II. 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 $N(v) - (U \cup \mathrm{Temp} \cup F)$;

(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 $N(v) - (U \cup \mathrm{Temp} \cup F)$;

(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 vertex $v \in G$,

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

Parts (a) and (b) ensure that we have many repeated colours in the neighbourhoods of those vertices and so when we come to colour them during Phase III, there will be many colours to choose from. Part (e) then ensures that the expected number of times that any particular colour is assigned to the neighbourhood of v during Phase III is $O(1)$.

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

Remark: In this statement of Lemma 3.1, we only list the properties required for the presentation in this extended abstract. In the full version of the paper, we have slightly stronger properties, and one additional property, to facilitate the (omitted) analysis of Phase II.

Proof. We apply the Lovasz Local Lemma, setting up a bad event for each of the possible situations that can violate (a)-(e). For example, for each vertex v, we define $A(v)$ to be the event that $\sum_{u \in N(v)\cap\mathrm{Temp}} \frac{1}{\max(|\mathrm{Out}_u|, \log^3\Delta)} > 299999$; i.e. that v violates (e). We show that each event has probability less that Δ^{-8}. It is easy to see that each event is mutually dependent of all but at most $O(\Delta^7)$ other events, and so the result follows.

The analysis for parts (a,b,c) is very much like that for the corresponding parts in [18, 12, 16]. Observe that, by Lemma 2.3(a), every $v \in S$ has $\Theta(\Delta^2)$ pairs of neighbours that are eligible to receive the same colour. Each such pair recieves the same colour with probability $\Theta(\frac{1}{\Delta})$. We prove that, with probability $\Theta(1)$, neither member of such a pair is in $U \cup \mathrm{Temp} \cup F$, and so the expected number of colours counted in part (a) is $\Theta(\Delta)$. A concentration argument proves that the bound in (a) holds with sufficiently high probability. Part (b) uses a similar argument based on Lemma 2.3(b). For part (c), note that Lemma 2.1(a) implies that $\mathbf{Exp}(|\mathrm{Temp}_i(a)|) < 1.5a$. Part (d) is easy using the calculations described in the introduction.

We include here only the details for part (e). The rest of the proof can be found in the full paper. Set

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

If $u \in N(v)$ is in Temp then at least one member of Out_u is assigned the same colour as u. So $\mathbf{Pr}(u \in \mathrm{Temp}) \le |\mathrm{Out}_u|/(\Delta + 1)$. Therefore

$$\mathbf{Exp}(Y) \le \sum_{u \in N(v)} \frac{|\mathrm{Out}_u|}{\Delta + 1} \times \frac{1}{|\mathrm{Out}_u|} < 1.$$

We apply McDiarmid's Inequality to show that Y is concentrated. 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 $\mathrm{Temp}' \subseteq \mathrm{Temp}$ defined to contain every vertex $u \in \mathcal{D}$ such that (1) u has a neighbour w which receives the same colour as u; (2) at most $\log \Delta$ 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 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: Fewer than Δ^2 vertices are 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

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

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

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

Note that $\mathbf{Exp}(Y_i') \le \mathbf{Exp}(Y) < 1$.

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 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 $1 + t_i \le 30{,}001$. McDiarmid's Inequality with $c = 2 \times 2^{-i} \log^{-2} \Delta$ and $r = 3 \times 2^{i+1} \log^3 \Delta$ yields:

$$\begin{aligned} & \mathbf{Pr}(Y_i' - \mathbf{Exp}(Y_i') > t_i) \\ \leq\ & 4e^{-t_i^2/(128\times(2\times2^{-i}\log^{-2}\Delta)^2\times3\times2^{i+1}\log^3\Delta(1+t_i))} \\ <\ & 4e^{-\frac{3\times10^8}{8\times128(1+t_i)}\log\Delta} \\ <\ & \Delta^{-9.5} \end{aligned}$$

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

$$Y' - \mathbf{Exp}(Y') < \sum_{i\geq 0} t_i = 30,000 \times \frac{1}{1 - 2^{-1/2}} < 299,999.$$

This implies:

$$\mathbf{Pr}(A(v)) \leq 2\Delta^{-9} < \Delta^{-8}.$$

□

4 Phase II: The kernels of the ornery dense sets

In this phase, we will recolour 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 will recolour any such vertex v by swapping its colour with an appropriate vertex in F_i, and then removing v from Temp. We prove that we can do this for every such vertex v so that no colour is swapped onto more than $20\log\Delta/\log\log\Delta$ neighbours of any one vertex.

We have to be careful about swaps involving colour classes of size 2, and in doing so we add more vertices to Temp; every vertex added to Temp will have more than $\log^3\Delta$ external neighbours (this is important in the proof of Lemma 5.1). We ensure that the bound in Lemma 3.1(e) still holds. The analysis is very much like a similar step in [18, 12, 16]. Details can be found in the full paper.

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.

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

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

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

$$\sum_{u\in N(v)\cap(U\cup\text{Temp})} \frac{1}{Q(u)} \leq \frac{3\times10^{14}}{\epsilon}.$$

Proof. Lemma 3.1(a) and the fact that only vertices in $\text{Temp} \cup F$ 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 u has $|\text{Out}_w| \leq |\text{Out}_u| < \log^3\Delta$. 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 part (b): since $Q(u) = \frac{\epsilon}{10^9}\Delta$ for every $u \in U$, we have $\sum_{u\in N(v)\cap U}\frac{1}{Q(u)} \leq \frac{10^9}{\epsilon}$. Lemma 3.1(e) (and the fact that its bound still held at the end of Phase II) imply that at the beginning of Phase III: $\sum_{u\in N(v)\cap\text{Temp}}\frac{1}{Q(u)} \leq \frac{10^9}{\epsilon} \times 299999$. □

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 $10^9/(\epsilon|\text{Out}_v|)$ for $v \in \text{Temp}$. Part (b) implies that for each $v \in G$, the expected number of neighbours of v to receive x is at most $O(1)$. 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 a similar bound for Phase II, 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. 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 yields 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 $v \in U \cup \text{Temp}$, we define $L_0(v)$ to be the set of colours not appearing on $N(v)$ at the beginning of Phase III.

Claim 1: For every $u \in (U \cup \text{Temp}) - N(v)$, 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 $(U \cup \text{Temp}) - N(v)$, it is straightforward to show that Claim 1 will imply the condition required for our application of the Lopsided Local Lemma (see the full paper for more explanation). To prove Claim 1, we start by proving:

Claim 2: Consider any independent set of vertices $w_1, ..., w_t \in U \cup \text{Temp}$ 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 $\mathcal{C} = (c(u) : u \in U \cup \text{Temp} - \{w_1, ..., w_t\})$. 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 that would still be 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 that would still be 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:

$$\begin{aligned}
& \frac{\rho(\alpha^*)}{\sum_{\alpha \in \Omega(\mathcal{C})} \rho(\alpha)} \\
\leq \quad & \frac{\prod_{v \in Y} 1/(\lambda'(v) - 1) \times \prod_{v \in U \cup \text{Temp} - Y} 1/\lambda'(v)}{\sum_{\alpha \in \Omega(\mathcal{C})} \prod_{v \in U \cup \text{Temp}} 1/\lambda'(v)} \\
= \quad & \frac{1}{|\Omega(\mathcal{C})|} \prod_{v \in Y} \frac{\lambda'(v)}{\lambda'(v) - 1} \\
\leq \quad & \frac{1}{\prod_{i=1}^{t} Q(w_i)} \times \prod_{v \in Y} \left(1 + \frac{1}{\lambda'(v) - 1}\right) \\
< \quad & \exp\left(\sum_{v \in Y} \frac{1}{\lambda'(v) - 1}\right) \times \prod_{i=1}^{t} \frac{1}{Q(w_i)}.
\end{aligned}$$

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:

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

This proves Claim 2.

We complete our proof by showing how Claim 2 implies Claim 1. There are $\Delta + 1$ choices for colour x, and at most $\binom{\Delta}{t}$ choices for 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 that 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,it is straightforward 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:

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

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:

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

□

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[2] in which he showed how to convert some applications of the Local Lemma into efficient algorithms.

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 [20]; the deterministic algorithm nearly follows from Theorem 3.1 of [13] (see also Chapter 25 of [15]). 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 (see the full paper for details).

For Phase III, we apply the Lopsided Local Lemma. The very large spread of dependency amongst our events prevents us from applying the theorems from [13] and [20] directly. Fortunately, Beck's technique works very well for this particular application of the Lopsided Local Lemma. 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.

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 6.1. *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 [15] - 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. 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 6.1 implies that we probably won't have to restart very often; in fact the expected number of runs is at most 2.

The rest of the algorithm and analysis is standard. We process the components of H' one-at-a-time. Running Stage 1 a second time reduces all component sizes to $O(\Delta^2 \log \log n)$. The remaining components are all small enough to be easily managed. See the full paper for more details. The resulting colouring is β-frugal for $\beta = (100 + 100 + 24) \log \Delta / \log \log \Delta < 250 \log \Delta / \log \log \Delta$.

References

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

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

[3] 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).

[4] 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).

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

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

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

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

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

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

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

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

[13] 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).

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

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

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

[17] S. Pemmaraju and A. Srinivasan, *The randomized coloring procedure with symmetry-breaking.* To appear in Proceedings of APPROX (2008).

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

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

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

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

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