Colouring Graphs When the Number of Colours is Almost the Maximum Degree*

Michael Molloy
Dept of Computer Science, University of Toronto
10 Kings College Rd.
Toronto, ON, Canada, M5S 3G4
molloy@cs.toronto.edu

Bruce Reed
School of Computer Science,McGill University
3480 University Street
McConnell Engineering Building, Room 318
Montreal, QB, Canada, H3A 2A7
breed@cs.mcgill.ca

April 14, 2009

Abstract

We consider the chromatic number of graphs with maximum degree Δ. For sufficiently large Δ, we determine the precise values of k for which the barrier to $(\Delta+1-k)$-colourability must be a local condition, i.e. a small subgraph. We also show that for Δ constant and sufficiently large, $(\Delta+1-k)$-colourability is either NP-complete or can be solved in linear time, and we determine precisely which values of k correspond to each case.

AMS Subject Classification: 05C15

*Short preliminary versions of this paper appeared in conference proceedings in [23] and [24].

1 An overview

1.1 The results

In 1852, Francis Guthrie asked if every map could be coloured with four colours so that each pair of countries which share a boundary receive different colours. This problem can be formulated in the language of graph theory (see [25] for an introduction to this language), as follows. We have a graph $G = (V, E)$ where V is the set of countries and two vertices are joined by an edge e of E if they have a common boundary. This graph is planar, i.e. it can be drawn in the plane so that edges intersect only at their common endpoints.

In this language, Guthrie's question is:

> Can the vertices of every planar graph be coloured with four colours so that no edge is monochromatic?

This question was finally answered in the affirmative in 1976, by Appel and Hakken[1] (see also [29]). It was the first of many problems which have been formulated in terms of graph colouring.

A *k-colouring* of a graph is a colouring of its vertices with k colours so that no edge is monochromatic. The chromatic number of G, denoted $\chi(G)$, is the minimum k for which such a colouring exists.

In general, the chromatic number of a graph can be forced to be large by its global structure rather than by small hard to colour subgraphs. Indeed, as shown by Erdos[12], for every $c > 2$ and N there is a graph which has chromatic number c but all of whose subgraphs with at most N vertices have chromatic number 2. A fundamental problem in graph theory is to characterize families of graphs for which the chromatic number is determined locally rather than globally.

The smallest graph with chromatic number c is a clique with c vertices (in a clique every two vertices are adjacent, and hence receive different colours). What can we say about the structure of c-colourable graphs on which we are forced to use c colours because they contain a clique of size c? Unfortunately, not very much. Given any c-colourable graph, we can take its (disjoint) union with a clique of size c to obtain a graph in this class. If, instead, we ask about the structure of graphs such that for every induced subgraph H, the chromatic number of H is equal to the size of the largest clique in H, then we obtain the celebrated *perfect graphs*. These graphs were defined by Berge[4], motivated by a question of Shannon in information theory. A perfect graph contains no odd cycle of length at least five as an induced subgraph, since these cycles have chromatic number three but contain no clique of size three. Furthermore, perfect graphs do not contain the complements of such cycles as induced subgraphs, since the complement of an odd cycle of length $2k+1 \geq 5$ has chromatic number $k+1$ but the size of its largest clique is k. Recently, Chudnovsky et al.[9] proved Berge's celebrated Strong Perfect Graph Conjecture which states that perfect graphs are precisely those graphs which satisfy these two conditions.

Another theorem yielding a family of graphs for which the chromatic number is determined by a clique is Brooks' Theorem[7], one of graph theory's foundational results. The degree of v, $\deg(v)$, is the number of vertices adjacent to v. A simple greedy colouring procedure (colour the vertices in any order, using on v a colour not yet used on any of its neighbours) yields that the chromatic number of a graph exceeds its maximum vertex degree by at most one. That is, using $\Delta(G)$ to denote the maximum vertex degree in G, we have: $\chi(G) \leq \Delta(G)+1$. Brooks' Theorem determines when this bound is tight. It states that for $\Delta(G) \geq 3$, if $\chi(G) = \Delta(G)+1$ then G contains a clique of order $\Delta(G)+1$; i.e. for such graphs, the chromatic number is determined by a clique. (The analagous result is not true for $\Delta(G) = 2$ as the odd cycles show).

Reed[28] settled a conjecture of Beutelspacher and Herring[5], by showing that for $\Delta(G)$ large enough, if $\chi(G) = \Delta(G)$ then G has a clique of size $\Delta(G)$. Borodin and Kostochka[6] conjectured that this is true as long as $\Delta(G)$ is at least nine.

It is natural to ask whether these results can be extended to show that for every k there is a Δ_k such that if $\Delta(G) \geq \Delta_k$ and $\chi(G) = \Delta(G)+1-k$ then G has a clique of size $\Delta(G)+1-k$. In fact this conjecture fails for the next value of k, two. To see this consider a graph obtained by adding all edges between a cycle of length five and a clique of size $\Delta - 4$. However, the results of this paper imply that for $\Delta(G)$ large enough, if $\chi(G) = \Delta(G)-1$ then G must contain either a clique of size $\Delta(G)-1$ or the graph described in the previous sentence (see [15]). So for such graphs, while $\chi(G)$ is not determined by a clique on $\chi(G)$ vertices, it is determined by something very close to a clique - by a subgraph on approximately $\chi(G)$ vertices.

This leads us to ask when the chromatic number of a graph is determined by a subgraph with approximately $\chi(G)$ vertices. Erdös' result shows that this is not true for general graphs. But the results of the past few paragraphs suggest that for large enough Δ, perhaps it is true whenever $\chi(G)$ is sufficiently close to its upper bound of $\Delta(G)+1$. A construction from Embden-Weinert et al[11] shows that it is not true for some graphs with $\chi(G) = \Delta - \lceil\sqrt{\Delta}\rceil$:

Definition: k_Δ is the maximum integer k such that $(k+1)(k+2) \leq \Delta$. Note that

$\sqrt{\Delta} - 1 > k_\Delta > \sqrt{\Delta} - 3$.

Theorem 1 *For every $\Delta \geq 2$ and $2 \leq c \leq \Delta - 1 - k_\Delta$, there are arbitrarily large graphs G with maximum degree Δ and with $\chi(G) = c + 1$ such that every proper subgraph $H \subset G$ satisfies $\chi(H) \leq c$.*

In this paper, we prove that for large enough Δ, if $\chi(G) > \Delta(G) + 1 - k_\Delta$, then $\chi(G)$ is indeed determined by a subgraph of size approximately $\chi(G)$; specifically by a subgraph of size at most $\Delta(G) + 1 < \chi(G) + \sqrt{\chi(G)}$:

Theorem 2 *There is an absolute constant Δ_0 such that for any $\Delta \geq \Delta_0$ and $c \geq \Delta + 1 - k_\Delta$, if G has maximum degree Δ and if $\chi(G) = c + 1$ then G contains a subgraph H such that*

(i) $|H| \leq \Delta + 1$;

(ii) $\chi(H) = c + 1$.

We conjecture that Theorem 2 in fact holds for every Δ (see Conjecture 6 below).

This leaves only the case $c = \Delta - k_\Delta$. For some values of Δ, this case is as in Theorem 1 and for all others it is as in Theorem 2 but with a weaker upper bound on $|H|$. We state this precisely in the next subsection, when we have developed the relevant machinery. Thus, for large Δ, we determine precisely how close c must be to $\Delta + 1$ for c-colorability to be a local rather than a global property. Conjecture 6 implies that this holds for every Δ.

We say a graph is ℓ-critical if its chromatic number is ℓ but all of its proper subgraphs can be $(\ell - 1)$-coloured. So Theorem 2 implies that for $c \geq \Delta + 1 - k_\Delta$ and Δ large, all $(c+1)$-critical graphs have bounded size while Theorem 1 says that for $c \leq \Delta - 1 - k_\Delta$ this is not true. When $c = \Delta - k_\Delta$ then, roughly speaking, for most values of Δ, all $(\Delta - k_\Delta + 1)$-critical graphs have size at most $O(\Delta^{5/2})$. For some values of Δ, the $(\Delta - k_\Delta + 1)$-critical graphs can have arbitrary size, but the large ones all have a fairly simple structure.

We give the simple short proof of Theorem 1 in the next subsection. It is Theorem 2, and the similar case $c = \Delta - k_\Delta$, on which we will have to spend the most effort.

1.2 Focussing on the threshold using reducers

Suppose G has a uniquely c-colourable subgraph D such that only one colour class in this unique colouring contains vertices with neighbours outside D. Clearly G is c-colourable if and only if the graph obtained from G by deleting D and adding a vertex adjacent to all the vertices of $G - D$ that have a neighbour in D is c-colourable. In such a situation, this new graph is the c-reduction of G via D and G is the c-expansion of the new graph via D.

We are interested in such subgraphs which have a special form.

Definition: A c-reducer consists of a clique C with $c-1$ vertices and a stable set S such that every vertex of C is adjacent to all of S but none of $V(G) - S - C$.

We often drop the "c" when its value is clear from context.

Definition: We say that a c-reducer $D = (C, S)$ is *deleteable* if there are fewer than c vertices of $V(G) - D$ with a neighbour in S.

Any c-colouring of $G - D$ for a deleteable c-reducer D of G, can be extended to a c-colouring of the c-reduction of G by D, since the only uncoloured vertex in this graph has fewer than c neighbours. Thus, G is c-colourable if and only if G with D deleted is c-colourable. This fact explains our choice of terminology.

We would like to iteratively apply c-expansion to build larger and larger $(c+1)$-critical graphs of maximum degree Δ. Suppose then, that we apply one such expansion and thereby construct a critical c-reducer $D = (C, S)$ in a $(c+1)$-critical graph of maximum degree Δ. The bound on the degrees of the vertices in C implies that $|S| \leq \Delta - c + 2$. The bound on the degrees of the vertices in S implies that each such vertex can have at most $\Delta - c + 1$ neighbours outside D. Thus, the degree of the vertex we expanded can be at most $(\Delta - c + 2)(\Delta - c + 1)$. Furthermore, we have:

Observation 3 *For $c \geq 3$: if G is a $(c+1)$-critical graph with maximum degree Δ, and v is a vertex in G with degree at most $(\Delta - c + 2)(\Delta - c + 1)$, then we can construct a $(c+1)$-critical c-expansion of G with maximum degree at most Δ.*

Proof First note that $\Delta \geq c$ in any $(c+1)$-critical graph. Now take a $(c+1)$-clique C and a stable set S of size $\min\{\Delta - c + 2, \deg(v)\}$, and join all of C to all of S. Delete v and join every neighbour of v to a vertex in S so that each vertex in S is joined to at least one and at most $T = \min\{\Delta - c + 1, c - 1\}$ neighbours of v. This is possible if $|S| \leq \deg(v) \leq |S| \times T$; the lower bound is trivial, and the upper bound follows from hypothesis when $T = \Delta - c + 1$.

When $T < \Delta - c + 1$, we have $c - 1 < \Delta - c + 1$ and so $c < \frac{1}{2}\Delta + 1$ and $|S| = \Delta - c + 2 > \frac{1}{2}\Delta$. This yields $|S| \times T = |S|(c-1) > \Delta \geq \deg(v)$ for $c \geq 3$.

It is an easy exercise to verify that the resulting expansion is $(c+1)$-critical. □

In the other direction, we have:

Observation 4 *Every c-reduction of a $(c+1)$-critical graph is $(c+1)$-critical.*

We omit the easy proof.

Observation 3 implies that we can build larger and larger $(c+1)$-critical graphs with maximum degree Δ by repeatedly expanding a $(c+1)$-clique provided $c \leq \Delta$ and $(\Delta + 2 - c)(\Delta + 1 - c) \geq \Delta$; i.e. $c \leq \Delta - 1 - k_\Delta$. This proves Theorem 1.

On the other hand, if $c \geq \Delta + 1 - k_\Delta$ then $(\Delta - c + 2)(\Delta - c + 1) \leq k_\Delta(k_\Delta + 1) < \Delta + 1 - k_\Delta < c$. It follows that for such values of c, we cannot build $(c+1)$-critical graphs via c-expansions since the c-reducer formed by any c-expansion is deleteable. This gives us some hope that Theorem 2 might hold.

Finally reducers allow us to characterize for which values of Δ there are arbitrarily large $(c+1)$-critical graphs of maximum degree Δ at the threshold $c = \Delta - k_\Delta$.

In order to do so, we apply the following strengthening of Theorem 2 which is the main result of this paper:

Theorem 5 *There is an absolute constant Δ_0 such that for any $\Delta \geq \Delta_0$ and $c \geq \Delta - k_\Delta$, if G has maximum degree at most Δ, $\chi(G) = c + 1$ and either*

(a) $c \geq \Delta + 1 - k_\Delta$; or

(b) G has no c-reducer

then there is some vertex v in G such that the subgraph induced by $\{v\} \cup N(v)$ has chromatic number $c + 1$.

We don't specify Δ_0; instead we implicitly take it to be large enough to satisfy the many inequalities arising in our proof. We think that it can be omitted; i.e. that Theorem 5 holds for every Δ:

Conjecture 6 *The condition "$\Delta \geq \Delta_0$" can be removed from Theorem 5.*

Before turning to what happens at the threshold, we note that Theorem 5 easily implies Theorem 2 where H is the subgraph induced by $\{v\} \cup N(v)$, since $|H| = \deg(v) + 1 \leq \Delta + 1$.

We now show how Theorem 3 gives us a better understanding of what happens at the threshold $c = \Delta - k_\Delta$. Every $(c+1)$-critical subgraph G with maximum degree Δ arises from a graph of size at most $\Delta+1$ via a series of c-expansions using c-reducers. If $(k_\Delta+1)(k_\Delta+2) = \Delta$ then this series can be unlimited. If $(k_\Delta + 1)(k_\Delta + 2) \leq \Delta - 1$ this is not the case, as we now explain.

Corollary 7 *There exists Δ_0 such that for every constant $\Delta \geq \Delta_0$ and every $0 \leq k \leq \Delta$:*

(a) For $c \geq \Delta + 1 - k_\Delta$, every $(c + 1)$-critical graph has size at most $\Delta + 1$.

(b) For $c \leq \Delta - 1 - k_\Delta$, there are arbitrarily large $(c + 1)$-critical graphs.

(c) For $c = \Delta - k_\Delta$, every $(c + 1)$-critical graph can be formed by starting with a $(c + 1)$-critical graph of size at most $\Delta + 1$ and then applying a sequence of c-expansions. Furthermore:

(i) if $(k_\Delta + 1)(k_\Delta + 2) \leq \Delta - 1$ then every $(c + 1)$-critical graph has size at most $(\Delta + 1)^2 k_\Delta$;

(ii) if $(k_\Delta + 1)(k_\Delta + 2) = \Delta$ then there are arbitrarily large $(c + 1)$-critical graphs.

Proof Parts (a,b) follow immediately from Theorems 2 and 1. For part (c): set $c = \Delta - k_\Delta$ and consider a $(c+1)$-critical graph G. By Theorem 5, either G has a c-reducer, or G has a vertex v with $\chi(\{v\} \cup N(v)) = c + 1$. In the former case, G is a c-expansion of a smaller $(c+1)$-critical graph, by Observation 4. In the latter case, since G is $(c+1)$-critical, $G = \{v\} \cup N(v)$ and so $|G| \leq \Delta + 1$. This proves the first assertion of part (c).

To prove part (ii), begin with any $(c + 1)$-critical graph of maximum degree Δ (eg. a $(c + 1)$-clique), and repeatedly apply Observation 3 to expand on any vertex and obtain a larger $(c + 1)$-critical graph. We can expand on a vertex so long as it has degree at most $(k_\Delta + 1)(k_\Delta + 2)$ and since this value is Δ, we can expand on any vertex. Thus, we can carry out this procedure an arbitrary number of times.

To prove part (i), set $c = \Delta - k_\Delta$ and consider a $(c + 1)$-critical graph G_0. We will prove a limit on how many times it can be expanded. Suppose G_1 is a c-expansion of G_0 replacing a vertex v by a reducer D with corresponding partition into a clique C and a stable set S joined by all possible edges. Because G_0 is $(c + 1)$-critical, it has minimum degree at least c. So there are $\deg_{G_0}(v) \geq c = \Delta - k_\Delta$ edges out of D in G_1. By the properties of a reducer, the number of such edges is at most $|S|(\Delta + 1 - c)$, and $|S| \leq \Delta + 2 - c$. This implies $|S| = \Delta + 2 - c$ since $(\Delta + 1 - c)(\Delta + 1 - c) = (k_\Delta + 1)^2 < \Delta - k_\Delta - 1 < c$. Therefore every vertex of C has degree $(|C| - 1) + |S| = (c - 2) + (\Delta + 2 - c) = \Delta$ in G_1. Furthermore, $\sum_{w \in S} \deg(w) = |C||S| - \deg_{G_0}(v)$. Therefore, defining the *deficiency* of v to be $\mathrm{def}(v) = \Delta - \deg(v)$, we have:

$$\sum_{w \in S} \mathrm{def}(w) = \Delta|S| - |C||S| - \deg_{G_0}(v) = (\Delta+1-c)(\Delta+2-c) - \deg_{G_0}(v) > \Delta - \deg_{G_0}(v) = \mathrm{def}(v).$$

Thus, if v has degree $\Delta - 1$, then every vertex of D has degree Δ in G_1 and no further expansions are possible in this part of the graph. More generally, a simple inductive argument yields that the number of vertices in the subgraph obtained from expanding v and then recursively expanding the nodes in the reducer thereby obtained is at most $(\Delta + 1)\mathrm{def}(v)$. Thus, if we start with a $(c+1)$-critical graph with at most $\Delta + 1$ vertices, since this graph has minimum degree at least c and hence total deficiency at most $(\Delta + 1)(\Delta - c) = (\Delta + 1)k_\Delta$, we can only expand until it has size $(\Delta + 1)^2 k_\Delta$. $\square$

In [15] (see also [14]) we use Corollary 7 to explicitly list all $(c + 1)$-critical graphs for $c \geq \Delta - 5$ and $\Delta \geq \Delta_0$.

The following observation, along with hypothesis (b) of Theorem 5, allows us to ignore reducers in proving that theorem.

Observation 8 *If $D = (C, S)$ is a non-deleteable reducer in a graph with maximum degree Δ, and if $c \geq \Delta - k_\Delta$ then we must have $c = \Delta - k_\Delta$ and $|S| = k_\Delta + 2$.*

Proof Since D is non-deleteable, there are at least c edges from S to $G - D$. Every vertex in S has at most $\Delta - |C| = \Delta - c + 1$ neighbours in $G - D$ and by considering the degree of a vertex in C, we have $|S| \leq \Delta - c + 2$. Therefore $(\Delta - c + 1)(\Delta - c + 2) \geq c$. If $c \geq \Delta - k_\Delta + 1$ then this yields $k_\Delta(k_\Delta + 1) \geq \Delta - k_\Delta + 1$ and so $(k_\Delta + 2)(k_\Delta + 1) \geq \Delta - k_\Delta + 1 + 2(k_\Delta + 1) > \Delta$ which violates the definition of k_Δ. Therefore $c = \Delta - k_\Delta$ and so $|S| \leq \Delta - c + 2 = k_\Delta + 2$. If $|S| \leq k_\Delta + 1$ then we must have $(k_\Delta + 1)(\Delta - c + 1) \geq c$ and so $(k_\Delta + 1)(k_\Delta + 1) \geq \Delta - k_\Delta$ and thus $(k_\Delta + 2)(k_\Delta + 1) \geq \Delta - k_\Delta + (k_\Delta + 1) > \Delta$ which again violates the definition of k_Δ. $\square$

1.3 Algorithmic Implications

Another corollary of our main theorem is that for every constant $\Delta \geq \Delta_0$, we determine (under the hypothesis that $P \neq NP$) the precise values of c for which one can test in polynomial time whether a graph of maximum degree Δ is c-colourable. This is well-known to be trivial for $c \leq 2$. Embden-Weinert et al[11] used their construction (see Section 1.2) to prove that for $3 \leq c \leq \Delta - k_\Delta - 1$, we cannot test for c-colourability of graphs with maximum degree Δ in polytime unless $P = NP$. On the other hand, Theorem 5 easily implies that for every constant $\Delta \geq \Delta_0$ and every $c \geq \Delta - k_\Delta$, there is a linear time deterministic algorithm to test whether graphs of maximum degree Δ are c-colourable. Furthermore, there is a polynomial time deterministic algorithm that will produce a c-colouring whenever one exists.

For the case where Δ is not constant, the threshold for polynomial testability of c-colouring is (probably) higher: at $\Delta - \Theta(\log \Delta)$. We will give the formal statements and proofs of these results in Section 11.

2 A Preliminary Proof Sketch

Our proof combines probabilistic arguments with a structural decomposition.

2.1 The Probabilistic Method

The Lovasz Local Lemma is a powerful tool which allows us to prove the existence of colourings whose local behaviour is that which we would expect from a random colouring. We introduce it in the next section and illustrate its power by proving the following result (which is a key part of the proof of our main theorem).

Lemma 9 *There exists a Δ_0 such that for $\Delta \geq \Delta_0$, if H has max degree Δ and each of its vertices either*

(i) has fewer than $\Delta - 3\sqrt{\Delta}$ neighbours; or

(ii) has at least $900\Delta^{3/2}$ non-adjacent pairs of neighbours

then there is a $\Delta - 2\sqrt{\Delta}$ colouring of H.

To prove this lemma, we actually prove the following result which easily implies it:

Lemma 10 *There exists a* Δ_0 *such that for every* $\Delta \geq \Delta_0$*, and every graph* H *of maximum degree* Δ *the following holds:*

There is a $\Delta - 2\sqrt{\Delta}$ *colouring of a subgraph of* H *such that for every vertex* v *of* H *which has at least* $900\Delta^{3/2}$ *non-adjacent pairs of neighbours:*

$$\text{there are at least } 2\sqrt{\Delta}+1 \text{ colours which appear on two neighbours of } v. \tag{1}$$

Lemma 10 yields Lemma 9 as follows:

Proof of Lemma 9 We use the same Δ_0 as in Lemma 10. We apply Lemma 10 to obtain a partial colouring of H. We then attempt to extend our partial colouring to a $\Delta - 2\sqrt{\Delta}$ colouring of H by greedily colouring the uncoloured vertices of H in any order. We claim that throughout this process at most $\Delta - 2\sqrt{\Delta} - 1$ colours appear in the neighbourhood of any vertex, and hence this greedy procedure will succeed. Our claim clearly holds for vertices of degree less than $\Delta - 3\sqrt{\Delta}$. It holds for the other vertices by the hypothesis of Lemma 9 and the properties of the partial colouring returned by Lemma 10. □

To prove Lemma 10, we analyze the partial colouring in which each vertex is randomly assigned a uniform colour from $\{1, 2, ..., \Delta - 2\sqrt{\Delta}\}$, where these choices are independent, and retains the colour provided it is assigned to none of its neighbours. We show that the probability that (1) fails for a specific v in this random colouring is very small. The Local Lemma then allows us to show that there is a colouring where this local property holds for every v.

We leave the details of the proof of Lemma 10 to the next section, closing this subsection with two remarks.

The colourings guaranteed to exist only look like a random colouring locally. In a truly random colouring of a large enough graph, although most vertices would satisfy (1), there would with high probability be some vertex on which (1) would fail. For this reason, we call both the colourings and the process used to create them pseudorandom.

To prove that (1) holds for v with high probability we consider the random variable counting the number of colours appearing twice in the neighbourhood of v. We compute the expected value of this random variable and then use a concentration inequality to show that the variable is close to its expected value with high probability. Such concentration inequalities form part of our probabilistic toolbox presented in more detail in Section 3.

2.2 The Structural Decomposition

Our key structural result implies that every large $(c+1)$-critical graph either has a subgraph which is quite similar to a deleteable c-reducer, or is easy to handle using the technique of the previous section. We start with two definitions:

Definition: D is a *c-near-reducer* if it consists of a clique C with $c-1$ vertices, and a stable set T with $\Delta-c+1$ vertices such that every vertex of C is joined to every vertex of T. Note that this implies that each vertex of C has at most one neighbour outside of D.

Definition: D is a *c-quasi-reducer* in G if it consists of a clique C with between c and $\Delta - 10^8\sqrt{\Delta}$ vertices and a set of $l \le c-|C|$ stable sets $T_1, ..., T_l$ such that for each T_i:

(a) the set of vertices outside D which have a neighbour in T_i is less than c;

(b) if $|T_i| > 2$ then every vertex of T_i sees all of C; and

(c) if $|T_i| = 2$ then either

  (i) there is $z \in C$ such that both vertices of T_i see all of $C - z$, or

  (ii) one vertex of T_i sees all of C, and the other vertex of T_i sees at least $\frac{2\Delta}{3}$ of the vertices of C.

Again, we often drop the "c" when its value is clear from context. Note that every c-reducer is a c-near-reducer and every c-near-reducer is a c-quasi-reducer.

Lemma 11 *There is a Δ_0 such that for $\Delta \ge \Delta_0$ and $c \ge \Delta - k_\Delta$ the following holds:*

If G is a $(c+1)$-critical graph of maximum degree at most Δ which has no vertex that is adjacent to all other vertices, then G can be partitioned into $X_1, ..., X_t, S$ such that

(a) each X_i induces a quasi-reducer, and

(b) each vertex of S satisfies hypothesis (i) or (ii) of Lemma 9.

Remark: Note that any $(c+1)$-critical graph with more than $\Delta+1$ vertices satisfies the conditions of Lemma 11.

We present this result in Section 4, where we strengthen it via a sequence of lemmas which imposes increasingly stronger conditions on the quasi-reducers of the decomposition. We close this section with a few more definitions.

Definition: An *internal neighbour* of $v \in X_i$ is a neighbour of v that is also in X_i. All neighbours of v in $G - X_i$ are *external neighbours*. For a stable set ρ of X_i, we say a vertex u is an *(internal/external) neighbour* of ρ if it is an (internal/external) neighbour of a vertex in ρ.

2.3 Putting It All Together

Our approach to proving Theorem 5 is straightforward. We fix $\Delta \geq \Delta_0$ and $c \geq \Delta - k_\Delta$ and let G be a minimum counterexample with respect to Δ, c; i.e. a counterexample with the smallest possible number of vertices. G clearly satisfies the conditions of Lemma 11, so we can apply that lemma to obtain a decomposition of G. Applying Lemma 9, we can c-colour the subgraph of G induced by S. We would now like to extend this colouring to the quasi-reducers, thereby completing it.

In doing so, we would like to mimic our approach for deleteable c-reducers from Section 1.2. As in that case, we contract each stable set into a single vertex and colour these vertices before we colour the cliques of the quasi-reducers.

The difficulty is that the vertices in the clique of a quasi-reducer may have neighbours outside the quasi-reducer. Indeed a quasi-reducer could simply consist of a clique C of size c. Each of its vertices then could have $\Delta - c$ external neighbours. If there was a colour i appearing on an external neighbour of every vertex of C, then we could not use this colour on C and hence would need at least $c + 1$ colours to complete the colouring.

As this specific example illustrates, in choosing our intial colouring of S, and in extending it, we need to take care that for each clique of a quasi-reducer, no colour which does not appear in the quasi-reducer is used on the external neighbourhood of a large majority of the vertices in the clique. In order to do so we:

(A) Massage the graph by adding edges between pairs of vertices that we want to get different colours and contracting a stable set into a vertex if we want its members to get the same colour. This results in a graph F, such that a c-colouring of F easily yields a c-colouring of G.

(B) Choose our initial colouring of S and its extension using the probabilistic method. We can thereby ensure that its local behaviour is what we would expect of a random colouring. Here, we focus on one specific aspect of the behaviour: for each clique of a quasi-reducer and colour i, how many vertices of the clique have an external neighbour of colour i?

(C) Interleave the colouring of the big stable sets of the quasi-reducers with the colouring of the cliques within them. The reason for doing so will become clear later. For the moment, the reader may choose to ignore this last complication. Doing so will only aid her intuition as to the structure of the proof.

More specifically, we prove Theorem 5 by combining two lemmas:

Lemma 12 *For any minimum counterexample G to Theorem 2, we can find a graph F of maximum degree $10^8\Delta$, whose c-colourability implies the c-colourability of G, and a partition of the vertices of F into $S, B, A_1, ..., A_t$ such that:*

(a) Every A_i is a clique with $c - 10^8\sqrt{\Delta} \leq |A_i| \leq c$.

(b) Every vertex of A_i has at most $10^8\sqrt{\Delta}$ neighbours in $F - A_i$.

(c) There is a set $\mathrm{All}_i \subseteq B$ of $c - |A_i|$ vertices which are adjacent to all of A_i. Every other vertex of $F - A_i$ is adjacent to at most $\frac{3\Delta}{4} + 10^8\sqrt{\Delta}$ vertices of A_i.

(d) Every vertex of S either has fewer than $\Delta - 3\sqrt{\Delta}$ neighbours in S or has at least $900\Delta\sqrt{\Delta}$ non-adjacent pairs of neighbours within S.

(e) Every vertex of B has fewer than $c - \sqrt{\Delta} + 9$ neighbours in $F - \cup_j A_j$.

(f) If a vertex $v \in B$ has at least $c - \Delta^{\frac{3}{4}}$ neighbours in $F - \cup_j A_j$, then there is some i such that: v has at most $c - \sqrt{\Delta} + 9$ neighbours in $F - A_i$ and every vertex of A_i has at most $30\Delta^{\frac{1}{4}}$ neighbours in $F - A_i$.

(g) For every A_i, every two vertices outside of $A_i \cup \mathrm{All}_i$ which have at least $2\Delta^{\frac{9}{10}}$ neighbours in A_i are joined by an edge of F.

Lemma 13 *Any graph F with a decomposition $\{S, B, A_1, ..., A_t\}$ as in Lemma 12, has a c-colouring.*

These yield our main theorem:

Proof of Theorem 5: We choose G to be a minimum counterexample to Theorem 5. Lemmas 12 and 13 imply that $\chi(G) \leq c$, which is a contradiction. □

In the next section, we will present our probabilistic tools. We give the proof of Lemma 12 in Sections 4 and 5. We prove Lemma 13 using a pseudorandom iterative colouring procedure in Sections 6 to 10. We close by discussing algorithmic implications in Section 11.

3 The Probabilistic Method and a First Application

3.1 Some Powerful Tools

The following lemma has a three line proof but nevertheless is extremely powerful and has had scores of applications. One of its applications is in this paper.

The Lovasz Local Lemma[13] *Let* $A_1, ..., A_n$ *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* $\mathbf{Pr}(\overline{A_1} \cup ... \cup \overline{A_n}) > 0$.

The lemma is powerful because it allows us to deduce global results via a local analysis. The best way to explain this is via an example. We give one in the next subsection, where we prove Lemma 10.

The binomial random variable $BIN(n, p)$ is the sum of n independent $0-1$ random variables each of which is equal to 1 with probability p. The following special case of Chernoff's original bound[8] can be found in [25]:

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

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

The following is a simple corollary of Hoeffding's Inequality[17] or Azuma's Inequality[2], as described in [25]:

Simple Concentration Bound: *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*.*

Then for any $0 \leq t \leq \mathbf{Exp}(X)$*, we have*

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

Talagrand's Inequality requires another condition, but often provides a stronger bound when $\mathbf{Exp}(X)$ is much smaller than n. Rather than providing Talagrand's original statement from [31], we present the following useful reworking, which we prove in an appendix:

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 $t \geq 0$, we have

$$\mathbf{Pr}(|X - \mathbf{Exp}(X)| > t + 20c\sqrt{r\mathbf{Exp}(X)} + 64c^2r) \leq 4e^{-\frac{t^2}{8c^2r(\mathbf{Exp}(X)+t)}}.$$

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. The derivation can also be found in the appendix. Talagrand[31] derived a similar result 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[22] *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 $t \geq 0$, we have

$$\mathbf{Pr}(|X - \mathbf{Exp}(X)| > t + 25c\sqrt{r\mathbf{Exp}(X)} + 128c^2r) \leq 4e^{-\frac{t^2}{32c^2r(\mathbf{Exp}(X)+t)}}.$$

Note that in both Talagrand's Inequality and McDiarmid's Inequality, if $t \geq 50c\sqrt{r\mathbf{Exp}(X)})+ 256c^2r$ then by substituting $t/2$ for t in the above bounds, we obtain the more concise:

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

and for McDiarmid's Inequality:

$$\mathbf{Pr}(|X - \mathbf{Exp}(X)| > t) \leq 4e^{-\frac{t^2}{128c^2r(\mathbf{Exp}(X)+t)}}. \tag{3}$$

Those are the bounds that we will usually use.

3.2 Proof of Lemma 10

Remark: The arguments in this section have appeared in several other papers. The reader who is familiar with, eg., [28, 25] may choose to skip to the next section.

We obtain a partial colouring of H as follows:

1. We activate each vertex of H with probability $\frac{9}{10}$.
2. Each activated vertex in H is assigned a uniformly random colour from $\{1, ..., c\}$. These choices are made independently.
3. We uncolour every vertex that has a neighbour with the same colour.

We will apply the Lovasz Local Lemma to prove that with positive probability the resultant colouring satisfies condition (1) of Lemma 10. For each $v \in H$ with at least $900\Delta^{3/2}$ non-adjacent pairs of neighbours, define $E_1(v)$ to be the event that v has fewer than $3\sqrt{\Delta}$ colours that appear at least twice in its neighbourhood. We will prove below that $\mathbf{Pr}(E_1(v)) \leq \Delta^{-10}$.

It is straightforward to check that each $E_1(v)$ is mutually independent of all the $E_1(w)$ for vertices w at distance more than 4 from v in H. Thus, each event is mutually independent of all but fewer than Δ^4 other events. Thus, our lemma follows from the Lovasz Local Lemma since $\Delta^{-10} \times \Delta^4 < \frac{1}{4}$.

We now fix a $v \in H$ and bound $\mathbf{Pr}(E_1(v))$. We let Ω be a collection of $900\Delta^{3/2}$ pairs of non-adjacent neighbours of v. We consider the random variable Y which counts the number of pairs in Ω which (i) are both assigned the same colour, (ii) both retain that colour, and (iii) are the only two vertices in $N(v)$ that are assigned that colour. Clearly Y is a lower bound on the number of colours appearing at least twice in $N(v)$.

The probability that some non-adjacent pair $u, w \in N(v)$ satisfies (i) is $\frac{9}{10} \times \frac{9}{10} \times \frac{1}{c}$. The total number of neighbours of v, u, w in H is at most 3Δ. Given that they satisfy (i), u, w also satisfy (ii) and (iii) if none of those vertices are activated and assigned the colour of u, v, and this occurs with probability at least $(1 - \frac{1}{c})^{3\Delta}$. Therefore,

$$\mathbf{Exp}(Y) \geq 900\Delta^{3/2} \times \frac{81}{100c} \times \exp(-3\Delta/c) > 4\sqrt{\Delta}.$$

So if $E_1(v)$ holds then Y must differ from its mean by at least $\sqrt{\Delta}$.

We will apply Talagrand's Inequality to show that Y is highly concentrated. To do so, we consider two related variables: Y_1 is the number of colours assigned to both members of at least one pair in Ω; Y_2 is the number of colours that are (i) assigned to both members of at least one pair in Ω and (ii) also assigned to one of their neighbours or to at least one other vertex in $N(v)$. Note that $Y = Y_1 - Y_2$. Thus, if $E_1(v)$ holds then either Y_1 or Y_2 must differ from its mean by at least $\frac{1}{2}\sqrt{\Delta}$. Note that:

$$\mathbf{Exp}(Y_2) \leq \mathbf{Exp}(Y_1) \leq c \times (900\Delta^{3/2}) \times \left(\frac{9}{10}\right)^2 \left(\frac{1}{c}\right)^2 < 1200\sqrt{\Delta}.$$

If $Y_1 \geq s$ then there is a set of at most $4s$ trials whose outcomes certify that $Y_1 \geq s$, namely the activation and colour assignment for s pairs of variables. Also, changing the outcome of any individual random trial can only affect Y_1 by at most 2 since at worse it affects whether Y_1 counts the old colour and the new colour of a vertex whose colour is changed. Therefore, applying (2) with $c = 2, r = 4$ yields:

$$\mathbf{Pr}(|Y_1 - \mathbf{Exp}(Y_1)| > \frac{1}{2}\sqrt{\Delta}) \leq 4\exp\left(-\tfrac{1}{4}\Delta/(32 \times 4 \times 4 \times 1201\sqrt{\Delta})\right) < \frac{1}{2}\Delta^{-10}.$$

Similarly, if $Y_2 \geq s$ then there is a set of at most $6s$ trials whose outcomes certify that $Y_2 \geq s$: the activation and colour assignment for s pairs of variables and for each pair the activation and assignment of the same colour to a neighbour of one member of the pair or to another member of $N(v)$. Also, changing the outcome of any individual random trial can only affect Y_2 by at most 2 for the same reason as for X_1. Therefore, applying (2) with $c = 2, r = 6$ yields:

$$\mathbf{Pr}(|Y_2 - \mathbf{Exp}(Y_2)| > \frac{1}{2}\sqrt{\Delta}) \leq 4\exp\left(-\tfrac{1}{4}\Delta/(32 \times 4 \times 6 \times 1201\sqrt{\Delta})\right) < \frac{1}{2}\Delta^{-10}.$$

Therefore, $|\mathbf{Pr}(Y - \mathbf{Exp}(Y)| > \sqrt{\Delta} \leq \Delta^{-10}$. □

4 The Structural Decomposition

In this section we prove Lemma 11. In addition, we show that in a minimum counterexample to Theorem 5, none of the X_i can be reducers or near-reducers. Thus we assume throughout that G is $(c+1)$-critical and contains no vertex that is adjacent to all the other vertices.

We start with a decomposition introduced in [28]. The following result is essentially Lemma 15.2 of [25], and actually holds for all graphs of maximum degree Δ, for any Δ.

Definition: We set $d = 10^6\sqrt{\Delta}$ and, following Section 15.2 of [25], call a vertex *sparse* if it has at most $\binom{\Delta}{2} - d\Delta$ edges in its neighbourhood.

Lemma 14 *There is a partition of $V(G)$ into sets $X_1, ..., X_t$ and S such that:*

(a) $\forall i$, $\Delta - 10^6\sqrt{\Delta} \leq |X_i| \leq \Delta + 10^6\sqrt{\Delta}$;

(b) $\forall i$, there are at most $10^7\Delta^{3/2}$ edges from X_i to $G - X_i$;

(c) $\forall i$, a vertex is adjacent to at least $\frac{3\Delta}{4}$ vertices of X_i iff it is in X_i;

(d) every vertex in S is sparse.

We note that every sparse vertex satisfies (i) or (ii) of Lemma 9. Thus, to prove Lemma 11 given Lemma 14, it is enough to prove that each X_i is a quasi-reducer. To begin, we show that under this hypothesis each X_i has a c-colouring.

Lemma 15 *For every $i, \chi(X_i) \leq c$.*

Proof Suppose X_i is not c-colourable. Since G is $(c+1)$-critical, we must have $G = X_i$. We let M be a maximum matching in $\overline{G}$, chosen subject to this to maximize the sum of the degrees (in G) of the vertices within it. M has fewer than $|G| - c$ edges as otherwise, we can colour G with c colours using the edges of M as the only non-singleton colour classes. Thus, by Lemma 14(a), M has at most $10^7\sqrt{\Delta}$ edges, and the clique C induced by the vertices of G not in M has at least $\Delta - 10^7\sqrt{\Delta}$ vertices.

By the maximality of M, for every edge of M, either (i) one endpoint is adjacent to every vertex in C or (ii) for some z in C both endpoints of M are adjacent to every vertex in $C - z$. Since G is $(c+1)$-critical, it has minimum degree at least c, and so each vertex in M is non-adjacent to at most $(|G| - c)$ vertices. It follows that for each edge of M, the endpoints have a total of at least $2|C| - (|G| - c)$ neighbours in C and so there are at least $2|C||M| - |M|(|G| - c)$ edges from C to M. Thus the average degree of the vertices in C is at least $|G| - 1 - \frac{|M|(|G|-c)}{|C|}$.

By our bounds on the size of M and G, if Δ is large enough, this is at least $|G| - 1 - 10^{15}$. Thus, $|G| \leq \Delta + 1 + 10^{15}$. This bound on $|G|$ implies that both $|M|$ and $|G| - c$ are at most $\Delta + 1 + 10^{15} - c \leq \sqrt{\Delta} + 10^{16}$. By exploiting the bound from the last paragraph on the number of edges from M to C, we obtain that for large enough Δ, the average degree of the vertices in C is at least $2|M| - \frac{|M|(|G|-c)}{|C|} > |G| - 3$. It follows that G has at most $\Delta + 2$ vertices. Since $|M| \leq |G| - (c+1) \leq \Delta + 1 - c \leq k_\Delta + 1$ this yields that $|C| \geq c + 1 - |M| \geq c - k_\Delta \geq \Delta - 2k_\Delta$.

We now count the non-edges between C and the vertices of M again, being a bit more careful. Consider some edge $e = uv$ of M. The number of edges from $\{u, v\}$ to C is at least $2|C| - 2$ in case (ii) above, but can be smaller in case (i) if u is adjacent to all of C and v is adjacent to as much of $G - C$ as possible. Since v is non-adjacent to u, v is adjacent to at least $\deg(v) - 2|M| + 2 \geq c - 2|M| + 2$ vertices in C. Hence there are at most $|C| + 2(|M| - 1) - c \leq |G| - c - 2$ non-edges between any edge of M and C. So, letting $\overline{E}(C, M)$ be the number of nonedges between C and $V(M)$, we have:

$$|\overline{E}(C,M)| \leq |M|(|G|-c-2) \leq (k_\Delta+1)k_\Delta \leq (k_\Delta+1)(k_\Delta+2) - 2k_\Delta - 2 \leq \Delta - 2k_\Delta - 2 < |C|.$$

Thus, some vertex $v \in C$ is adjacent to all of $V(M)$ and so v is adjacent to all vertices in G (other than v), thus contradicting the hypothesis of Lemma 11. □

We next prove that in any c-colouring of an X_i, the majority of the colour classes are singletons which together form a clique.

Lemma 16 *For each i, $\overline{G[X_i]}$ has no matching of size $\lceil 10^2\sqrt{\Delta}\rceil$.*

Proof Suppose that M is a matching of size $\lceil 10^2\sqrt{\Delta}\rceil$ in $\overline{G[X_i]}$. Let R be the unmatched vertices in X_i; by Lemma 14(a), $\Delta - 10^6\sqrt{\Delta} < |R| < \Delta + 10^6\sqrt{\Delta}$. For each pair u, v that are matched in M, the number of neighbours of u in X_i plus the number of neighbours of v in X_i is at least $3\Delta/2$, by Lemma 14(c). Thus there are at least $3\Delta/2 - 4|M| - |R| > |R|/3$ vertices in R that are adjacent to both of u, v. So on average, a vertex of R is adjacent to both members of at least $|M|/3$ pairs. This implies that at least $|R|/5 > \Delta/10$ members of R are adjacent to both members of at least $|M|/10$ pairs. Let Z be $\Delta/10$ such vertices in R.

Any vertex of $R - Z$ that is adjacent to less than half of Z must have at least $\Delta - (|X_i| - \frac{1}{2}|Z|) > \Delta/25$ neighbours outside of X_i. Thus, Lemma 14(b) implies that there are at least $\Delta/2$ vertices in $R - Z$ which are adjacent to at least half of Z. Let Y be a set of $\Delta/2$ such vertices.

Since G is $(c+1)$-critical, we can c-colour $G - X_i$. We will extend that c-colouring to X_i greedily as follows, thus obtaining a contradiction.

1. Colour the vertices of M, assigning the same colour to both members of each matched pair. This is possible because, by Lemma 14(c), each pair has at most $\Delta/2 + |M| < c$ previously coloured neighbours.

2. Colour the vertices of $X_i - Y - Z - M$. This is possible since each such vertex has at most $\Delta/4$ neighbours outside of X_i (by Lemma 14(c)) and at most $|X_i| - |Z| - |Y| < \Delta/2$ previously coloured neighbours in X_i.

3. Colour the vertices of Y. This is possible since each vertex of Y has at least $\frac{1}{2}|Z| \geq \Delta/20$ uncoloured neighbours and hence at most $19\Delta/20 < c$ coloured neighbours.

4. Colour the vertices of Z. This is possible since each vertex of Z has at least $|M|/10 \geq 10\sqrt{\Delta}$ colours that appear twice in its neighbourhood, and thus has at most $\Delta - 10\sqrt{\Delta} < c$ colours appearing in its neighbourhood.

□

We now describe how to construct a particular c-colouring of each X_i, which will demonstrate that X_i is a quasi-reducer. There are two cases.

Construction

Case 1: $\overline{G[X_i]}$ has a matching of size at least $|X_i| - c$.

Let M be a maximum matching in $\overline{G[X_i]}$. Each edge of M forms a colour class of size 2. The remaining vertices are singleton colour classes.

Case 2: $\overline{G[X_i]}$ has no matching of size at least $|X_i| - c$.

Note that $|X_i| > c$ or else we would trivially be in Case 1. This, plus the fact that $\chi(G[X_i]) \leq c$, means that we can take a colouring of $G[X_i]$, which uses *exactly* c colours. For each j let λ_j be the number of colour classes of size j. Note that, since $|X_i| \leq \Delta + 10^6\sqrt{\Delta}$ (by Lemma 14(a)) and since $c \geq \Delta - k_\Delta > \Delta - \sqrt{\Delta} + 1$, we can assume that $\lambda_j = 0$ for each $j > (10^6 + 1)\sqrt{\Delta}$.

We choose our colouring to be optimal with respect to the following criteria amongst all colourings using exactly c colours.

1. First, λ_1 is as small as possible.
2. Second, subject to 1, the sequence $\lambda_{10^8\sqrt{\Delta}}, ..., \lambda_2$ is lexicographically minimum.

End of Construction

Definition: Let C_i denote the set of colour classes in X_i of size 1.

We now use these colourings to show that each X_i is a quasi-reducer with clique C_i and whose stable sets are the colour classes of size at least two. We will need:

Observation 17 *If our colouring was constructed via Case 2, then there was at least one colour class of size greater than 2.*

Proof If every colour class has size at most 2, then the colour classes of size 2 form a matching in $\overline{G[X_i]}$ of size exactly $|X_i| - c$; this would put us in Case 1. □

Observation 17 implies:

Lemma 18 *C_i is a clique in G of size at least $\Delta - 2 \times 10^6\sqrt{\Delta}$.*

Proof First we show that C_i is a clique in G. This is true if we carry out the reduction described in Case 1, since M is a maximum matching and $|M| \leq 10^2\sqrt{\Delta}$ by Lemma 16. So we can assume that we carry out the reduction of Case 2. By Observation 17, there is at least one colour class ρ of size at least 3. Suppose that $u, v \in C_i$ are not adjacent. Then by making $\{u, v\}$ a colour class and by splitting ρ into 2 colour classes, we obtain a colouring using exactly c colours and with fewer classes of size 1. Therefore C_i is a clique.

Note that if we performed the reduction of Case 2 then since all classes outside of C_i have size at least two, and since we have exactly c colour classes, we have $C_i + 2(c - |C_i|) \leq |X_i| \leq \Delta + 10^6\sqrt{\Delta}$ (by Lemma 14(a)). So $C_i \geq 2c - (\Delta + 10^6\sqrt{\Delta}) > \Delta - 2 \times 10^6\sqrt{\Delta}$. On the other hand, if we performed the reduction of Case 1 then by Lemma 16, $|C_i| \geq |X_i| - 2 \times 10^2\sqrt{\Delta} > \Delta - 2 \times \lceil 10^6\sqrt{\Delta}\rceil$ (by Lemma 14(a)). □

We also have:

Lemma 19 *Every vertex in C_i is adjacent to every vertex in any colour class of X_i that has size at least 3.*

Proof Suppose that $v \in C_i$ is not adjacent to $u \in \rho$ where ρ is a colour class of size at least 3. Thus we applied Case 2 of our construction to X_i. By making $\{u, v\}$ a colour class and $\rho - u$ a colour class, we obtain a colouring using exactly c colours and with fewer classes of size 1, thus contradicting our choice of the colouring of X_i. □

Lemma 20 *Suppose that $\{u, v\}$ is a colour class. Then either*

(i) there is some $x \in C_i$ such that u, v are both adjacent to all of $C_i - x$; or

(ii) one of $\{u, v\}$ is adjacent to all of C_i, and the other is adjacent to all but at most $\frac{\Delta}{4} + 10^7\sqrt{\Delta}$ vertices of C_i.

Proof *Case A:* One of the vertices is adjacent to all of C_i; w.l.o.g. assume it is u. By Lemma 14(c), v is adjacent to at least $\frac{3\Delta}{4}$ vertices in X_i. At most $|X_i| - |C_i| < 10^7\sqrt{\Delta}$ (by Lemmas 14(a) and 18) of them are not in C_i. This yields (ii), since $|C_i| \leq \Delta$.

Case B: There are $x, y \in C_i$ so that u is not adjacent to x and v is not adjacent to y. If $x \neq y$ then we know that we did not perform Case 1 of the reduction since replacing uv by ux and vy would yield a matching larger than M. So by Observation 17, there must be a colour class ρ of size at least three in the colouring. But by making $\{u, x\}$ and $\{v, y\}$ two colour classes, and splitting ρ in two, we obtain a colouring using c colours with fewer classes of size 1. Thus we must have $x = y$ and so (i) holds. □

Lemmas 18, 19 and 20 establish that X_i satisfies most of the definition of a quasi-reducer. All that remains is to bound the number of external neighbours of the classes of size at least two. We will need the following two lemmas.

Lemma 21 *Consider any two colour classes ρ_1, ρ_2 where $|\rho_1| - 2 \geq |\rho_2| \geq 2$. Then every vertex in ρ_1 has at least one neighbour in ρ_2.*

Proof Suppose that $v \in \rho_1$ has no neighbour in ρ_2. Then by moving v from ρ_1 to ρ_2, we obtain a colouring using exactly c colours in which the sequence $\lambda_{(10^6+1)\sqrt{\Delta}}, ..., \lambda_2$ is lexicographically smaller, thus contradicting our choice of the colouring in our construction. □

Recall from our construction that λ_ℓ is the number of colour classes of size ℓ in X_i.

Lemma 22 *For any X_i that has a colour class of size at least 3:*

$$\sum_{\ell \geq 3} (\ell - 1)\lambda_\ell \leq \Delta - c + 1 - \frac{2}{3}\lambda_2.$$

Proof By Lemma 17, we know that $|C_i| \geq \Delta - 2 \times 10^6\sqrt{\Delta}$. So, by Lemma 20, The number of edges between C_i and a colour class of size 2 is at least $|C_i|(2 - \frac{\Delta}{4} - 2 \times 10^6\sqrt{\Delta}) > \frac{5}{3}|C_i|$. By Lemma 19, every vertex of C_i is adjacent to every member of every colour class of size greater than 2. Lemma 18 implies that every vertex of C_i is also adjacent to every other vertex of C_i. Combining these results implies that the average degree of the vertices of C_i in $G[X_i]$ is at least $|C_i| - 1 + \frac{5}{3}\lambda_2 + \sum_{\ell \geq 3} \ell\lambda_\ell \leq \Delta$, since the average degree of a vertex in C_i is at most Δ.

Since there is a colour class of size at least 3, we are in Case 2 of the construction. So $|C_i| + \sum_{\ell \geq 2} \lambda_\ell = c$ The lemma follows by subtracting this inequality from the previous one. □

Lemmas 21 and 22 allow us to bound the sizes of the colour classes:

Lemma 23 *For any X_i for which we applied Case 2 of the construction, there is no colour class of size greater than $\Delta - c + 2$. Furthermore, if there is a colour class of size $\Delta - c + 1$ then X_i is a near-reducer and if there is a colour class of size $\Delta - c + 2$ then X_i is a reducer.*

Proof If there is a colour class of size at least $\Delta - c + 2$, then since there are c colour classes in total, applying Lemma 22, we obtain that this is the only colour class of size exceeding 1. By Lemma 19 this implies that X_i is a c-reducer.

Assume there is a colour class ρ of size $\Delta - c + 1$. By the preceding paragraph, we can assume that there is no colour class of size at least $\Delta - c + 2$. So since we applied Case 2 of the construction, we must have $|\rho| \geq 3$. By Lemma 22 either this is the only colour class of size exceeding 1 or there is one other such colour class and it has size 2. By Lemma 19, if X_i is not a near-reducer then there is such a colour class ρ' of size 2. By Lemma 20, and the fact that $|C_i| \geq \frac{2\Delta}{3}$ (by Lemma 18), there is a set Z of at least $\frac{\Delta}{4}$ vertices of C_i which see all of ρ'. By Lemma 19, these vertices also see all of ρ. If $|\rho| > 3$ then Lemmas 19 and 21 imply that every vertex in ρ has at least $c-1$ neighbours in X_i and hence at most $\Delta - c + 1$ neighbours outside X_i. It follows that there are at most $(\Delta-c+1)(\Delta-c+1) \leq (k_\Delta+1)^2 \leq \Delta-k_\Delta-1 < c$ vertices of $G - X_i$ adjacent to vertices of ρ. If $|\rho| = 3$ then Lemma 19 implies that every vertex in ρ has at least $c-2$ neighbours in X_i, and so there are at most $3(\Delta - c + 2) < c$ vertices of $G - X_i$ adjacent to vertices of ρ.

Since G is $(c+1)$-critical, we can c-colour $G - X_i$. We extend this to a c-colouring of G as follows: (a) First colour all of ρ with a colour not appearing on any of the fewer than c vertices of $G - X_i$ adjacent to ρ. (b) Then colour the two vertices of ρ' with a colour not appearing on any of their at most $2 \times \frac{\Delta}{4}$ neighbours in $G - X_i$ (by Lemma 14(c)) nor on ρ. (c) Next colour all of $C_i - Z$ which is possible since each vertex in this set has $\frac{\Delta}{4}$ uncoloured neighbours in Z. (d) Finally colour Z which is possible since each vertex of Z sees all the $\Delta - c + 3$ vertices of $\rho \cup \rho'$ and those are coloured using two colours. This contradiction implies the desired result. □

And finally we bound the size of the external neighbourhoods of the colour classes of size at least two. Parts (a,b) of the following lemma are enough to prove that X_i is a quasi-reducer. The remaining parts will be used in the proof of Lemma 12.

Lemma 24 *(a) If a colour class is not the unique largest colour class in X_i, then it has at most $\frac{\Delta}{2} + 10\sqrt{\Delta}$ external neighbours.*

(b) If X_i is not a reducer or near-reducer then every colour class ρ of X_i, has at most $c - \sqrt{\Delta} + 3$ external neighbours.

(c) If X_i is not a reducer or near-reducer and there is a colour class ρ of X_i with more than $c - 10^8\sqrt{\Delta}$ external neighbours then $|C_i| \geq c - 2 \times 10^8$ and each vertex of C_i has at most 3×10^8 external neighbours.

(d) If X_i is not a reducer or near-reducer and there is a colour class ρ of X_i with more than $c - 2\sqrt{\Delta} + 3$ external neighbours then $|C_i| = c - 1$ and each vertex of C_i has at most 5 external neighbours.

(e) *If X_i is not a reducer or near-reducer and there is a colour class ρ of X_i with more than $c - 2\Delta^{3/4}$ external neighbours then $|C_i| \geq c - 5\Delta^{1/4}$ and each vertex of C_i has at most $8\Delta^{1/4}$ external neighbours.*

Proof We first prove (a). Every vertex v which is a singleton colour class has at most $\frac{\Delta}{4}$ external neighbours by Lemma 14(c) and hence satisfies (a).

If $|\rho| = 2$ then by Lemma 20, at least one member of ρ has at least $|C_i| - 1 > \Delta - 10^7\sqrt{\Delta}$ (by Lemma 18) neighbours within X_i and thus at most $10^7\sqrt{\Delta}$ external neighbours. By Lemma 14(c), the other member of ρ has at most $\frac{\Delta}{4}$ external neighbours. Thus they have a total of at most $\frac{\Delta}{4} + 10^7\sqrt{\Delta}$ external neighbours which is less than $\frac{\Delta}{2}$ for Δ sufficiently large. So any such colour class satisfies (a).

So we assume $|\rho| \geq 3$. By Lemma 19, every vertex of ρ sees all of C_i. So we know that there are at most $|\rho|10^7\sqrt{\Delta}$ edges from ρ to vertices outside X_i. So we are done unless $|\rho| > \frac{1}{2\times 10^7}\sqrt{\Delta}$. Suppose that there are j colour classes (including ρ) of size at least $|\rho| - 1$. By Lemma 22, $j(|\rho| - 2) \leq \Delta - c + 1$ and so $|\rho| \leq (\Delta - c + 1 + 2j)/j$ and $j \leq \Delta - c + 1 \leq \sqrt{\Delta}$. By Lemmas 19 and 21 each vertex in ρ is adjacent to a member of each of the $c - j$ colour classes of size less than $|\rho| - 1$ and so has at most $\Delta - c + j$ external neighbours. Therefore, the number of external neighbours of ρ is at most $(\Delta - c + j)(\Delta - c + 1 + 2j)/j$. For j between 2 and $\sqrt{\Delta}$, this is easily seen to be at most $\frac{1}{2}\Delta + 10\sqrt{\Delta}$ since $c \geq \Delta - \sqrt{\Delta} + 2$. If $j = 1$ then ρ is the unique largest colour class, and so this finishes the proof of part (a).

We turn now to (b), (c), and (d). Since $\frac{1}{2}\Delta + 10\sqrt{\Delta} \leq c - 10^8\sqrt{\Delta}$, by the above remarks we can restrict our attention to unique largest colour classes ρ such that $|\rho| \geq 3$ and $j = 1$.

Since X_i is not a reducer or near-reducer, Lemma 23 implies $|\rho| \leq \Delta - c$.

We first focus on the case $|\rho| \leq \Delta - c - 1$. Since $j = 1$, we know by Lemma 21 that every vertex of ρ has at least $c - 1$ neighbours in X_i and hence at most $\Delta - c + 1$ neighbours outside X_i. So, using the facts that $c \geq \Delta - k_\Delta$, $(k_\Delta + 1)(k_\Delta + 2) \leq \Delta$ and $k_\Delta > \sqrt{\Delta} - 3$, the number of external neighbours of ρ is at most

$$|\rho|(\Delta - c + 1) \leq (k_\Delta - 1)(k_\Delta + 1) < (k_\Delta + 2)(k_\Delta + 1) - (k_\Delta + 1) - 2(\sqrt{\Delta} - 2) \leq c - 2\sqrt{\Delta} + 3.$$

Thus (b,d) hold for this case. Furthermore, if $|\rho| \leq \Delta - c - 10^8$ then ρ has at most $|\rho|(\Delta - c + 1) \leq (k_\Delta - 10^8)(k_\Delta + 1) < c - 10^8\sqrt{\Delta}$ external neighbours and so (c) holds.

Suppose that $|\rho| > \Delta - c - 10^8$. By Lemma 22, $\frac{2}{3}\lambda_2$ plus the sum of $|\rho'| - 1$ over all colour classes $\rho' \neq \rho$ with $|\rho'| \geq 3$ is at most $\Delta - c + 1 - (|\rho| - 1) \leq 10^8 + 1$. Each colour class of size at least 2, other than ρ, contributes at least $\frac{2}{3}$ to that sum. Thus, there are at most

$\frac{3}{2} \times (10^8+1)+1 < 2\times 10^8 - 1$ colour classes of size at least 2 and so $|C_i| \geq c - 2\times 10^8 + 1$. By Lemmas 18 and 19, each vertex of C_i has at least $|C_i| - 1 + |\rho| \geq \Delta - c - 10^8 + c - 2\times 10^8 = \Delta - 3\times 10^8$ neighbours in X_i and hence has at most 3×10^8 neighbours outside of X_i. Thus (c) holds (even when $|\rho| = \Delta - c$).

We now turn to the remaining case for (b,d): $|\rho| = \Delta - c$. The same argument used twice already above yields that ρ has at most $|\rho|(\Delta - c + 1) \leq k_\Delta(k_\Delta + 1) < c - \sqrt{\Delta} + 1$ external neighbours, and so (b) holds. Furthermore, applying Lemma 22 as in the previous paragraph yields that there are at most $\frac{3}{2}(\Delta - c + 1 - (|\rho| - 1)) + 1 = 4$ colour classes of size at least 2. Hence $|C_i| \geq c - 4$. By the same argument as the previous paragraph, each vertex of C_i has at least $|C_i| - 1 + |\rho| \geq c - 5 + \Delta - c = \Delta - 5$ neighbours in X_i and hence has at most 5 neighbours outside of X_i. This proves (d).

The proof of (e) is similar. Again, we can assume that ρ is the unique largest colour class, and that there is no colour class of size $\rho - 1$. If $|\rho| \leq \Delta - c - 3\Delta^{1/4}$ then ρ has at most $|\rho|(\Delta - c + 1) \leq (k_\Delta - 3\Delta^{1/4})(k_\Delta + 1) < c - 2\Delta^{3/4}$ external neighbours. So we can assume $|\rho| \geq \Delta - c - 3\Delta^{1/4}$. Thus there are at most $\frac{3}{2}(\Delta - c + 1 - (|\rho| - 1)) + 1 < 5\Delta^{1/4} - 1$ colour classes of size at least two and so $|C_i| \geq c - 5\Delta^{1/4} + 1$ and each vertex of C_i has at least $|C_i| - 1 + |\rho| \geq \Delta - 8\Delta^{1/4}$ internal neighbours. This proves (e). □

Proof of Lemma 11: Consider a graph G as in Lemma 11, and the decomposition of G that is ensured by Lemma 14. The sparseness of the vertices in S easily implies condition (b) of Lemma 11.

If X_i is a reducer or near-reducer then it is easy to verify that the unique stable set in X_i has at most c external neighbours. Otherwise, Lemma 24(a,b) imply that every colour class of size at least two in X_i has fewer than c external neighbours. This, along with Lemmas 18, 19 and 20 prove that X_i is a quasi-reducer. □

We strengthen Lemma 11 by showing that for minimum counterexamples to Theorem 5, the X_i can be neither reducers nor near-reducers. Recall that we fix $\Delta \geq \Delta_0$ and $c \geq \Delta - k_\Delta$ and say that G is a minimum counterexample to Theorem 5 if it is a counterexample with the smallest possible number of vertices for those values of Δ, c.

Observation 25 *In a minimum counterexample to Theorem 5, no X_i is a reducer.*

Proof If G is a minimum counterexample to Theorem 5, then G is $(c+1)$-critical. Therefore, G cannot contain a reducible c-reducer and so by Observation 3 we have $c = \Delta - k_\Delta$. Thus, G contains no reducers by the hypothesis of Theorem 2. □

Lemma 26 *In a minimum counterexample to Theorem 5, no X_i is a near-reducer.*

Remark: In the future, when we apply Lemma 24, we implicitly exploit Observation 25 and Lemma 26 to remove from that lemma the conditions that X_i is not a reducer or near-reducer.

Proof of Lemma 26: Suppose X_i is a near-reducer, with corresponding clique K_i and stable set S_i. By the minimality of G, there is a c-colouring of $G - X_i$; we will extend this c-colouring to G. If some vertex $v \in K_i$ has no neighbour in $G - X_i$, then we can extend any such colouring, as we now show.

Since $|S_i|+|K_i| = \Delta$, every vertex of K_i has at most one neighbour outside of X_i. We first colour all of the vertices of S_i using a colour not appearing on the at most $|S_i|(\Delta - |K_i|) = (\Delta - c+1)(\Delta - c+1) \leq (k_\Delta+1)(k_\Delta+1) \leq \Delta - (k_\Delta+1) < c$ neighbours of S_i in $G - X_i$. We then colour all of the vertices of $K_i - v$, which is possible because at most $c-2$ colours are forbidden because they are used on X_i and at most 1 is forbidden by a neighbour outside of X_i. Finally we colour v with the unique colour not yet appearing in X_i.

So we can assume that every vertex of K_i has a neighbour in $G - X_i$. We want to colour $G - X_i$ so that these neighbours do not all get the same colour. To this end consider the set Z of vertices of $G - X_i$ with neighbours in K_i. If any two vertices of Z are adjacent, then any colouring of $G - X_i$ will suffice since those two vertices will have different colours; so we can assume that Z is a stable set. If there is a vertex z of Z with fewer than $c-1$ neighbours outside of X_i, then fix a colouring of $G - X_i$ and check if there is some colour appearing on all of the elements of Z. If there is then recolour z with some other colour which appears on none of its neighbours, to produce the desired colouring of $G - X_i$.

Otherwise, $|Z| \geq \lceil \frac{|K_i|}{\Delta-(c-1)} \rceil = \lceil \frac{c-1}{\Delta+1-c} \rceil \geq k_\Delta + 1$. For each pair of vertices x and y in Z, consider the graph obtained from $G - X_i$ by adding the edge xy. Since G is a minimum counterexample to Theorem 5, $G - X_i + xy$ has a c-colouring unless in this graph either (i) xy is contained in a reducer, or (ii) there is a vertex v such that $x, y \in \{v\} \cup N(v)$ and the graph induced by $\{v\} \cup N(v)$ is not c-colourable. In either case, x and y are both dense vertices in $G - X_i$ with at least $\Delta - 2\sqrt{\Delta}$ common neighbours. Such a c-colouring of $G - X_i + xy$ would be a c-colouring of $G - X_i$ in which Z is not monochromatic, as desired. It follows that if we cannot find such a colouring of $G - X_i$ then there is some j such that $Z \subseteq X_j$ and in every c-colouring of X_j, all the vertices of Z receive the same colour; thus, Z is contained in a colour class of X_j. But there are $c-1$ edges from this colour class to X_i, so applying Lemma 24(b), we obtain that X_j is a near-reducer with $Z \subseteq S_j$. Simply counting shows that all of the edges leaving S_j go to K_i.

We construct a directed graph on the X_i which are near-reducers by adding an edge from

i to j if all the edges from S_i to $G - X_i$ go to X_j. Clearly this graph has maximum outdegree 1. So, it either has a vertex of indegree zero or is a set of directed cycles. In the latter case, we consider some such cycle J. The X_i that are on J form a component of G so by the minimality of G they must span G. For each X_l on J, we colour all the vertices in S_l with colour 1, and colour K_l using the colours $2, ..., c$. This yields a c-colouring of G and thereby a contradiction.

So we can assume that there is a near-reducer X_j with indegree zero in this graph, which by the argument above implies that we can colour $G - X_j$ so that there are two vertices v and w in K_j whose neighbours outside of X_j are coloured with different colours. We extend our colouring by first giving all of S_j the same colour. We then colour w, giving preference to the colour assigned to the neighbour of v outside X_j. The only reason we fail to use this colour is if we have already used it on the vertices of S_j. In either case, we colour the remaining vertices of K_j, colouring v last. Since every colour incident to v appears in X_j, we will be able to colour v, thereby obtaining a c-colouring of G. □

5 Massaging the Decomposition

In this section we prove Lemma 12 which we restate for the reader's convenience:

Lemma 12 *We can find a graph F of maximum degree $10^9\Delta$, whose c-colourability implies the c-colourability of G, and a partition of the vertices of F into $S, B, A_1, ..., A_t$ such that:*

(a) Every A_i is a clique with $c - 10^8\sqrt{\Delta} \leq |A_i| \leq c$.

(b) Every vertex of A_i has at most $10^8\sqrt{\Delta}$ neighbours in $F - A_i$.

(c) There is a set $\mathrm{All}_i \subseteq B$ of $c - |A_i|$ vertices which are adjacent to all of A_i. Every other vertex of $F - A_i$ is adjacent to at most $\frac{3\Delta}{4} + 10^8\sqrt{\Delta}$ vertices of A_i.

(d) Every vertex of S has degree at most Δ and either has fewer than $\Delta - 3\sqrt{\Delta}$ neighbours in S or has at least $900\Delta\sqrt{\Delta}$ non-adjacent pairs of neighbours within S.

(e) Every vertex of B has fewer than $c - \sqrt{\Delta} + 9$ neighbours in $F - \cup_j A_j$.

(f) If a vertex $v \in B$ has at least $c - \Delta^{\frac{3}{4}}$ neighbours in $F - \cup_j A_j$, then there is some i such that: v has at most $c - \sqrt{\Delta} + 9$ neighbours in $F - A_i$ and every vertex of A_i has at most $30\Delta^{\frac{1}{4}}$ neighbours in $F - A_i$.

(g) *For every* A_i*, every two vertices outside of* $A_i \cup \mathrm{All}_i$ *which have at least* $2\Delta^{\frac{9}{10}}$ *neighbours in* A_i *are joined by an edge of* F*.*

The starting point for the proof of Lemma 12 is the decomposition into $X_1, ..., X_t, S$ given by Lemma 11. The first step in our construction of F will be to contract the non-singleton stable sets of each X_i into vertices as we did with reducers. If we were to let B be the set of these contracted vertices, let A_i be the clique formed by the singleton colour classes of X_i, and leave S unchanged then (a) and (b) would follow from the definition of a quasi-reducer, (d) would follow from Lemma 11(b), (e) would follow from Lemma 24(b), and (f) would follow from Lemma 24(b,e) (letting A_i be the clique from the quasi-reducer X_i containing v).

Conditions (c) and (g) were added to give us better control over the pseudorandom colouring process. In particular, these conditions will help us ensure that for most colours α, not too many vertices of A_i have a neighbour of colour α. (Thus our approach is slightly more nuanced than the sketch in Section 2.3 suggests. In fact, we can let up to $c - |A_i|$ colours appear on external neighbours of all vertices in A_i without precluding a colouring; these will be the colours on the vertices of All_i.)

To ensure that conditions (c) and (g) hold we will need to massage the graph obtained in the first step of the proof by contracting vertices and adding edges, while ensuring that the other properties continue to hold. This may increase the maximum degree, but it will remain $O(\Delta)$.

We describe a three step process for constructing F in the next three subsections. We then prove that F has the properties claimed in the lemma.

5.1 Constructing G'

We construct a graph G' as follows:

1. For each X_i, if our colouring of the X_i used $r < c$ colours then we add $c - r$ vertices to C_i adding edges so they are adjacent to all of X_i, and
2. We contract each stable set of X_i into a vertex and add edges so that the vertices of X_i form a clique D_i of size c.

Clearly, if we can c-colour G', then we can c-colour G: simply give each vertex of X_i the colour that it, or the vertex that it was contracted to, received in D_i. The way in which we selected the vertices to be contracted ensures that this will be a valid colouring.

Note that G' may have maximum degree greater than Δ. Note also that the analogue of Lemma 14(c) does not hold here; in particular, perhaps $v \in D_i$ has more than $\frac{3}{4}\Delta$ neighbours in some D_j. Nevertheless, we do have the following properties:

Lemma 27 *For each i:*

(a) $|D_i| = c$;

(b) every vertex in C_i has at most $2 \times 10^6\sqrt{\Delta}$ neighbours in $G' - D_i$;

(c) every vertex in S has degree at most Δ in G' and is adjacent to at most $\frac{3\Delta}{4}$ vertices in D_i;

(d) $|D_i - C_i| \leq 3 \times 10^6\sqrt{\Delta}$.

(e) every vertex in D_i has degree at most 2Δ in G'.

Proof Our construction implies that (a) holds. Lemma 18 then yields (b) and (d). We obtain (c) by combining Lemma 14(c) and the fact that every edge from a vertex $v \in S$ to D_i corresponds to an edge from v to X_i.

For part (e): Lemma 24(b) implies that every stable set of X_i has fewer than Δ external neighbours. (In fact, this is also implied by the definition of a quasi-reducer.) Part (a) implies that the vertex it is contracted to has $c - 1 < \Delta$ neighbours in D_i for at total of fewer than 2Δ neighbours in G'. Parts (a,b) imply that each vertex in C_i has degree at most $c + 2 \times 10^6\sqrt{\Delta} < 2\Delta$ in G'. □

Definition: As with X_i, an *internal neighbour* of $v \in D_i$ is a neighbour of v that is also in D_i and all neighbours of v in $G - D_i$ are *external neighbours*.

5.2 The First Modification

We now present our first modification, which enforces Lemma 12(c). Before doing so, we remind the reader why we want that condition to hold.

Since each D_i is a c-clique, we must ensure that, in our c-colouring, every colour appears on a vertex of D_i. One difficulty in doing so is that a vertex u in some D_j might be adjacent to almost all of some D_i. If we are not careful, we could end up colouring all the non-neighbours in D_i of such a vertex u with colours different from that appearing on u. This would exclude the use of the colour appearing on u in D_i thereby preventing a c-colouring of G'. The way we deal with this issue is to plan in advance which member of D_i will be assigned the colour of each external vertex that is adjacent to too much of D_i. We will take "too much" to mean "more than $\frac{3\Delta}{4}$ vertices in C_i".

We will define, for some of the vertices $v \in D_i$, a set R_v which contains vertices outside of D_i that have more than $\frac{3\Delta}{4}$ neighbours in C_i. We note that, by Lemma 27(b,c), these vertices are all in $D_j - C_j$ for some $j \neq i$. In our colouring, we will ensure that v and all members of R_v receive the same colour.

We will define R_i to be the vertices $v \in D_i$ for which R_v is defined.

We begin with those $b \in \cup_{j \neq i}(D_j - C_j)$ that have at least $|C_i| - \Delta^{3/4}$ neighbours in C_i. Note that, by Lemma 27(b), there are at most $2 \times 10^6 \sqrt{\Delta}|C_i|/(|D_i| - \Delta^{3/4}) < 10^7\sqrt{\Delta}$ such vertices. Note furthermore that, by Lemma 24(a), there is at most one such vertex in each D_j.

We process these vertices one-at-a-time. We say that $v \in D_i$ is *eligible* for $b \in D_j$ if (i) v has fewer than $\frac{1}{2}\Delta + 10\sqrt{\Delta}$ external neighbours and (ii) b is not adjacent to v nor to any members of R_v. By Lemma 24(a), condition (i) rules out at most one vertex. By Lemma 24(b), and the fact that no other vertex in D_j can be in any R_v, condition (ii) rules out at most $c - \sqrt{\Delta} + 3$ vertices. So, at least $(|D_i| - 1) - (c - \sqrt{\Delta} + 3) = \sqrt{\Delta} - 4$ vertices of D_i are eligible for b. Amongst all those eligible vertices, we select a v for which $|R_v|$ is (currently) smallest and we place b into R_v. Since we only do this for at most $10^7\sqrt{\Delta}$ vertices b, and since each has at least $\sqrt{\Delta} - 4$ eligible vertices to choose from, this guarantees that every R_v contains at most 2×10^7 vertices.

Next we process those b with between $\frac{3\Delta}{4}$ and $|C_i| - \Delta^{3/4}$ neighbours in C_i. By Lemma 27(b), the total number of vertices we process in the two phases is at most $2 \times 10^6 \Delta\sqrt{\Delta}/\frac{3\Delta}{4} < 10^7\sqrt{\Delta}$. So, when we come to process a vertex b in this phase, the number of vertices eligible for b will be at least the number of non-neighbours of b in C_i minus the number of $v \in C_i$ with $R_v \neq \emptyset$, which is at least $\Delta^{3/4} - 10^7\sqrt{\Delta} \geq \frac{1}{2}\Delta^{3/4}$. Thus, we can process these vertices one-at-a-time, adding at most one of them to each R_v. This guarantees that $|R_v| \leq 2 \times 10^7 + 1 < 10^8$.

We now note some consequences of the way in which we constructed each R_v:

Lemma 28 *For each D_i, and each $v \in R_i$:*

(a) $\cup_{w \in R_i} R_w$ is the set of vertices not in D_i which have more than $\frac{3}{4}\Delta$ neighbours in C_i;

(b) Every vertex of R_v is in $D_j - (C_j \cup R_j)$ for some j, and there is no $v \neq u$ with $R_v \cap R_u \neq \emptyset$;

(c) $v \cup R_v$ is an independent set;

(d) $|R_i| \leq 10^7\sqrt{\Delta}$;

(e) $|R_v| < 10^8$;

(f) the external degree of v plus the sum over each $b \in R_v$ of the number of external neighbours of b outside of D_i is at most $\frac{3}{4}\Delta + 10^8\Delta^{3/4}$;

Before proving this lemma, we show how we use it to modify G':

Modification 1: For every $v \in \cup_{i=1}^t R_i$, we contract R_v into v. We define $A_i = C_i \backslash R_i$. So the D_i's may now intersect, but by Lemma 28(b), the A_i's are still disjoint. We will set $\text{All}_i = D_i - A_i$. Note that after Modification 1, each D_i is still a clique and so every vertex in All_i is adjacent to all of C_i (as required in Lemma 12(c)).

The effect of this contraction is equivalent to enforcing a rule that all vertices in R_v must get the same colour as v. Recall that this was our goal: $\cup_{v\in R_i} R_v$ are all those vertices $u \notin D_i$ for which we had to select in advance a $v \in D_i$ that would have the same colour as u.

Note that Lemma 28(b) implies that the modification is well-defined in that no vertex is contracted into two different vertices and no member of $\cup_{i=1}^t R_i$ is contracted into another vertex. Note also that Lemma 28(c) implies that a c-colouring of the modified graph yields a c-colouring of G. And note that the contracted vertices may have degree higher than Δ, but Lemmas 28 and 27(e) ensure that none have degree higher than $2 \times 10^8\Delta$.

And finally, note that Lemma 28(f) implies Lemma 12(c).

Proof of Lemma 28 Parts (a,b,c) are guaranteed by the way we choose R_v, the fact that, by Lemma 24(b), no vertex of G' can have $\frac{3\Delta}{4}$ neighbours in two different D_i's, and Lemma 27(b,c) which implies that we need not place any vertices in $S \cup (\cup_j C_j)$ into R_v. For part (d): by Lemma 27(b), at most $2 \times 10^6\Delta\sqrt{\Delta}/(\frac{3}{4}\Delta) < 10^7\sqrt{\Delta}$ vertices lie in the union of all the R_v's. Part (e) is described above. For part (f): in the first phase of the construction, we add at most 2×10^7 vertices to R_v and by Lemma 24(b) each such vertex has at most $c - \sqrt{\Delta} + 3 - (|C_i| - \Delta^{3/4}) < 2\Delta^{3/4}$ external neighbours outside of D_i. In the second phase, we add at most one vertex and by Lemma 24(b) it has at most $c - \sqrt{\Delta} + 3 - \frac{3}{4}\Delta < \frac{1}{4}\Delta$ external neighbours outside of D_i. By condition (i) of the definition of "eligible", the external degree of v is at most $\frac{1}{2}\Delta + 10\sqrt{\Delta}$. So we have a total of at most $\frac{1}{2}\Delta + 10\sqrt{\Delta} + 2 \times 10^7 \times 2\Delta^{3/4} + \frac{1}{4}\Delta < \frac{3}{4}\Delta + 10^8\Delta^{3/4}$. □

5.3 Our Second Modification

As we said above, every colour must appear on D_i, since it is a c-clique. Thus every colour not on $D_i - A_i$ must appear on A_i. This will be impossible if there is some such colour that appears in the external neighbourhoods of every vertex of A_i.

We started to deal with this problem in the previous subsection with Modification 1. In particular, we ensured that there is no one vertex outside of D_i that is adjacent to too many vertices of A_i. However, we still might have the problem that a vertex appears on several vertices outside of D_i, which between them are adjacent to all of A_i. In this section, we present some further definitions, results and a modification which will help us to prevent that from occurring.

Definition: Big_i is the set of vertices not in D_i with at least $\Delta^{9/10}$ neighbours in A_i, after Modification 1 is applied. Each pair of members of some Big_i are said to be *big-neighbours*.

Modification 2: We add an edge between every pair of big-neighbours, thus turning each Big_i into a clique.

Thus, for each colour α, at most one vertex outside of D_i with more than $\Delta^{9/10}$ neighbours in D_i can be assigned α. So α cannot be forbidden from all of A_i because it appears on only a few vertices which between them are adjacent to all of A_i. The only remaining thing to worry about is that α appears on *many* vertices which between them are adjacent to all of A_i. When we construct our colouring, we will see that we can prevent this using our probabilistic tools.

The following lemma with respect to Big_i will be useful when we analyze the effects of Modification 2.

Lemma 29 *In the graph formed by applying Modification 1 to G':*

(a) $|\mathrm{Big}_i| \leq 10^8\Delta^{3/5}$.

(b) No vertex has more than $10^{16}\Delta^{7/10}$ *big-neighbours.*

(c) If $v \in A_i$ then v is not in any Big_j.

(d) If $v \in S$ is in Big_i *for at least one i, then the number of neighbours v has in S plus the number of big-neighbours v has in S is at most* $\Delta - \frac{1}{2}\Delta^{9/10}$.

Note that adding these edges may increase some vertex degrees, but Lemma 29(b) ensures that these increases are only $o(\Delta)$.

Proof (a) The number of edges from D_i to $G' - D_i$ before Modification 1 is at most the number of edges from X_i to $G - X_i$, which is at most $10^7\Delta^{3/2}$ by Lemma 14(b). In the proof of Lemma 28(d), we showed that at most $10^7\sqrt{\Delta}$ vertices were contracted into vertices of R_i during Modification 1. Each of these vertices has less than 2Δ neighbours outside of D_i (by Lemma 24(b)). So the total number of edges from D_i to $G' - D_i$ after Modification 1 is at most $3 \times 10^7\Delta^{3/2}$ and so $|\mathrm{Big}_i| \leq 3 \times 10^7\Delta^{3/2}/\Delta^{9/10} < 10^8\Delta^{3/5}$.

(b) By Lemma 28(e), no vertex is the contraction of more than 10^8 vertices of G'. Lemma 24(b) implies that, before Modification 1, each vertex had degree less than 2Δ. So after Modification 1, no vertex has degree greater than $2 \times 10^8\Delta$. Therefore no vertex can belong to Big_i for more than $2 \times 10^8\Delta/\Delta^{9/10} = 2 \times 10^8\Delta^{1/10}$ cliques D_i. So by part (a), each vertex has at most $2 \times 10^8\Delta^{1/10} \times 10^7\Delta^{3/5} < 10^{16}\Delta^{7/10}$ big-neighbours.

(c) By Lemma 28(b), v did not gain any new neighbours during Modification 1 so by Lemma 27(b), v has at most $2 \times 10^6\sqrt{\Delta}$ external neighbours. Thus v has fewer than $\Delta^{9/10}$ neighbours in each A_j.

(d) Let ℓ be the number of A_i with $v \in \mathrm{Big}_i$. Then v has at least $\ell \times \Delta^{9/10}$ neighbours outside of S. By part (a), v has at most $\ell \times 10^8\Delta^{3/5} < \frac{1}{2}\ell\Delta^{9/10}$ big-neighbours. So the number of neighbours of v in S plus the number of big-neighbours of v in S is at most $\Delta - \ell \times \Delta^{9/10} + \ell \times \frac{1}{2}\Delta^{9/10} \leq \Delta - \frac{1}{2}\Delta^{9/10}$. □

5.4 Proof of Lemma 12

F is the graph formed by applying Modifications 1 and 2 to G'. We close this section by proving the properties of F that are listed in Lemma 12.

Proof of Lemma 12: We define All_i to be $D_i - A_i$. If any member of $D_i - A_i$ was contracted into some v during Modification 1, then we include v in All_i. Thus we may have $\mathrm{All}_i \cap \mathrm{All}_j \neq \emptyset$. We define B to be $\cup_{i=1}^{t}\mathrm{All}_i$.

We begin by bounding the maximum degree in F. Lemma 27(c,e) implies that the maximum degree in G' is at most 2Δ. Lemma 28(e) ensures that the maximum degree after Modification 1 is at most $2 \times 10^8\Delta$. And Lemma 29(b) ensures that no vertex gains more than $10^{16}\Delta^{7/10}$ new neighbours during Modification 2. So the maximum degree in F is at most $2 \times 10^8\Delta + 10^{16}\Delta^{7/10} < 10^9\Delta$.

Part (a): $A_i \subseteq C_i$ so it is a clique of size at most $|D_i| = c$. Lemmas 27(d) and 28(d) yield that its size is at least $c - 3 \times 10^6\sqrt{\Delta} - 10^7\sqrt{\Delta} > c - \frac{1}{2} \times 10^8\sqrt{\Delta}$.

Part (b): Consider any $v \in A_i$. By the calculations of part (a), v has fewer than $\frac{1}{2}10^8\sqrt{\Delta}$ neighbours in $D_i - A_i$. Lemma 27(b) implies that, in G', v has at most $2 \times 10^6\sqrt{\Delta}$ neighbours outside of D_i. Since $v \in A_i = C_i \backslash R_i$, Lemma 28(b) implies that v did not gain any new neighbours from Modification 1. Lemma 29(c) implies that v did not gain any new neighbours from Modification 2. So v has at most $\frac{1}{2}10^8\sqrt{\Delta} + 2 \times 10^6\sqrt{\Delta} < 10^8\sqrt{\Delta}$ neighbours outside of A_i.

Part (c): $|\text{All}_i| = |D_i| - |A_i| = c - |A_i|$. Lemma 28(a) implies that any vertex outside of $A_i \cup \text{All}_i = D_i$ that had more than $\frac{3}{4}\Delta$ neighbours in A_i was contracted into a vertex of All_i during Modification 1, and Lemma 28(f) implies that Modification 1 did not cause any other vertices outside of D_i to have more than $\frac{3}{4}\Delta + 10^8\sqrt{\Delta}$ neighbours in A_i. Lemma 29(c) implies that no vertex gained any new neighbours in A_i from Modification 2.

Part (d): Consider any $v \in S$. If v is not in any Big_i then its degree in F is at most its degree in G and hence at most Δ. Otherwise, its degree in F is less than Δ by Lemma 29(d).

Now suppose that, in F, v has more than $\Delta - 3\sqrt{\Delta}$ neighbours in S. Note that any edges in S that are in the graph F but not in the graph G, were added by Modification 2. By Lemma 29(d), v is not in any Big_i.

For each $u \in N(v) \cap S$, if $u \in \text{Big}_i$ for some D_i then, in the graph G, u has at most $\Delta - \Delta^{9/10}$ neighbours in S and hence is non-adjacent to at least $\Delta^{9/10} - 3\sqrt{\Delta}$ members of $N(v) \cap S$. By Lemma 29(b), u is non-adjacent to at least $\Delta^{9/10} - 3\sqrt{\Delta} - 10^{16}\Delta^{7/10} \geq \frac{1}{2}\Delta^{9/10}$ members of $N(v) \cap S$. It follows that if at least $3600\Delta^{6/10}$ neighbours of v gained new edges during Modification 2 then v has, in the graph F, at least $900\Delta\sqrt{\Delta}$ pairs of non-adjacent neighbours in S.

So suppose that at most $3600\Delta^{6/10}$ neighbours of v in S gained new edges during Modfication 2. Then the total number of edges added to $N(v) \cap S$ is at most $\binom{3600\Delta^{6/10}}{2} < \Delta\sqrt{\Delta}$. So by Lemma 14(d), in the graph F, $N(v) \cap S$ has at most $\binom{\Delta}{2} - (10^5 + 1)\Delta\sqrt{\Delta}$ edges. Straightforward calculations imply that, in the graph F, v has at least $900\Delta\sqrt{\Delta}$ pairs of non-adjacent neighbours in S.

Part (e): This is a simple corollary of part (f).

Part (f): We will first establish that v did not gain any new neighbours during Modifications 1 and 2. Suppose that vertices were contracted into v during Modification 1; i.e.

$v \in R_j$ for some D_j. The number of neighbours that v has in $F - \cup A_i$ is at most the sum of: (i) the external degree of v plus the sum over each $b \in R_v$ of the number of external neighbours of b outside of D_j; (ii) $|D_j - A_j|$ plus the sum of $|D_\ell - A_\ell|$ over all D_ℓ that contain a member of R_v; and (iii) the total number of big-neighbours of v and all members of R_v. By Lemma 28(f), (i) is at most $\frac{3}{4}\Delta + 10^8\sqrt{\Delta}$. By Lemma 28(e) and by part (a), (ii) is at most $10^8 \times (10^8\sqrt{\Delta})$. By Lemma 29(b), (iii) is at most $10^{16}\Delta^{7/10}$. This yields a total of less that $c - \Delta^{3/4}$ thus contradicting the hypothesis of part (f).

Suppose that $v \in \text{Big}_j$ for some D_j; i.e. v has at least $\Delta^{9/10}$ neighbours in D_j. By part (a), v has at most $|D_j - A_j| < 10^8\sqrt{\Delta}$ neighbours in $D_j - A_j$, and so v has at least $\frac{1}{2}\Delta^{9/10}$ in A_j. By Lemma 29(b), v gained at most $10^{16}\Delta^{7/10}$ new neighbours during Modification 2. By the previous paragraph, v was not contracted in Modification 1. So v has at most $\Delta - \frac{1}{2}\Delta^{9/10} + 10^{16}\Delta^{7/10} < \Delta - \Delta^{3/4}$ neighbours in $F - \cup_i D_i$ thus contradicting the hypothesis of part (f).

Since $v \in B$ we have that $v \in D_i - A_i$ for some i. By Lemma 24(e), $|D_i - C_i| < 5\Delta^{1/4}$ and every vertex of C_i has at most $8\Delta^{1/4}$ neighbours in $G' - D_i$. Thus, the total number of edges from C_i to $G' - D_i$ is at most $8\Delta^{1/4}|C_i| < 8\Delta^{5/4}$. So by the construction of R_i, at most $8\Delta^{5/4}/(\frac{3}{4}\Delta) < 11\Delta^{1/4}$ vertices are in $\cup_{u \in R_i} R_u$ and so $|R_i| < 11\Delta^{1/4}$. Thus, no vertex of A_i has more than $|R_i| + |D_i - C_i| < 16\Delta^{1/4}$ neighbours in $D_i - A_i$. By Lemmas 28(b) and 29(c), no vertex of A_i gained any new neighbours during Modifications 1 and 2. So every vertex of A_i has at most $8\Delta^{1/4} + 16\Delta^{1/4} < 30\Delta^{1/4}$ neighbours in $F - A_i$.

Now we bound the number of neighbours that v has in $F - A_i$. Let ρ be the colour class of X_i corresponding to v. The number of neighbours that v has in $F - A_i$ is the number of external neighbours of ρ plus $|D_i - A_i| - 1$. The former term is at most $c - \sqrt{\Delta} + 3$, by Lemma 24(b) (and Lemma 26). By part (a), $|D_i - A_i| \leq 10^8\sqrt{\Delta}$. So if ρ has at most $c - 10^8\sqrt{\Delta}$ external neighbours, then v has fewer than $c - \sqrt{\Delta}$ neighbours in $F - A_i$. Else, by Lemma 24(c), $|C_i| \geq c - 2 \times 10^8$ and each vertex of C_i has at most 3×10^8 external neighbours. Every member of $\cup_{u \in R_i} R_u$ is adjacent to at least $\frac{3}{4}$ vertices of C_i and so $|R_i| \leq 3 \times 10^8 |C_i|/(\frac{3}{4}\Delta) \leq 4 \times 10^8$. Therefore, $|D_i - A_i| \leq 2 \times 10^8 + 4 \times 10^8 < 10^9$. Thus, if ρ has at most $c - \sqrt{\Delta} - 10^9$ external neighbours then v has at most $c - \sqrt{\Delta}$ neighbours in $F - A_i$. If ρ has more than $c - \sqrt{\Delta} - 10^9$ external neighbours then applying Lemma 24(d) in the same way yields $|C_i| = c - 1$ and $|R_i| \leq \lfloor 5|C_i|/(\frac{3}{4}\Delta) \rfloor \leq 6$. So $|D_i - A_i| \leq 7$ and v has at most $\Delta - \sqrt{\Delta} + 9$ neighbours in $F - A_i$.

Part (g): By Lemma 29 (b), if a vertex outside of $A_i \cup \text{All}_i = D_i$ has at least $2\Delta^{9/10} > \Delta^{9/10} + 10^{16}\Delta^{7/10}$ neighbours in A_i after Modification 2 then it must be in Big_i. So any two such vertices were joined during Modification 2.

6 The Pseudorandom Process: A High Level Overview

The rest of the paper is devoted to proving Lemma 13; i.e. to showing that F can be c-coloured.

We begin with a sketch of the proof. We will colour the graph in stages, at each stage being careful so as to ensure that we can extend our colouring to the whole graph in later stages. We have to be most careful about extending the colouring to the cliques A_i. When we come to colour A_i, we will use the at least $c - |\mathrm{All}_i| \geq |A_i|$ colours that do not appear on All_i (see Lemma 12(c)). It turns out that it is sufficient to ensure that for each uncoloured A_i, every colour which does not appear on a vertex of All_i appears on neighbours of at most $\frac{4}{5}\Delta$ vertices of A_i.

In each stage, we will choose our colouring pseudorandomly. For ease of discussion, suppose that we do so in an (oversimplified) manner in which each vertex gets a uniformly random colour. By Lemma 12(b), every vertex of A_i has $O(\sqrt{\Delta})$ neighbours outside of $A_i \cup \mathrm{All}_i$ and so for any fixed colour x, the expected number of vertices of A_i with a neighbour outside of $A_i \cup \mathrm{All}_i$ that gets colour x is $|A_i| \times O(\frac{\sqrt{\Delta}}{c}) = O(\sqrt{\Delta})$. Of course we cannot prove that this random variable is typically near its expectation, as there may be a vertex $v \notin A_i \cup \mathrm{All}_i$ with a large number, say $\frac{\Delta}{2}$, neighbours in A_i. In that case, those $\frac{\Delta}{2}$ vertices of A_i will each have a neighbour (namely v) in $F - A_i \cup \mathrm{All}_i$ assigned one colour (namely the colour of v). However, Lemma 12(g) (i.e. Modification 2) ensures that at most one vertex $b \notin A_i \cup \mathrm{All}_i$ with more than $2\Delta^{9/10}$ neighbours in A_i will have colour x, and Lemma 12(c) (i.e. Modification 1) ensures that b has at most $\frac{3}{4}\Delta$ neighbours in A_i. To make use of these properties, we define:

Definition:

- Big_i^+ is the set of vertices not in $A_i \cup \mathrm{All}_i$ with at least $2\Delta^{9/10}$ neighbours in A_i.
- $\mathrm{Notbig}(i, x)$ is the number of vertices of A_i which have a neighbour $v \notin A_i \cup \mathrm{All}_i \cup \mathrm{Big}_i^+$ with colour x.

We will always be able to ensure that $\mathrm{Notbig}(i, x)$ is $o(\Delta)$. So at any point in our colouring process, at most $\frac{3}{4}\Delta + o(\Delta) < \frac{4}{5}\Delta$ vertices of A_i will have a neighbour outside of $A_i \cup \mathrm{All}_i$ with colour x.

We begin by colouring S so that each $\mathrm{Notbig}(i, x)$ is $o(\Delta)$. We do so by applying an iterative colouring algorithm to S. The first iteration is different from the others. We extend the proof of Lemma 10 to show that we can partially colour S so that every vertex has at most $\frac{19\Delta}{20}$ colours appearing in its neighbourhood and if it has degree exceeding $\Delta - 3\sqrt{\Delta}$ then it has at least 3Δ repeated colours. In addition, we can do so while keeping every Notbig(i, x) small. We then finish the colouring, exploiting the fact that for every vertex v, the number of colours available at v exceeds the number of uncoloured vertices in $N(v)$ by at least $\sqrt{\Delta}$.

We would now like to colour the vertices of B, using the same technique in order to keep Notbig(i, x) bounded for each i and x. By Lemma 12(e), we know that for each vertex v of B, given the colouring of the rest of $B \cup S$, there will be at least $\sqrt{\Delta} - 9$ colours which do not appear on its neighbours in v, so if we picked an extension of our colouring to B at random, then the probability v gets any particular colour is at most $\frac{1}{\sqrt{\Delta}-9}$. This bounds the expectation of the increase in Notbig(i, x) during this phase by $|A_i| \times O(\frac{\sqrt{\Delta}}{\sqrt{\Delta}-9}) = O(\Delta)$; unfortunately, this bound is too large for our purposes. To deal with this problem, we will actually interleave the colourings of the A_i with our colourings of the vertices in B.

Definition:

- B_H is the set of those vertices in B which have degree at most $\Delta - \Delta^{3/4}$ in $F - \cup A_i$;
- $B_L = B - B_H$;
- A_L is the set of those A_i such that every vertex of A_i has at most $30\Delta^{1/4}$ neighbours outside of $A_i \cup \text{All}_i$
- $A_H = A - A_L$.

We will first colour the vertices of B_H, then those A_i in A_H, then the vertices in B_L then those A_i in A_L. When colouring B_H, each vertex will have at most $\Delta - \Delta^{3/4}$ coloured neighbours, and so the probability that a particular colour x appears on a particular member of B_H is at most $\Delta^{-3/4}$. Thus the expected increase in Notbig(i, x) while colouring B_H is at most $|A_i| \times O(\frac{\sqrt{\Delta}}{\Delta^{3/4}}) = o(\Delta)$.

When colouring B_L we have already coloured the members of A_H, so we only need bound the increase Notbig(i, x) for each $A_i \in A_L$. Lemma 12(f) implies that each $b \in B_L$ has at most $c - \sqrt{\Delta}$ neighbours outside of A_L, and it follows as above that the probability that a particular colour x appears on a particular member of B_L is at most $\Delta^{-1/2}$. Thus for any $A_i \in A_L$, the expected increase in Notbig(i, x) while colouring B_L is at most $|A_i| \times O(\frac{\Delta^{1/4}}{\sqrt{\Delta}}) = o(\Delta)$.

Furthermore, we will be able to prove that these increases in Notbig(i, x) are concentrated enough to enable us to colour all of B whilst keeping Notbig(i, x) sufficiently small for all uncoloured A_i.

Our algorithm has three main subroutines.

The first, which we apply only once on S to ensure that there are repeated colours in the neighbourhoods of high degree vertices, is a simple modification of that used in the proof of Lemma 10.

The second is used in colouring S, B_H, and B_L. It takes an uncoloured subgraph H with an upper bound on the number of neighbours any uncoloured A_i has in H, and lower bounds on the number colours available to the vertices of H. It colours H so that for every colour x and every uncoloured A_i, the increase in Notbig(i, x) is $o(\Delta)$. (There is no need to bound this increase for any coloured A_i, since the only purpose of the bound is to help us colour A_i.) This is a simple iterative procedure.

The third procedure is used to colour some of the A_i's. It requires that for each A_i that we wish to colour, and for every colour not already used in All_i, there are not too many vertices of A_i with a neighbour of this colour; say at most $\frac{4\Delta}{5}$. It assigns a random permutation of the colours not used on All_i to A_i. It then recolours any vertex which receives a colour also appearing on an external neighbour. Each $v \in A_i$ has at most $O(\sqrt{\Delta})$ external neighbours (by Lemma 12(b)) and the probability of v receiving one of the $O(\sqrt{\Delta})$ colours from those neighbours is at most $O(\sqrt{\Delta})/|A_i| = O(\Delta^{-1/2})$, so the expected number of vertices in A_i that need to be recoloured is at most $|A_i| \times O(\Delta^{-1/2}) = O(\sqrt{\Delta})$. This is a manageably small number, and because we insisted that at the start of this procedure, no more than $\frac{4}{5}\Delta$ vertices in A_i had neighbours with any given colour, we will be able to recolour by having each of the vertices in conflict switch colours with a suitable vertex in A_i. More strongly, this hypothesis ensures that there will be at least $\Delta/20$ suitable vertices from which to pick, and so we can show that the probability that a vertex v ends up with a particular colour α is at most $O(\Delta^{-1})$. This ensures that when colouring A_H, for every $A_i \in A_L$ and colour x, the expected increase in Notbig(i, x) is at most $O(\sqrt{\Delta})$. By applying the Local Lemma we can show that with positive probability, at the end of the procedure every such Notbig(i, x) still has size at most $O(\Delta^{19/20})$.

We describe the iterative colouring procedure used on S, B_L, and B_H in Section 8. We describe our procedure for colouring the A_i in Section 9. We then go on to describe how we apply these procedures in each phase of our process to prove Lemma 13. First however, we present a lemma which we will need in analyzing both of these procedures.

7 Bounding Notbig(i, x)

Here we present a lemma which will be used several times to prove that Notbig(i, x) does not grow too large.

Suppose that we have a collection of at most Δ subsets of $V(F)$. Each set contains at most Q vertices. No vertex lies in more than $2\Delta^{9/10}$ sets. We conduct a random experiment where each vertex is marked with probability at most $1/(Q \times \Delta^{1/5})$. The vertices are not necessarily marked independently, but the experiment has the following property:

(P7.1) For any set of $\ell \geq 1$ vertices, the probability that all are marked is at most $1/(Q \times \Delta^{1/5})^{\ell}$.

Lemma 30 *The probability that at least* $\Delta^{37/40}$ *sets contain at least one marked vertex is at most* $\exp(-\Delta^{1/40})$.

When we use this to bound Notbig(i, x), the sets will be the external neighbourhoods of the vertices in A_i, after removing the vertices of Big_i^+ from those neighbourhoods. Thus, there are $|A_i| \leq \Delta$ such sets and no vertex lies in more than $2\Delta^{9/10}$ sets as otherwise it would be in Big_i^+. Typically, a vertex is marked if it receives the colour x. So the increase in Notbig(i, x) will be the number of sets that contain at least one marked vertex. Our goal is roughly to keep that number less than $\Delta^{19/20}$. We have no hope of showing that, with exponentially low probability, this number will be less than $2\Delta^{9/10}$, as it might become that high with the marking of a single vertex. So we choose $\Delta^{37/40}$ which is asymptotically between $2\Delta^{9/10}$ and $\Delta^{19/20}$.

Proof For each $1 \leq i \leq 9$, let T_i be the set of vertices lying in between $\Delta^{(i-1)/10}$ and $\Delta^{i/10}$ sets. Let E_i denote the event that at least $\frac{1}{9}\Delta^{37/40}$ sets contain a marked member of T_i. Note that if at least $\Delta^{37/40}$ sets contain at least one marked vertex, then at least one E_i must hold.

Since the sizes of the sets total at most ΔQ, $|T_i| \leq \Delta Q/\Delta^{(i-1)/10}$. If E_i holds, then at least $\frac{1}{9}\Delta^{37/40}/\Delta^{i/10}$ members of T_i must be marked. Therefore, applying (P7.1), we have:

$$
\begin{aligned}
\mathbf{Pr}(E_i) &\leq \binom{\Delta Q/\Delta^{(i-1)/10}}{\frac{1}{9}\Delta^{37/40}/\Delta^{i/10}} \left(\frac{1}{Q\Delta^{1/5}}\right)^{\frac{1}{9}\Delta^{37/40}/\Delta^{i/10}} \\
&\leq \left(\frac{e\Delta Q/\Delta^{(i-1)/10}}{\frac{1}{9}(\Delta^{37/40}/\Delta^{i/10}) \times Q\Delta^{1/5}}\right)^{\frac{1}{9}\Delta^{37/40}/\Delta^{i/10}}
\end{aligned}
$$

$$= \left(\frac{9e}{\Delta^{1/40}}\right)^{\frac{1}{9}\Delta^{37/40}/\Delta^{i/10}}.$$

Since $\frac{1}{9}\Delta^{37/40}/\Delta^{i/10} > \frac{1}{9}\Delta^{1/40}$ and $9e/\Delta^{1/40} < \frac{1}{2e}$, this yields that $\mathbf{Pr}(E_i) < \frac{1}{9}\exp(-\Delta^{1/40})$. So the probability that at least one E_i holds is at most $\exp(-\Delta^{1/40})$. □

8 An Iterative Colouring Procedure

In this section, we describe the technique that we use to colour S, B_L, and B_H. In order to have it apply to all three situations, we use preconditions that are somewhat general, although not nearly as general as possible.

In this setting, F may be partially coloured, but every A_i is either completely coloured or completely uncoloured. We are given an uncoloured subgraph H of $F - \cup_i A_i$ (and so H has maximum degree Δ by Lemma 12(d,e)). We wish to extend our partial colouring to H. For each vertex $u \in H$ we are given an initial list $L(u)$ of the colours available to u. We have values $X \geq \frac{1}{2}\sqrt{\Delta}$ and $U \geq \Delta^{1/4}$ such that:

(P8.1) every uncoloured vertex in each A_i has at most U neighbours in $H - \text{All}_i$;

(P8.2) for all $u \in H$, $|L(u)|$ is at least $5U \times \Delta^{1/5}$ and is at least $\deg_H(u) + X$.

The main thrust of (P8.2) is that each vertex in H initially has at least $5U \times \Delta^{1/5}$ available colours and throughout the procedure will always have at least X available colours.

Lemma 31 *We can extend our partial colouring to H such that for every uncoloured $A_i \in A$ and every colour x,* $\text{Notbig}(i, x)$ *increases by at most $\Delta^{19/20}$.*

We will colour H using a two step pseudorandom procedure.

Step 1: We fix a small constant $\frac{1}{100000} > \epsilon > 0$ and carry out $I = \lceil 2\Delta^{\epsilon} \log \Delta \rceil$ iterations. In each iteration, we will analyze the following random colouring procedure:

1. We activate each uncoloured $u \in H$ with probability $\alpha = \Delta^{-\epsilon}$.

2. We assign each activated u a uniformly random colour from $L(u)$.

3. If two activated neighbours receive the same colour, we uncolour them both.

4. Each activated u that is still coloured is uncoloured with probability $q(u)$ where $q(u)$ is defined so that u has probability exactly $\frac{1}{2}\alpha$ of being activated and retaining a colour.

5. For each vertex u which retains a colour x, we remove x from $L(v)$ for each $v \in N_H(u)$.

Note that it is possible to define $q(u)$ as desired since, defining $N_1(u)$ to be the set of uncoloured neighbours of u in H, the probability of an activated u being uncoloured in the third step is at most:

$$
\begin{aligned}
& \sum_{x\in L(u)} \mathbf{Pr}(u \text{ is assigned } x) \times \sum_{u'\in N_1(u)} \alpha \mathbf{Pr}(u' \text{ is assigned } x) \\
= \ & \frac{\alpha}{|L(u)|} \sum_{u'\in N_1(u)} \sum_{x\in L(u)} \mathbf{Pr}(u' \text{ is assigned } x) \\
\leq \ & \frac{\alpha}{|L(u)|} \sum_{u'\in N_1(u)} 1 \\
= \ & \frac{\alpha |N_1(u)|}{|L(u)|} \\
< \ & \alpha,
\end{aligned}
$$

since $|L(u)| > |N_1(u)|$ (by (P8.2)). Thus, the probability of being activated and not being uncolored in the third step is at least $\alpha(1-\alpha) > \frac{1}{2}\alpha$. So $q(u)$ is well-defined as a function of the lists of u and $N_1(u)$.

Our analysis allows us to prove the following:

Lemma 32 *After I iterations, with positive probability:*

(a) Each $u \in H$ has at most $\Delta^{200\epsilon}$ uncoloured neighbours in H.

(b) Each uncoloured vertex in $\cup A_i$ has at most $\Delta^{200\epsilon}$ uncoloured neighbours in H.

(c) For every A_i and colour x, $\text{Notbig}(i,x) \leq \frac{1}{2}\Delta^{19/20}$.

We choose a partial colouring satisfying (a)-(c) of this lemma.

Step 2: To finish the colouring we analyze a different procedure. Note that for every $v \in H$, $|L(v)| \geq X$ by (P8.2) and the fact that at most $\deg_H(v)$ colours have been removed from $L(v)$.

1. For each uncoloured $v \in H$, we choose a uniformly random subset $L'(v) \subset L(v)$ of size $2\Delta^{200\epsilon}$.
2. We colour all such v from their sublists, one-at-a-time.

Of course, the second step is possible because of Lemma 32(a).

Lemma 33 *With positive probability, for every uncoloured A_i and colour x, at most $\frac{1}{2}\Delta^{19/20}$ vertices of A_i have neighbours outside of Big_i^+ with colour x in their sublists.*

We choose a set of lists as guaranteed by Lemma 33.

Proof of Lemma 31: By Lemma 32(c), for each i, at most $\frac{1}{2}\Delta^{19/20}$ vertices of A_i have neighbours outside of Big_i^+ that were given colour x during Step 1. By Lemma 33, for each i, at most $\frac{1}{2}\Delta^{19/20}$ vertices of A_i have neighbours outside of Big_i^+ that were given colour x during Step 2, since in order for a vertex to be given colour x, x must appear on its sublist. $\square$

8.1 Proofs of Lemmas 32 and 33

To prove Lemma 32, we will recursively obtain a bound U_i on the maximum number of uncoloured external neighbours of a vertex of $\cup_j A_j$ after the i^{th} iteration. We also obtain upper and lower bounds $d_i^+(v)$ and $d_i^-(v)$ on the number of uncoloured neighbours of a vertex $v \in H$ after the i^{th} iteration. These bounds are defined as follows:

$$U_0 = U; \text{ for each } i > 0, U_i = \left(1 - \frac{1}{2}\Delta^{-\epsilon}\right) \times U_{i-1} + U_{i-1}^{49/50}$$

and for each vertex $v \in H$,

$$d_0^+(v) = \deg_H(v); \text{for each } i > 0, d_i^+ = \left(1 - \frac{1}{2}\Delta^{-\epsilon}\right) \times d_{i-1}^+ + (d_{i-1}^+)^{49/50}$$
$$d_0^-(v) = \deg_H(v); \text{for each } i > 0, d_i^- = \left(1 - \frac{1}{2}\Delta^{-\epsilon}\right) \times d_{i-1}^- - (d_{i-1}^-)^{49/50}$$

Some standard analysis shows that the small order terms in these definitions do not accumulate substantially, and so we obtain:

Lemma 34 *(a) if $U_i \geq \Delta^{150\epsilon}$, then*

$$U_i \leq 2\left(1 - \frac{1}{2}\Delta^{-\epsilon}\right)^i U;$$

(b) if $d_i^-(v) \geq \Delta^{150\epsilon}$, then

$$\frac{1}{2}\left(1 - \frac{1}{2}\Delta^{-\epsilon}\right)^i \deg_H(v) \leq d_i^-(v) \leq d_i^+(v) \leq 2\left(1 - \frac{1}{2}\Delta^{-\epsilon}\right)^i \deg_H(v).$$

We will defer the proof until the end of this subsection. For now, note the immediate corollary:

Corollary 35 *(a) if $d_i^-(v) \geq \Delta^{150\epsilon}$, then $d_i^-(v) \geq \frac{1}{4}d_i^+(v)$;*

(b) if $d_i^-(v), U_i \geq \Delta^{150\epsilon}$, then $d_i^-(v) \geq \frac{1}{4}U_i \times (\deg_H(v)/U)$.

Proof of Lemma 32 We will apply the Lovasz Local Lemma to each iteration of the procedure to prove inductively that with positive probability, after $i \leq I$ iterations:

(8.1) If $U_i \geq \frac{1}{2}\Delta^{200\epsilon}$ then every uncoloured vertex in $\cup_j A_j$ has at most U_i uncoloured external neighbours in H.

(8.2) For every vertex $v \in H$, if $d_i^-(v) \geq \frac{1}{8}\Delta^{200\epsilon}$ then v has between $d_i^-(v)$ and $d_i^+(v)$ uncoloured neighbours in H.

(8.3) For every uncoloured A_j and colour x, $|\mathrm{Notbig}(j,x)|$ increases by at most $\Delta^{19/20\ -\epsilon}/4\log\Delta$ during iteration i.

These will establish Lemma 32 as follows. $(1 - \frac{1}{2}\Delta^{-\epsilon})^I\Delta < 1 < \frac{1}{2}\Delta^{150\epsilon}$, so by Lemma 34, $U_I < \Delta^{150\epsilon} < \frac{1}{2}\Delta^{200\epsilon}$ and $d_i^-(v) < \Delta^{150\epsilon} < \frac{1}{8}\Delta^{200\epsilon}$ for all v (since $\deg_H(v) \leq \Delta$ by Lemma 12(d,e) as $H \subseteq B \cup S$). Furthermore, these parameters decrease by less than half, and so there are $i_1, i_2(v) < I$ such that $\frac{1}{2}\Delta^{200\epsilon} < U_{i_1} < \Delta^{200\epsilon}$ and, if $\deg_H(v) > \Delta^{200\epsilon}$, $\frac{1}{8}\Delta^{200\epsilon} < d_{i_2(v)}^-(v) < \frac{1}{4}\Delta^{200\epsilon}$. Thus (8.1) applied at iteration i_1 establishes part (a) and (8.2) applied at iteration $i_2(v)$ establishes part (b) as $d_{i_2(v)}^+(v) < 4d_{i_2(v)}^-(v) < \Delta^{200\epsilon}$ (by Corollary 35(a)). Part (c) follows from (8.3) since the number of iterations is $I = \lceil 2\Delta^\epsilon \log\Delta \rceil$.

Our statements hold trivially for $i = 0$. Consider some larger value of $i \le I$ and suppose that the statements hold for all smaller values of i. We will prove that with positive probability, the random choices made during iteration i will result in the statements holding for this value i.

For each uncoloured vertex $v \in A$, define $E_1(v)$ to be the event that v violates (8.1). For each vertex $v \in H$, define $E_2(v)$ to be the event that v violates (8.2). For each uncoloured $A_j \in A$ and colour x define $E_3(j,x)$ to be the event that A_j, x violate (8.3). We will prove that each of these events holds with probability at most Δ^{-10}.

For each vertex v, define $\mathcal{D}(v)$ to be the set consisting of: (i) $E_1(u)$ and $E_2(u)$ for every vertex u of distance at most 4 from v; and (ii) $E_3(j,x)$ for every colour x and every uncoloured set A_j containing a vertex of distance at most 4 from v. For each uncoloured set A_j, define $\mathcal{D}(j) = \cup_{v \in A_j} \mathcal{D}(v)$. It is straightforward to check that the random choices which determine whether $E_1(v)$ holds have no affect on whether any events outside of $\mathcal{D}(v)$ hold; it follows that $E_1(v)$ is mutually independent of all events outside of $\mathcal{D}(v)$. Similarly, $E_2(v)$ and $E_3(j,x)$ are mutually independent of all events outside of $\mathcal{D}(v)$ and $\mathcal{D}(j)$ respectively. Since F has maximum degree at most $10^9\Delta$, each $\mathcal{D}(v)$ has size less than $3 \times 10^9 c\Delta^4$ and so each $\mathcal{D}(j)$ has size less than $3 \times 10^9 c\Delta^5$ (as $|A_j| \le \Delta$). Thus with positive probability, none of these events hold since $3 \times 10^9 c\Delta^5 \times \Delta^{-10} < \frac{1}{4}$.

Bounding $\mathbf{Pr}(E_3(j,x))$*:*

We begin with $\mathbf{Pr}(E_3(j,x))$ as it is the shortest of the three arguments. Consider any uncoloured A_j and colour x. We will apply Lemma 30 with $Q = \max(U_{i-1}, \frac{1}{2}\Delta^{200\epsilon})$ to bound $\mathbf{Pr}(E_3(j,x))$. To do so, we will show that for every vertex $v \in H$, at the beginning of iteration i we have $|L(v)| > Q\Delta^{1/5}$.

At most $\deg_H(v)$ colours are removed from $L(v)$. So by (P8.2), $|L(v)| \ge X \ge \frac{1}{2}\sqrt{\Delta}$. This establishes $|L(v)| > Q\Delta^{1/5}$ for the case where $Q < \frac{1}{2}\Delta^{3/10}$ and so we can assume $Q = U_{i-1}$ (since $\epsilon < \frac{3}{2000}$) and $U_{i-1} \ge \frac{1}{2}\Delta^{3/10}$. By (P8.2), $L(v)$ initially has size at least $5U\Delta^{1/5}$, so if $\deg_H(v) < 4U\Delta^{1/5}$ then $L(v)$ will always have size at least $U\Delta^{1/5} \ge U_{i-1}\Delta^{1/5}$. Suppose that $\deg_H(v) \ge 4U\Delta^{1/5}$. Using the facts that $U_{i-1} \ge \frac{1}{2}\Delta^{3/10} >> \Delta^{150\epsilon}$ and that $d^-_{j-1}(v) > \frac{1}{2}d^-_j(v)$, a simple inductive application of Corollary 35(b) implies that $d^-_j(v) \ge \frac{1}{4}U_j \times (4\Delta^{1/5}) >> \Delta^{150\epsilon}$ for every $j \le i-1$; in particular, $d^-_{i-1}(v) \ge U_{i-1}\Delta^{1/5} = Q\Delta^{1/5}$. Since $|L(v)|$ is at least the number of uncoloured neighbours of v in H, (8.2) establishes that $|L(v)| > Q\Delta^{1/5}$.

By applying induction, and using (8.1), we know that at the beginning of iteration i, every vertex in A_j has at most Q uncoloured neighbours in H. Also, the probability that a vertex $v \in H$ receives colour x is at most $1/L(v) > 1/(Q\Delta^{1/5})$. Furthermore, these colour assignments are independent and so (P7.1) holds. Therefore, since $\Delta^{19/20}/\log^2\Delta > \Delta^{37/40}$, Lemma 30 implies that $\mathbf{Pr}(E_3(j,x)) < \exp(-\Delta^{1/40}) < \Delta^{-10}$, as required.

Bounding $\mathbf{Pr}(E_1(v))$*:*

Next, we turn to $\mathbf{Pr}(E_1(v))$. Consider any uncoloured vertex $v \in \cup_j A_j$. Since $H \subset F - \cup_j A_j$, every neighbour of v in H is external. So by induction, we can assume that at the beginning of the ith iteration, v has at most U_{i-1} neighbours in H that are not coloured. Let Y be the number of these neighbours which have colours at the end of iteration i. The probability of any particular uncoloured vertex becoming coloured in iteration i is exactly $\frac{1}{2}\Delta^{-\epsilon}$, and so $\mathbf{Exp}(Y) = \frac{1}{2}\Delta^{-\epsilon}U_{i-1}$. Thus, if $E_1(v)$ holds, then Y must differ from its mean by more than $U_{i-1}^{49/50}$.

Let Y_1 be the number of neighbours of v that get activated, and let Y_2 be the number that get activated and have their colours removed. Note that $Y = Y_1 - Y_2$ and so if $E_1(v)$ holds then either Y_1 or Y_2 differs from its mean by more than $\frac{1}{2}U_{i-1}^{49/50}$. Y_1 is a binomial variable. So the Chernoff Bound implies that the probability of Y_1 differing from its mean by that much is at most $2\exp(-U_{i-1}^{49/25}/(12U_{i-1}\Delta^{-\epsilon})) < \frac{1}{2}\Delta^{-10}$, since $U_{i-1} > \frac{1}{2}\Delta^{150\epsilon}$. So we turn our attention to Y_2.

Rather than showing directly that Y_2 is concentrated, it will be more convenient to deal with Y_2' which we define to be the number of uncoloured neighbours of v that are activated and (i) have their colours removed, or (ii) are assigned a colour that is assigned to at least $\log\Delta$ neighbours of v. Clearly $Y_2 \le Y_2'$. Furthermore, it is straightforward to show that, with high probability, $Y_2 = Y_2'$ as follows:

For each vertex u, we will use $\tilde{d}(u)$ to denote the number of neighbours of u in H that do not have colours at the beginning of iteration i.

If $Y_2 \neq Y_2'$ then some colour x must be assigned to at least $\log\Delta$ neighbours of v. Applying (8.1), (8.2) and Lemma 34 we have: $\tilde{d}(v) \le U_{i-1} \le 2(1-\frac{1}{2}\Delta^{-\epsilon})^{i-1}U$, and for every $u \in N_H(v)$, $\tilde{d}(u) \ge d_{i-1}^-(u) \ge \frac{1}{2}(1-\frac{1}{2}\Delta^{-\epsilon})^{i-1}\deg_H(u)$. By (P8.2) the number of colours available for $u \in N_H(v)$ is at least

$$\begin{aligned}
& \max\{\deg_H(u)+X, 5U\times\Delta^{1/5}\} - (\deg_H(u) - d_{i-1}^-(u)) \\
\ge\ & \max\{\deg_H(u)+X, 5U\times\Delta^{1/5}\} - \left(1-\frac{1}{2}\left(1-\frac{1}{2}\Delta^{-\epsilon}\right)^{i-1}\right)\deg_H(u) \\
\ge\ & \max\{\deg_H(u)+X, 5U\times\Delta^{1/5}\} - \left(1-\frac{1}{2}\left(1-\frac{1}{2}\Delta^{-\epsilon}\right)^{i-1}\right)\max\{\deg_H(u)+X, 5U\times\Delta^{1/5}\}
\end{aligned}$$

$$\geq \frac{1}{2}\left(1-\frac{1}{2}\Delta^{-\epsilon}\right)^{i-1} \times 5U\Delta^{1/5}$$
$$\geq \frac{5}{4}\Delta^{1/5}\tilde{d}(v).$$

Therefore,

$$\mathbf{Pr}(Y_2 \neq Y_2') \leq c \times \binom{\tilde{d}(v)}{\log\Delta}\left(\frac{5}{4}\Delta^{1/5}\tilde{d}(v)\right)^{-\log\Delta} \leq c \times \left(\frac{4e/5}{\Delta^{1/5}\log\Delta}\right)^{\log\Delta} < \tfrac{1}{4}\Delta^{-10}.$$

Note that, since $|Y_2 - Y_2'| \leq \Delta$, this implies that $|\mathbf{Exp}(Y_2) - \mathbf{Exp}(Y_2')| < \frac{1}{4}\Delta^{-9} = o(1)$.

We wish to show that $\mathbf{Pr}(|Y_2 - \mathbf{Exp}(Y_2)| > \frac{1}{2}U_{i-1}^{49/50}) < \frac{1}{2}\Delta^{-10}$. By the preceding bound on $\mathbf{Pr}(Y_2 \neq Y_2')$ and on $|\mathbf{Exp}(Y_2) - \mathbf{Exp}(Y_2')|$, it will suffice to show that $\mathbf{Pr}(|Y_2' - \mathbf{Exp}(Y_2')| > \frac{1}{4}U_{i-1}^{49/50}) < \frac{1}{4}\Delta^{-10}$.

We will apply Talagrand's Inequality. It will be helpful to consider each $v \in H$ to be involved in two random trials. The first one combines Steps 1 and 2: v is assigned the label "unactivated" or "activated with colour x" for some $x \in L(v)$; the first label is given with probability $1 - \Delta^{-\epsilon}$, and each of the other labels with probability $\Delta^{-\epsilon}/|L_v|$. In the second random trial, v is labelled "uncoloured" with probability $q(v)$, even if v was unactivated or lost its colour in Step 3. Of course, from the view of the algorithm, it is pointless to carry out the second step on vertices that don't have colours at the beginning of Step 4, but doing so has the technical advantage that all our random trials are independent in the sense that the outcome of one trial does not affect whether another is carried out.

If $Y_2' \geq s$ then there is a set of at most $(\log\Delta) \times s$ random trials that certify this fact, namely for each of s vertices counted by Y_2': the activation and colour assignment of the vertex and either the choice to uncolour that vertex in Step 4, or the activation and assignment of the same colour to either a neighbour of the vertex or to $\log\Delta - 1$ other vertices in $N(v)$. Also, changing the outcome of one of the random trials can only affect Y_2' by at most $\log\Delta$: in the extreme cases, changing the colour of u from Red to Blue either (i) causes u and fewer than $\log\Delta$ of its neighbours with colour Blue to be uncoloured, (ii) causes u and $\log\Delta - 1$ other vertices in $N(v)$ with colour Blue to be counted by Y_2', or (iii) reduces the number of vertices in $N(v)$ with colour Red to below $\log\Delta$ thus causing them all to not be counted by Y_2'. (Remark: this is the benefit of working with Y_2' rather than Y_2; a single colour assignment could potentially affect Y_2 by a much larger amount.) Also note that $Y_2 \leq U_{i-1}$ and so $\mathbf{Exp}(Y_2) \leq U_{i-1}$. Therefore, applying (2) with $c = r = \log\Delta$ yields:

$$\mathbf{Pr}(|Y_2' - \mathbf{Exp}(Y_2')| > \tfrac{1}{4}U_{i-1}^{49/50}) < 4\exp(-\left(\tfrac{1}{4}U_{i-1}^{49/50})^2/(32\log^3\Delta(U_{i-1} + \tfrac{1}{4}U_{i-1}^{49/50}))\right) < \tfrac{1}{4}\Delta^{-10},$$

since $U_{i-1} \geq \Delta^{200\epsilon}$. This implies $\mathbf{Pr}(E_1(v)) < \Delta^{-10}$.

Bounding $\mathbf{Pr}(E_2(v))$.

The proof that $\mathbf{Pr}(E_2(v)) < \Delta^{-10}$ is very similar to that for $E_1(v)$. The main difference is that for each $u \in N_H(v)$, $\deg_H(u)$ may be a lot bigger than $\deg_H(v)$, and this makes it more difficult to bound the analogue of $\mathbf{Pr}(Y_2' \neq Y_2)$.

For each vertex u, we will again use $\widetilde{d}(u) \leq d_{i-1}^{+}(u)$ to denote the number of neighbours of u in H that do not have colours at the beginning of iteration i. We use L_u to denote the set of colours that are available for u at the beginning of iteration i. By (P8.2), $|L_u| \geq \widetilde{d}(u) + X > \widetilde{d}(u)$.

Consider any $v \in H$ with $d_i^{-}(v) \geq \frac{1}{8}\Delta^{200\epsilon}$. We define Z_i analogously to Y_i. So we let Z_1 be the number of neighbours of v that get activated, and let Z_2 be the number that get activated and then have their colours removed. Again, it suffices to prove that with high probability, neither Z_1 nor Z_2 differs from its mean by more than $\frac{1}{2}(d_{i-1}(v)^{-})^{49/50}$. The Chernoff Bound implies that the probability of Z_1 differing from its mean by that much is at most

$$2\exp\left(-(d_{i-1}^{-}(v))^{49/25}/(12 d_{i-1}^{-}(v)\Delta^{-\epsilon})\right) < \frac{1}{2}\Delta^{-10},$$

since $d_{i-1}^{-}(v) \geq d_i^{-}(v) \geq \frac{1}{8}\Delta^{200\epsilon}$. So we turn our attention to Z_2.

We partition the vertices of $N_H(v)$ that are not coloured at the beginning of iteration i into $N_C \cup N_D$ where N_C contains those vertices u with $\widetilde{d}(u) \geq \widetilde{d}(v)^{3/4}$ and N_D contains those with $\widetilde{d}(u) < \widetilde{d}(v)^{3/4}$. We let Z_C, Z_D be the number of vertices in N_C, N_D respectively that get activated and uncoloured during this iteration. Thus $Z_2 = Z_C + Z_D$.

A similar argument to that used for Y_2 shows that Z_C is concentrated: Let Z_C' be the number of vertices in N_C that get activated and (i) are uncoloured or (ii) are assigned a colour that is assigned to at least $\widetilde{d}(v)^{3/10}$ members of N_C. Since $|N_C| \leq \widetilde{d}(v)$ and each vertex $u \in N_C$ has $L_u \geq \widetilde{d}(u) \geq \widetilde{d}(v)^{3/4}$, the probability that $Z_C \neq Z_C'$ is at most

$$c \times \binom{\widetilde{d}(v)}{\widetilde{d}(v)^{3/10}} (\widetilde{d}(v)^{3/4})^{-\widetilde{d}(v)^{3/10}} < \Delta \times \left(\frac{e\widetilde{d}(v)}{\widetilde{d}(v)^{3/10}\widetilde{d}(v)^{3/4}}\right)^{\widetilde{d}(v)^{3/10}} < \tfrac{1}{8}\Delta^{-8},$$

since $\widetilde{d}(v) \geq d_{i-1}^{-}(v) > \frac{1}{8}\Delta^{200\epsilon}$. As $|Z_C - Z_C'| \leq \Delta$, this implies $|\mathbf{Exp}(Z_C) - \mathbf{Exp}(Z_C')| = o(1)$.

By the same arguments as for Y_2': If $Z_C' \geq s$ then there are $\widetilde{d}(v)^{3/10}s$ trials whose outcomes certify this fact. Each trial can affect Z_C' by at most $\widetilde{d}(v)^{3/10}$. So applying (2) with

$c = r = \widetilde{d}(v)^{3/10}$ yields

$$\begin{aligned} & \mathbf{Pr}(|Z'_C - \mathbf{Exp}(Z'_C)| > \tfrac{1}{4}(d^-_{i-1}(v))^{49/50}) \\ \leq\ & 4\exp\left(-\tfrac{1}{16}(d^-_{i-1})^{49/25}/(128(d^-_{i-1}(v))^{6/10}(d^-_{i-1}(v))^{3/10}(d^-_{i-1}(v) + \tfrac{1}{4}(d^-_{i-1}(v))^{49/50}))\right) \\ <\ & \tfrac{1}{8}\Delta^{-10} \end{aligned}$$

and so

$$\mathbf{Pr}(|Z_C - \mathbf{Exp}(Z_C)| > \tfrac{1}{4}(d^-_{i-1}(v))^{49/50}) \leq \tfrac{1}{8}\Delta^{-10}.$$

Z_D requires a bit more care. We first expose the assignments to all vertices other than N_D - we denote this assignment by $\mathcal{H}$. We will then focus on the conditional distribution of Z_D. Our first step is to show that, with high probability, the conditional mean does not differ from the mean by very much:

Claim: $\mathbf{Pr}(|\mathbf{Exp}(Z_D|\mathcal{H}) - \mathbf{Exp}(Z_D)| > \frac{1}{2}(d^-_{i-1}(v))^{49/50}) < \frac{1}{8}\Delta^{-10}$.

Proof: For any assignment $\mathcal{H}$ to the vertices of $H - N_D$, we use $\mu_{\mathcal{H}}$ to denote the conditional expectation $\mathbf{Exp}(Z_D|\mathcal{H})$. Note that, the expected value of $\mu_{\mathcal{H}}$ over the space of random colourings of $H - N_D$ is equal to the expected value of Z_D over the space of random colourings of H. So our claim simply says that $\mu_{\mathcal{H}}$ is concentrated.

For each $u \in N_D$, we define $F_u = F_u(\mathcal{H}) \subseteq L_u$ to be the set of colours in L_u that are assigned by $\mathcal{H}$ to vertices in $N_H(u) - N_D$. We start by using Talagrand's Inequality to show that $|F_u|$ is highly concentrated.

F_u is determined by the independent colour assignments to the vertices of $H - N_D$. If $|F_u| \geq s$ then there is a set of s assignments that certifies this fact; namely the assignments of different colours to s vertices. The assignment to one vertex can affect $|F_u|$ by at most one, since it contributes at most one new colour to F_u. Since $|F_u| \leq |L_u|$ we have $\mathbf{Exp}(F_u) < |L_u|$. Therefore, applying (2) with $c = r = 1$ yields:

$$\mathbf{Pr}(|F_u| - \mathbf{Exp}(|F_u|) > \Delta^{-1/5}|L_u|) < 4\exp\left(-\Delta^{-2/5}|L_u|^2/(32(|L_u| + \Delta^{-1/5}|L_u|))\right) < \tfrac{1}{8}\Delta^{-10},$$

since by (P8.2) we have $|L_u| \geq X \geq \frac{1}{2}\sqrt{\Delta}$.

Therefore, the probability that there is at least one vertex $u \in N_D$ for which $|F_u|$ differs from its mean by more than $\Delta^{-1/5}|L_u|$ is at most $|N_D| \times \frac{1}{8}\Delta^{-11} \leq \frac{1}{8}\Delta^{-10}$. So we assume that there is no such u, and show that this implies $|\mu_{\mathcal{H}} - \mathbf{Exp}(\mu_{\mathcal{H}})| < \frac{1}{2}(d^-_{i-1}(v))^{49/50}$.

Given a particular assignment $\mathcal{H}$ to $H - N_D$, and colour $\gamma \in L_u$, the probability that u keeps its colour if it is assigned γ is 0 if $\gamma \in F_u$ and otherwise is $(1-q(u))\prod_w(1 - \frac{1}{|L_w|})$ where the product is over all vertices $w \in N(u) \cap N_D$ with $\gamma \in L(w)$. Noting that the latter product is at most 1 and is not a function of F_u, it follows that changing whether $\gamma \in F_u$ affects the probability that u retains its colour by at most $1/|L_u|$. Therefore, since $|F_u|$ differs from its mean by at most $\Delta^{-1/5}|L_u|$ the conditional probability that u is uncoloured differs from its mean by at most $\Delta^{-1/5}$. $\mu_{\mathcal{H}}$ is the sum over all $u \in N_D$ of these probabilities and so

$$|\mu_{\mathcal{H}} - \mathbf{Exp}(\mu_{\mathcal{H}})| \leq \Delta^{-1/5}|N_D| < \frac{1}{2}(d_{i-1}^{-}(v))^{49/50},$$

since $|N_D| \leq d_{i-1}^{+}(v) < 4d_{i-1}^{-}(v)$ (by Corollary 35(b)), and $d_{i-1}^{-}(v) \leq \Delta$. This proves our claim. QED

We define Z'_D to be the number of vertices in N_D that get activated and have their colours removed because (i) they are assigned the same colour as a neighbour outside of N_D or (ii) they are assigned the same colour as a neighbour $w \in N_D$ and that colour is assigned to fewer than $\widetilde{d}(v)^{3/10}$ vertices in $N(w) \cap N_D$. Note that this definition is a bit different than the definition of Z'_C. If $Z_D \neq Z'_D$, then some $u \in N_D$ gets the same colour as at least $\widetilde{d}(v)^{3/10}$ of its neighbours. Since each such u has at most $\widetilde{d}(v)^{3/4}$ neighbours, and every vertex w has $L_w \geq X \geq \frac{1}{2}\sqrt{\Delta}$ (by (P8.2), for every choice of $\mathcal{H}$ the conditional probability that $Z_D \neq Z'_D$ is at most:

$$|N_D| \times \binom{\widetilde{d}(v)^{3/4}}{\deg_H(v)^{3/10}}(\frac{1}{2}\sqrt{\Delta})^{-\deg_H(v)^{3/10}} < \Delta \times \left(\frac{2e\Delta^{3/4}}{\Delta^{3/10}\sqrt{\Delta}}\right)^{\widetilde{d}(v)^{3/10}} < \tfrac{1}{8}\Delta^{-10}.$$

Since $|Z_D| - |Z'_D| \leq \Delta$, for every choice of $\mathcal{H}$ we have $|\mathbf{Exp}(Z_D|\mathcal{H}) - \mathbf{Exp}(Z'_D|\mathcal{H})| = o(1)$.

After conditioning on $\mathcal{H}$, Z'_D is determined by $|N_D| \leq \widetilde{d}(v)$ assignments and each assignment can affect Z'_D by at most $\widetilde{d}(v)^{3/10}$. Note that $\widetilde{d}(v) \leq d_{i-1}^{+}(v) < 4d_{i-1}^{-}(v)$ (by Corollary 35(a)). So the Simple Concentration Bound with $c = (4d_{i-1}^{-}(v))^{3/10}$ yields that for any choice of $\mathcal{H}$ we have

$$\begin{aligned} & \mathbf{Pr}(|Z'_D - \mathbf{Exp}(Z'_D|\mathcal{H})| > \tfrac{1}{4}(d_{i-1}^{-}(v))^{49/50}|\mathcal{H}) \\ < \ & 2\exp\left(-\tfrac{1}{16}(d_{i-1}^{-}(v))^{49/25}/(2 \times (4d_{i-1}^{-}(v))^{3/5} \times 4d_{i-1}^{-}(v))\right) \\ < \ & \tfrac{1}{8}\Delta^{-10}, \end{aligned}$$

since $d_{i-1}^{-}(v) \geq \frac{1}{8}\Delta^{200\epsilon}$. This, along with our claim and our bound on $\mathbf{Pr}(Z_D \neq Z'_D)$, implies that $\mathbf{Pr}(|Z_D - \mathbf{Exp}(Z_D)| > \frac{1}{2}(d_{i-1}^{-}(v))^{49/50}) < \frac{1}{2}\Delta^{-10}$. Along with the analogous bound for Z_C, we have

$$\mathbf{Pr}(E_2(v)) < \Delta^{-10}.$$

□

Proof of Lemma 33 We will apply the Lovasz Local Lemma. For each uncoloured A_i and colour x, define $E(i,x)$ to be the event that more than $\frac{1}{2}\Delta^{19/20}$ vertices of A_i have neighbours outside of Big_i^+ with colour x in their sublists. We will apply Lemma 30 with $Q = \Delta^{200\epsilon}$ to bound the probability of $E(i,x)$. Each vertex $v \in A_i$ has at most $\Delta^{200\epsilon}$ uncoloured external neighbours (by Lemma 32(b)). Each such neighbour u chooses x for its sublist with probability at most $2\Delta^{200\epsilon}/|L(u)| < 1/(\Delta^{200\epsilon}\Delta^{1/5})$, since by (P8.2), $|L(u)| \geq X \geq \frac{1}{2}\sqrt{\Delta}$ and $\epsilon < \frac{3}{40000}$. Furthermore, these colour assignments are made independently and so (P7.1) holds. Since $\Delta^{19/20} > \Delta^{37/40}$, applying Lemma 30 yields $\mathbf{Pr}(E(i,x)) < \exp(-\Delta^{1/40}) < \Delta^{-10}$.

For each uncoloured A_i, define $\mathcal{D}(i)$ to be the set of events consisting of $E(j,x)$ for every colour x and uncoloured A_j such that A_j contains a vertex within distance 2 of some vertex in A_i. It is straightforward to check that the random choices which determine whether $E(i,x)$ holds have no affect on whether any events outside of $\mathcal{D}(i)$ hold; it follows that $E(i,x)$ is mutually independent of all events outside of $\mathcal{D}(i)$. Since F has maximum degree at most $10^9\Delta$, each $\mathcal{D}(i)$ has size less than $10^9 c\Delta^3$. Thus with positive probability, none of these events hold since $c\Delta^3 \times \Delta^{-10} < \frac{1}{4}$. □

We now close this subsection with our deferred proof:

Proof of Lemma 34:

Part (a): Set $\theta = 1 - \frac{1}{2}\Delta^{-\epsilon}$. We use induction to prove the stronger statement:

$$U_i \leq \theta^i U + (\theta^i U)^{99/100}.$$

The base case is trivial as $U_0 = U$. Assuming the statement holds for $i-1$, we have $U_{i-1} \leq \theta^{i-1}U + (\theta^{i-1}U)^{99/100} < 2\theta^{i-1}U < 4\theta^i U$ (since $\theta > \frac{1}{2}$). Since $U_{i-1} > U_i \geq \Delta^{150\epsilon}$, this implies

$$\theta^i U > \tfrac{1}{4}\Delta^{150\epsilon}. \tag{4}$$

We also have:

$$\begin{aligned}
U_i &= \theta U_{i-1} + U_{i-1}^{49/50} \\
&\leq \theta\left(\theta^{i-1}U + (\theta^{i-1}U)^{99/100}\right) + \left(\theta^{i-1}U + (\theta^{i-1}U)^{99/100}\right)^{49/50} \\
&< \theta^i U + \theta^{\frac{99}{100}i + \frac{1}{100}}U^{99/100} + 2^{49/50}\theta^{\frac{49}{50}i - \frac{49}{50}}U^{49/50} \\
&< \theta^i U + \theta^{\frac{99}{100}i + \frac{1}{100}}U^{99/100} + 2\theta^{\frac{49}{50}i}U^{49/50} \qquad (\text{since } 2^{\frac{1}{50}} > \theta^{-\frac{49}{50}}) \\
&< \theta^i U + \theta^{\frac{99}{100}i}U^{99/100}(\theta^{1/100} + 2(\theta^i U)^{-1/100})
\end{aligned}$$

$$
\begin{aligned}
&< \theta^i U + (\theta^i U)^{99/100}((1 - \frac{1}{2}\Delta^{-\epsilon})^{1/100} + 2(\tfrac{1}{4}\Delta^{150\epsilon})^{-1/100}) \qquad \text{by (4)} \\
&< \theta^i U + (\theta^i U)^{99/100}(1 - \tfrac{1}{200}\Delta^{-\epsilon} + 3\Delta^{-1.5\epsilon}) \\
&< \theta^i U + (\theta^i U)^{99/100}.
\end{aligned}
$$

Part (b): It is trivial that $d_i^-(v) \le d_i^+(v)$. The rightmost inequality follows from a nearly identical proof to that of part (a); we omit the repetitive details. So we focus on the leftmost inequality - the lower bound on $d_i^-(v)$. We will use induction to prove a stronger statement:

$$d_i^-(v) \ge \theta^i \deg_H(v) - (\theta^i \deg_H(v))^{99/100}.$$

The argument is very close to that of part (a), but has a few slight differences. Again, the base case is trivial and we have $\theta^i \deg_H(v) > \frac{1}{4}\Delta^{150\epsilon}$ by the same reasoning as for (4). We also have (by the aforementioned omitted proof) the upper bound

$$d_{i-1}^-(v) \le d_{i-1}^+(v) \le \theta^{i-1}\deg_H(v) + (\theta^{i-1}\deg_H(v))^{99/100} < 2\theta^{i-1}\deg_H(v) < 4\theta^i \deg_H(v),$$

which we will use to obtain the second line below.

$$
\begin{aligned}
d_i^-(v) &= \theta d_{i-1}^-(v) - d_{i-1}^-(v)^{49/50} \\
&\ge \theta\left(\theta^{i-1}\deg_H(v) - (\theta^{i-1}\deg_H(v))^{99/100}\right) - \left(4\theta^i \deg_H(v)\right)^{49/50} \\
&> \theta^i \deg_H(v) - \theta^{\frac{99}{100}i}\deg_H(v)^{99/100}(\theta^{1/100} + 4(\theta^i \deg_H(v))^{-1/100}) \\
&> \theta^i \deg_H(v) - (\theta^i \deg_H(v))^{99/100},
\end{aligned}
$$

where the last line follows by the same argument as in part (a). □

9 Colouring Cliques

In this section, we present the technique that we use for colouring the vertices of $\cup_i A_i$.

In this setting, F is partially coloured so that no two neighbours have the same colour. We have a collection of cliques $A' \subset \{A_1, ..., A_t\}$ where no vertices in the cliques of A' are coloured. We have the following condition on the colours appearing outside of the cliques. (Recall that for $v \in A_i$, an *external neighbour* of v is a neighbour of v that is not in $A_i \cup \text{All}_i$.)

(P9.1) For each $A_i \in A'$ and each colour x, $\text{Notbig}(i, x) \le 10\Delta^{19/20}$.

(P9.2) Every uncoloured vertex in some $A_i \notin A'$ has at most $30\Delta^{1/4}$ external neighbours.

Lemma 36 *We can extend our partial colouring to all vertices of the cliques in A' such that for every uncoloured $A_i \notin A'$ and every colour x,* $\mathrm{Notbig}(i,x)$ *increases by at most $\Delta^{19/20}$.*

We use a two step procedure. In the first step, we assign a random permutation of colours to the vertices of A_i. Possibly some vertices will need to be recoloured because they receive a colour that is on an external neighbour; those vertices will be recoloured more carefully in the second step.

Step 1: For each $A_i \in A'$, we choose uniformly at random a set of $|A_i|$ colours that don't appear on All_i and we assign a random permutation of those colours onto the vertices of A_i. Note that by Lemma 12(c), at least $|A_i|$ such colours are available. We define:

- Temp_i is the set of vertices $v \in A_i$ that receive a colour which is also on an external neighbour of v.

We remark that that colour might have been on the external neighbour of v in the given partial colouring, or it might have been assigned to the external neighbour during Step 1.

Lemma 37 *With positive probability:*

(a) For each $A_i \in A'$, $|\mathrm{Temp}_i| \le 10^8\sqrt{\Delta}$.

(b) For each uncoloured $A_i \notin A'$ and each colour x, at most $\frac{1}{2}\Delta^{19/20}$ vertices in A_i have an external neighbour not in Big_i^+ that receives x.

We prove this lemma in Subsection 9.1. We choose a colouring satisfying conditions (a) and (b) of Lemma 37. For each vertex v, we use $\gamma(v)$ to denote the colour of v.

Step 2: For each $A_i \in A'$ and each vertex $v \in \mathrm{Temp}_i$, we define $\mathrm{Swappable}_v$ to be the set of vertices $u \in A_i - \mathrm{Temp}_i$ that can swap colours with v. More specifically, $u \in A_i$ is in $\mathrm{Swappable}_v$ if:

(a) $u \notin \mathrm{Temp}_i$;

(b) $\gamma(u)$ does not appear on any external neighbour of v;

(c) $\gamma(v)$ does not appear on any external neighbour of u.

By Lemma 37(a), at most $10^8\sqrt{\Delta}$ vertices $u \in A_i$ violate (i). By Lemma 12(b), there are at most $10^8\sqrt{\Delta}$ colours appearing on external neighbours of v and so at most $10^8\sqrt{\Delta}$ vertices violate (b). Since Big_i^+ is a clique of F (by Lemma 12(g)) at most one vertex of Big_i^+ has colour $\gamma(v)$. So Lemma 12(c) and the fact that $\gamma(u), \gamma(v)$ do not appear on All_i imply that at most $\frac{3}{4}\Delta + 10^8\sqrt{\Delta}$ vertices u violate (c) because $\gamma(v)$ appears on an external neighbour of u in Big_i^+. At the beginning of this procedure, at most $10\Delta^{19/20}$ vertices $u \in A_i$ had $\gamma(v)$ on an external neighbour not in Big_i^+ (by (P9.1)), and by Lemma 37(b), at most $\frac{1}{2}\Delta^{19/20}$ vertices $u \in A_i$ had such an external neighbour receive $\gamma(v)$ during Part 1. So in total, fewer than $\frac{3}{4}\Delta + 11\Delta^{19/20}$ vertices $u \in A_i$ violate (c). Therefore, since $|A_i| \geq \Delta - 10^8\sqrt{\Delta}$ (by Lemma 12(a)), we have:

$$|\mathrm{Swappable}_v| \geq |A_i| - 10^8\sqrt{\Delta} - 10^8\sqrt{\Delta} - \frac{3}{4}\Delta - 11\Delta^{19/20} > \Delta/10.$$

For each $A_i \in A'$ and each $v \in \mathrm{Temp}_i$, we choose 10^4 uniformly random members of $\mathrm{Swappable}_v$. We call these vertices *candidates* of v. We say that a *candidate* u of v is *bad* if either:

(i) u is a candidate of some other vertex;

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

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

(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' and $\gamma(w') = \gamma(v)$.

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

Lemma 38 *With positive probability:*

(a) For each $A_i \in A'$, every $v \in \mathrm{Temp}_i$ has a good candidate.

(b) For each uncoloured $A_i \notin A'$ and each colour x, at most $\frac{1}{2}\Delta^{19/20}$ vertices in A_i have a neighbour not in Big_i^+ that either has a candidate of colour x or is a candidate for some vertex of colour x.

We prove this lemma in Subsection 9.1.

We choose candidates satisfying conditions (a) and (b) of Lemma 38. For each $v \in \text{Temp}_i$, we swap the colour of v with that of one of its good candidates. Our definition of *good* ensures that we have a proper partial colouring of F.

Proof of Lemma 36: Consider the extension of the colouring to A' obtained using the process above. The fact that $\text{Notbig}(i, x)$ does not increase by more than $\Delta^{19/20}$ follows from Lemmas 37(b) and 38(b), noting that a vertex can only be given the colour x in Step 2 if it has a candidate with colour x or if it is a candidate for a vertex with colour x. □

9.1 Proofs of Lemmas 37 and 38

We say that A_i, A_j are *adjacent* if some $u \in A_i$ is adjacent to some $v \in A_j$. We say they are at distance at most two if they are identical, or adjacent, or some A_k is adjacent to A_i and to A_j.

Proof of Lemma 37 We will use the Lovasz Local Lemma. For each $A_i \in A'$, we define $E_1(i)$ to be the event that $|\text{Temp}_i| > 10^8\sqrt{\Delta}$. For each uncoloured $A_i \notin A'$ and colour x, we define $E_2(i, x)$ to be the probability that more than $\frac{1}{2}\Delta^{19/20}$ vertices in A_i have an external neighbour not in Big_i^+ that receives x.

Define $\mathcal{D}(i)$ to be the set of events $E_1(j)$ and $E_2(j, x)$ defined for any colour x and any A_j at distance at most two from A_i. It is straightforward to check that the random choices which determine whether $E_1(i)$ holds have no affect on whether any events outside of $\mathcal{D}(i)$ hold; it follows that $E_1(i)$ is mutually independent of all events outside of $\mathcal{D}(i)$. Similarly, $E_2(i, x)$ is also mutually independent of all events outside of $\mathcal{D}(i)$. Since F has maximum degree at most 10^9, each A_j is at distance at most two from fewer than $10^9\Delta^4$ other sets and so each $\mathcal{D}(i)$ has size less than $2 \times 10^9 c\Delta^4$. Thus with positive probability, none of these events hold since $2 \times 10^9 c\Delta^4 \times \Delta^{-10} < \frac{1}{4}$.

We actually bound the conditional probability of $E_1(i)$ given the colour assignments for all cliques other than A_i, and the choice of $|A_i|$ colours to be used on A_i. Summing over all possible choices gives us the bound on the unconditional probability of this event. So, our random experiment will be only the choice of the permutation of those colours onto the vertices of A_i. For any vertex $v \in A_i$, the probability that $v \in \text{Temp}_i$ is at most $10^8\sqrt{\Delta}/|A_i|$ since v has at most $10^8\sqrt{\Delta}$ external neighbours by Lemma 12(b). So $\mathbf{Exp}(|\text{Temp}_i|) \le 10^8\sqrt{\Delta}$.

An easy application of McDiarmid's Inequality shows that, given our conditioning, $|\text{Temp}_i|$

is highly concentrated. $|\text{Temp}_i|$ is determined by the random permutation from the $|A_i|$ colours to the vertices of A_i. It will be more convenient for us to view this step as a permutation from the vertices to the colours, rather than vice versa.

If $|\text{Temp}_i| \geq s$ then the colours of s members of Temp_i certify that fact. Switching the colours of two vertices in A_i only affects whether those two vertices are in Temp_i. So applying (3) with $c = r = 1$ yields:

$$\mathbf{Pr}(E_1(i)) \leq \mathbf{Pr}(|\text{Temp}_i| - \mathbf{Exp}(|\text{Temp}_i|) > \sqrt{\Delta}) < 4\exp\left(-(\sqrt{\Delta})^2/(128 \times (10^8+1)\sqrt{\Delta})\right) < \Delta^{-10}.$$

We apply Lemma 30 with $Q = 30\Delta^{1/4}$ to bound $E_2(i,v)$. For each A_i and colour x, each vertex $v \in A_i$ has at most $30\Delta^{1/4}$ external neighbours (by (P9.2)) and each such neighbour in some A_j is assigned x with probability at most $\frac{1}{|A_j|} < \frac{2}{\Delta} < 1/(30\Delta^{1/4} \times \Delta^{1/5})$. Furthermore, at most one vertex in each A_i is assigned x and the random permutations for different cliques are independent; so (P7.1) holds. Thus applying Lemma 30 yields: the probability that more than $\frac{1}{2}\Delta^{19/20}$ vertices $v \in A_i$ have an external neighbour not in Big_i^+ that is assigned x is at most $\exp(-\Delta^{1/40}) < \Delta^{-10}$. □

Proof of Lemma 38 We will again use the Lovasz Local Lemma. For each v in some Temp_i, we define $E_1(v)$ to be the event that v does not have a good candidate. For each uncoloured $A_i \notin A'$ and colour x, we define $E_2(i,x)$ to be the probability that more than $\frac{1}{2}\Delta^{19/20}$ vertices in A_i have an external neighbour not in Big_i^+ that either has a candidate of colour x or is a candidate of a vertex with colour x.

As in the previous proof, we define $\mathcal{D}(i)$ to be the set of events (i) $E_1(v)$ where $v \in A_j$ for some A_j at distance at most two from A_i and (ii) $E_2(j,x)$ for any colour x and any A_j at distance at most two from A_i. For each $v \in A_i$, it is straightforward to check that the random choices which determine whether $E_1(v)$ holds have no affect on whether any events outside of $\mathcal{D}(i)$ hold; it follows that $E_1(v)$ is mutually independent of all events outside of $\mathcal{D}(i)$. Similarly, $E_2(i,x)$ is also mutually independent of all events outside of $\mathcal{D}(i)$. Since F has maximum degree at most 10^9, each A_j is at distance at most two from fewer than $10^9\Delta^4$ other sets and each A_j contains at most c vertices. So each $\mathcal{D}(i)$ has size less than $2\times 10^9 c\Delta^4$. Thus with positive probability, none of these events hold since $2 \times 10^9 c\Delta^4 \times \Delta^{-10} < \frac{1}{4}$.

It will be useful to note that for any vertex u which chooses candidates, the probability that a particular vertex $w \in \text{Swappable}_v$ is chosen is $10^4/|\text{Swappable}_v| \leq 10^5/\Delta$.

To bound $\mathbf{Pr}(E_1(v))$ for some $v \in \text{Temp}_i$, we will choose the candidates in two rounds. In the first round, we choose the candidates for all vertices but v; in the second round, we choose the candidates for v.

Let Y be the number of vertices $u \in \text{Swappable}_v$ that meet conditions (iv) or (v) of the definition of *bad*; note that Y is determined by the candidates selected in the first round. We will use Lemma 30 to show that, with high probability, Y is not too large. For each vertex $u \in \text{Swappable}_v$, we define θ_u to be the set of neighbours of u in $A' - A_i$. Thus, since every vertex in A' has at most $10^8\sqrt{\Delta}$ external neighbours, each set θ_u has size at most $10^8\sqrt{\Delta} < 2\Delta^{9/10}$ and no vertex lies in more than $Q = 10^8\sqrt{\Delta}$ of these sets. We consider a vertex in $\cup_{u\in\text{Swappable}_v}\theta_u$ to be *marked* if it chooses a candidate with colour $\gamma(v)$ or if it is chosen as a candidate for a vertex of colour $\gamma(v)$. Each of these vertices has at most one potential candidate with colour $\gamma(v)$ and can be chosen as a candidate for at most one vertex with colour $\gamma(v)$ (and both cannot occur). So the probability that a vertex is marked is at most $10^5/\Delta < 1/(Q \times \Delta^{1/5})$. Furthermore, it is easy to see that (P7.1) holds. Therefore, Lemma 30 implies that:

$$\mathbf{Pr}(Y > \Delta^{39/40}) \leq \exp\left(-\Delta^{1/40}\right) < \frac{1}{2}\Delta^{-10}.$$

Now we analyze the second round. By Lemma 37(a), at most $10^4 \times 10^8\sqrt{\Delta}$ members of Swappable_v meet condition (i) of the definition of *bad*. By Lemma 12(b), v has at most $10^8\sqrt{\Delta}$ external neighbours, and each has at most 10^4 candidates; this, along with the fact that every colour γ appears on at most one member of Swappable_v, implies that at most $10^4 \times 10^8\sqrt{\Delta}$ members of Swappable_v violate (ii) or (iii). So if $Y \leq \Delta^{39/40}$ then the number of bad members of Swappable_v is at most $\Delta^{39/40} + 10^4 \times 10^8\sqrt{\Delta} + 10^4 \times 10^8\sqrt{\Delta} < 2\Delta^{39/40}$, and so the probability that v does not choose a good candidate during the second round is at most

$$\left(\frac{2\Delta^{39/40}}{\Delta/10}\right)^{10^4} < \frac{1}{2}\Delta^{-10}.$$

Therefore, $\mathbf{Pr}(E_1(v)) \leq \frac{1}{2}\Delta^{-10} + \frac{1}{2}\Delta^{-10} = \Delta^{-10}$.

For each uncoloured A_i and colour x, each vertex $v \in A_i$ has at most $30\Delta^{1/4}$ external neighbours by (P9.2). Each such neighbour in some A_j chooses a candidate with colour x or is chosen as a candidate by the at most one vertex in its clique with colour x with probability at most $10^4/\Delta < 1/(30\Delta^{1/4} \times \Delta^{1/5})$, and it is easy to check that (P7.1) holds. So applying Lemma 30 with $Q = 30\Delta^{1/4}$ yields that $\mathbf{Pr}(E_2(i,x)) \leq \exp(-\Delta^{1/40}) < \Delta^{-10}$. □

10 The Proof At Last

We now prove Lemma 13 using a five phase procedure to colour the graph F from Lemma 12.

10.1 Phase I

In this phase, we colour S.

Step 1: Here we obtain a colouring of some of S by applying the following lemma.

Lemma 39 *There is a colouring of the subgraph of F induced by S such that:*

(a) For every $v \in S$ either $d_H(v) < \Delta - 3\sqrt{\Delta}$, or there are at least $3\sqrt{\Delta}$ colours that appear at least twice in $N(v) \cap S$;

(b) For every i, x, $\mathrm{Notbig}_{i,x} \leq \Delta^{19/20}$.

(c) Every vertex in S has at most $\frac{19\Delta}{20}$ coloured neighbours.

We give the proof of this lemma below.

Step 2 We colour the remainder of S using the procedure from Section 8.

We set H to be the subgraph of F induced by the uncoloured vertices of S. For each $u \in H$ we initialize $L(u)$ to be the set of colours that do not appear on neighbours of u in F. By Lemma 12(b) we can take $U = 10^8\sqrt{\Delta}$. By Lemma 39(a), $L(u) \geq \deg_H(u) + 3\sqrt{\Delta}$ and so we can take $X = 3\sqrt{\Delta}$. Lemma 39(c) implies that $|L(u)| \geq c - \frac{19\Delta}{20} > 5\Delta^{1/5}U$, thus satisfying (P8.2).

Therefore, Lemmas 39(b) and 31 imply that we can extend our colouring to the vertices of S so that no two neighbours have the same colour and:

$$\text{For every } i, x, \mathrm{Notbig}_{i,x} \leq 2\Delta^{19/20}. \tag{5}$$

We close this subsection with:

Proof of Lemma 39 The proof is very similar to that of Lemma 10. We will apply the Local Lemma to the same process studied there. For each $v \in S$ with more than $\Delta - 3\sqrt{\Delta}$ neighbours in H, define $E_1(v)$ to be the event that v has fewer than $3\sqrt{\Delta}$ colours that appear at least twice in its neighbourhood. For each i, x, define $E_2(i, x)$ to be the event that $|\mathrm{Notbig}_{i,x}| > \Delta^{19/20}$. For each $v \in S$, define $E_3(v)$ to be the event that v has more than $19\Delta/20$ coloured neighbours. We will prove below that the probability of each of these events is at most Δ^{-10}.

Since every vertex in S has degree at most Δ (by Lemma 12(d)), it is straightforward to check that each event is mutually independent of all but fewer than $3c\Delta^5$ other events. Thus, our lemma follows from the Lovasz Local Lemma since $\Delta^{-10} \times 3c\Delta^5 < \frac{1}{4}$.

We gave already the bound on the probability of $E_1(v)$ in the proof of Lemma 10. We consider now $E_2(i,v)$. For each A_i and colour x, each vertex $v \in A_i$ has at most $10^8\sqrt{\Delta}$ external neighbours (by Lemma 12(b)) and each such neighbour is activated and assigned x with probability at most $\frac{9}{10} \times \frac{1}{c} < \frac{1}{\Delta} < 1/(10^8\sqrt{\Delta} \times \Delta^{1/5})$. Furthermore, these assignments are made independently and so (P7.1) holds. Thus applying Lemma 30 with $Q = 10^8\sqrt{\Delta}$ yields: the probability that more than $\Delta^{19/20}$ vertices $v \in A_i$ have an external neighbour not in Big_i^+ that is assigned x in Step 2 is at most $\exp(-\Delta^{1/40}) < \Delta^{-10}$. This is clearly an upper bound on $\mathbf{Pr}(E_2(i,x))$.

Finally, the number of neighbours of v that are activated is distributed like $BIN(\deg(v), \frac{9}{10})$. The probability that this number is at least $\frac{19}{20}\Delta$ is maximized when $\deg(v) = \Delta$, and the Chernoff Bound implies that probability is at most $2e^{-(\Delta/20)^2/3\Delta(9/10)} < \Delta^{-10}$. This is clearly an upper bound on $\mathbf{Pr}(E_3(i,x))$. □

10.2 Phase II

In this phase, we colour all of B_H, using the procedure from Section 8. We start by setting H to be the subgraph of F induced by B_H. For each $v \in H$ we initialize $L(v)$ to be the set of colours not appearing on any neighbours of v. By Lemma 12(b) we can take $U = 10^8\sqrt{\Delta}$. By the definition of B_H, each $v \in H$ has at most $\Delta - \Delta^{3/4} > c - \frac{1}{2}\Delta^{3/4}$ neighbours in $S \cup B_H$. Therefore $|L(v)| \geq \deg_H(v) + X$ where $X = \frac{1}{2}\Delta^{3/4}$ and so $|L(v)| \geq X \geq 5U \times \Delta^{1/5}$ and (P8.2) holds. Thus, Lemma 31 and (5) imply that we can extend our colouring to B_H such that no two neighbours have the same colour, and:

$$\text{For every } i, x, \mathrm{Notbig}(i,x) \leq 3\Delta^{19/20}. \tag{6}$$

10.3 Phase III

In this step, we colour each $A_i \in A_H$, using the procedure from Section 9. We set $A' = A_H$. By (6) we satisfy condition (P9.1). By definition, every vertex in a clique $A_i \notin A'$ has at most $30\Delta^{1/4}$ neighbours in $F - A_i \cup \mathrm{All}_i$ and so we satisfy condition (P9.2). Lemma 36 and (6) imply that we can extend our colouring to the vertices of A_H such that:

$$\text{For every colour } x \text{ and every uncoloured } A_i, \text{Notbig}(i,x) \le 4\Delta^{19/20}. \quad (7)$$

10.4 Phase IV

In this phase, we extend our colouring to B_L. Once again, we apply the procedure from Section 8.

We start by setting H to be the subgraph of F induced by the vertices of B_L. For each $u \in H$ we initialize $L(u)$ to be the set of colours not appearing on any neighbours of u. All remaining uncoloured A_i are in A_L and thus we can take $U = 30\Delta^{1/4}$. For each $v \in H$, Lemma 12(f) guarantees that there is some $A_i \in A_L$ such that v has at most $c - \sqrt{\Delta} + 9$ neighbours outside of A_i. So $|L(v)| \ge \deg_H(v) + X$ where $X = \sqrt{\Delta} - 9$ and so $|L(v)| \ge X \ge 5U \times \Delta^{1/5}$ and we satisfy (P8.2). Thus, Lemma 31 and (7) imply that we can extend our colouring to B_L such that no two neighbours have the same colour and:

$$\text{For every colour } x \text{ and every uncoloured } A_i, \text{Notbig}(i,x) \le 5\Delta^{19/20}. \quad (8)$$

10.5 Phase V

In this step, we colour each $A_i \in A_L$, using the procedure from Section 9. We set $A' = A_L$. By (8) we satisfy condition (P9.1). No set $A_i \notin A'$ is uncoloured and so we trivially satisfy condition (P9.2). So Lemma 36 implies that we can extend our colouring to the vertices of A_L.

This completes the colouring of F and hence completes the proof of Lemma 13, and hence of our main theorem.

11 The Algorithms

We close this paper by presenting the algorithmic implications of this work.

As a corollary of Theorem 5, we can determine for every constant $\Delta \ge \Delta_0$, the precise values of c for which one can test in polynomial time whether a graph of maximum degree Δ is c-colourable (under the hypothesis that $P \ne NP$). This is well-known to be trivial for

$c \leq 2$. Embden-Weinert et al[11] used their construction that we presented in Section 1.2 to prove that for $3 \leq c \leq \Delta - k_\Delta - 1$, we cannot test this in polytime unless $P = NP$:

Theorem 40 *For every constant Δ and every $3 \leq c \leq \Delta - k_\Delta - 1$, it is NP-hard to test whether graphs of maximum degree Δ are c-colourable.*

Theorem 5 easily implies that, for sufficiently large Δ, Theorem 40 is tight. Furthermore, the proof of Theorem 5 yields a deterministic polynomial time algorithm that will actually produce the colouring that it gurantees to exist. So we have the following complement to Theorem 40:

Theorem 41 *For every constant $\Delta \geq \Delta_0$ and every $c \geq \Delta - k_\Delta$, there is a linear time deterministic algorithm to test whether graphs of maximum degree Δ are c-colourable. Furthermore, there is a polynomial time deterministic algorithm that will produce a c-colouring whenever one exists.*

For the case where Δ is not constant, the threshold for polynomial testability of c-colouring is (probably) higher: at $\Delta - \Theta(\log \Delta)$.

Theorem 42 *(a) For any constant a and any function $c : Z^+ \to Z^+$ with $c(\Delta) \geq \Delta - a \log \Delta$, there is a polynomial time deterministic algorithm to test whether a graph G is $c(\Delta(G))$-colourable so long as $\Delta(G) \geq \Delta_0$. Furthermore, there is a polynomial time randomized algorithm that will produce a $c(\Delta(G))$-colouring whenever one exists.*

(b) Consider any function $\gamma : Z^+ \to Z^+$ such that:

(i) $\gamma(n) = o(n)$;

(ii) $\lim_{n\to\infty} \gamma(n)/\log n = \infty$; *and*

(iii) $\gamma(n)$ *can be computed in* $\mathrm{poly}(n)$ *time.*

If there is a polynomial time algorithm to test whether a graph G is $(\Delta(G) - \gamma(\Delta(G)))$-colourable so long as $\Delta(G) \geq \Delta_0$, then there is a subexponential time algorithm to test whether any graph is 3-colourable.

Since it is NP-hard to determine the chromatic number of a graph[19], it is widely believed that this cannot be done in subexponential time, and so Theorem 42(b) indicates that for unbounded Δ, we probably cannot test for c-colourability if $\Delta - c$ is asymptotically larger than $O(\log \Delta)$.

Proof of Theorem 41: We are given a graph G of maximum degree $\Delta \geq \Delta_0$ where Δ is a constant, and we wish to determine whether it is c-colourable for some $c \geq \Delta - k_\Delta$. We will present a linear time algorithm to do so.

If $c = \Delta - k_\Delta$ then we carry out an initial step in which we find all of the reducers. Each time we find a reducer D, we take the c-reduction via D; i.e., we remove the clique C and we contract the vertices of the stable set S into a single vertex. It is easily seen that the reduced graph is c-colourable iff G is (see the discussion in Section 1.2). To carry this step out in linear time, we loop through every vertex and examine its neighbourhood to determine whether it is in a reducer. For each vertex, this check takes constant time since the size of the neighbourhood is Δ which is constant. If we find a reducer, then the reduction takes constant time as we are only removing and contracting a constant number of vertices.

Theorem 5 now implies that G is c-colourable iff the subgraph induced by $\{v\} \cup N(v)$ can be c-coloured for every v. For each v, we can check this in constant time since the subgraph has constant size. So it takes linear time to check all of these subgraphs and determine whether G is c-colourable.

Next, we describe how to actually produce a c-colouring in polynomial time. Basically, we work through the proof of Theorem 5, showing how to make it constructive. But we have to be a bit careful since that proof assumes that G is a minimum counterexample, and this might not be the case for a general input. That assumption is only used in the proofs of Lemmas 15, 16, 23, 26, and Observation 25; it will be straightforward to handle the case where G violates those.

First we carry out the same initial step as above, to obtain a reduction of G that contains no reducers. We do this even if $c > \Delta - k_\Delta$, and so we know that Observation 25 holds for the resulting graph.

Next we carry out a similar step to remove all near-reducers. We search the graph for any near-reducers. If we find one then there are a few possible ways to deal with it, corresponding to different arguments from the proof of Lemma 26. Let K, S denote the clique and stable set of the near-reducer X. If there is a vertex of K that has no neighbours outside of X, then we remove X from G. Let $Z = N(K) - X$; i.e., the set of vertices in $G - X$ that are adjacent to vertices in K. If Z contains any edges then we remove X from G. If any vertex in Z has fewer than $c - 1$ neighbours in $G - X$, then we remove X from G. Otherwise, for every pair $x, y \in Z$ we consider the graph $G - X + xy$; i.e. the graph obtained by removing

X and adding the edge xy. We check whether, in this graph, x, y lie in either a reducer or in a non-c-colourable subgraph induced by $\{v\} \cup N(v)$ for some v. If they do not, then we replace G by $G - X + xy$. We repeat this step iteratively, until the graph contains no near-reducers that we can remove in this manner. The arguments in the proof of Lemma 26 show that if we can c-colour the resulting graph, then we can modify that into a c-colouring of G. This modification is easily done in polynomial time, as is this iterative step.

As described in the proof of Lemma 26: if any connected component of the graph still contains near-reducers, then that component must consist of a cycle of near-reducers. Such a component is easily c-coloured, as described in that proof, and thus we can remove it. We iterate until the remaining graph has no near-reducers, and thus satisfies Lemma 26.

Our next step is to decompose the graph into $X_1, X_2, ..., X_t, S$ as described in Lemma 14. The way to produce this decomposition is given in [28], and it can easily be implemented in linear time. Since G is c-colourable, so is each X_i; i.e. Lemma 15 holds. For each set X_i, we check to see whether $\overline{G[X_i]}$ has a matching of size $\lceil 10^2\sqrt{\Delta}\rceil$. (This takes constant time since $|X_i|$ is bounded by a constant.) If it does then we remove X_i. The proof of Lemma 16 explains how to extend any c-colouring of the resulting graph to X_i.

So at this point, we have a graph that is decomposed as in Lemma 14 such that Lemmas 15, 16, 23, 26, and Observation 25 all hold. All other lemmas also hold, since they do not rely on the graph being a minimum counterexample.

Next we construct the c-colouring for each X_i as described in Section 4; this takes constant time for each X_i since $|X_i|$ is bounded by a constant. Then we construct G' and carry out Modifications 1 and 2 as described in Section 5, thus forming the graph F from Lemma 12. Again, those Modifications require constant time for each X_i and so F is constructed in linear time.

Thus far, the algorithm has been straightforward. The remaining work is to produce the colourings whose existence was proven in Sections 8, 9 and 10.1. Those colourings were proven to exist using the Lovasz Local Lemma. For each of those proofs, the main theorem of Moser and Tardos[27] implies that a very simple randomized algorithm will produce the colouring in polynomial expected time. Furthermore, the fact that Δ is constant implies that the maximum degree in the underlying dependency graph is bounded by a constant. This is enough to allow us to use the technique from [26] to derandomize the algorithm and thus obtain a deterministic algorithm.

Note: the main theorem of [27] is stated in terms of a more general version of the Lovasz Local Lemma than what we use in this paper. To convert our applications into their terms, we use a standard substitution such as $x_i = 2p$ (see Section 19.3 of [25]). $\square$

Remark: The algorithmic technique introduced by Beck[3] would also apply here, but the newer approach of [27] (see also [26]) is simpler.

And now we turn to the case where Δ is not constant. The proof of Theorem 42(a) is much like that of Theorem 41.

Proof of Theorem 42(a): We can assume that $a \geq 1$ as this implies the statement for smaller values of a.

We start with the decision algorithm; i.e. we wish to determine whether the graph is c-colourable where $c = \Delta - a \ln \Delta$. The main difficulty here is that we have no constant bound on the size of $\{v\} \cup N(v)$ and so we cannot check whether the graph it induces is c-colourable in constant time. Note that Lemma 15 is the only place in the proof where we use the fact that, in a counterexample, $\{v\} \cup N(v)$ is c-colourable for each v. So our proof actually shows that G is c-colourable iff every X_i is c-colourable.

The fact that we have $c \geq \Delta - a \ln \Delta$ allows us to revise the parameters in our lemmas somewhat. We only outline the differences as the proofs remain the same. We say that a vertex is *nearsparse* if it has at most $\binom{\Delta}{2} - 10^6 a\Delta \ln \Delta$ edges in its neighbourhood. We take a decomposition like that of Lemma 14, except that for each i: $\Delta - 10^7 a\Delta \ln \Delta \leq |X_i| \leq \Delta + 10^7 a\Delta \ln \Delta$, there are at most $10^7 a\Delta \ln \Delta$ edges from X_i to $G - X_i$, and every vertex in S is nearsparse. Nearsparseness is sufficient for arguments nearly identical to those in Section 10.1, since $c \geq \Delta - a \ln \Delta$. The other conditions are even stronger than those in Lemma 14 and so the arguments in the other sections still apply. In particular, it is still true that G is c-colourable iff every X_i is c-colourable.

So it suffices to check whether each X_i is c-colourable. We use Edmond's algorithm[10] to find a maximum matching M in $\overline{X_i}$, the complement of X_i. If M contains more than $2 \times 10^7 a \ln \Delta$ edges, then we can easily colour X_i with $|X_i| - 2 \times 10^7 a \ln \Delta < c$ colours by treating each edge of M as a colour class of size two.

For the case when M has fewer than $2 \times 10^7 a \ln \Delta$ edges: Let C denote the vertices of X_i that are not matched in M. Since M is maximum, C is a clique, and $|C| = |X_i| - 2|M| \geq \Delta - 3 \times 10^7 a \ln \Delta$. We colour the vertices of C using colours $\{1, ..., |C|\}$. Every vertex $v \in X_i - C$ now has a list L_v of colours that do not appear on $N_{X_i}(v) \cap C$. X_i can be c-coloured iff $X_i - C$ can be list-coloured using these lists. The latter condition can be tested using dynamic programming (see below) in time poly$(|G|) \times 2^{2|X_i - C|}$. Since $|X_i - C| = 2|M| = O(\ln \Delta)$ (as a is constant), this is polynomial in $|G|$.

To test whether $X_i - C$ can be list-coloured, we do the following: For each subgraph

$H \subseteq X_i - C$, and each $1 \leq i \leq c$, we test whether H can be list-coloured using only the colours from $\{1, ..., i\}$; i.e. if every vertex $v \in H$ has list $L_v \cap \{1, ..., i\}$. We carry out these tests in increasing order of $|H|$. There are only $c \times 2^{|X_i - C|}$ tests that have to be completed. For any particular H, we check every $H' \subset H$ to see whether H' could be the set of vertices coloured i. That is, we check whether (1) H' is a stable set; (2) $i \in L_v$ for each $v \in H'$; (3) $H - H'$ can be list-coloured using the colours from $\{1, ..., i-1\}$. To test H, we only have to check $2^{|H|} \leq 2^{|X_i - C|}$ sets H' and each H' is easily checked in polytime (since the test for $H - H'$ is already done). So the overall running time is $c \times \text{poly}(|X_i - C|) \times 2^{2|X_i - C|} = \text{poly}(|G|) \times 2^{2|X_i - C|}$ as required.

The algorithm to find an actual c-colouring of G follows much like that for the case when Δ is fixed. The main differences are: (1) During the initial steps, when we eliminate reducers and near-reducers, some of the checks cannot be done in constant time; and (2) we can no longer construct the c-colouring of X_i or carry out Modifications 1 and 2 in constant time.

To deal with (1): It is an easy exercise to find all reducers and near-reducers in a graph in polynomial time. If we find a near-reducer then we process it using one of the methods by which we processed near-reducers in the proof of Theorem 41. But we have to carry out a few tests to determine which method to use. Most of the tests are easily done in polynomial time. The only subtle one is checking whether a pair of vertices $x, y \in Z$ lie in a non-c-colourable subgraph induced by some $\{v\} \cup N(v)$. Do do so, we check every vertex v with $x, y \in \{v\} \cup N(v)$ and use the algorithm described above to see whether $\{v\} \cup N(v)$ is c-colourable. This yields a graph for which Observation 25 and Lemma 26 hold. We ensure that Lemmas 15, 16, 23 hold in the same way that we did for the proof of Theorem 41, except that in this case the revised parameters described above ensure a tighter version of Lemma 16: $\overline{G[X_i]}$ has no matching of size $\lceil 10^2 a \ln \Delta \rceil$.

To deal with (2): We first try to colour X_i as in Case 1 of the colouring construction from Section 4; i.e. we apply Edmonds' algorithm to find a maximum matching in $\overline{G[X_i]}$, and if it has size at least $|X_i| - c$ then we fix our colouring by letting the edges of the matching be colour classes of size 2. If it has size smaller than $|X_i| - c$ then we find any c-colouring of X_i using dynamic programming as described above; we ensure that every colour is used at least once by recolouring some vertices if necessary. This colouring might not be optimal with respect to the criteria given in Case 2 of the construction, and hence it might not satisfy Lemmas 19, 20 and 21. We iteratively check whether the colouring satisfies those lemmas, and if it does not, then we modify it as described in the proofs of those lemmas; each such iteration is easily done in polytime. Each time we do this for Lemma 19 or 20, we reduce the size of C_i and so we make at most $|X_i|$ such modifications. Each time we do this for

Lemma 21, we reduce $\sum \lambda_i^2$, and so we make (far) fewer than $|X_i|^2$ such modifications. So within polynomial time we will have a colouring satisfying those three lemmas. The proofs of those lemmas are the only places in which we used the fact that the colouring is optimal with respect to the criteria for Case 2 of the construction. Thus, the colouring of X_i that we obtain in this manner is sufficient to yield the graph F of Lemma 12.

Finally, we obtain the colouring of F using the randomized algorithm from [27], just as we did in the proof of Theorem 41. Because Δ is not bounded, it is not clear whether the algorithm can be derandomized. □

Finally, we prove the corresponding negative result:

Proof of Theorem 42(b): Suppose that we have a function $\gamma(n)$ as in the theorem statement. Since $\gamma(n) = o(n)$ and $\lim_{n\to\infty} \gamma(n)/\log n = \infty$, there is a function $t : Z^+ \to Z^+$ such that: (i) $t(n) - n - 2 \geq \gamma(t(n)) \geq n$; (ii) $t(n) \leq 2^{o(n)}$; and (iii) $t(n)$ can be easily computed in poly(n) time using binary search, since $\gamma(n)$ can be computed in poly(n) time.

For any graph H and $c \leq |H|$, we can test the c-colourability of H as follows:

First compute $t = t(|H|)$. Next, choose a graph X on $t - |H| + 1$ vertices with $\chi(X) = t - \gamma(t) - c$, and such that X has at least one vertex that is adjacent to every other vertex in X. This is straightforward since $t - c - 2 \geq t - |H| - 2 \geq \gamma(t) \geq |H|$ and so $2 \leq t - \gamma(t) - c < |X|$. Then we form G by joining every vertex of X to every vertex in H. $\chi(G) = \chi(H) + \chi(X) = \chi(H) + t - \gamma(t) - c$. Also, $\Delta(G) = |H| + |X| - 1 = t$. Therefore, testing the c-colourability of H is equivalent to testing the $(\Delta(G) - \gamma(\Delta(G))$-colourability of G which, by hypothesis, can be done in poly($|G|$) = poly(t) time - a running time that is subexponential in $|H|$. □

References

[1] K. Appel and W. Hakken. *Every Planar Map Is Four Colorable.* Bull. Amer. Math. Soc. **82** (1976), 711–712.

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

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

[4] C. Berge. *Perfect Graphs.* In: Six Papers on Graph Theory, Indian Statistical Institute, Calcutta, (1963) 1–21.

[5] A. Beutelspacher and P. Hering. *Minimal Graphs for Which the Chromatic Number Equals the Maximal Degree.* Ars Combinatorica **18** (1984), 201–216

[6] O. Borodin and A. Kostochka. *On an Upper Bound on a Graph's Chromatic Number, Depending on the Graphs's Degree and Density.* J. Combinatorial Th. (B) **23** (1977), 247–250

[7] R. Brooks. *On Colouring the Nodes of a Network.* Proc. Cambridge Phil. Soc. **37** (1941), 194–197

[8] 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).

[9] M. Chudnovsky, N. Robertson, P. Seymour and R. Thomas. *The strong perfect graph theorem.* Ann. of Math. (2) **164** (2006), 51 - 229.

[10] J. Edmonds. *Paths, Stars and Flowers.* Canadian J. Math. **17** (1965), 449 - 467.

[11] S. Embden-Weinert, S. Hougardy and B. Kreuter. *Uniquely Colourable Graphs and the Hardness of Colouring Graphs of Large Girth.* Combinatorics Probability & Computing **7** (1998), 375 - 386.

[12] P. Erdős. *Graph Theory and Probability.* Canadian J. Math. **11** (1959), 34 - 38.

[13] 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 (1975), 609 - 627.

[14] B. Farzad. *When the Chromatic Number is Close to the Maximum Degree.* M.Sc. thesis, Dept. of Computer Science, University of Toronto, (2001).

[15] B. Farzad, M. Molloy and B. Reed. $(\Delta - k)$*-critical graphs.* J. Comb. Th. (B), **93** (2005), 173 - 185.

[16] M. Furer. *Improved Hardness Results for Approximating the Chromatic Number.* Proceedings of the Thirty-Sixth Annual IEEE Symposium on Foundations of Computer Science (1995), 414 - 421.

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

[18] T. Jensen and B. Toft. *Graph Coloring Problems.* Wiley (1995).

[19] R. Karp. *Reducibility among Combinatorial Problems.* In: Complexity of Computer Computations, Plenum Press (1972), 85 - 103.

[20] C. McDiarmid. *On the Method of Bounded Differences.* In: Surveys in Combinatorics, Proc. 14th British Combinatorial Conference, Cambridge University Press (1989), 148 - 188.

[21] C. McDiarmid. *Concentration.* In: Probabilistic Methods for Algorithmic Discrete Mathematics, (Habib M., McDiarmid C., Ramirez-Alfonsin J., Reed B., Eds.), Springer (1998), 195 - 248.

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

[23] M. Molloy and B. Reed. *Colouring Graphs Whose Chromatic Number Is Near Their Maximum Degree.* Proceedings of Latin American Theoretical Informatics (1998).

[24] M. Molloy and B. Reed. *Colouring Graphs When the Number of Colours is Almost the Maximum Degree.* Proceedings of the 33rd ACM Symposium on Theory of Computing (2001), 462 - 470.

[25] M. Molloy and B. Reed. *Graph Colouring and the Probabilistic Method.* Springer (2002).

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

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

[28] B. Reed. $\omega, \Delta,$ *and* χ. J. Graph Th. **27** (1998), 177 - 212.

[29] N. Robertson, D. Sanders, P. Seymour and R. Thomas. *The four-colour theorem.* J. Combin. Th. (B) **70** (1997), 2 - 44.

[30] P. Tait. *Remarks on the Colouring of Maps.* Proceedings of the Royal Edinborough Society, **10** (1878–1880), 501 - 503, 729.

[31] M. Talagrand. *Concentration of Measure and Isoperimetric Inequalities in Product Spaces.* Institut Des Hautes Études Scientifiques, Publications Mathématiques **81** (1995), 73 - 205.

12 Appendix: Talagrand's Inequality

Here we prove the version of Talagrand's Inequality that we gave in Section 3. The original inequality provided by Talagrand is much more general than the one we stated there, but it does not apply as directly in the setting of this paper. We start by stating it:

Consider any n independent random trials $T_1, ..., T_n$, and let $\mathcal{A}$ be the set of all the possible sequences of n outcomes of those trials. For any subset $A \subseteq \mathcal{A}$, and any real ℓ, we define $A_\ell \subseteq \mathcal{A}$ to be the subset of sequences which are within a distance ℓ of some sequence in A with regards to an unusual measure. In particular, we say that $x = (x_1, ..., x_n) \in A_\ell$ iff for every set of reals $b_1, ..., b_n$, there exists at least one $y = (y_1, ..., y_n) \in A$ such that

$$\sum_{x_i \neq y_i} b_i < \ell \left(\sum_{i=1}^{n} b_i^2 \right)^{1/2}.$$

Setting each $b_i = 1$ (or in fact, setting each $b_i = c$ for any constant $c > 0$), we see that if $y \in A_\ell$ then there is an $x \in A$ such that x and y differ on at most $\ell\sqrt{n}$ trials. Furthermore, if $y \in A_\ell$, then no matter how we weight the trials with b_i's, there will be an $x \in A$ such that the total weight of the trials that x and y differ on is small.

Talagrand's Original Inequality[31]: *For any n independent trials $T_1, ..., T_m$, any set $A \subseteq \mathcal{A}$ and any real ℓ,*

$$\mathbf{Pr}(A) \times \mathbf{Pr}(\overline{A}_\ell) \leq \mathrm{e}^{-\ell^2/4}.$$

Recall our reworking of Talagrand's Inequality from Section 3: *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 $t \geq 0$, we have

$$\mathbf{Pr}(|X - \mathbf{Exp}(X)| > t + 20c\sqrt{r\mathbf{Exp}(X)} + 64c^2r) \leq 4e^{-\frac{t^2}{8c^2r(\mathbf{Exp}(X)+t)}}. \tag{9}$$

Remark: The $64c^2r$ term only arises from Lemma 43 below. It can be eliminated in settings where one can show $\mathbf{Exp}(X)$ differs from the median by less than, say, $\frac{1}{2}t$; eg. when $t > 2\mathbf{Exp}(X)$.

We start by proving that X is highly concentrated around its *median*, $\mathbf{Med}(X)$.

Define $A = \{x : X(x) \geq \mathbf{Med}(X)\}$, and $C = \{y : X(y) < \mathbf{Med}(X) - t\}$. Our first step is to prove that $A \subseteq \overline{C_\ell}$, where $\ell = t/(c\sqrt{r\mathbf{Med}(X)})$. So consider some $x \in A$. Let I be the set of indices of the at most $r\mathbf{Med}(X)$ trials whose outcomes certify that $X(x) \geq \mathbf{Med}(X)$. For each i, we set $b_i = c$ if $i \in I$ and $b_i = 0$ if $i \notin I$. Note that

$$\sum_{i=1}^{m} b_i^2 = c^2|I| \leq rc^2\mathbf{Med}(X) = (t/\ell)^2.$$

Now consider any $y \in C$. Define y' to be the outcome which agrees with x on all indices of I, and with y on all other indices. Since I certifies that $X(x) \geq \mathbf{Med}(X)$, we also have $X(y') \geq \mathbf{Med}(X)$. Since y and y' differ only on trials in I on which x and y differ, and since changing the outcome of any one $T_i \in I$ can affect X by at most $b_i = c$, we have that $X(y) \geq X(y') - \sum_{x_i \neq y_i} b_i \geq \mathbf{Med}(X) - \sum_{x_i \neq y_i} b_i$. Thus,

$$\sum_{x_i \neq y_i} b_i > t \geq \ell \left(\sum_{i=1}^{m} b_i^2 \right)^{1/2}.$$

Since this is true for every $y \in C$, we have $x \notin C_\ell$.

Therefore $A \subseteq \overline{C_\ell}$, and so $\mathbf{Pr}(\overline{C_\ell}) \geq \mathbf{Pr}(A) \geq \frac{1}{2}$, by the definition of median. Therefore, by Talagrand's Inequality:

$$\mathbf{Pr}(C) \leq 2\mathrm{e}^{-\ell^2/4} < 2\mathrm{e}^{-\frac{t^2}{4c^2r\mathbf{Med}(X)}}.$$

Next, set $A' = \{x : X(x) \leq \mathbf{Med}(X)\}$, $C' = \{y : X(y) > \mathbf{Med}(X) + t\}$ and $\ell = t/(c\sqrt{r(\mathbf{Med}(X) + t)})$. By a nearly identical argument, we obtain that $\overline{A'_\ell} \supseteq C'$, and so

$$\mathbf{Pr}(C') \leq \mathbf{Pr}(\overline{A'_\ell}) \leq 2\mathrm{e}^{-\frac{t^2}{4c^2r(\mathbf{Med}(X)+t)}}.$$

Therefore

$$\mathbf{Pr}(|X - \mathbf{Med}(X)| > t) \le \mathbf{Pr}(C \cup C') \le 4\mathrm{e}^{-\frac{t^2}{4c^2 r(\mathbf{Med}(X)+t)}} \tag{10}$$

To show that this implies concentration around the *mean*, we prove that the mean and median do not differ by very much:

Lemma 43 *Under the preconditions of (9),* $|\mathbf{Exp}(X) - \mathbf{Med}(X)| \le 20c\sqrt{r\mathbf{Exp}(X)} + 64c^2 r.$

Proof First, observe that $\mathbf{Exp}(X) - \mathbf{Med}(X) = \mathbf{Exp}(X - \mathbf{Med}(X))$. We will bound the absolute value of this latter term by partitioning the positive real line into the intervals $I_i = (i \times c\sqrt{r\mathbf{Med}(X)}, (i+1) \times c\sqrt{r\mathbf{Med}(X)}]$, defined for each integer $i \ge 0$. Clearly, $|\mathbf{Exp}(X - \mathbf{Med}(X))|$ is at most the sum over all I_i of the maximum value in I_i times the probability that $|X - \mathbf{Med}(X)| \in I_i$, which is

$$\begin{aligned} & \sum_{i \ge 0} (i+1) \times c\sqrt{r\mathbf{Med}(X)} \times \mathbf{Pr}(|X - \mathbf{Med}(X)| \in I_i) \\ = & \sum_{i \ge 0} c\sqrt{r\mathbf{Med}(X)} \times \mathbf{Pr}(|X - \mathbf{Med}(X)| \in \cup_{j \ge i} I_j). \end{aligned}$$

Setting $t = i \times c\sqrt{r\mathbf{Med}(X)}$ in (10) yields a bound on the probability that $|X - \mathbf{Med}(X)| \in \cup_{j \ge i} I_j$ of:

$$4\mathrm{e}^{-\frac{t^2}{4c^2 r(\mathbf{Med}(X)+t)}} < 4\mathrm{e}^{-\frac{t^2}{8c^2 r \max\{\mathbf{Med}(X), t\}}} < 4(\mathrm{e}^{-\frac{t^2}{8c^2 r\mathbf{Med}(X)}} + \mathrm{e}^{-\frac{t^2}{8c^2 rt}}) = 4(\mathrm{e}^{-\frac{i^2}{8}} + \mathrm{e}^{-\frac{i\sqrt{\mathbf{Med}(X)}}{8c\sqrt{r}}}).$$

Therefore:

$$|\mathbf{Exp}(X - \mathbf{Med}(X))| < c\sqrt{r\mathbf{Med}(X)} \times \sum_{i \ge 0} 4(\mathrm{e}^{-\frac{i^2}{8}} + \mathrm{e}^{-\frac{i\sqrt{\mathbf{Med}(X)}}{8c\sqrt{r}}}).$$

It is straightforward to bound $\sum_{i \ge 0} 4\mathrm{e}^{-\frac{i^2}{8}} < 12.5$. For the second summation, we use the identity $\sum_{i \ge 0} x^i = (1-x)^{-1}$ for $x < 1$ and the bound $\mathrm{e}^{-y} < 1 - \frac{1}{2}y$ for $y < 1.5$. We set $y = \frac{\sqrt{\mathbf{Med}(X)}}{8c\sqrt{r}}$ and $x = \mathrm{e}^{-y}$. For the case $y \ge 1.5$, the second summation is less than $(1 - \mathrm{e}^{-1.5})^{-1} < 1.3$ and so the total is less than $14c\sqrt{r\mathbf{Med}(X)}$. For $y < 1.5$ we have:

$$\sum_{i \ge 0} 4\mathrm{e}^{-\frac{i\sqrt{\mathbf{Med}(X)}}{8c\sqrt{r}}} = \frac{4}{1-x} < \frac{4}{\frac{1}{2}y} = 64c\sqrt{r}/\sqrt{\mathbf{Med}(X)}.$$

This, with our bounds above, yields:

$$\begin{aligned}|\mathbf{Exp}(X) - \mathbf{Med}(X)| &< 14c\sqrt{r\mathbf{Med}(X)} + 64c^2 r \\ &< 20c\sqrt{r\mathbf{Exp}(X)} + 64c^2 r,\end{aligned}$$

where the last inequality makes use of the fact that $X \geq 0$ and so $\mathbf{Exp}(X) \geq \frac{1}{2}\mathbf{Med}(X)$. □

We apply again the fact that $X \geq 0$ and so $\mathbf{Exp}(X) \geq \frac{1}{2}\mathbf{Med}(X)$ to obtain: $\mathbf{Med}(X) + t \leq 2(\mathbf{Exp}(X) + t)$. By Lemma 43, if $|X - \mathbf{Exp}(X)| \geq t + 20c\sqrt{r\mathbf{Exp}(X)}) + 64c^2 r$ then $|X - \mathbf{Med}(X)| \geq t$. Therefore:

$$\begin{aligned}\mathbf{Pr}(|X - \mathbf{Exp}(X)| > t + 20c\sqrt{r\mathbf{Exp}(X)} + 64c^2 r) &\leq \mathbf{Pr}(|X - \mathbf{Med}(X)| > t) \\ &\leq 4\mathrm{e}^{-\frac{t^2}{4c^2 r(\mathbf{Med}(X)+t)}} \\ &\leq 4\mathrm{e}^{-\frac{t^2}{8c^2 r(\mathbf{Exp}(X)+t)}}\end{aligned}$$

as required. □

McDiarmid's Inequality[22] extends Talagrand's original inequality to the setting of independent permutations, with a minor loss in the constant in an exponent (see Corollary 2.8 of [22]). Theorem 1.1 of [22] shows that, under the conditions of our statement from Section 3, (10) holds, after replacing a "4" with a "16", i.e.:

$$\mathbf{Pr}(|X - \mathbf{Med}(X)| > t) \leq \mathbf{Pr}(C \cup C') \leq 4\mathrm{e}^{-\frac{t^2}{16c^2 r(\mathbf{Med}(X)+t)}}$$

The same arguments used to prove Lemma 43 yield a similar bound again, just with somewhat worse constants:

$$|\mathbf{Exp}(X) - \mathbf{Med}(X)| \leq 25c\sqrt{r\mathbf{Exp}(X)} + 128c^2 r.$$

Our reworking of McDiarmid's Inequality from Section 3 then follows in the same manner as our reworking of Talagrand's Inequality.