Colouring graphs when the number of colours is nearly the maximum degree

[Extended Abstract]

Michael Molloy*
Dept of Computer Science
University of Toronto
Toronto, ON, Canada
molloy@cs.toronto.edu

Bruce Reed
Equipe Combinatoire
CNRS
Paris, France.
reed@ecp6.jussieu.fr

ABSTRACT

We consider for graphs of maximum degree Δ, the problem of determining whether $\chi(G) > \Delta - k$ for various values of k. We obtain sharp theorems characterizing when the barrier to $\Delta - k$ colourability must be a local condition, i.e. a small subgraph, and when it can be global. We also show that for large fixed Δ, this problem is either NP-complete or can be solved in linear time, and we determine precisely which values of k correspond to each case

Keywords

graph colouring, probabilistic method

1. INTRODUCTION

Graph colouring is one of the most fundamental and notoriously difficult problems in graph theory. For example, it was one of the original six problems that Karp showed to be NP-complete in his seminal paper[10]. Since then, the complexity of various types of constrained colouring has been a very active ongoing research area. In this paper, we consider the complexity of colouring graphs with fixed maximum degree. We completely settle this issue whenever the maximum degree is sufficiently large. We do this using a graph theoretic theorem which characterizes non-c-colourable graphs whenever c is close to the maximum degree.

We are interested in determining when the chromatic number of a graph is the maximum of the chromatic number of its small subgraphs and when it is instead determined by more global factors. Obviously, this question has algorithmic implications since if "small" means that the size of the subgraph is bounded a constant t then in the first case, we can determine the chromatic number of G in polynomial time by calculating the chromatic number of each of the $O(n^t)$ such small subgraphs using brute force.

There have always been close ties between the maximum degree of a graph and bounds on its chromatic number. Brooks' Theorem, one of the earliest and most cited results in graph theory, is of this flavour. It states that for a graph G of maximum degree $\Delta > 2$, if $\chi(G) = \Delta + 1$ then G contains a clique of size $\Delta + 1$. (This is not true for $\Delta = 2$, consider an odd cycle of length at least 5). Now, every graph has a $(\Delta + 1)$-colouring because we can just colour the vertices one at a time, at each point choosing a colour which does not appear on a neighbour of the vertex under consideration. Further, if G has a clique of size $\Delta + 1$ then its chromatic number is at least $\Delta + 1$. Thus, Brooks' Theorem says that for $\Delta > 2$, $\chi(G) = \Delta + 1$ if and only if G contains a $\Delta + 1$ clique. Such a clique will be a component of G so we can test if G is Δ-colourable in linear time. (In fact, we can also test this in the case $\Delta = 2$, as then we are just asking if G is bipartite.)

It is natural to ask whether Brooks' Theorem extends further. Eg., for what k can we say that if $\chi(G) > \Delta - k$ then G has a $(\Delta - k + 1)$-clique? Reed[18] showed that this holds for $k = 1$ when Δ is large, thus settling a conjecture of Beutelspacher and Herring[2]. It is conjectured by Borodin and Kostochka [3] that this is true for $\Delta \geq 9$; counterexamples are known for each $\Delta \leq 8$ (this is Problem 4.8 in [8]).

This pattern does not extend to $k = 2$. For any $\Delta \geq 4$, consider the graph formed by adding all edges between a $(\Delta - 4)$-clique and a 5-cycle. It is easy to see that this graph has chromatic number $\Delta - 1$ but its largest clique has size $\Delta - 2$. However, using the results of this paper, Farzad, Molloy and Reed[7] show that for Δ sufficiently large, if $\chi(G) > \Delta - 2$ then G contains either a $(\Delta - 1)$-clique, or the graph described above. Similar results are shown for $\chi(G) > \Delta - 3$ and $\chi(G) > \Delta - 4$; in these cases, G must contain one of 5 and 38 graphs respectively. Recall that a t-*critical* graph is a graph with chromatic number t such that every proper subgraph can be $(t-1)$-coloured. So the results mentioned above characterize the $(\Delta - k + 1)$-critical graphs for large Δ and $k \leq 3$.

For any $\Delta \geq 2$, we define k_Δ to be the maximum k s.t. $(k+1)(k+2) \leq \Delta$. Thus, $k_\Delta \approx \sqrt{\Delta} - 2$.

Our main theorem, Theorem 1, says that for Δ sufficiently large, and for $k < k_\Delta$, something very similar to Brooks' Theorem holds. If $\chi(G) > \Delta - k$ then G must contain a

*Supported in part by an NSERC Research Grant and a Sloan Research Fellowship.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
STOC'01, July 6-8, 2001, Hersonissos, Crete, Greece.
Copyright 2001 ACM 0-58113-349-9/01/0007 ...$5.00.

subgraph H which is very close to a $(\Delta - k + 1)$-clique, in that H has small size and cannot be $(\Delta - k)$-coloured. Theorem 2 and, more specifically, Theorem 7 show that our main theorem is tight.

Theorem 1 is graph theoretical, but it's proof yields a very interesting result about the complexity of graph colouring, namely that for $k \leq k_\Delta$ we can test $(\Delta - k)$-colourability for graphs with large constant maximum degree Δ in linear time; this is Theorem 4. Theorem 3 shows that this is tight (unless P=NP). Theorems 5 and 6 give similar results when Δ is allowed to grow with the size of the input graph.

THEOREM 1. *There is an absolute constant Δ_0 such that for any $\Delta \geq \Delta_0$ and any $k < k_\Delta$, if G has maximum degree Δ and if $\chi(G) > \Delta - k$ then G contains a subgraph H such that*

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

(ii) $\chi(H) = \Delta - k + 1$.

We contrast this with:

THEOREM 2. *For every $\Delta \geq 2$ and $k > k_\Delta$, there are arbitrarily large $(\Delta - k + 1)$-critical graphs with maximum degree Δ.*

So Theorem 1 implies that for $k < k_\Delta$ and Δ large, all $(\Delta - k + 1)$-critical graphs have bounded size while Theorem 2 says that for $k > k_\Delta$ this is not true. We leave the case $k = k_\Delta$ until Section 2, as it is less easily described. Roughly speaking, for most values of Δ, all $(\Delta - k_\Delta + 1)$-critical graphs have size at most $O(\Delta^2)$. 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.

Theorem 2 is derived from the proof of the following theorem due to Embden-Weinert, Hougardy and Kreuter[5] which is even stronger (unless P=NP).

THEOREM 3. *For any constant Δ and $\Delta - 2 > k > k_\Delta$, $(\Delta - k)$-colourability of graphs with maximum degree Δ is NP-hard.*

(Of course, this theorem won't hold for $k \geq \Delta - 2$ since then the problem is no harder than 2-colourability.)

An immediate corollary of Theorem 1 is that for constant $\Delta \geq \Delta_0$ and $k < k_\Delta$, $(\Delta - k)$-colourability is in P. This is because, as we described earlier, for constant Δ a simple brute force algorithm can test every subgraph of size at most $\Delta + 1$ for $(\Delta - k)$-colourability in polytime. In fact, we will see that we don't have to search *every* subgraph of size $\Delta + 1$, and we can test for $(\Delta - k)$-colourability in linear time. If $k = k_\Delta$ then we will see that any large $(\Delta - k + 1)$-critical subgraphs are easily recognizable, and any small ones can be found quickly; so again we have a linear time algorithm.

THEOREM 4. *For any constant $\Delta \geq \Delta_0$ and $k \leq k_\Delta$, there is a simple linear time algorithm for $(\Delta - k)$-colourability of graphs with maximum degree Δ.*

When Δ can grow with the size of the graph, our proof of Theorem 1 will yield the following, which we showed in [15].

THEOREM 5. *For any constant a, there is a polytime algorithm which tests c-colourability of graphs with maximum degree $\Delta \geq \Delta_0$ for any $c \geq \Delta - a \log \Delta$.*

This is (probably) sharp since it is not hard to show that

THEOREM 6. *Consider any function $\gamma(n)$ which tends to infinity with n. Suppose that there is a polytime algorithm which tests c-colourability of graphs with maximum degree $\Delta \geq \Delta_0$, for any $c \geq \Delta - \gamma(\Delta) \log \Delta$. Then there is a subexponential time algorithm for colourability.*

The requirement $\Delta \geq \Delta_0$ in Theorems 1, 4 is somewhat unpalatable, since the constant Δ_0 is very large. Unfortunately, this type of requirement is an artifact of our proof technique - a complicated probabilistic argument. It is very common for applications of the probabilistic method to require this condition, and it is often obfuscated by the fact that the results of the theorems are usually at best asymptotically tight - for example, Kahn's seminal proof[9] that the list chromatic index of a graph is $\Delta + o(\Delta)$. Theorem 1 is one of the relatively few applications of this sort where the chromatic number of the graphs considered is determined exactly.

We conjecture that the condition $\Delta \geq \Delta_0$ can be dropped from our theorems (we state this more concretely as Conjecture 8 in Section 2.) If true, this would completely characterize the complexity of $(\Delta - k)$-colourability for every fixed pair (Δ, k).

In [15], we proved that Theorem 1 was true for $c > \Delta - \epsilon\sqrt{\Delta}$ for a small constant $\epsilon > 0$. The present situation is more delicate, and requires a much more complicated argument (see the discussion surrounding Lemma 16 in section 7.)

Theorems 1 and 4 follow from Theorem 12 in section 4. We will discuss the proofs of Theorems 2 and 3 in section 2. Theorem 5 first appeared in [15] and Theorem 6 is easy. Their proofs appear in the appendix.

Throughout this paper, we will use $c = \Delta - k$ to denote the number of colours being considered.

2. FOCUSSING ON THE THRESHOLD

In this section we discuss the proof of Theorem 2. The key lies in what we call reducers. We say that a subgraph H is a (k, Δ)-*reducer* if it consists of a $(\Delta - k - 1)$-clique, an independent set of size $k + 2$, and all possible edges joining them. When k, Δ are clear, we just call it a reducer.

If G has a reducer, then in every $(\Delta - k)$-colouring of G, all vertices in the independent set must have the same colour. Thus, if we replace H by a single vertex, and join that vertex to every vertex in $G - H$ which had a neighbour in the independent set, then the resulting graph is $(\Delta - k)$-colourable iff G is $(\Delta - k)$-colourable. This operation is called *reducing* the reducer to a vertex. The inverse of the operation is called *expanding* a vertex to a reducer.

Note also that within H, every vertex in the clique has degree Δ and every vertex in the independent set has degree $\Delta - k - 1$. Therefore, each vertex in the independent set might have up to $k + 1$ neighbours outside of H, and so if we reduce H to a vertex, then that vertex could have degree up to $(k + 1)(k + 2)$.

Moreover, if $(k + 1)(k + 2) = \Delta_1 > \Delta$ then consider any graph G with maximum degree at most Δ_1. By expanding every vertex of degree greater than Δ into a reducer, we can transform it into a graph G' with maximum degree Δ such that G is $(\Delta - k)$-colourable iff G' is. In fact, by using two rounds of expansions, we can do the same for any graph of maximum degree at most $\Delta_2 = (\Delta_1 - \Delta + k + 1)(\Delta_1 - \Delta + k + 2) > \Delta_1$. And with enough rounds of expansion, we can transform any graph G, regardless of its maximum degree, into a graph G' such that G is $(\Delta - k)$-colourable iff G' is.

Furthermore, G is $(\Delta - k + 1)$-critical iff G' is. Therefore, by starting with any arbitrarily large $(\Delta - k + 1)$-critical graph G, we can construct an arbitrarily large $(\Delta - k + 1)$-critical graph G' with maximum degree Δ. This proves Theorem 2.

(This is essentially the same construction that was used in [5] to prove Theorem 3.)

If $(k + 1)(k + 2) = d \leq \Delta$, then the same operation will allow us to transform any graph G with maximum degree d to a larger graph G' with maximum degree Δ such that G is $(\Delta - k)$-colourable iff G' is. If $d \leq \Delta - k - 1$, this operation is not very interesting since every graph with maximum degree at most d is $(\Delta - k)$-colourable. It is easy to check that $d \leq \Delta - k - 1$ when $k \leq k_\Delta - 1$, but for $k = k_\Delta$, possibly d can be as high as Δ. This is why $k = k_\Delta$ is treated as a special case, and the situation for it is somewhere between that in Theorem 1 and that in Theorem 2.

THEOREM 7. *For any $\Delta \geq \Delta_0$, and for any graph G with maximum degree Δ and $\chi(G) > \Delta - k_\Delta$:*

(a) When $(k_\Delta + 1)(k_\Delta + 2) < \Delta$: *Then G contains a subgraph H such that*

  (i) $|H| \leq O(\Delta^2)$;

  (ii) $\chi(H) = \Delta - k_\Delta + 1$.

(b) When $(k_\Delta + 1)(k_\Delta + 2) = \Delta$: *There are arbitrarily large $(\Delta - k_\Delta + 1)$-critical graphs with maximum degree Δ. But, if G has no reducer, then G must have a subgraph H as in (a).*

We are now able to state our conjecture:

CONJECTURE 8. *For every Δ, there exists $f(\Delta)$ such that: For any $k < k_\Delta$, if G has maximum degree Δ and if $\chi(G) > \Delta - k$ then G contains a subgraph H such that*

(i) $|H| \leq f(\Delta)$;

(ii) $\chi(H) = \Delta - k + 1$.

For $k = k_\Delta$, the same is true if G has no reducer.

If true, this conjecture would imply that the condition $\Delta \geq \Delta_0$ can be dropped from Theorem 4, but we would only have a polytime bound on $(\Delta - k)$-colouring, rather than a linear time bound.

Note that Brooks' Theorem implies Conjecture 8 for the case $\Delta \geq 2, k = 0$, since if $\chi(G) > \Delta - k$ then $\chi(G) = \Delta + 1$. If $\Delta \geq 3$ then G contains a $(\Delta + 1)$-clique, H, and if $\Delta = 2$ then a reducer is an induced P_3, i.e. an induced subgraph isomorphic to the path on 3 vertices, and so every odd cycle contains a reducer.

3. GRAPHS WITH NO NEAR-REDUCERS

As we saw in the last section, for $k > k_\Delta$, reducers can be used to create arbitrarily large $(\Delta - k + 1)$-critical graphs. In order to prove Theorem 1 when k is not much less than k_Δ, we are going to have to be very careful in dealing with subgraphs which are close to being reducers.

Note that every vertex v in a reducer has at least $\binom{\Delta}{2} - k\Delta$ edges in the subgraph induced by its neighbourhood. In this section, we shall consider graphs in which every vertex has somewhat fewer edges in the subgraph induced by its neighbourhood. In this sense, there are no subgraphs which are too near to being reducers. As we will see, this makes our task much easier.

More precisely, we say that a vertex v of G is *d-sparse* if the subgraph induced by its neighbourhood contains fewer than $\binom{\Delta}{2} - d\Delta$ edges. Otherwise, v is *d-dense* (we often simply use dense or sparse if d is clear from the context).

In this section, we will consider graphs with no $2000\sqrt{\Delta}$-dense vertices. Recall that every vertex in a reducer is k-dense, and so in this sense, our density parameter is within a constant multiple of what is required to forbid reducers. This is why we say that a graph with no $2000\sqrt{\Delta}$-dense vertices has no subgraphs which are near reducers.

Kahn[9] introduced the following simple procedure to find a partial c-colouring of a graph G.

SPARSE:

1. For each vertex v, assign to v a uniformly random colour from amongst the c colours.

2. For each vertex v which is assigned the same colour as a neighbour u, we uncolour v.

Clearly, this will produce a valid partial colouring, i.e. an assignment of colours to a (possibly empty) subset of the vertices such that no two adjacent vertices have the same colour. The procedure is simple; the substance is in its application. Essentially, one proves that with positive probability, the partial colouring obtained will satisfy certain properties which ensure that it can be completed to a c-colouring of the entire graph. Most commonly, but not always, this completion is obtained via several iterations of the procedure, each time assigning random colours to the remaining uncoloured vertices. We say "with positive probability" versus "with high probability" because we always use the Lovasz Local Lemma [6], and so the probability of success might be exponentially low. Nevertheless, this is enough to prove the existence of a c-colouring; an efficient method to produce a c-colouring is a different matter (see Section 11).

A common goal is to show that, with positive probability, every vertex has many repeated colours in its neighbourhood. For example, in [12] we essentially prove the following (see [16] for a full proof):

LEMMA 9. *Consider any $d \geq \log^2 \Delta$, any $c \geq \frac{1}{2}\Delta$ and any graph G with maximum degree Δ in which every vertex is d-sparse. With positive probability, SPARSE will produce a partial colouring such that:*

For every vertex v of degree at least $\frac{1}{2}\Delta$, there will be at least $\frac{d}{1000}$ colours which appear at least twice in $N(v)$.

Note that Lemma 9 immediately implies that for $d \geq \log^2 \Delta$, any graph in which every vertex is d-sparse has chromatic number at most $\Delta + 1 - \frac{d}{1000}$: Start with the partial colouring guaranteed by the lemma, and try to complete it in a greedy manner, one vertex at a time. When we come to vertex v, the number of colours not available for v is at most the number of coloured neighbours minus the number of colours which appear at least twice in $N(v)$, which is at most $\Delta - \frac{d}{1000}$. Therefore, there will always be at least one colour available for v.

Thus, as a corollary, we see that every graph with no $2000\sqrt{\Delta}$-dense vertices can be $(\Delta - k)$-coloured, for any $k < k_\Delta$ and so, of course, our main theorems hold for such graphs.

4. DENSE DECOMPOSITIONS

In the last section, we saw that for graphs in which every vertex is sufficiently sparse, Theorem 1 follows quickly from Lemma 9. So the meat of our problem is in dealing with dense vertices. In this section, we will describe how to seperate these vertices from the rest of the graph so that they will be easier to handle.

We will partition $V(G)$ into sets $D_1, ..., D_l$ and S where the D_i are dense clique-like sets which contain all the dense vertices. A dense set D_i will be like a Δ-clique in that:

(i) it has about Δ vertices,

(ii) most of the pairs of vertices within it are adjacent,

(iii) there are only a few edges between D_i and $G - D_i$

(iv) a vertex is in D_i if and only if most of its neighbours are.

DEFINITION 10. *Let d be a non-negative integer and let G be a graph of maximum degree Δ. We say $D_1, ..., D_l, S$ form a d-*dense decomposition *of G if:*

(a) $D_1, ..., D_l, S$ are disjoint and partition $V(G)$;

(b) every D_i has between $\Delta+1-8d$ and $\Delta+1+4d$ vertices;

(c) there are at most $8d\Delta$ edges between D_i and $V - D_i$;

(d) a vertex is adjacent to at least $\frac{3\Delta}{4}$ vertices of D_i if and only if it is in D_i;

(e) every vertex in S is d-sparse.

We shall call $D_1, ..., D_l$, *d-dense sets*, or simply dense sets.

This decomposition was introduced in [17]. It was also used in [13] and [15].

A neighbour u of a vertex v in D_i is *internal* if u is also in D_i and *external* otherwise. The internal and external neighbourhoods of v are the sets of its internal and external neighbours, respectively. The external neighbourhood of a subset X of D_i is the union of the external neighbourhoods of its elements.

This next key lemma follows in the same lines as the proof of (3.2) - (3.7) in [17] (which essentially establishes Lemma 11 for $d = \frac{1}{500}\Delta$). It's full proof appears in [16].

LEMMA 11. *For every Δ and $d \leq \frac{\Delta}{100}$, every graph G of maximum degree Δ has a d-dense decomposition.*

Furthermore, this d-dense decomposition can be found using a simple linear time algorithm. See [16] or [17] for details.

To prove Theorem 1, we prove the following stronger theorem:

THEOREM 12. *There is an absolute constant Δ_0 such that for any $\Delta \geq \Delta_0$ and any $k < k_\Delta$: Consider any graph G of maximum degree Δ and any d-dense decomposition of G with $d = 2000\sqrt{\Delta}$. G is $(\Delta - k)$-colourable iff every D_i is $(\Delta - k)$-colourable. For $k = k_\Delta$, the same is true if G has no reducer.*

Theorem 12 implies Theorem 1 where the bound on $|H|$ is weakened from $\Delta + 1$ to $\Delta + O(\sqrt{\Delta})$. This is because if $\chi(G) > \Delta - k$, i.e. if G is not $(\Delta - k)$-colourable, then there is at least one dense set D_i which is not $(\Delta - k)$-colourable. So we can set H to be a $(\Delta - k + 1)$-critical subgraph of D_i. Below, we will see how this yields the bound of $\Delta+1$ required for Theorem 1. If $k = k_\Delta$, then we repeatedly reduce all reducers until none remain, and then apply Theorem 12. With a small amount of calculation, this yields Theorem 7.

Theorem 12 also implies Theorem 4. We repeatedly reduce all reducers until none remain. Then we take a dense decomposition, and test whether each D_i is $(\Delta-k)$-colourable. Since $|D_i| = O(1)$, each such test can be done in constant time.

The proof of Theorem 12 is very long and complicated. There is no way to present even a thorough outline of it in this extended abstract. So instead, we discuss some of the main ideas and prove a few key lemmas. Our aim is to get across a feeling for the type of arguments used in the proof and an intuitive idea of why it works.

We close this section by describing how Theorem 12 yields the sharper bound on $|H|$ required for Theorem 1.

LEMMA 13. *Consider any graph G with a dense decomposition as in Theorem 12. Suppose some D_i is not $(\Delta - k)$-colourable for some $k < k_\Delta$. Then D_i has a $(\Delta - k + 1)$-critical subgraph on at most $\Delta + 1$ vertices.*

Proof Outline: Let H be any $(\Delta - k + 1)$-critical subgraph of D_i. We know that $|H| \leq |D_i| \leq \Delta + 8000\sqrt{\Delta} + 1$, and of course, $|H| \geq \Delta - k + 1$. Consider a maximum matching M in $\overline{H}$, the complement of the graph induced by H. Suppose that M has at least $10000\sqrt{\Delta}$ edges. Then we could assign a colour to each pair of vertices in this matching, thus leaving at most $\Delta - 12000\sqrt{\Delta} + 1$ uncoloured vertices. Therefore, the chromatic number of H would be at most $\Delta - 12000\sqrt{\Delta} + 1 + 10000\sqrt{\Delta} < \Delta - k + 1$. Therefore, $|M| < 10000\sqrt{\Delta}$.

The vertices of H not covered by M form a clique, since if any two of them were non-adjacent, then they could be added to M. Therefore, H has a clique A of size at least $(\Delta - k + 1) - 2|M| > \Delta - 21000\sqrt{\Delta}$. We set $T = V(H) - A$; then $|T| \leq |H| - |A| < 29000\sqrt{\Delta}$.

Since H is $(\Delta - k + 1)$-critical, each vertex $v \in H$ has degree at least $\Delta - k$ in H (as otherwise any $(\Delta-k)$-colouring of $H - v$ can easily be completed to v), and so it is non-adjacent to at most $9000\sqrt{\Delta}$ vertices in H. In particular, every vertex in T has at least $|A| - 9000\sqrt{\Delta}$ neighbours in A and so $E(A, T)$, the set of edges from A to T, has size at least $|T|(|A| - 9000\sqrt{\Delta})$. Therefore at least one vertex in $v \in A$ has at least $|T|(|A| - 9000\sqrt{\Delta})/|A| \geq |T| - 9000 \times 29000$.

Since v has degree at most Δ and is adjacent to every member of A, and all but 9000×29000 members of T, we have that $|H| \leq \Delta + 9000 \times 29000$. This allows us to repeat the same argument with sharper numbers, and this narrows the bound to $\Delta + 3$. A more careful version of the same argument yields $\Delta + 1$ (we omit the details). □

5. MODIFYING THE COLOURING PROCEDURE.

As we saw in section 2, reducers are small $(\Delta-k)$-colourable subgraphs which force global non-$(\Delta-k)$-colourability when

$k > k_\Delta$. This is because they have a unique $(\Delta - k)$-colouring, and if they are joined in an inauspicous manner, we cannot choose a $(\Delta - k)$-colouring of each of them which yields a $(\Delta - k)$-colouring of the whole graph. Similarly, the most problematic dense sets D_i, i.e. those which are closest to being reducers, each have a very small number of $(\Delta - k)$-colourings. This makes it difficult to choose $(\Delta - k)$-colourings for each of them which can be extended to a $(\Delta - k)$-colouring of the whole graph.

In particular, consider applying SPARSE to a graph G as in Theorem 12 using $c = \Delta - k$ colours. It is quite likely that SPARSE will produce a partial colouring which does not extend to a c-colouring even within some D_i. To avoid this problem, we must guide the procedure somewhat. We fix a partition of each dense set into c stable sets, and modify SPARSE so that the random colour classes chosen yield the same partition of D_i into stable sets.

So consider a graph G, with a dense decomposition as in Theorem 12, and consider a particular c-colouring of each dense set D_i. This yields a partition of D_i into c stable sets (possibly some are empty); we refer to these stable sets as the *partition classes* of D_i.

A simple way to modify SPARSE for graphs as in Theorem 12 is as follows:

DENSE:

1. For each vertex $v \in S$, assign to v a uniformly random colour from amongst the c colours.

2. For each dense set D_i, take a uniformly random permutation of the c colours to the partition classes of D_i.

3. For each vertex $v \in S$ which is assigned the same color as a neighbour u, we uncolour v.

4. For each vertex $v \in D_i$ which is assigned the same color as a neighbour u, we uncolour v and all other vertices in the same partition class as v.

One nice property of DENSE is that the analogue of Lemma 9 holds for it, with the modification that we only guarantee many repeats in the neighbourhood of each $v \in S$. So, at least intuitively, completing our colouring on the vertices of S is not a difficult issue. Completing the colouring on the dense sets requires most of the work. However, to get some idea as to why this might be achievable, note that a vertex in D_i never gets assigned the same colour as an internal neighbour, so condition (c) of Definition 10 ensures that we don't expect many vertices in D_i to be uncoloured. Furthermore, the small number of edges from D_i to $G - D_i$ allows us to complete the colouring of each D_i almost independently of the rest of the graph. Almost, but not quite, and this "not quite" is what causes most of the work in our proof.

To deal with the issue of completing the colouring on the dense sets, we must use a procedure which, while at heart is inspired by DENSE, is actually much more complicated. So complicated, in fact, that we do not even formally define it in this extended abstract. Instead, we just discuss some pieces of it over the next few sections.

6. ANOTHER LOOK AT REDUCERS

At this point, it is worth taking a closer look at how the presence of a reducer will prevent us from proving Theorem 12 for $k \geq k_\Delta$, It is easily seen that if G has a reducer, then that reducer will be a dense set D_i, and its c partition classes will be the $c-1$ vertices of the clique along with the independent set. After obtaining an initial partial colouring, we need to complete that colouring to the uncoloured vertices of G, including any uncoloured colour classes of D_i. Suppose that the independent set is an uncoloured partition class. Recall that the number of external neighbours of this partition class is at most $(k + 1)(k + 2)$. If this number is at least c, then conceivably, each of the c colours could already appear on a neighbour of this partition class, and so it would be impossible to assign a colour to it. (When extending our partial colouring, we generally attempt to assign the same colour to the vertices of each partition class, but sometimes we allow ourselves to break this rule. In the case of a reducer, however, this rule obviously can't be broken.)

More generally, partition classes with a large number of external neighbours are difficult to deal with, since if they are uncoloured in our partial colouring, then when we try to colour them they might have few available colours. We will explore this theme more closely in the next section.

The reader may note that the same problem can occur because of a large number of internal neighbours as well. However, internal neighbours are not nearly as problematic. The reason is that, when recolouring the partition classes of a dense set, we can readjust some colours within the set without too much difficulty since this keeps the recolouring local. However, trying to readjust the colours on neighbours outside of the dense set leads to the recolouring being spread throughout the graph, and this makes any analysis hopeless.

7. CHOOSING A GOOD PARTITION OF A DENSE SET

When using the procedure DENSE, it is important to make a good intial choice of our partition of each dense set into c stable sets. The first step is to gain a better understanding of these dense sets. We list a few key lemmas here, and briefly discuss their proofs. These proofs are presented more fully in the appendix.

Our first lemma says that we can assume that the complement of a dense set does not contain a matching with more than $O(\sqrt{\Delta})$ edges:

LEMMA 14. *If G is a minimal counterexample to Theorem 12, then no $\overline{D_i}$ has a matching of size greater than* $10000\sqrt{\Delta}$.

Proof Outline: By minimality, there is a $(\Delta-k)$-colouring of $G - D_i$. If D_i had such a matching, then we could complete that colouring to D_i, thus obtaining a $(\Delta-k)$-colouring of G; a contradiction. To complete the colouring, we let every pair in M be a colour class of size 2. Then we greedily colour the remaining vertices, in decreasing order of their external degrees. The first ones coloured will always have an available colour since most of their neighbours haven't been coloured yet. The last ones coloured will be adjacent to many pairs of vertices in M; thus they will have many colours repeated in their neighbourhoods, so there will be an available colour. □

This immediately yields the following corollary, by the same short argument as in the proof of Lemma 13:

COROLLARY 15. *If G is a minimal counterexample to Theorem 12, then every dense set D_i has a clique of size at least $\Delta - 21000\sqrt{\Delta}$.*

So from now on, we will assume that every D_i contains a clique A_i of size at least $\Delta - 21000\sqrt{\Delta}$.

In the previous section we described how a dense set in which a partition class has too many external neighbours is particularly troublesome. The next lemma is the main reason why the situation was much simpler in [15], when we were only concerned with $k \leq \epsilon\sqrt{\Delta}$ for a small constant $\epsilon > 0$. There, instead of using density parameter $d = 2000\sqrt{\Delta}$, we could take $d = 2000\epsilon\sqrt{\Delta}$.

LEMMA 16. *Consider a graph G of maximum degree $\Delta \geq \Delta_0$ and suppose that $c \geq \Delta - \frac{1}{100000}\sqrt{\Delta}$. Take any d-dense decomposition of G, where $d = \frac{1}{50}\sqrt{\Delta}$, and consider a dense set D_i. If $\chi(D_i) \leq c$ then D_i has colouring with at most c colours in which no colour class has more than $\frac{1}{2}\Delta$ external neighbours.*

Proof Outline: We look at a c-colouring which maximizes the number of colour classes of size 2. We see that in such a colouring, every vertex in a class larger than 2 must have many internal neighbours. So many, in fact, that it's total external neighbourhood is small. □

In our present situation, things are more delicate, and instead of Lemma 16 we have the following.

LEMMA 17. *Each c-colourable D_i has a partition in which no partition class has more than $c - 2$ external neighbours.*

Proof Outline: The proof follows along the same lines as that of Lemma 16. This time we choose a c-colouring which first minimizes the number of colour classes of size 1, and then lexicographically minimizes the sequence of the sizes of the other classes. □

Recall from the previous section that for $k > k_\Delta$, our proof breaks down because of dense sets with partition classes which have at least c external neighbours. Thus, our situation here is almost, but not quite, fatal. By again considering a minimal counterexample, we can provide ourselves with a little more breathing room.

LEMMA 18. *If G is a minimal counterexample to Theorem 12 then each D_i has a partition in which no partition class has more than $c - \sqrt{\Delta} + 2$ external neighbours.*

Proof Outline Any dense set D_i which has such a partition class, has a very constrained structure, and it is very close to a reducer. We use this to show that at least one c-colouring of $G - D_i$ can be completed to D_i. □

In general, the partitions which yield parts with the highest possible outdegree are very close to reducers in that they have one large partition class, and the remaining parts small. In order to simplify our discussion, for the remainder of this paper we will assume that every dense set has the following reducer-like structure which captures the nature of the problematic sets.

Simplifying Assumption 1 *Each dense set D_i consists of a $(c-1)$-clique A_i, a stable set B_i and all edges in between them. The maximum outdegree of vertices in A_i is $t < \Delta - c + 2$ and $|B_i| = \Delta - (c - 2 + t)$.*

(Note that if we had $t \geq \Delta - c + 2$, then B_i would be empty.)

Of course, the general situation is more complicated in that there can be a relatively small number (of order $O(\sqrt{\Delta})$) of additional vertices in D_i, and dealing with such vertices adds a great deal of additional complication to our proof. But even under this assumption, we get many of our key ideas across.

8. A TWO STEP PROCESS

We now break the problem up into two steps. We first colour the graph containing each B_i and all of S and then we finish off by colouring all of the A_i's.

In the first step, we will give all vertices of each B_i the same colour so we actually obtain a colouring of the graph F formed by contracting every B_i into a single vertex, b_i whose neighbourhood corresponds to the external neighbourhood of B_i, and by deleting the vertices in each A_i. If the D_i were actually reducers then this would correspond to colouring the graph to which we reduced.

It is easy to see that F has a c-colouring as we can use SPARSE to c-colour S and since each B_i has degree less than c by Lemma 18, we can greedily complete this to a c-colouring of F.

The problem is that there are many c-colourings of F which cannot be completed to a c-colouring of G because of conflicts between the vertices of A_i and their external neighbours. So, in the first step we need to ensure that we choose a c-colouring of F that can be completed to G. To do this, we need to colour F so that no colour appears on the external neighbourhood of too many vertices in some A_i. The reason that this is neccessary is because we will eventually need to assign each colour (except for that appearing on B_i) to some vertex of A_i. So we will try to find a colouring of F such that the following property holds:

(*) For each color γ and set A_i, fewer than $\frac{1}{2}\Delta$ vertices of A_i have an external neighbour coloured γ.

This is a difficult property to enforce if some B_j has too many neighbours in A_i, because the colour appearing on B_j will be on the external neighbourhood of each of B_j's neighbours in A_i. So for now, we will make another simplifying assumption:

Simplifying Assumption 2: *For each A_i, no set B_j, $j \neq i$ has more than $\Delta^{9/10}$ neighbours in A_i. Also, no vertex outside of D_i has more than $\Delta^{9/10}$ neighbours in A_i.*

We will see that obtaining a colouring satisfying (*) is relatively easy provided we make one final assumption.

Simplifying Assumption 3: *F is $\Delta^{2/3}$-sparse.*

If this simplifying assumption holds, then we can colour F using SPARSE. In section 3, we used SPARSE to show that Theorem 1 holds when G is sufficiently sparse. There, we considered a single iteration of SPARSE, and then completed the partial colouring greedily. Here, our situation is more complicated. (*) will hold for the partial colouring produced by SPARSE, but if we complete with a greedy colouring, then we might violate (*). So instead, we apply several iterations of SPARSE, until every vertex in F is coloured. (Actually, we have to do something a little different to finish the last few vertices, but we omit such technical details here.)

By Lemma 9, after the first iteration of SPARSE, with positive probability each vertex in F will have at least $\Theta(\Delta^{2/3})$ repeated colours in its neighbourhood. Since $c \geq \Delta - O(\sqrt{\Delta})$, this means that for any vertex v, no matter how future iterations go, there will always be at least $O(\Delta^{2/3})$

colours which are available for v. It follows that for any colour γ, the probability that v is coloured γ by this procedure is at most $O(\Delta^{-2/3})$.

This fact enables us to prove that with positive probability (*) holds. Consider any set A_i and colour γ. The number of external edges from A_i is at most $O(\Delta\sqrt{\Delta})$, and the external endpoint of any such edge recieves γ with probability at most $O(\Delta^{-2/3})$. Therefore, if $X_{i,\gamma}$ is the number of vertices in A_i with an external neighbour of colour γ, then $\mathbf{Exp}(X_{i,\gamma}) \leq O(\Delta^{5/6})$. We apply Talagrand's Inequality[19] to show that $X_{i,\gamma}$ is highly concentrated, and so $\mathbf{Pr}(X_{i,\gamma} \geq \frac{1}{2}\Delta) < \mathrm{e}^{-\Delta/20}$. The Local Lemma then implies that with positive probability, $X_{i,\gamma} < \frac{1}{2}\Delta$ for *every* pair A_i, γ.

We omit the details, only remarking that Simplifying Assumption 2 is crucial in showing that $X_{i,\gamma}$ is concentrated. That is because when we apply Talgrand's Inequality, we require that no single colour assignment can affect $X_{i,\gamma}$ by very much. This follows from the fact that if a vertex u retains the colour γ, then this can increase $X_{i,\gamma}$ by at most $\Delta^{9/10}$.

In Section 10 we will discuss how to obtain a colouring satisfying (*) even if Simplifying Assumption 3 does not hold. First, we will see how such a colouring can be completed to the entire graph.

9. COMPLETING THE COLOURING

Suppose that we have a colouring of F satisfying (*). We will extend the colouring to a partial colouring of G satisfying the following properties:

(P1) Each A_i has at most $20000\sqrt{\Delta}$ uncoloured vertices.

(P2) For each color c and set A_i, fewer than $\frac{3}{4}\Delta$ vertices of A_i have an external neighbour coloured c.

We can obtain this partial colouring using a single iteration of the obvious modification of DENSE, i.e. where instead of permuting the c colours onto D_i, we permute the $c-1$ colours not assigned to B_i onto A_i.

Consider any particular A_i. Property (c) of Definition 10 implies that the expected number of uncoloured vertices in A_i is at most $17000\sqrt{\Delta}$. We can prove that this variable is strongly concentrated, and so with very high probability, (P1) holds for A_i.

For any vertex v and colour γ, the probability that γ is assigned to v is at most $\frac{1}{c-1} \approx \frac{1}{\Delta}$. A very similar analysis to that used to enforce (*) in the previous section, applies here to show that with high probability no vertex in A_i has more than $\frac{1}{4}\Delta$ external neighbours which are assigned γ; this means that (P2) holds. Here, Talagrand's Inequality does not apply, and so we use McDiarmid's Inequality[11], which was commissioned by us specifically for this type of application.

An application of the Local Lemma now implies that with positive probability, (P1) and (P2) hold throughout the graph. We omit the details for this extended abstract.

Finally, we describe how a colouring with properties (P1) and (P2) can be completed to the entire graph G.

Consider any A_i. There are $r \leq O(\sqrt{\Delta})$ uncoloured vertices in A_i, we call this set U_i, and r colours which do not yet appear in D_i. Start with an arbitrary assignment of these r colours to these r vertices. For each $v \in U_i$, we will swap its colour with that of a vertex in $A_i - U_i$. We will only consider swapping the colours on v and u if the colour on v does not appear on any external neighbours of u and vice versa. For any v, the outdegree of v is at most $\Delta - c + 2 < 2\sqrt{\Delta}$, by Simplifying Assumption 1. Thus there are at least $c - 2\sqrt{\Delta}$ colours which do not appear on external neighbours of v. Also, by property (P2), the colour assigned to v does not appear on the external neighbourhood of more than $\frac{3}{4}\Delta$ vertices in A_i. Therefore, there will be at least $\frac{1}{5}\Delta$ vertices u which v can swap with.

Any one such swap will be successful, in that both u and v will be coloured legally. But of course, we must take care to ensure that the cummulation of all swaps does not create some conflicting neighbours. We do so as follows. For each vertex $v \in U_i$ we randomly select 20 candidates from amongst the at least $\frac{1}{5}\Delta$ vertices which v can swap with. We prove that with very high probability, there is at least one candidate u such that swapping with u will not create any conflicts, regardless on which candidates all the other uncoloured vertices throughout the graph choose to swap with. We omit the details from this extended abstract, commenting only that this is very similar to an argument used in [13].

This completes the colouring of each A_i and thus of the entire graph.

We close this section by briefly discussing Simplifying Assumptions 1, 2. Thus far, our colouring procedures have been fairly simple. In the last section, we applied several iterations of SPARSE; in this section, we applied a single iteration of DENSE followed by a colour swapping routine to finish off the few remaining vertices. If Assumption 1 or 2 does not hold, then we must tinker with these procedures somewhat. For example, if a vertex $u \in F$ has more than $\Delta^{9/10}$ neighbours in A_i, then we pick one of its at least $\frac{1}{4}\Delta$ non-neighbours $v \in A_i$ and deterministically assign v the same colour as u. Note that by Property (c) of Definition 10, there are at most $O(\Delta^{3/5})$ such vertices in F and so, at least intuitively, this should not have too much of an effect on our analysis.

Similarly, if some of the c partition classes in D_i, other than B_i, have size greater than 1, then we must treat them a little more carefully in the second step of our 2 step process. However, since $|D_i| \leq \Delta + O(\sqrt{\Delta})$, there are at most $O(\sqrt{\Delta})$ such partition classes, and so again, this should not have a significant effect on our analysis.

On the other hand, if Simplifying Assumption 3 does not hold, then things are much more complicated.

10. WHEN F IS DENSE

The reader may feel that the whole analysis thus far has been a cheat. We reduce the problem of c-colouring G to the problem of c-colouring the auxiliary graph F, complicating our task by requiring that the colouring also satisfy (*). We then make the simplifying assumption that F is sparse. It seems that we could just have easily assumed that G is sparse and then used the straightforward one iteration analysis given in Section 3.

However, a moment's reflection shows that F, although not always sparse is never very dense either. In particular, by Lemma 18, it consists of the sparse subgraph of G induced by S together with some vertices (the b_i) which have maximum degree $c - \sqrt{\Delta} + 2$. So consider any dense set D in a d-dense decomposition of F, and c-colour $D \cap S$ using colour classes of size 1 and 2 (this is easy as $F \cap S$ is sparse).

If we then greedily add the vertices of $D-S$ to the colouring, by Lemma 18 we will have a choice of $\sqrt{\Delta}-2$ colour classes to which we can assign each vertex. This implies that we can c-colour D so that no partition class has more than $\frac{c}{2}$ external neighbours.

We then construct $F^{'}$ from F in the same way that we constructed F from G. Each new vertex $b^{'}$ will have maximum degree at most $\frac{c}{2}$ and so $F^{'}$ will be very sparse. Thus $F^{'}$ can be coloured using several iterations of SPARSE as at the end of section 8, and this colouring can be completed to $F-F^{'}$ using the technique described in Section 9. This is sufficiently random for us to ensure that when it finishes (*) will hold, and so we can color $G-F$ again using the techniques of Section 9. There are of course many complications, readers should be grateful that we have no room for them here.

11. MORE ON ALGORITHMS

We look back at the proof of Theorem 4. Note that all we do is show that a simple algorithm checks to see whether every dense set D_i can be c-coloured. If they can, then Theorem 12 implies that G can be c-coloured, so we have solved the decision problem of c-colourability. However, we have not provided any algorithm to actually obtain the guaranteed colouring of G.

The proof of Theorem 12 involves many applications of the Local Lemma. So long as $\Delta = O(1)$ as in Theorem 4, each of these applications meets the requirements of Theorem 3.1 in [14] which implies that the technique introduced by Beck[1] yields a polytime algorithm to find the object (usually an extension of our colouring), guaranteed by the application. Combining all of these algorithms yields a polytime algorithm to find the colouring of G.

12. REFERENCES

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

[2] A.Beutelspacher and P.R.Hering, *Minimal graphs for which the chromatic number equals the maximum degree.* Ars. Comb. **18** (1984), 201 - 216.

[3] O. Borodin and A. Kostochka, *On an upper bound on a graph's chromatic number, depending on the graphs's degree and density* J.Comb.Th.(B) **23** (1977), 247 - 250.

[4] R.L. Brooks, *On colouring the nodes of a network*, Proc. Cambridge Phil. Soc. **37** (1941), 194 - 197.

[5] T. Emden-Weinert, S. Hougardy, and B. Kreuter, *Uniquely colourable graphs and the hardness of colouring graphs of large girth*, Prob. Comb. and Comp. **7** (1998), 375 - 386.

[6] P. Erdős and L. Lovász, *Problems and results on 3-chromatic hypergraphs and some related questions*, in: "Infinite and Finite Sets" (A. Hajnal et. al. Eds), Colloq. Math. Soc. J. Bolyai 11, North Holland, Amsterdam, 1975, 609-627.

[7] B. Farzad, M. Molloy and B. Reed, in preparation.

[8] T. Jensen and B. Toft, *Graph Coloring Problems*, Wiley, 1995.

[9] J. Kahn, *Asymptotically good list-colorings*, J. Combinatorial Th. (A), **73** (1996), 1 - 59.

[10] R. Karp, *Reducibility among combinatorial problems*, In Complexity of Computer Computations, Plenum Press (1972), 85 - 103.

[11] C. McDiarmid, manuscript.

[12] M. Molloy and B. Reed, *A bound on the strong chromatic index of a graph*, J.Comb.Th. (B) **69** (1997), 103 - 109.

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

[14] M. Molloy and B. Reed, *An algorithmic version of the Lovasz Local Lemma*, Proceedings of STOC 1998, 524 - 529.

[15] M. Molloy and B. Reed, *Colouring graphs whose chromatic number is near their maximum degree*, Proceedings of Latin American Theoretical Informatics, 1998.

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

[17] B. Reed, $\chi, \Delta,$ *and* ω, J.Graph Th., **27** (1998), 177 - 212.

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

[19] M. Talagrand, *Concentration of measure and isoperimetric inequalities in product spaces.* Institut Des Hautes Études Scientifiques, Publications Mathématiques **81** (1995), 73 - 205.

APPENDIX

A. A FEW MORE PROOF DETAILS

Proof of Lemma 14 Assume that $\overline{D_i}$ has a matching M with more than $10000\sqrt{\Delta}$ edges. We will see that this violates the minimality of G.

Since G is a minimal counterexample, the graph $G-D_i$ can be c-coloured; consider any c-colouring of $G-D_i$. We will extend this c-colouring to D_i. We begin by assigning a colour to each of $10000\sqrt{\Delta}$ of the vertex-pairs in M. This is always possible, since each such pair has at most $\frac{1}{2}\Delta$ colours forbidden by external neighbours and at most $10000\sqrt{\Delta}$ colours forbidden by internal neighbours. Next, we assign a colour to each remaining vertex with outdegree at least $20\sqrt{\Delta}$. By condition (c) of Definition 10, there are at most $\frac{1}{2}\Delta$ such vertices and so each vertex will have at most $\frac{1}{4}\Delta$ colours forbidden by external neighbours and at most $\frac{1}{2}\Delta + 10000\sqrt{\Delta}$ by internal neighbours. Finally, we colour all remaining vertices in D_i. Each has at most $20\sqrt{\Delta}$ colours forbidden by external neighbours and at most $|D_i| - 10000\sqrt{\Delta} \leq \Delta - 2000\sqrt{\Delta}$ by internal neighbours. Thus we will always be successful since $c > \Delta - \sqrt{\Delta}$. $\square$

Proof of Lemma 16 First note that any colour class of size at most 2 has at most $\frac{1}{2}\Delta$ external neighbours since every vertex has at most $\frac{1}{4}\Delta$ external neighbours. As in the proof of Lemma 15, we focus on a maximum matching M in $\overline{D_i}$. If $|M| \geq \frac{1}{10}\sqrt{\Delta}$ then since $|D_i| \leq \Delta+1+4d < c+|M|$, we can colour D_i with at most c colours, using each of the pairs in M as a colour classe of size 2 and using all other vertices as colour classes of size 1.

So we assume that $|M| < \frac{1}{10}\sqrt{\Delta}$. Now take any c-colouring of D_i which maximizes the number of colour classes of size 2. Thus, the classes of size 1 must form a clique, and furthermore, each vertex v in that clique is adjacent to each vertex u in any colour class of size greater than 2. This latter fact is true because otherwise u could be recoloured with the same colour as v, thus increasing the number of colour classes of size 2.

Since $|M| < \frac{1}{10}\sqrt{\Delta}$, there are at least $c - \frac{1}{10}\sqrt{\Delta} > \Delta - \frac{1}{9}\sqrt{\Delta}$ colour classes of size 1. Since each u in any colour class of size greater than 2 is adjacent to each of these vertices, it has at most $\frac{1}{9}\sqrt{\Delta}$ external neighbours. Furthermore, no colour class has size greater than $\frac{1}{5}\sqrt{\Delta}$ as otherwise within that class we could find a matching of size greater than $\frac{1}{10}\sqrt{D_i}$ in $\overline{D_i}$. Therefore, every class of size greater than 2 has fewer than $\frac{1}{5}\sqrt{\Delta} \times \frac{1}{9}\sqrt{\Delta} < \frac{\Delta}{2}$ external neighbours. □

Proof Outline of Lemma 17 The proof follows along the same lines as that of Lemma 16. We show that we can choose a colouring in which (i) every colour class of size 1 is adjacent to every vertex in a colour class of size greater than 2. We also show that (ii) most colour classes of size 1 are adjacent to at least one vertex in every colour class of size 2; and (iii) every vertex in any class of size greater than 2 has at least one neighbour in every class with at least 2 fewer vertices.

(i) and (ii) imply that no colour class has size greater than $\Delta - c + 2 = k + 2$. Also, if the two largest classes differ in size by at most 1, then neither has size greater than $\frac{1}{2}(\Delta - c + 2)$. (iii) implies that if the largest class exceeds the second largest by more than 1 then every vertex in the largest class has at most $\Delta - c + 1 = k + 1$ external neighbours. It follows that in the extremal case, the largest class has $k+2$ vertices and a total of $(k+2)(k+1)$ external neighbours, and no other classes have nearly this number.

For $k \le k_\Delta - 1$, it is easily seen that $(k+2)(k+1) < c - 2$. If $k = k_\Delta$, it can be shown that this extremal case only occurs if D_i is a reducer. We show that the extremal case for non-reducers has a largest class of size $k+1$, and a total of $(k+1)(k+1) \le c-2$ external neighbours. □

Proof Outline of Lemma 18 As we saw in the proof of Lemma 17, the largest possible number of external neighbours for a partition class is at most $(k+1)(k+2)$ which is easily checked to be at most $c - \sqrt{\Delta} + 2$ if $k \le k_\Delta - 1$. When $k = k_\Delta$, then D_i cannot be a reducer and the largest possible number of external neighbours is at most $(k+1)(k+1)$. This can only be achieved if D_i contains a $(c-1)$-clique A_i in which every vertex has at most one external neighbour, and an independent set of size $\Delta + 1 - c$. Suppose we have such a D_i.

By the minimality of G, we can c-colour $G - D_i$. We show that we can complete this colouring to D_i, thus obtaining a contradiction. To do this, we first colour all vertices in the independent set the same colour. This is possible because that set has at most $c-2$ colours forbidden by external neighbours. Next we colour the vertices of A_i one at a time. Before we reach the last vertex, we will always succeed since at most $c-2$ colours are forbidden by neighbours in D_i and at most 1 by an external neighbour. By choosing the last vertex and the c-colouring of $G - D_i$ carefully, we can guarantee that the colour appearing on the external neighbour of that vertex will also be assigned to a vertex in D_i. Thus, the total number of colours forbidden ot the last vertex is at most $c-1$ and so we will succeed in coluring it as well.

Therefore, G cannot contain such a D_i. This implies that for $k = k_\Delta$, the largest possible number of external neighbours is at most $k(k+1) < c - \sqrt{\Delta} + 2$ □

Proof outline of Theorem 5: When $c = \Delta - a \log \Delta$, the same argument as in Lemma 15 shows that we can assume each D_i has a clique of size $t \ge \Delta - a_1 \log \Delta$, for a constant $a_1 = a_1(a)$. Examining the proof of Lemma 13, we see that by repeatedly removing vertices with degree less than c in D_i, we reach a subgraph H of size $\Delta + O(1)$. Therefore, to test whether D_i can be c-coloured is equivalent to testing whether the $r \approx a_1 \log \Delta$ non-clique vertices can be $c - t$ coloured. This can be done using dynamic programming in $O(2^r)$ time, which is in P. □

Proof of Theorem 6: Suppose that for some $\gamma(n) \to \infty$ and for all $\Delta \ge \Delta_0$ and $k \ge \Delta - \gamma(\Delta)\log\Delta$, we have a polytime algorithm, $\mathcal{A}$ for k-colourability of graphs with maximum degree Δ, where Δ may grow with the size of the input. We show that this yields a subexponential time algorithm for colourability.

By adding a sufficiently large clique to a graph G, we can reduce general k-colourability to the case where $k \ge \frac{1}{2}|G|$. So, consider any graph G on n vertices, and integer $\frac{1}{2}n \le k \le n$; we wish to test whether G is k-colourable. Again, we reduce by adding a large clique to G.

Choose an integer h s.t. $\frac{1}{2}(n-k) \le \log h \times \gamma(h) \le n - k$; thus h, and in particular, $\gamma(h)$ tends to infinity with n. We can assume that $\gamma(h) \to \infty$ slowly and so $h >> n$. Construct G' by taking an h-clique, C, a copy of G, and all the edges between them. Note that G' has maximum degree $\Delta = h + n - 1$ and that G is k-colourable iff G' is $(h+k)$-colourable. Since

$$\Delta - \gamma(\Delta)\log\Delta \le \Delta - \gamma(h)\log h < h + k,$$

we can use $\mathcal{A}$ to test for k-colourability of G in poly$(|G'|)$ time, which is subexponential in n. □