Colouring Graphs whose Chromatic Number Is Almost Their Maximum Degree

Michael Molloy and Bruce Reed*

November 2, 2000

1 The Result

Determining the minimum number of colours required to colour the vertices of a given graph so that every pair of adjacent vertices receive different colours is a notoriously difficult problem. In fact, determining whether three colours will suffice was one of the original six problems that Karp reduced Satisfiability to in his seminal 1972 paper [6].

However, it is easy to see that if the maximum degree of G is Δ, then $\chi(G)$ is at most $\Delta+1$. In fact there is a simple recursive greedy procedure for colouring a graph G with $\Delta(G)+1$ colours. Having coloured $G-x$ we simply colour x with a colour not appearing on the (at most) Δ neighbours of $G-x$. Of course, this colouring may not be optimal. In fact, in 1941, Brooks[3] proved that $\chi(G)=\Delta+1$ precisely if some component of G is a $\Delta+1$-clique or $\Delta=2$ and G is not bipartite. As these two condition are easy to test, we see that we can determine if G has chromatic number $\Delta+1$ in polynomial (in fact, linear) time. It is natural to wonder if one can develop polynomial time algorithms for determining if G has chromatic number k for k close to $\Delta+1$.

One negative result in this area was provided by Maffray and Preissman[7], who proved that determing if $\chi(G)$ is Δ is NP-complete, even for graphs with $\Delta=4$ which are triangle-free. Furthermore, Emden-Weinert, Hougardy, and Kreuter[4] have shown that determining if a graph G with maximum degree Δ has a colouring using $k=\lfloor\Delta-\sqrt{\Delta}+1\rfloor$ colours is NP-complete provided $k\geq 3$ (i.e. $\Delta\geq 4$). In this paper we present two companion positive results:

Theorem 1 *There is a positive ϵ such that for each fixed Δ, and each k between $\lfloor\Delta+1-\epsilon\sqrt{\Delta}\rfloor$ and $\Delta+1$, we can determine if $\chi(G)=k$ in polynomial time.*

*This work was supported by NATO Collaborative Research Grant #CRG950235. The work of the first author is supported by an NSERC Research Grant. The work of the second author was partially supported by a FAPESP grant

Theorem 2 *There is a positive ϵ such that for any graph G of maximum degree Δ and each k between $\lfloor \Delta+1-\epsilon \log \Delta \rfloor$ and $\Delta+1$, we can determine if $\chi(G) = k$ in polynomial time.*

We remark that these theorems are only relevant for graphs for which Δ is at least ϵ^{-2} which is why they do not contradict the result of Maffray and Preissman above.

We note that as a consequence of our technique, we obtain the following interesting non-algorithmic result.

Theorem 3 *There is a positive ϵ such that if G is a graph of maximum degree Δ and chromatix number at least $\Delta + 1 - \epsilon\sqrt{\Delta}$ then the chromatic number of G is the maximum of the chromatic numbers of its subgraphs with at most $\Delta + \sqrt{\Delta}$ vertices.*

Our approach is to examine a naive colouring procedure which consists of choosing a random colour for each vertex with each possibility equally likely, uncolouring those involved in a conflict, and then completing the colouring, essentially by greedily extending the partial colouring. We shall see that, suprisingly, analyzing a slight variant of this procedure using simple but powerful probabilistic tools allows us to develop an algorithm which performs as specified above. See [8],[9],[10], [12] [13], [14] for related results, in particular [12] contains a survey of a large variety of problems solved using this and similar techniques.

We note that our result requires the application of a recent result of the authors[11], that enables us to develop algorithms from applications of the Lovasz Local Lemma([5], see also [1]). It also requires the use of Azuma's inequality[2] and a relatively new concentration inequality due to Talleygrand[15].

However, in this short summary, we do not discuss these aspects. Rather we focus on a structural decomposition used in the proof and point out why $\Delta - \sqrt{\Delta}$ is (approximately) the boundary at which things become difficult.

2 A Proof Sketch

One simple result which is crucial to the algorithm, is the following:

Lemma 4 *Any partial $\Delta - r$ colouring of G such that for every vertex v there are at least $r+1$ colours appearing twice in the neighbourhood of v, can be extended to a $\Delta - r$ colouring of G in linear time.*

Proof Simply colour the uncoloured vertices one at a time. When we come to colour v, it has at most $\Delta - r - 1$ colours appearing in its neighbourhood and hence we can use some colour not appearing on its neighbourhood. $\square$

Now, we will attempt to find such a partial colouring with $\Delta - r$ at most the maximum of $\lfloor \Delta - c\sqrt{\Delta} \rfloor$ and the chromatic number of G by analyzing the probabilistic procedure described above. That is, by considering a partial colouring obtained by colouring each vertex with a uniformily independently chosen colour and then uncolouring vertices which are involved in conflicts.

Of course, our technique will not always yield repeated colours in every neighbourhood. For example if the neighbourhood of a vertex is a clique then it will contain no repeated colours. More strongly, if there are very few non-edges in the neighbourhood of a vertex, then we cannot expect to say very much. It turns out that if every vertex has enough non-edges in its neighbourhood then G has a partial $\Delta - \epsilon\sqrt{\Delta}$ colouring in which there are $\epsilon\sqrt{\Delta}$ repeated colours in each neighbourhood. Furthermore, we can find such a colouring easily using the probabilistic method sketched above. To begin, we consider such sparse graphs. To ease our exposition during these introductory remarks, we assume that G is Δ regular. We also assume that each vertex of G has at least $e^6\epsilon\Delta^{\frac{3}{2}}$ non-edges in its neighbourhood and analyze a uniformily random partial $\frac{\Delta}{2}$ colouring $\mathcal{C}$ of the vertices of G.

Our first step is to examine the number of repeated colours we expect in N_v. We let Z_v be the number of pairs of vertices of $N(v)$ which receive the same colour and such that this colour is assigned to no other vertex of $N(v)$ (this second condition is to avoid double counting). There are at least $e^6\epsilon\Delta^{\frac{3}{2}}$ pairs of non-adjacent vertices in N_v. For any particular such pair, $\{u, w\}$, both vertices receive the same colour with probability $\frac{2}{\Delta}$. If u, w both receive the same colour, then they will form one of the pairs counted by Z_v provided they both retain that colour in $\mathcal{C}$ and no other colour in $N(v)$ is assigned that colour. I.e. iff no other vertex in $(N(v) - u - w) \cup N(u) \cup N(w)$ receives that colour. The probability that no such vertex receives that colour is $(1 - \frac{2}{\Delta})^{|N(v)-u-w) \cup N(u) \cup N(w)|} \geq (1 - \frac{2}{\Delta})^{3\Delta} \geq \mathrm{e}^{-6}$. It follows that $\mathbf{Exp}(Z_v) \geq \epsilon^6\epsilon\Delta^{\frac{3}{2}} \times \frac{2}{\Delta} \times \mathrm{e}^{-6} = 2\epsilon\sqrt{\Delta}$.

Thus, for each vertex, the expected number of colours which appear twice in its neighbourhood in $\mathcal{C}$ is high, and so it seems hopeful that with positive probability, $\mathcal{C}$ will have the desired property. To prove this requires two more steps: First, we would like to show that for each v, Z_v is highly concentrated, i.e. that the probability that Z_v differs from its expected value by a significant amount is small. This will establish that for any one particular v, with high probability N_v will contain many repeated colours. We could do this using the Concentration Inequality of Talleygrand. The second step is to strengthen this statement, showing that with positive probability *every* vertex will have many such pairs in its neighbourhood. This second step is a straightforward application of the Lovasz Local Lemma.

However, we will not prove this result about sparse graphs, we have just sketched it to indicate the main lines of our approach. It shows that the crux of the problem is to deal with those vertices with dense neighbourhoods.

We call a vertex *dense* if there are more than $\binom{\Delta}{2} - e^4\epsilon\Delta^{\frac{3}{2}}$ edges in the subgraph induced by its neighbourhood. We note that this implies that there are at most $2e^4\epsilon\Delta^{\frac{3}{2}}$ edges between $N(v)$ and $V - N(v)$. We shall show that we can partition the vertex set of G into sets $S_1, ..., S_l$ and a set L (for Light vertices) such that:

- every dense vertex is in some S_i,
- each dense set contains at most $\Delta + 4e^4\epsilon\Delta^{\frac{3}{2}}$ vertices.
- there are at most $12e^4\epsilon\Delta\frac{3}{2}$ edges between S_i and $V - S_i$,s and
- for each S_i, every vertex in S_i has at least $\frac{3}{4}\Delta$ neighbours in S_i and every vertex outside S_i has fewer than $\frac{3}{4}\Delta$ neighbours in S_i.

Now, the light vertices are easy to colour using $\Delta - \epsilon\sqrt{\Delta}$ colours, as discussed above. Furthermore, each S_i can essentially be dealt with individually, as the conditions above ensure that it is isolated from the rest of the graph. Thus, we can produce an optimal colouring of S_i. Combining these colourings, we obtain a $max(\Delta + \epsilon\sqrt{\Delta}, max\{\chi(S_i)|1 \leq i \leq l\})$. colouring of G. Thus, this colouring is a $max(\Delta + \epsilon\sqrt{\Delta}, \chi(G))$ colouring of G. We remark that if Δ is bounded, then we can efficently find the colourings of each S_i because they are of bounded size. As we shall see, if $\chi(S_i)$ is within $b \leq \epsilon\sqrt{\Delta}$ of Δ then to find an optimal colouring of each S_i, we need only find an optimal colouring for some set X_i of at most $50b$ vertices of S_i. This will permit us to find an optimal colourings provided $b \leq \epsilon log\Delta$.

Of course, we cannot colour each S_i completely independently of the rest of the graph. Consider, for example, the situation in which the dense set S_j with maximum chromatic number is a $\Delta - 2$ clique each of whose elements has 3 neighbours outside S_j. If all these neighbours receive colour 1 then we will have to use $\Delta - 2$ new colours in colouring S_j, and thereby obtain a colouring using one more colour than desired. To avoid such a problem, we need to consider colouring all the S_i at once at the same time as we perform the partial colouring of the light vertices. Essentially, we fix some optimal colouring of each S_i and then randomly permute the names of the colour classes (the permutation for distinct S_i are chosen independently). As with the earlier random colouring we have to uncolour vertices involved in conflicts. However, because there are only a few edges out of each dense set, there will be very few uncoloured vertices in each dense set. Carefully choosing the original optimal colourings of the S_i, and the order in which we colour the uncoloured vertices, will allow us to complete the colouring.

In order to do so, we need to specify how we construct the dense sets, examine their structure, and discuss how we choose the original colour classes for each S_i. To begin, we discuss constructing the dense sets. We present a slightly more general result then that needed here, with a view to future applications (see [13] for a precursor of this result).

3 A Structural Decomposition

Consider any integer p. For any graph G of maximum degree Δ at least $100p$, we say that a vertex v of G is $p-sparse$ if the subgraph induced by its neighbourhood contains fewer than $\binom{\Delta}{2} - p\Delta$ vertices (note v can have a clique as a neighbourhood and still be p-sparse, provided this clique has fewer than $\Delta - p$ vertices). Otherwise, v is $p - dense$. Note that if v is p-dense then there are at most $2p\Delta$ edges between $N(v)$ and $V - N(v)$. We now define a $p - dense$ *decomposition* of G.

For each p-dense vertex v of G, we define a set S_v as follows.

1. Step 0. Set $S_v = N(v) + v$.
2. Step 1. While there is a vertex y in S_v with $|N(y) \cap S_v| < \frac{3\Delta}{4}$ delete some such vertex from S_v.
3. Step 2. While there is a vertex y outside of S_v with $|N(y) \cap S_v| \geq \frac{3\Delta}{4}$ add some such vertex to S_v.

It is easy to see that the set S_v is uniquely defined.

We note that:

5 *every vertex in S_v has at least $\frac{3\Delta}{4}$ neighbours in S_v.*

6 *every vertex outside of S_v has less than $\frac{3\Delta}{4}$ neighbours in S_v.*

7 *The neighbourhood of any p-dense vertex has at least $\Delta - 2p$ elements.*

8 *Since $|E(N(v))| \geq \binom{\Delta}{2} - p\Delta$, fewer than $10p$ vertices of $N(v)$ are deleted from S_v in Step 1, and $|E(S_v, G - S_v)| \leq 12p\Delta$ (Note that in Step 2, the number of edges out of S_v always decreases).*

9 *Since $|E(N(v), G - N(v))| \leq 2p\Delta$, fewer than $4p$ vertices of $G - N(v)$ are added to S_v in Step 2.*

Thus, we have:

10 $\Delta - 12p \leq |S(v)| \leq \Delta + 4p$.

Note that (8), (7) and (6) together imply that $v \in S_v$. Because S_v is so close to N_v, we can prove:

11 *If x and y are special vertices and S_x intersects S_y then y is in S_x and x is in S_y.*

By (11), for some integer l, we can construct a sequence of special vertices $x_1, ..., x_l$ and corresponding disjoint sets $S_1, ..., S_l$ with $S_i = S_{x_i}$ such that every special vertex is in $\cup_1^l S_i$. We shall call $S_1, ..., S_l$, *p-dense sets*. We call $S_1, ..., S_l, L = V - \cup_1^l S_i$, a *$p$-dense decomposition* Recapping, we have:

12 *If $S_1, ..., S_l$ is a p-dense decomposition then:*

- *every dense set has between $\Delta - 12p$ and $\Delta + 4p$ vertices*
- *there are at most $12p\Delta$ edges between S_i and $V - S_i$*
- *a vertex is adjacent to $\frac{3\Delta}{4}$ vertices of S_i if and only if it is in S_i*
- *every vertex in L is p-sparse*

Remark 13 *There are many variants of this structural decomposition. In particular, one could choose the dense set S_i to be S_v for the p-dense vertex v in $V - \cup_1^{i-1} S_i$ with the most edges in its neighbourhood. Furthermore, for small p the constants 4 and 12 can be improved.*

4 Partioning the dense sets

We turn now to the proof of Theorem 3. An analysis of this proof provides the algorithms mentioned in Theorems 1 and 2. We shall show that there is some Δ_0 such that Theorem 3 holds for $\Delta \geq \Delta_0$ setting $\epsilon = \frac{1}{10^6}$. The general result follows by setting $\epsilon = max(\frac{1}{\Delta_0+1}, \frac{1}{10^6})$. We shall not compute Δ_0 precisely, we only insist that it satisfies certain inequalities scattered throughout the paper. So, now let G be a graph of maximum degree $\Delta \geq D_0$, let $\epsilon = \frac{1}{10^6}$, and let $p = \epsilon e^6 \sqrt{\Delta}$. Let $S_1, ..., S_l, L$ be a p-dense decomposition of G.

We say a dense set is *matchable* if their is a matching with $6p$ edges in its complement. We call a dense set which is not matchable a *near clique*. We note that:

14 *a near clique S_i has chromatic number at least $\Delta - 30p$. Furthermore, it has an optimal colouring in which the singleton colour classes form a clique with at least $\Delta - 36p$ elements, all of whose vertices see all of every colour class which has more than two elements.*

Proof Since S_i has at least $\Delta - 12p$ elements, in any $\Delta - 30p$ colouring of S_i, there must be at least $18p$ vertices in non-singleton colour classes. But each non singleton colour class U contains a matching in the complement of S_i with $\lfloor \frac{|U|}{2} \rfloor$ edges. The first statement follows. Since S_i is a near-clique, at most $6p$ of the colour classes in an optimal colouring are non-singleton. Furthermore, in an optimal colouring of any graph, the singleton colour classes form a clique. By choosing an optimal colouring with the maximal number of two element colour classes, we can ensure that there is no edge of $\overline{S_i}$ joining a singleton colour class and a colour class with more than three elements. Thus, the second statement also holds. □

Now, for each matchable S_i, we let $k_i = |S_i| - 6p < \Delta - \epsilon\sqrt{\Delta}$ and we take a k_i colouring of S_i: $U_1^i, ..., U_{k_i}^i$ consisting of $6p$ colour classes with 2 elements and $|S_i| - 6p$ singleton colour classes. For each near clique S_i, we let k_i be the chromatic number of S_i, and pick a k_i-colouring of S_i: $\{U_1^i, ..., U_{k_i}^i\}$ with the properties specified in (14). We say a colour class is *compound* if it is not a singleton. We may refer to a vertex forming a singleton colour class as a *singleton vertex*. We let $k = max(\lfloor \Delta - \epsilon\sqrt{\Delta} \rfloor, max\{k_i | 1 \leq i \leq l\})$.

The *core* of S_i is the set of vertices of S_i with fewer than p neighbours outside of S_i. A vertex in the core of a dense set S_i is a *core vertex*. Every vertex in a dense set which is not a core vertex is a *peripheral vertex*.

We let E_j^i be the union of the external neighbourhoods of the vertices of U_j^i. We let F_j^i be the union of the external neighbourhoods of the core vertices elements of U_j^i.

We can prove the following easily.

15 *For a colouring of a near-clique as in (14), Every E_j^i has at most $\frac{\Delta}{2}$ elements and every F_j^i has at most $\frac{\Delta}{100}$ elements.*

We say that a vertex is *highly adjacent* to S_i if it is a vertex with at least $\frac{\Delta}{log^3\Delta}$ neighbours in S_i or if it is a core vertex in some compound colour class U_j^r such that F_j^r has more than $\frac{\Delta}{log^3\Delta}$ elements in S_i. We let HA_i be the vertices highly adjacent to S_i.

5 Why Is The Boundary at The Square Root of Delta?

We pause to note that it is this last property which explains why $k = \Delta - \sqrt{\Delta} + 1$ is (approximately) the boundary at which k colouring becomes difficult.

To see this consider $\Delta = r^2$, and consider a dense set S which consists of $\Delta - r$ vertices inducing a clique, and a stable set U with $r + 1$ elements all joined to every vertex of the clique. Then, in any $k = \Delta - r + 1$ colouring of G, all the vertices of U must receive the same colour. Note that each vertex in U may have up to r neighbours outside S. Thus, U can be thought of as a pseudo vertex whose degree may be $r^2 + r = \Delta + r$. So, we can determine if a graph H with maximum degree $\Delta + r$ has a k colouring by replacing each vertex of G by a dense set $S(v)$ and corresponding $U(v)$. Repeated applications of this procedure allow us to widen the gap between k and the maximum degree of the graph we are trying to colour even further. This permits us to reduce k colouring to $\Delta - \sqrt{\Delta} + 1$ colouring.

6 The Random Partial Colouring

We shall consider the random colouring process in which we:

- assign each vertex in L a random uniformily chosen colour between 1 and k, with all these choices independent
- choose, for each dense set S_i, a permutation of 1 to k, where these choices are all independent of each other and the colours assigned to the light vertices, then for $1 \leq j \leq k_i$, we give U_j^i the jth colour in the permutation for S_i.
- uncolour any vertex which is assigned the same colour as one of its neighbours.
- uncolour any vertex in HA_i which is coloured with a colour that no longer appears on any singleton vertex of S_i which has less than $\frac{\Delta}{log^2\Delta}$ external neighbours (either because it was not assigned to any such vertex or because the vertex it was assigned to was uncoloured in the last step; this is to avoid having a colour which appears on an external neighbour of every vertex of S_i.).
- uncolour any vertex in a compound colour class in which we have uncoloured a core vertex.

By analyzing this procedure, we can show that:

16 *There is a partial k colouring of G such that:*

1. *(P1) For every U_j^i, either all its vertices receive the same colour, none of its vertices are coloured, or all of its core vertices receive the same colour and all of its peripheral vertices are coloured with this colour or uncoloured.*
2. *(P2) For every vertex v in L, at least $1,5\epsilon\sqrt{\Delta}$ colours appear on exactly two neighbours of v.*
3. *(P3) For every peripheral vertex v, at least $1,5\epsilon\sqrt{\Delta}$ colours appear on exactly two neighbours of v.*
4. *(P4) At most $log\ \Delta$ compound colour classes of any S_i contain an uncoloured core vertex.*
5. *(P5) At most $13p$ vertices of any S_i which are singleton colour classes in the partition of S_i are uncoloured.*
6. *(P6) There are at least $\frac{2\Delta}{5}$ vertices v of each S_i satisfying: every colour appearing in the external neighbourhood of v also appears in S_i.*

We show now that if there is only one dense set S_1, these properties ensure that we can extend our partial k-colouring to a colouring of G. Properties (P2) and (P3) indicate that we need really only worry about the core vertices.

To begin we complete the colouring of the core vertices of the compound colour classes of S_1. We shall choose a colour for each of the at most $\log\Delta$ compound colour classes which contain an uncoloured core vertex and assign it to all of these vertices. When choosing a colour for such a U_j^1, we forbid all the colours which appear on E_j^1, and all colours which appear on a compound colour class of S_1. We may however choose a colour which appears on some singleton vertex in S_1 and if we do so, we uncolour this vertex. We do not permit ourselves to colour the uncoloured core of some compound U_j^i with a colour appearing on the external neighbourhood of more than $50p$ vertices of S_1. Now, we know that every E_j^i has at most $\frac{\Delta}{100}$ elements. Also, there are fewer than $18p$ vertices in the union of the compound colour classes, so we forbid at most $18p$ colours because they appear in some compound colour class. Furthermore, since there are at most $12p\Delta$ edges out of S_1, we forbid fewer than $\frac{\Delta}{4}$ colours because they appear in the external neighbourhood of too many vertices of S_1. It follows that there are at least $\frac{\Delta}{3}$ possible colours to assign to the core of any uncoloured compound colour class. Thus, we can choose a distinct colour satisfying our conditions for each of the at most $\log\Delta$ compound colour classes with uncoloured cores. We note that after colouring the uncoloured cores of compound colour classes with these colours, (P1) still holds. Furthermore, (P2) and (P3) hold if we replace the $1.5\epsilon\sqrt{\Delta}$ by $1.5\epsilon\sqrt{\Delta} - log\Delta$. Finally, (P4) holds if we replace $\frac{2\Delta}{5}$ by $\frac{9\Delta}{25}$.

We now colour the core vertices not in compound colour classes. If S_1 is matchable then this is easy. Every core vertex sees at most p vertices outside of S_1. If it has fewer than $\Delta - \frac{p}{2} < k$ neighbours then obviously we can choose a colour for it which appears on none of its neighbours and hence with which it can be coloured. Otherwise, it must see at least $\Delta - \frac{3p}{2}$ of the at most $\Delta + 4p$ vertices of S_1 and hence sees both vertices of at least $\frac{p}{2}$ of the $6p$ compound colour classes. Thus, there is a colour available with which to colour it.

If S_i is a near clique, things are more difficult. We will chose a set of $3p$ vertices satisfying the conditions of (P6) and uncolour them. We want to do so in such a way that their are still at least $\epsilon\sqrt{\Delta} + 1$ repeated colours in the neighbourhood of each peripheral or light vertex and so that at a set R of least p of the $3p$ vertices continue to satisfy the condition of (P6), (i.e. non of their external neighbours were coloured with a colour used on one of the $3p$ vertices we have just uncoloured). To help to ensure that the second condition holds, we only pick from the at least $\frac{4\Delta}{25}$ of our $\frac{9\Delta}{25}$ possibilities which are coloured with a colour appearing in the external neighbourhoods of fewer than $60p$ vertices of S_1 (here we again use the fact that there are at most $12p\Delta$ edges out if the dense set). If we make such a restricted choice then it is easy to check that with high probability, the two conditions we want to satisfy will hold (we will not do this simple computation; we note only that each vertex is uncoloured with probability less than $\frac{1}{\sqrt{\Delta}}$ so we can see that we do not expect this step to have huge effects). Now, we colour every core singleton vertex not in R, which we can do because such a vertex has p uncoloured neighbours in R. Next, we colour all the vertices of R, which we can do because, as (P1) holds, if there is an uncoloured vertex v left in R then there are fewer than $k_1 - 1 \leq k - 1$ colours appearing in S_1. Hence, because every colour appearing in the neighbourhood of v also appears in S_1, we know that there is a colour available with which to colour v.

Having coloured the core vertices, we can greedily finish off the peripheral and light vertices because for each such vertex there are at least $\epsilon\sqrt{\Delta}+1$ colurs which are repeated in its neighbourhood.

Now, if there is more than one dense set, things get a bit more complicated. We can't colour them sequentially because completing one dense set will affect what happens in the others. So, we will colour all the uncoloured cores of compound colour classes at the same time, using a random procedure. We then choose a set R_i of each near-clique S_i which will form the last p core vertices of S_i coloured. Again, we need to choose these R_i randomly and all at once. Having performed these two steps, we can then finish off the colouring in a simply greedy fashion.

The details get a little hairier, and the technicalities are quite lengthy. We save them for a longer version of this paper. We hope that this short account has enabled the reader to grasp how our structural decomposition can be used in such proofs and why the complexity boundary is at $\Delta - \sqrt{\Delta}$.

References

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

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

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

[4] T. Emden-Weinert, S. Hougardy, and B. Kreuter, *Uniquely colourable graphs and the hardness of colouring graphs with large girth*, Probability, Combinatorics, and Computing, to appear.

[5] 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.

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

[7] F. Maffray and M. Preissmann, *On the NP-completeness of the k-colourability problem for triangle-free graphs*, Discrete Mathematics **162** (1996), 313-317.

[8] M. Molloy and B. Reed, *A bound on the strong chromatic index of a graph*, Journal of Combinatorial Theory (B), to appear.

[9] M. Molloy and B. Reed, *A bound on the total chromatic number*, submitted.

[10] M. Molloy and B. Reed, *Assymptotically better list colourings*, manuscript.

[11] M. Molloy and B. Reed, *An algorithmic version of the Lovasz Local Lemma*, in preparation.

[12] M. Molloy and B. Reed, *Graph Colouring via the Probabilistic Method* , to appear in a book edited by A. Gyarfas and L. Lovasz.

[13] B. Reed, $\chi, \Delta,$ *and* ω, Journal of Graph Theory, to appear.

[14] B. Reed, *A strengthening of Brook's Theorem*, in preparation.

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