Near-Optimal List Colourings

Michael Molloy*
Department of Computer Science
University of Toronto
Toronto, Canada

Bruce Reed*
CNRS, Paris, France
IME-USP, Sao Paulo, Brasil

July 31, 2000

Abstract We show that a simple variant of a naive colouring procedure can be used to list colour the edges of a k-uniform linear hypergraph of maximum degree Δ provided every vertex has a list of at least $\Delta + c(\log\Delta)^4\Delta^{1-\frac{1}{k}}$ available colours (where c is a constant which depends on k). We can extend this to colour hypergraphs of maximum codegree $o(\Delta)$ with $\Delta + o(\Delta)$ colours. This improves earlier results of Kahn and our techniques are quite similar. We also develop efficient algorithms to obtain such colourings when Δ is constant.

1 A Colouring Problem and a Colouring Procedure

A *hypergraph* H consists of a set $V(H)$ (or simply V) of vertices and a set $E(H)$ (or simply E) of edges, each of which is a subset of $V(H)$. The *chromatic index* of a hypergraph is the minimum number of colours needed to colour its edges so that no two edges which intersect receive the same colour. Given a list of colours for each edge of H, we say that an edge colouring is *acceptable* if every edge is coloured with a colour on its list. The *list chromatic index* of a hypergraph H is the minimum r which satisfies: if every list has at least r members then there is an acceptable colouring. Obviously, both of these numbers can be no smaller than the maximum degree of H, denoted $\Delta(H)$.

A hypergraph is k-*uniform* if every edge contains k vertices. It is k-*bounded* if every edge contains at most k vertices. We shall discuss list colouring the edges of k-bounded hypergraphs. For ease of exposition we consider only k-uniform

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

hypergraphs, the results extend easily (we simply add dummy vertices to each edge).

As we describe below, the technique used in this paper is one that has been used many times in various graph colouring problems. In order to see more clearly how our current problem fits into the general set of problems amenable to this approach, we will reformulate it as a vertex-colouring problem. To this end, we define the *chromatic number* of a graph G, denoted $\chi(G)$ to be the minimum number of colours needed to colour its vertices so that no two adjacent vertices receive the same colour. Given a list of colours for each vertex of G, we say that a vertex colouring is *acceptable* if every vertex is coloured with a colour on its list. The *list chromatic number* of G, denoted $\chi_\ell(G)$ is the minimum r which satisfies: if every list has at least r members then there is an acceptable colouring. Obviously, both of these numbers can be no smaller than the size of the largest clique (set of pairwise adjacent vertices) of G.

In order to cast hypergraph edge-colouring problems in terms of colouring the vertices of a graph we need to consider line graphs of hypergraphs. The line graph $L(H)$ of a hypergraph H has a vertex corresponding to each edge of H. Two vertices of $L(H)$ are adjacent if and only if the corresponding two edges of H share a vertex. We want to calculate the list chromatic number of line graphs of certain k-uniform hypergraphs. Note first that the maximum degree of a line graph of a k-uniform hypergraph of maximum degree Δ may be as large as $k(\Delta-1)$. Furthermore, the chromatic index of H may be as large as $k(\Delta-1)$. To see this, consider $k=p+1$ for some prime p and let H be a hypergraph whose edges and vertices correspond to the lines and points of a projective plane generated by Z_p. That is we have a set of k^2-k+1 edges, every pair of which meet in exactly one vertex. Now, we let $\Delta = kd$ for some integer d and take d copies of each edge of H to obtain a new hypergraph H' whose line graph is a clique.

Note that in this example, every pair of vertices is contained in many edges. The *codegree* of a pair $\{x,y\}$ of vertices in a hypergraph H, denoted $d_H(x,y)$ or simply $d(x,y)$, is the number of edges that contain both the vertices.

We will study the list chromatic index of a k-uniform hypergraph in which the maximum codegree is $o(\Delta)$. To begin, we consider linear hypergraphs, that is, hypergraphs such that no two edges intersect in more than one vertex. Kahn[13] proved that if a k-uniform linear hypergraph has maximum degree Δ then it has list chromatic index at most $\Delta + o(\Delta)$ (see also [18] for an earlier proof of the same bound on the chromatic index). We show:

Theorem 1 *For all k, there is a constant c_k (depending on k) such that any k-uniform linear hypergraph has list chromatic index at most $\Delta + c_k\Delta^{1-\frac{1}{k}}(\log\Delta)^4$.*

We remark that this bound on the list chromatic index of a k-uniform linear hypergraph is also the best known asymptotic bound on the chromatic index of such a hypergraph.

This result also improves a result of Häggkvist and Janssen[11] who proved that when $k = 2$, i.e. for graphs, the list chromatic index is bounded by $\Delta + O(\Delta^{\frac{2}{3}}\text{poly}(\log\Delta))$. It has been conjectured by many people (probably first by Vizing in 1975, see [12]) that the list chromatic index of a graph H equals its chromatic index and hence is at most $\Delta + 1$.

We then turn to k-uniform hypergraphs with larger bounds on the maximum codegree. We show:

Theorem 2 *For all k, there is a constant c_k (depending on k) such that any k-uniform hypergraph of maximum codegree B has list chromatic index at most $\Delta + c_k(B^{\frac{1}{k}}\Delta^{1-\frac{1}{k}}(\log\frac{\Delta}{B})^4)$.*

Again, this improves the main result of Kahn[13].

Remark: In Theorems 1 and 2, by being a little more careful these techniques can certainly reduce the exponent in the poly(log Δ) terms, and in fact, if $k \geq 3$ or if B is moderately large, say at least $\log^{10}\Delta$, the poly(log Δ) terms can probably be eliminated. We have made no attempt to do so.

In this paper we use the naive colouring procedure which consists of choosing a colour for each vertex, uncolouring those involved in a conflict, and recursively repeating this process on the uncoloured vertices. Suprisingly, analyzing this procedure using simple but powerful probabilistic tools allows us to find optimal or at least asymptotically optimal bounds on the chromatic number and/or list chromatic number of various classes of graphs. This procedure falls under the wider class of techniques commonly referred to as the "semi-random method" or the "Rödl nibble". [1] and [20] are usually cited as the first of many uses of this technique, and in fact, the first application of this particular procedure by Kahn[13] has a lineage traceable directly to [20]. See [17] for a survey of applications of the technique.

A typical application of the semi-random method consists of two parts: first, one must decide upon an appropriate colouring procedure which will form the basis of the proof; second, one must analyze that procedure, the bulk of this analysis usually consisting of expected value computations and proofs of concentration. Our new contributions are not in the first part - the naive colouring procedure that we use is exactly that suggested by Kahn[13] (with one change to accomodate the colouring of the final edges - the use of Save_v in section 5). It is in the second part that this paper differs from [13] - by using a different approach to the analysis, we are able to get much tighter control on the procedure which allows us to replace Kahn's $o(\Delta)$ term by $\sqrt{\Delta}\text{poly}(\log\Delta)$.

We believe the results in this paper to be of interest for a number of reasons. The first is that determining the list chromatic index of graphs and hypergraphs is an important open problem and any improvement on the upper bound is of interest. The second is that this paper, which provides a slightly different analysis of Kahn's original colouring method, shows more clearly how his procedure fits into the general schema given in [17]. Furthermore, our treatment simplifies some expected value computations (via inclusion-exclusion), and significantly simplifies the concentration details, by appealing to Talagrand's Inequality and thus avoiding the introduction of new martingale inequalities. We hope that this may make it more accessible than Kahn's treatment to some readers. Furthermore, the improved bounds on the chromatic index of linear hypergraphs are of interest in other fields, see e.g. [14] where our results have been applied to bound the maximum possible size of a collection of r-uniform sets which do not contain a Δ-subsystem (see [14] for definitions). Finally, as discussed in Section 7, we can use recent results on algorithmic versions of the Local Lemma to obtain efficent algorithms for finding acceptable colourings for sets of lists which satisfy our bounds in the case where Δ is constant.

We shall need the following two standard tools:

The Lovász Local Lemma[9]: *Let $\mathcal{E}$ be a set of events in a probability space. Suppose that each $A \in \mathcal{E}$ is mutually independent of all but at most d other events and occurs with probability at most p. If $ep(d+1) < 1$ then with positive probability, none of the events in $\mathcal{E}$ occur.*

A Simple Observation[8]: *The vertices of any graph G of maximum degree d can be coloured using $d+1$ colours using a greedy procedure which takes $O(d|V(G)|)$ time. We simply consider the vertices in any order, colouring each vertex with a colour that does not appear on its coloured neighbours. The same approach allows to find an acceptable colouring in linear time given a list of at least $d+1$ acceptable colours for each vertex.*

We will make use of the following lemma from [19]:

A List Colouring Lemma *Given a graph G and a list L_v of acceptable colours for each vertex in G let Poss_i be the set of vertices on which i is acceptable. Suppose:*

$$min\{|L_v| \ s.t. \ v \in V(G)\} > 2\mathrm{e} \times \max\{|N(v) \cap \mathrm{Poss}_i| \ s.t. \ v \in \mathrm{Poss}_i\}$$

Then there is an acceptable colouring for G and the given lists.

The proof of the List Colouring Lemma is a simple application of the Local Lemma. For fixed Δ, the proof can be made constructive, thus yielding an efficient algorithm for finding the acceptable colouring, by a somewhat complicated application of the techniques in [16].

Throughout the paper, where there is no possibility of confusion, we omit the $\lceil\rceil$ and $\lfloor\rfloor$ signs. All logarithms are natural logarithms.

2 The Chromatic Index of Linear Hypergraphs

To begin, we consider the chromatic number of the line graphs of linear hypergraphs. We will then generalize to list colouring such linegraphs.

So, consider a graph G which is the line graph of a k-uniform Δ-regular linear hypergraph. This implies that there is a set $\mathcal{C}$ of cliques of G each containing Δ vertices such that every vertex is in exactly k elements of $\mathcal{C}$ and each pair of vertices is contained in at most one element of $\mathcal{C}$. We denote by $\mathcal{C}_v$ the subset of $\mathcal{C}$ containing v.

As remarked earlier, we are going to colour the vertices of G using a naive colouring procedure whereby we repeatedly *assign* to each uncoloured vertex a random colour which appears on none of its coloured neighbours and then colour a vertex with the colour it is assigned if this colour is not assigned to any of its (uncoloured) neighbours. If a vertex is successfully coloured during an iteration, then we say that the vertex *retains* its colour. For each vertex v, at any time in the algorithm we will have a list L_v of colours which do not appear on any coloured neighbour of v and hence can be assigned to v. Similarily for each clique $C \in \mathcal{C}$, we will have a list L_C of colours which appear on no vertex of C. Note that L_v is simply the intersection of $\{L_C | C \in \mathcal{C}_v\}$. Finally, for each $C \in \mathcal{C}$ and colour $i \in L_C$, we have a subset $S_{C,i}$ of $V(C)$ consisting of those vertices of C for which i is available, i.e. $S_{C,i} = \{v | v \in C, i \in L_v\}$. Of course, whether v is in $S_{C,i}$ or not depends on the colours of the vertices in the other cliques of $\mathcal{C}_v$. For any set of vertices T, at any point during our procedure, we define U_T (Uncoloured$_T$) to be the vertices of T which are not coloured.

To begin, we attempt to colour G with Δ colours. Thus, initially $|L_v| = \Delta$ for all $v \in V(G)$, $|L_C| = \Delta$ for all $C \in \mathcal{C}$, and $|S_{C,i}| = \Delta$ for all $C \in \mathcal{C}$ and each colour i.

Consider what happens during the first iteration of the naive colouring procedure in which each uncoloured vertex v chooses a colour uniformly from those on $L_v = \{1, ..., \Delta\}$ with these choices mutually independent. Each vertex v has $k\Delta - k$ neighbours each of which is assigned the same colour as v with probability $\frac{1}{\Delta}$. Thus the probability that v is coloured is $(1 - \frac{1}{\Delta})^{k\Delta - k}$ which for Δ large is close to e^{-k}. So, for each $C \in \mathcal{C}$, we have that $\mathbf{Exp}(|L_C|)$ is about $(1 - \frac{1}{e^k})\Delta$.

Now, if the choices of the new L_C were all independent, then for each v which remains uncoloured we would have $\mathbf{Exp}(|L_v|)$ is about $(1 - \frac{1}{e^k})^k\Delta$. Similarily, for a colour i not used in C the probability that an uncoloured vertex in $S_{C,i}$ does not have a neighbour coloured i in any of the other $k-1$ cliques containing it would be about $(1 - \frac{1}{e^k})^{k-1}$ and since we expect a vertex to be uncoloured with probability about $(1 - \frac{1}{e^k})$, we have that for each C and i remaining in L_C the expected number of vertices of $S_{C,i}$ is about $(1 - \frac{1}{e^k})^k\Delta$.

Of course, the choices made on the various cliques are not independent - the cliques intersect. However, the linearity of the hypergraph can be used to show that, for the random variables under consideration, the expected values computed above are never far wrong. We can then show (using a concentration inequality, introduced below), that the probability that a particular such variable is far from its expected value is quite small. Finally, an application of the Local Lemma allows us to prove the existence of an assignment of colours where the size of every L_v,L_C, and $S_{C,i}$ is close to its expected value.

At this point we have:

1 *Each* L_C *has about* $(1-\frac{1}{e^k})\Delta$ *elements,*

2 *Each* L_v *has about* $(1-\frac{1}{e^k})^k\Delta$ *elements,*

3 *Each* $S_{C,i}$ *has about* $(1-\frac{1}{e^k})^k\Delta$ *elements.*

Of course $|L_C|$ is simply the number of uncoloured vertices in C, so we have essentially reduced our problem to a list colouring problem on a subhypergraph whose degree is about $(1-\frac{1}{e^k})\Delta$. Unfortunately, the lists are no longer identical. Nevertheless conditions (2) and (3) are enough to allow us to iterate our analysis, obtaining a sequence of iterations so that after the jth iteration, we have:

4 *Each* L_C *has about* $(1-\frac{1}{e^k})^j\Delta$ *elements,*

5 *Each* L_v *has about* $(1-\frac{1}{e^k})^{kj}\Delta$ *elements,*

6 *Each* $S_{C,i}$ *has about* $(1-\frac{1}{e^k})^{kj}\Delta$ *elements.*

So, for example, the sizes of each L_v and each $S_{C,i}$ should remain roughly the same.

Unfortunately, as the number of iterations increases, "about" becomes a looser and looser term and we eventually lose control of the whole process. However, with care, we can continue until the L_v are almost empty, at which point by (**5**) $(1-\mathrm{e}^k)^j$ is approximately $\Delta^{-\frac{1}{k}}$ and so by (**4**) each L_C has around $\Delta^{1-\frac{1}{k}}$ elements. In fact, we stop when the maximum order of an L_C is $O((\log\Delta)^4\Delta^{1-\frac{1}{k}})$. Since the degree of any vertex in the uncoloured graph is at most k times this maximum, we can complete the colouring using $O(k(\log\Delta)^4\Delta^{1-\frac{1}{k}})$ new colours (by our Simple Observation). Forthwith the details.

2.1 A Detailed Description

Here we formalize the argument described above. The main difficulty lies in retaining control of the various parameters as the number of iterations increases.

We define $U_1 = S_1 = \Delta, L_1 = \Delta + \Delta^{1-\frac{1}{k}}\log^4\Delta$ and recursively define

- $U_{i+1} = U_i((1-\frac{1}{L_i})^{kS_i-k}) + \sqrt{U_i}\log^2\Delta$;
- $L_{i+1} = L_i \times (1-\frac{S_i}{L_i}\times(1-\frac{1}{L_i})^{kS_i-k})^k - \sqrt{L_i}\log^2\Delta$;
- $S_{i+1} = S_i\times(1-\frac{S_i}{L_i}\times(1-\frac{1}{L_i})^{kS_i-k})^{k-1}\times(1-(1-\frac{1}{L_i})^{kS_i-k}) + \sqrt{S_i}\log^2\Delta$.

Throughout the execution of our colouring procedure, we will enforce that at the beginning of each iteration i:

- $|L_v| \geq L_i$ for every uncoloured vertex v
- $|U_C| \leq U_i$ for every clique $C \in \mathcal{C}$
- $|S_{C,\gamma}| \leq S_i$ for every clique $C \in \mathcal{C}$ and colour $\gamma \in L_C$.

In order to retain control of the process, it will be very helpful if all the lists L_v are of the same size. Thus, at the beginning of iteration i, we delete $|L_v| - L_i$ colours from each L_v.

It will also be helpful if all the sets $S_{C,\gamma}$ are of the same size. To ensure this, we will introduce some *artificial vertices* to the graph. A vertex which is not artificial is a *real vertex*. Each artificial vertex v has an *associated colour* $\gamma(v)$ which is the only colour that it can receive. An artifical vertex v lies in exactly one clique $C(v)$ of $\mathcal{C}$, and is considered to belong to $S_{C(v),\gamma(v)}$. During iteration i of our procedure, each artificial vertex v receives $\gamma(v)$ with probability $\frac{1}{L_i}$. If v receives $\gamma(v)$ and if no other vertex in $C(v)$ receives $\gamma(v)$, then v retains $\gamma(v)$ with probability $(1-\frac{1}{L_i})^{(k-1)(S_i-1)}$ (see the remark following Lemma 3). If an artificial vertex v receives $\gamma(v)$ then every other vertex in $C(v)$ which receives $\gamma(v)$ is uncoloured, and if v retains $\gamma(v)$ then $\gamma(v)$ is removed from $L_{C(v)}$ and thus is removed from L_u for each vertex $u \in C(v)$. At the beginning of iteration i, for each clique $C \in \mathcal{C}$ and colour $\gamma \in L_C$, we add $S_i - |S_{C,\gamma}|$ artificial vertices to C, each with associated colour γ. This simulates the property that $|S_{C,\gamma}| = S_i$. At the end of each iteration, we delete all artificial vertices, and so in particular, at the beginning of each iteration each set $S_{C,\gamma}$ contains only real vertices.

It turns out that it is not important to simulate that each set U_C has size exactly U_i.

In summary, each iteration i proceeds as follows:

1. For each vertex v, remove $|L_v| - L_{i+1}$ colours from L_v (these colours can be chosen arbitrarily).
2. For each clique $C \in \mathcal{C}$ and colour $\gamma \in L_C$, add $S_i - |S_{C,\gamma}|$ artificial vertices to C, each with associated colour γ.
3. (a) For each uncoloured real vertex v, assign to v a colour chosen uniformly at random from L_v.
   (b) For each artificial vertex v, assign to v its associated colour with probability $\frac{1}{L_i}$.

   If a vertex v is assigned a colour γ then we say v *receives* γ.
4. (a) Uncolour every (real or artificial) vertex which receives the same colour as a (real or artificial) neighbour.
   (b) If an artificial vertex v received its associated colour and did not lose that colour in step 4(a), then uncolour v with probability $1-(1-\frac{1}{L_i})^{(k-1)(S_i-1)}$.

   If a vertex is not uncoloured, then we say that it *retains* its colour.

5. For each (real or artificial) vertex v which retained a colour γ, and each clique $C \in \mathcal{C}$ with $v \in C$, remove γ from L_C, and from L_u for any real vertex $u \in C$.

6. Delete all artificial vertices from the graph.

7. If one of the following conditions fails to hold, then HALT.

   - $|L_v| \geq L_{i+1}$ for every uncoloured vertex v
   - $|U_C| \leq U_{i+1}$ for every clique $C \in \mathcal{C}$
   - $|S_{C,\gamma}| \leq S_{i+1}$ for every clique $C \in \mathcal{C}$ and colour $\gamma \in L_C$.

We consider each iteration of our procedure to be a separate random experiment. In what follows, we assume that are about to carry out iteration i, having already carried out all previous iterations without halting. All probabilities are in terms of the random choices made during iteration i.

The first step of our analysis is to determine the probability that a vertex retains its colour, and the probability that a list keeps a colour.

Lemma 3 *(a) For every uncoloured vertex v and colour $\gamma \in L_v$, the probability that v keeps γ, conditional on v receiving γ, is precisely $(1-\frac{1}{L_i})^{kS_i-k}$.*

(b) For every clique $C \in \mathcal{C}$ and colour $\gamma \in L_C$, the probability that γ is removed from L_C during this iteration is at most $\frac{S_i}{L_i}(1-\frac{1}{L_i})^{kS_i-k}$.

Proof (a) If v is a real vertex, then this is the probability that no neighbour of v, real or artificial, receives γ. There are exactly $kS_i - k$ such neighbours which can receive γ and each receives γ with probability exactly $\frac{1}{L_i}$. Furthermore, the colour assignments are all independent. Therefore the probability that no neighbour of v receives γ is $(1-\frac{1}{L_i})^{kS_i-k}$.

If v is an artificial vertex, then this is the probability that no other vertex in $C(v)$ receives γ, times the probability that v does not lose its colour in Step 4(b). This product is simply $(1-\frac{1}{L_i})^{kS_i-k}$.

(b) This is the probability that at least one vertex in C, real or artificial, retains γ. There are S_i vertices $u \in C$ which can retain γ, and each retains γ with probability precisely $\frac{1}{L_i} \times (1-\frac{1}{L_i})^{kS_i-k}$. Therefore the probability that at least one of these vertices retains γ is at most $\frac{S_i}{L_i}(1-\frac{1}{L_i})^{kS_i-k}$. □

Inspired by Lemma 3, we define the following for each iteration i:

- $\text{Retain}_i = (1-\frac{1}{L_i})^{kS_i-k}$
- $\text{Keep}_i = 1 - \frac{S_i}{L_i} \times \text{Retain}_i$

Retain_i is the probability that a vertex v retains a colour $\gamma \in L_v$, conditional on v receiving that colour, and Keep_i is a lower bound on the probability that a list L_C keeps a colour γ which is in L_C at the beginning of the iteration.

Remark This explains the choice of probability with which we uncolour an artificial vertex in Step 4(b) - the probability was chosen such that an artificial vertex which receives a colour retains that colour with probability Retain_i.

Our main goal here is to keep track of how $L_v, U_C, S_{C,\gamma}$ change from iteration to iteration. Consider any particular iteration i. Through a slight abuse of notation, we will use $L_v, S_{C,\gamma}$ and U_T for any $T \subseteq V(G)$, to refer to the contents of those sets at the beginning of the iteration and we will use $L_v^{'}, U_C^{'}, S_{C,\gamma}^{'}$ to refer to their contents at the end of the iteration.

Unfortunately, the size of $S_{C,\gamma}^{'}$ is difficult to control, mainly because its value can fluctuate greatly with the change of the colour assigned to a single vertex. For example, if one vertex in C receives and retains γ then $S_{C,\gamma}^{'} = \emptyset$, but by changing the colour assigned to that vertex, we might change $S_{C,\gamma}^{'}$ to a fairly large set. (The readers who are accustomed to concentration proofs will immediately see why this is a problem. The other readers might see this after reading section 4.) To overcome this problem, we will examine the size of a different, but very closely related set:

$n(C,\gamma)$ is the number of vertices $u \in C$ such that (1) u does not retain its colour during this iteration and (2) no neighbour of u *outside of* C retains γ. Note that if γ remains in L_C then $S_{C,\gamma} = n(C,\gamma)$, while if γ is removed from L_C then we are no longer interested in $S_{C,\gamma}$. This justifies our focussing on $n(C,\gamma)$.

Our first step in controlling the sizes of these sets at the end of the iteration is to compute the expected values of these sizes.

Lemma 4 *If $L_i \geq S_i \geq \frac{1}{2}L_i$ then for each real vertex v, clique $C \in \mathcal{C}$, subset $T \subseteq C$, and colour $\gamma \in L_C$:*

(a) $\mathbf{Exp}(|U_T^{'}|) = |U_T|(1 - \text{Retain}_i)$;

(b) $\mathbf{Exp}(|L_v^{'}|) \geq |L_v| \times \text{Keep}_i^k - \text{O}(1)$;

(c) $\mathbf{Exp}(n(C,\gamma)) \leq |S_{C,\gamma}| \times \text{Keep}_i^{k-1} \times (1 - \text{Retain}_i) + \text{O}(1)$.

We defer the proof of this lemma until the next section. Next, we must show that $|L_v^{'}|, |U_C^{'}|, n(C,\gamma)$ are strongly concentrated:

Lemma 5 *For every real vertex v, clique $C \in \mathcal{C}$, subset $T \subseteq C$ and colour $\gamma \in L_C$, such that $|L_v|, |U_T|, |S_{C,\gamma}| \geq \log^7 \Delta$, we have*

(a) $\mathbf{Pr}(||U_T^{'}| - \mathbf{Exp}(|U_T^{'}|)| > \mathbf{Exp}(|U_T^{'}|)^{1/2} \log^2 \Delta) < \Delta^{-\log \Delta}$;

(b) $\mathbf{Pr}(||L_v^{'}| - \mathbf{Exp}(|L_v^{'}|)| > \frac{1}{2}\mathbf{Exp}(|L_v^{'}|)^{1/2} \log^2 \Delta) < \Delta^{-\log \Delta}$;

(c) $\mathbf{Pr}(|n(C,\gamma)) - \mathbf{Exp}(n(C,\gamma))| > \frac{1}{2}\mathbf{Exp}(n(C,\gamma)^{'})^{1/2} \log^2 \Delta) < \Delta^{-\log \Delta}$.

Again, we defer the proof of Lemma 5 until later. Upon proving this lemma, a straightforward application of the Local Lemma yields:

Lemma 6 *With positive probability, for every i such that $U_i, L_i, S_i \geq \log^7 \Delta$ and such that $L_j > S_j > \frac{1}{2} L_i$ for each $j \leq i$, we do not HALT during iteration i. In other words, for each real vertex v, clique $C \in \mathcal{C}$ and colour $\gamma \in L_C$, at the end of iteration i we have:*

(a) $|U_C| \leq U_{i+1}$;

(b) $|L_v| \geq L_{i+1}$;

(c) $|S_{C,\gamma}| \leq S_{i+1}$.

Proof Consider any iteration i, and assume that we have not yet halted.

For each vertex v, clique $C \in \mathcal{C}$ and colour $\gamma \in L_C$, let $A_v, B_C, C_{C,\gamma}$ denote the events that C violates (a), v violates (b), and C, γ violates (c), respectively. Observing that $\mathbf{Exp}(|U_C^{'}|) \leq |U_C|, \mathbf{Exp}(|L_v^{'}|) \leq |L_v|$ and $\mathbf{Exp}(n(C,\gamma)^{'}) \leq |S_{C,\gamma}|$, Lemmas 4 and 5 imply that the probability of any one of these events is at most $\Delta^{-\log \Delta}$. It is straightforward to show that each event E is mutually independent of all events which do not involve any vertices of distance at most 4 from the vertices involved in E, and so E is mutually independent of all but at most $\mathrm{O}(\Delta^5)$ other events. Since $\mathrm{e} \times \Delta^{-\log \Delta} \times \mathrm{O}(\Delta^5) < 1$ for large Δ, the lemma follows from the Local Lemma. $\square$

Our final step will be to show that we can continue until the size of each U_i is sufficiently small. In the next section we will prove the following:

Lemma 7 *Define $i_0 = \min\{i : L_i < 3 \log^7 \Delta\}$. Then $U_{i_0}, L_{i_0}, S_{i_0} > \log^7 \Delta$, $U_{i_0} \leq \Delta^{1-\frac{1}{k}} \log^4 \Delta$ and $L_i > S_i > \frac{1}{2} L_i$ for each $i \leq i_0$.*

We can now prove a weaker version of Theorem 1 where we restrict our attention to conventional colouring rather than list colouring.

Theorem 8 *For fixed k, if H is a k-uniform linear hypergraph with maximum degree Δ, then $\chi^{'}(H) \leq \Delta + \mathrm{O}(\Delta^{1-\frac{1}{k}} \log^4 \Delta)$.*

Proof We apply our procedure to the line graph of H for i_0 iterations, or until our procedure halts prematurely. By Lemmas 6 and 7, with positive probability we do not halt prematurely, and we have a partial proper colouring using $\Delta + \Delta^{1-\frac{1}{k}} \log^4 \Delta$ colours such that each clique $C \in \mathcal{C}$ contains at most $\mathrm{O}(\Delta^{1-\frac{1}{k}} \log^4 \Delta)$ uncoloured vertices. Therefore, the uncoloured vertices induce a graph with maximum degree at most $\mathrm{O}(\Delta^{1-\frac{1}{k}} \log^4 \Delta)$, and so it can be coloured using $\mathrm{O}(\Delta^{1-\frac{1}{k}} \log^4 \Delta)$ new colours. $\square$

3 Most of the Proofs

We begin by introducing two tools which will prove useful. Our first can be looked at as showing that the events that different vertices retain particular colours, are nearly independent, except for the obvious dependency between the events that 2 particular adjacent vertices retain the same colour.

Lemma 9 *For any constant t and set of (real or artificial) vertices $v_1, ..., v_t$ along with colours $\gamma_1 \in L_{v_1}, ..., \gamma_t \in L_{v_t}$, such that if v_j and v_r are adjacent then $\gamma_j \neq \gamma_r$, the probability that each v_j retains γ_j, conditional on each v_j receiving γ_j is $(\frac{1}{L_i}\mathrm{Retain}_i)^t(1 + O(\frac{1}{L_i}))$.*

Proof Let t_1, t_2 denote the number of real and artificial vertices respectively amongst the t vertices being considered.

For each (real or artificial) vertex v_j, denote by N_j the set of (real or artificial) neighbours of v_j which have γ_j in their lists, i.e. the union over each C containing v_i of $S_{C,\gamma_i} - v_i$. Each S_{C,γ_i} has size exactly S_i and so by the linearity of H, $|N_j| = kS_i - k$ if v_j is real and $|N_j| = S_i - 1$ if v_j is artificial. For each colour $\gamma \in \{\gamma_1, ..., \gamma_t\}$ we let $n_1(\gamma), n_2(\gamma)$ be the number of real, artificial vertices v_j such that $\gamma_j = \gamma$, and we let $N(\gamma)$ be the union over all these $n_1(\gamma) + n_2(\gamma)$ vertices v_i of N_i. If $\gamma_j = \gamma_s = \gamma$ then v_j is not adjacent to v_s and so by the linearity of H, these two vertices have at most k^2 common neighbours. Therefore, $|N(\gamma)| = (kn_1(\gamma) + n_2(\gamma))S_i(1 - O(\frac{1}{S_i}))$.

For any vertex $u \in \cup_\gamma N(\gamma)$, we let $n(u)$ denote the number of sets $N(\gamma)$ to which u belongs (note that if u is artificial then $n(u) = 1$). We have $n(u) \leq t$ for all u, and $\sum_u n(u) = (kt_1 + t_2)S_i(1 - O(\frac{1}{S_i}))$. The probability that for each real v_j, no vertex in N_j receives γ_j is $\prod_u(1 - \frac{n(u)}{L_i}) = (1 - \frac{1}{L_i})^{\sum n(u)}(1 + O(\frac{1}{L_i})) = \mathrm{Retain}_i^{t_1}(1 - \frac{1}{L_i})^{t_2(S_1 - 1)}(1 + O(\frac{1}{L_i}))$. The probability that the t_2 artificial vertices do not lose their colours in step 4(b) is $(1 - \frac{1}{L_i})^{t_2(k-1)(S_1-1)}$ and the result follows. □

Our next tool is a simple but powerful counting tool whose proof is so short that we include it (this tool is a simple special case of the much more powerful, but equally easy to prove inclusion-exclusion principle):

Theorem 10 *Let X be a set and let $\mathcal{S}$ be a multiset of (possibly empty) subsets of X. Then the number of non-empty elements of $\mathcal{S}$ equals*

$$\sum_{S \in \mathcal{S}} \sum_{j=1}^{|S|} (-1)^{j+1} |\{U | U \subseteq S, |U| = j\}|.$$

Proof We know that $(1 - x)^k = \sum_{i=0}^k \binom{k}{i} x^i$. So, in particular, $(1 - 1)^k = \sum_{i=0}^k \binom{k}{i}(-1)^i$. Thus, for any nonempty set S with k elements we have:

$0 = \sum_{i=0}^{k} |\{U|U \subseteq S, |U| = i\}|(-1)^i$. Hence, $-1 = \sum_{j=1}^{k} |\{U|U \subseteq S, |U| = j\}|(-1)^j$, that is $1 = \sum_{j=1}^{k} |\{U|U \subseteq S, |U| = j\}|(-1)^{j+1}$. The desired result follows immediately. □

Corollary 11 *Let S be a (not neccessarily uniform) random set chosen from a finite ground set X. Then,* $\mathbf{Pr}(S \neq \emptyset) = \sum_{j=1}^{|X|}(-1)^{j+1}\mathbf{Exp}(|\{U|U \subseteq S, |U| = j\}|)$.

Proof To begin, we assume that for each subset Y of X, $\mathbf{Pr}(S = Y)$ is rational (as is the case in all our applications of the corollary), the general result follows by successive approximations. So, by taking multiple copies of each subset of Y we can generate a multiset $\mathcal{S}$ of subsets of X where S is equally likely to be any element of $\mathcal{S}$. Suppose $\mathcal{S}$ has m elements. Now, $\mathbf{Pr}(S \neq \emptyset) = \frac{|\{S \in \mathcal{S}, S \neq \emptyset\}|}{m}$. Furthermore, $\mathbf{Exp}(|\{U|U \subseteq S, |U| = j\}|) = \frac{\sum_{S \in \mathcal{S}} |\{U|U \subseteq S, |U|=j\}|}{m}$. The result follows immediately. □

Armed with these tools, we can now prove our bounds on the expected values of our parameters:

Proof of Lemma 4: (a) This follows immediately from the fact that a vertex retains its colour with probability precisely Retain_i.

(b) For each colour $\gamma \in L_v$, we need to compute the probability that no neighbour of v retains colour γ. To do so, we will apply Theorem 10 to calculate the probability of the complementary event that at least one neighbour of v retains colour γ.

We let S be the (random) set of (real or artificial) neighbours of v which retain γ and calculate:

$$\mathbf{Pr}(S \neq \emptyset) = \sum_{j=1}^{k}(-1)^{j+1}\mathbf{Exp}(|\{U|U \subseteq S, |U| = j\}|).$$

We first calculate for each j, the number of stable sets of size j in $N(v)$ such that each vertex has γ in its list. To begin we fix the $\binom{k}{j}$ cliques containing v which contain a member of the stable set. Having done so we arbitrarily order these cliques, and pick the vertices of the stable set. In the first clique C, we are free to pick any vertex of $S_{C,\gamma}$. In any later clique C', we must avoid picking a neighbour of any vertex already picked. However, by the linearity of the hypergraph, each vertex not in C' has at most k neighbours in C', so in total at most $k(k-1)$ vertices are forbidden, and we have between S_i and $S_i - k(k-1)$ choices for the next vertex of the stable set. Thus, the number of such stable sets is $\binom{k}{j}S_i^j(1 - \mathrm{O}(\frac{1}{S_i}))$.

By Lemma 9, the probability that every vertex in such a stable set receives and retains γ is $(\frac{1}{L_i}\text{Retain}_i)^j(1 - O(\frac{1}{L_i}))$. Therefore, $\mathbf{Pr}(S \neq \emptyset) =$

$(1-O(\frac{1}{L_i}))\sum_{j=1}^{k}\binom{k}{j}(-1)^{j+1}(\frac{S_i}{L_i}\text{Retain}_i)^j$. But this sum is simply $1-(1-\frac{S_i}{L_i}\text{Retain}_i)^k$ (to see this we use the expansion of $(1-x)^k$ multiplied through by -1). Therefore, the probability that no neighbour of v retains γ is $(1-O(\frac{1}{L_i}))\text{Keep}_i^k$.

Therefore, $\mathbf{Exp}(|L_v^{'}|)=|L_v|\times(1-O(\frac{1}{L_i}))\text{Keep}_i^k$, as required.

(c) For any real vertex $u\in S_{C,\gamma}$ let A_u be the event that u does not retain its colour and let B_u be the event that γ remains in $L_{C'}$ for every $C^{'}\neq C$ which contains u. $\mathbf{Pr}(A_u)=1-\text{Retain}_i$ and using the same computations as in part (b), we obtain that $\mathbf{Pr}(B_u)=\text{Keep}_i^{k-1}$ and so if these two events were independent then the result would follow. Of course, they are not independent, but what we will show is that their dependency is limited. In particular, conditional, on u receiving some colour $\alpha\neq\gamma$ they are very nearly independent, and since u receives such a colour with probability $1-O(\frac{1}{L_i})$ the strong dependency created by conditioning on v receiving γ is negligible.

For any vertex x and colour α we denote by "$x\leftarrow\alpha$" the event that x receives α. Consider any colour $\alpha\neq\gamma$ in L_u. We start by computing the probability, conditional on $u\leftarrow\alpha$, that u retains α and that no neighbour of u outside of C retains γ. Applying the same computations as in part (b), we obtain that this probability is $\text{Retain}_i\text{Keep}_i^{k-1}(1-O(\frac{1}{L_i}))$ (the details are nearly identical to those in (b) and we omit them). Since this holds for each of the L_i-1 such colours in L_u, the probability that u retains some colour and that no neighbour of v outside of C retains γ is $\text{Retain}_i\text{Keep}_i^{k-1}(1-O(\frac{1}{L_i}))$. As mentioned earlier, the probability that no neighbour of u outside of C retains γ is $\text{Keep}_i^{k-1}(1-O(\frac{1}{L_i}))$ (this time the details are even more identical to those in (b) and are also omitted). Therefore, the probability that u does not retain a colour and that no neighbour of u outside of C retains γ, i.e. that u remains in $n(C,\gamma)$, is $(1-\text{Retain}_i)\text{Keep}_i^{k-1}(1-O(\frac{1}{L_i}))$ as required. □

Proof of Lemma 7: Recursively define U_i^*, L_i^*, S_i^* as follows. $U_1^*=L_1^*=S_1^*=\Delta$, and for $i\geq 1$,

- $\text{Keep}_i^*=1-\frac{S_i^*}{L_i^*}\times\text{Retain}_i$;
- $U_{i+1}^*=U_i^*(1-\text{Retain}_i)$;
- $L_{i+1}^*=L_i^*\times(\text{Keep}_i^*)^k$;
- $S_{i+1}^*=S_i^*\times(\text{Keep}_i^*)^{k-1}\times(1-\text{Retain}_i)$.

Remark: It turns out that future calculations are simpler if we use Retain_i here instead of defining Retain_i^* in the obvious manner and using it.

First observe that a straightforward induction yields that $L_i^*=S_i^*$ and $\text{Keep}_i^*=1-\text{Retain}_i$ for all $i\geq 1$. Therefore,

- $U_i^*=\Delta\prod_{j=1}^{i-1}(1-\text{Retain}_j)$;

- $L_i^* = S_i^* = \Delta \prod_{j=1}^{i-1}(1 - \text{Retain}_j)^k$;

and so when $L_i^* \approx \log^7 \Delta$, we have $U_i^* \approx \Delta^{1-\frac{1}{k}} \log^{7/k} \Delta$.

We will show that U_i, L_i, S_i are not very far from U_i^*, L_i^*, S_i^* respectively. To do so, we will verify the following inductively for each $1 \le i \le i_0$:

i. $U_i \le U_i^* + \sqrt{U_i^*} \log^{2.5} \Delta$;

ii. $L_i \le L_i^* + \sqrt{L_i^*} \log^{2.5} \Delta$;

iii. $L_i \ge S_i + (\log^4 \Delta - 4i \log \Delta) L_i^{1-\frac{1}{k}}$.

Note that trivially, $U_i > U_i^*$ and $S_i > S_i^* = L_i^*$. Note also that since $L_i > S_i > L_i - L_i^{1/4}$, and $\text{Retain}_i < 1$, we have $\text{Keep}_i + L_i^{-3/4} > \text{Keep}_i^* > \text{Keep}_i$.

Assume that our inductive claims hold for some $i < i_0$:

Part (i):

$$\begin{aligned}
U_{i+1} &= U_i(1 - \text{Retain}_i) + \sqrt{U_i} \log^2 \Delta \\
&\le (U_i^* + \sqrt{U_i^*} \log^{2.5} \Delta)(1 - \text{Retain}_i) + \sqrt{U_i} \log^2 \Delta \\
&\le U_{i+1}^* + \sqrt{U_{i+1}^*} \log^{2.5} \Delta \sqrt{1 - \text{Retain}_i} + \sqrt{U_i} \log^2 \Delta \\
&\le U_{i+1}^* + \sqrt{U_{i+1}^*} \log^{2.5} \Delta.
\end{aligned}$$

where the last inequality follows from the fact that $\sqrt{1 - \text{Retain}_i} = 1 - \theta(1)$.

Part (ii):

$$\begin{aligned}
L_{i+1} &= L_i \text{Keep}_i^k - \sqrt{L_i} \log^2 \Delta \\
&< L_i (\text{Keep}_i^*)^k + O(L_i^{1/4}) - \sqrt{L_i} \log^2 \Delta \\
&< L_i (\text{Keep}_i^*)^k \\
&< (L_i^* + \sqrt{L_i^*} \log^{2.5} \Delta)(\text{Keep}_i^*)^k \\
&< L_{i+1}^* + \sqrt{L_{i+1}^*} \log^{2.5} \Delta
\end{aligned}$$

Part (iii):

This last part is the most delicate, particularly when $k = 2$, because of the fact that $\text{Keep}_i - \text{Keep}_i^*$ is large enough to cause some difficulties.

It will be convenient to define $\zeta_i = L_i - S_i$. At one point we will use the fact that $\text{Keep}_i = 1 - \text{Retain}_i + \frac{\zeta_i}{L_i} \text{Retain}_i$.

$$\begin{aligned}
L_{i+1} &= L_i \times \text{Keep}_i^k - \sqrt{L_i} \log^2 \Delta \\
&> (S_i + \zeta_i) \text{Keep}_i^{k-1} (1 - \text{Retain}_i + \frac{\zeta_i}{L_i} \text{Retain}_i) - \sqrt{L_i} \log^2 \Delta
\end{aligned}$$

$$
\begin{aligned}
&> && S_i \times \mathrm{Keep}_i^{k-1}(1-\mathrm{Retain}_i) + \zeta_i \times \mathrm{Keep}_i^{k-1}(1-\mathrm{Retain}_i) \\
& && \quad +\zeta_i \times \mathrm{Keep}_i^{k-1} \times \mathrm{Retain}_i \times \frac{S_i}{L_i} - \sqrt{L_i}\log^2\Delta \\
&= && S_{i+1} + \zeta_i \times \mathrm{Keep}_i^{k-1} - 2\sqrt{L_i}\log^2\Delta - \frac{\zeta_i^2}{L_i} \\
&> && S_{i+1} + \zeta_i \times \mathrm{Keep}_i^{k-1} - 3L_i^{1-\frac{1}{k}}\log^2\Delta,
\end{aligned}
$$

where the last inequality follows because $\zeta_i^2/L_i = O(\log^5\Delta)$ by part (ii) and because $L_i^{1-\frac{1}{k}}\log^2\Delta = \theta(\log^{5.5}\Delta)$ since $i < i_0$. Therefore,

$$
\begin{aligned}
\zeta_{i+1} &\geq \zeta_i \times \mathrm{Keep}_i^{k-1} - 3L_i^{1-\frac{1}{k}}\log^2\Delta \\
&\geq (\log^4\Delta - 4i\log^2\Delta)L_i^{1-\frac{1}{k}}\mathrm{Keep}_i^{k-1} - 3L_i^{1-\frac{1}{k}}\log^2\Delta \\
&\geq (\log^4\Delta - 4i\log^2\Delta)L_{i+1}^{1-\frac{1}{k}} - 4L_{i+1}^{1-\frac{1}{k}}\log^2\Delta \\
&= (\log^4\Delta - 4(i+1)\log^2\Delta)L_{i+1}^{1-\frac{1}{k}}
\end{aligned}
$$

as required, where the last inequality follows from the easily verified fact that $\mathrm{Keep}_i \approx (1-\frac{1}{e^k})$, and so $3L_i^{1-\frac{1}{k}} < 4L_{i+1}^{1-\frac{1}{k}}$.

Now $\mathrm{Retain}_i \approx \mathrm{e}^{-k}$ and so $i_0 = O(\log\Delta)$. Therefore (iii) implies $L_i \geq S_i$. The remainder of the lemma follows immediately. □

4 The Concentration Results

We turn now to the proof of Lemma 5. We will need the following tool:

A Concentration Lemma: *Let R be a random variable defined as a function of a set of independent random variables $X_1, ..., X_m$. Let β, D be positive reals. Suppose that for each outcome $x = (x_1, ..., x_m)$, we can choose values $b_1, ..., b_m$ all between 0 and β so that*

(a) $\sum_{i=1}^m b_i \leq D$; and

(b) for any other outcome $y = (y_1, ..., y_m)$, $R(x) \leq R(y) + \sum_{x_i \neq y_i} b_i$.

Then there exists T which is a function only of β such that for all $t \geq T\sqrt{D}$,

$$\mathbf{Pr}(|R - \mathbf{Exp}(R)| > t) \leq 4e^{-\frac{t^2}{5\beta D}}.$$

This tool is a minor variation on a standard useful corollary of an inequality of Talagrand[23], and we prove it in the final section of this paper. With the tool in hand we can concentrate on constructing the proof.

Proof of Lemma 5:

Part (a): Set $R = |U_T^{'}|$. Clearly, R depends only on the colours assigned to the vertices in U_T, and to the set W of (real or artificial) vertices not in U_T but adjacent to a vertex in U_T with which they share an available colour. We enumerate $W \cup U_T$ as $v_1, ..., v_m$ for some m.

We consider a fixed assignment X of colours to the vertices in $W \cup U_T$, and the corresponding set $U_T^{'}$ of uncoloured vertices. (A colour assignment to an artificial vertex v is an assignment of either it's associated colour or of no colour.) Obviously, each vertex in $U_T^{'}$ is adjacent to a vertex assigned the same colour. For each vertex v_i of $U_T^{'}$ which does not have a neighbour in $U_T^{'}$ which is assigned the same colour as v_i, we choose a single witness $v_j \in W$ which is adjacent to v_i and has the same colour as v_i. We note that a vertex can be a witness for at most one vertex of $U_T^{'}$ because U_T is a clique. For each v_i in $U_T^{'}$, we set $b_i = 2$ and for each v_i in W we set b_i to be one if v_i is a witness and 0 otherwise. We claim that these b_i satisfy the conditions of the concentration lemma with $\beta = 2$ and $D = 3|U_T|$.

Obviously, the b_i sum to at most $2|U_T|+|U_T^{'}| \leq 3D$, and so the first condition holds.

To verify the second condition, consider some other colour assignment Y which yields d fewer vertices in $U_T^{'}$ than X does. Let $S_1 \subseteq U_T$ be the set of vertices which have the same assignment in X as in Y and yet do not retain their colours in X but retain them in Y. Let S_2 be the set of vertices of U_T which have a different assignment in Y than in X. Let S_3 be the set of witnesses which have a different assignment in Y than in X. Now, $d \leq |S_1| + |S_2|$. We observe that $|S_1| \leq |S_2| + |S_3|$, because since U_T is a clique, for each vertex $u \in S_1$ of colour γ, there is no other vertex in S_1 which receives γ under Y and so either there is a vertex of S_2 which receives γ under X or the witness for u is in S_3. Hence, $d \leq 2|S_2| + |S_3| \leq \sum_{v_j \in S_2 \cup S_3} b_j$.

Applying the Concentration Lemma, along with the fact that $|U_T^{'}| \leq |U_T|$ we obtain that for $t >> \sqrt{|U_T|}$:

7 $\mathbf{Pr}(||U_T^{'}| - \mathbf{Exp}(|U_T^{'}|)| > t) \leq 4e^{-\frac{t^2}{30|U_T|}}$.

Setting $t = (\log \Delta)\sqrt{|U_T|}$, this yields Lemma 5(a).

Part (b): Next, we consider the concentration of $|L_v^{'}|$. We find it easier to consider some related but easier to analyze variables.

Namely, for $0 \leq b \leq a \leq k$, we define $\mathcal{M}(a,b)$ to be the set of colours in L_v which are assigned to at least 1 (real or artificial) vertex in *at least* a of the k cliques in $\mathcal{C}_v$ and are removed from a (real or artificial) vertex in *at least* b of the k cliques in $\mathcal{C}_v$. We let $M(a,b) = |\mathcal{M}(a,b)|$. It will prove convenient to extend this definition to $a = k+1$, in which case $M(k+1,b) = 0$ for all b.

It will be fairly straightforward to apply our Concentration Inequality to prove that for all a, b, $M(a,b)$ is highly concentrated. As we shall see next, R

can easily be expressed as a linear combination of $M(a,b)$ terms, and so it will follow that R is also highly concentrated.

First observe that if a colour is assigned to at least one vertex in a clique, and it is removed from at least one vertex in a clique, then it is removed from every (real or artificial) vertex that it was assigned to in that clique. Therefore, $M(a,b)$ counts that number of colours which were assigned to at least one vertex in at least a cliques, and purged entirely from at least b of those cliques.

For $0 \leq b \leq a \leq k$, the number of colours which are assigned to at least 1 vertex in *exactly* a cliques, and are removed from at least b of those cliques is $M(a,b) - M(a+1,b)$. Similarly, the number of colours which are assigned to at least 1 vertex in *exactly* a cliques, and are retained by a vertex in at least one of those cliques is $(M(a,0) - M(a+1,0)) - (M(a,a) - M(a+1,a))$. Thus

$$R = |L_v| - \sum_{a=1}^{k} (M(a,a) - M(a+1,a)) - (M(a,0) - M(a+1,0)).$$

Now, we shall prove:

8 $\mathbf{Pr}(|M(a,b) - \mathbf{Exp}(M(a,b))| > t) \leq 4e^{-\frac{t^2}{10kL_i}}$.

It immediately follows that:

$$\mathbf{Pr}(|R - \mathbf{Exp}(R)| > 4kt) \leq 16ke^{-\frac{t^2}{10kL_i}}.$$

Setting $t = (\log^2 \Delta)\sqrt{L_i}$, this yields Lemma 5(b).

So, let X be a fixed colour assignment on all vertices of distance at most 2 from v, $\{v_1, ..., v_m\}$. along with the decisions of whether or not to uncolour each of the artificial such vertices in step 4(b). (In order to ensure that X is chosen from a product space, we make these decisions in advance at the same time as we assign colours, even though for some artificial vertices these decisions may not have to be made.) For each $\gamma \in \mathcal{M}(a,b)$ we consider b cliques $C_1, .., C_j$ of $\mathcal{C}_v$ from which γ was removed and $a-b$ other cliques $K_1, .., K_{l-j}$ of $\mathcal{C}_v$ in which γ was assigned to at least one vertex. We choose two witnesses corresponding to each C_ℓ and one witness in each K_ℓ. The witnesses for C_ℓ consist of a (real or artificial) vertex coloured γ in C_ℓ and a witness which demonstrates that the first witness lost its colour. This second witness can either be a neighbour of the first witness which is also assigned colour γ, or in the case that the first witness is an artificial vertex, the decision that this vertex would lose its colour. The witness for K_ℓ consists of a vertex coloured γ. We set $b_j = 1$ if v_j is a witness and $b_j = 0$ otherwise. We note that $\sum b_j \leq 2kM(a,b)$. Clearly, we can apply the concentration lemma with $\beta = 1, D = 2kL_i$ and obtain:

$$\mathbf{Pr}(|M(a,b) - \mathbf{Exp}(M(a,b))| > t) \leq 4e^{-\frac{t^2}{10kL_i}},$$

as required.

Part (c): The proof is very similar to that of part (b). Set $R = n(C, \gamma)$.

This time, for $0 \le b \le a \le k-1$, we define $\mathcal{M}(a,b)$ to be the set of real vertices $u \in S_{C,\gamma}$ such that

i. u does not retain its colour;

ii. γ is assigned to at least 1 (real or artificial) vertex in *at least* a of the $k-1$ cliques in $\mathcal{C}_u$ other than C; and

iii. γ is removed from at least 1 (real or artificial) vertex in *at least* b of those cliques of $\mathcal{C}_u$.

Again we let $M(a,b) = |\mathcal{M}(a,b)|$, and we extend this definition to $a = k$, setting $M(k,b) = 0$ for all b. Also, let U be the set of real vertices in $S_{C,\gamma}$ which do not retain their colours.

As before,

$$R = |U| - \sum_{a=1}^{k-1} (M(a,0) - M(a+1,0)) - (M(a,a) - M(a+1,a)).$$

That $|U|$ is highly concentrated follows from (7), which yields that for $t >> \sqrt{S_i}$:

9 $\mathbf{Pr}(|U| - \mathbf{Exp}(|U|)| > t) \le 4e^{-\frac{t^2}{30S_i}}$,

We will prove that for such t:

10 $\mathbf{Pr}(|\mathcal{M}(a,b)| - \mathbf{Exp}(\mathcal{M}(a,b))| > t) \le 4e^{-\frac{t^2}{10k^2(2k+3)S_i}}$,

which, along with (9), will imply part (c) in the same manner as (8) implied part (b).

To prove (10), let X be a fixed colour assignment on the set of vertices of distance at most 3 from C, $\{v_1, ..., v_m\}$ along with the decisions of whether or not to uncolour each of the artificial such vertices in step 4(b). For each vertex u in $\mathcal{M}(a,b)$, if u has a colour different from that of every other vertex of $\mathcal{M}(a,b)$ then we choose a *hostile witness* for u which is a (real or artificial) neighbour of u assigned the same colour. Note that each vertex can be a hostile witness for at most one element of $\mathcal{M}(a,b)$. We also consider b distinct cliques $C_1, .., C_b \in \mathcal{C}_v - C$ each containing some (real or artificial) vertex other than u from which γ was removed and $a-b$ other cliques $K_1, .., K_{a-b} \in \mathcal{C}_v - C$ in which γ was assigned to at least one (real or artificial) vertex other than u. For each K_ℓ, we choose one *direct witness* which is a vertex in K_ℓ of colour γ. For each C_ℓ we choose one *direct witness*, i.e. a vertex w in C_ℓ of colour γ, and one *corroborating witness*, which, in the case that w is a real vertex, is a neighbour of w that is also assigned colour γ, and in the case that w is an artificial vertex, is either such a neighbour or the choice to uncolour w in Step 4(b) of our procedure. It is possible that this corroborating witness is a direct witness for a different clique, or in fact for a different member of $\mathcal{M}(a,b)$.

By the linearity of the underlying hypergraph, it is straightforward to verify that no vertex can be a direct witness more than k times. However, it is quite possible for a vertex to serve as a corroborating witness many times. In order to avoid this problem, we must be careful when we choose our corroborating witnesses. First we choose all direct witnesses for all members of $\mathcal{M}(a,b)$. Whenever a particular clique of $\mathcal{C}$ contains more than one direct witness, then for each of these direct witnesses that requires a corroborating witness, we choose for it one of the other direct witnesses in that clique. Note that we can easily do this such that no direct witness serves as a corroborating witness for more than one direct witness in that clique. If a direct witness for a clique C_ℓ is not adjacent to any other direct witness, then we choose for it a corroborating witness which is not a direct witness. This ensures that no vertex serves as a corroborating witness for more than one direct witness in the same clique of $\mathcal{C}$, and so serves as a corroborating witness for at most k direct witnesses in each of k different cliques for a total of at most k^2 times. Obviously, the choice to uncolour an artificial vertex in Step 4(b) of our procedure can be the corroborating witness for at most 1 direct witness, and so can serve as a corroborating witness at most k times.

For each x which is either a vertex not in $\mathcal{M}(a,b)$, or a choice to uncolour an artificial vertex we let x be the number of vertices for which x is a witness. For each x which is a vertex in $\mathcal{M}(a,b)$, we let b_x be two more than the number of vertices for which x is a witness. We note that $\sum b_x \leq (2k+3)M(a,b)$. We observe that if for some other colour assignment Y, a vertex u is in $M(a,b)$ under X but not under Y then either some witness for u is assigned a different colour in the two colour assignments, or u has a different colour under the two assignments, or all the vertices of $M(a,b)-u$ which received the same colour as u under assignment X have different colours under the two assignments. It follows that the conditions of the Concentration Lemma hold with $\beta = 2k^2, D = (2k+3)S_i$, so we obtain that for $t >> \sqrt{S_i}$:

$$\mathbf{Pr}(|M(a,b) - \mathbf{Exp}(M(a,b))| > t) \leq 6e^{-\frac{t^2}{10k^2(2k+3)S_i}}$$

which is (10). □

Remark: It will be useful in Section 6 to note that the only point where we require that the underlying hypergraph be linear is in the proof of part (c).

We close this section by presenting a second very similar concentration inequality which will prove useful in section 6. The proofs of the two inequalities are nearly identical.

Concentration Lemma II: *Let R be a random variable defined as a func-*

tion of a set of independent random variables $X_1, ..., X_m$. Let β, D be positive reals. Suppose that for each outcome $x = (x_1, ..., x_m)$, we can choose values $b_1, ..., b_m$ all between 0 and β so that

(a) $\sum_{i=1}^{m} b_i \leq R(x)$*; and*

(b) for any other outcome $y = (y_1, ..., y_m)$, $R(x) \leq R(y) + \sum_{x_i \neq y_i} b_i$.

Then there exists T which is a function only of β such that for all t with $T\sqrt{\mathbf{Exp}(R)} \leq t \leq \mathbf{Exp}(R)$,

$$\mathbf{Pr}(|R - \mathbf{Exp}(R)| > t) \leq 4e^{-\frac{t^2}{10\beta \mathbf{Exp}(R)}}.$$

We remark that the first concentration bound applies well in cases where R has an absolute upper bound which is a constant multiple of $\mathbf{Exp}(R)$. When this does not hold, we may require the second bound. (See for example the proof of Lemma 13 later.) If one wished, these two bounds could obviously be combined into a single bound, and in fact Talagrand's Inequality yields even more powerful bounds of this flavour.

5 The List Chromatic Index of Linear Hypergraphs

We turn now to the list colouring problem. We can again assume that Δ is at least Δ_k because any graph of smaller degree can be coloured with $k\Delta_k$ colours which we can add to the constant c_k in Theorem 1.

Our procedure for list edge colouring linear hypergraphs is similar to that given above. Of course, the initial L_v are given as input to the problem, we only insist that each have at least $\Delta + \lceil 3k(\log \Delta)^4 \Delta^{1-\frac{1}{k}} \rceil$ elements and then truncate them so they all have exactly this many elements. The only complication arises at the end of the procedure, as we cannot simply introduce a new set of colours to finish things off. Instead, at the beginning of the procedure we will choose a subset $\mathrm{Save}_v \subseteq L_v$ for each v which will be used to colour the vertices still uncoloured after our general iterative procedure has run out of steam.

The preprocessing procedure takes as input a hypergraph with $\Delta + \lceil 3k(\log \Delta)^4 \Delta^{1-\frac{1}{k}} \rceil$ colours on each list. It chooses a set Save_C of colours for each clique C which will not be used on C during the main colouring procedure but rather are saved for the final phase. During the final phase, a vertex v of the line-graph will be coloured with a colour from $\mathrm{Save}_v = L_v \cap (\cap_{c \in \mathcal{C}_v} \mathrm{Save}_C)$. Of course, during the main procedure, v can only be coloured with a colour in $\mathrm{Avail}_v = L_v - (\cup_{c \in \mathcal{C}_v} \mathrm{Save}_C)$. For each $C \in \mathcal{C}$ and $\gamma \in \mathrm{Save}_C$ we define $T_{C,\gamma} = \{u \in C : \gamma \in \mathrm{Save}_u\}$.

Lemma 12 *We can choose a set of colours* Save_C *for each* $C \in \mathcal{C}$*, such that for each vertex* v*, clique* $C \in \mathcal{C}$ *and colour* $\gamma \in \mathrm{Save}_C$*:*

(a) $|L_v \cap (\cup_{C \in \mathcal{C}_v} \mathrm{Save}_C)| \leq 2k\Delta^{1-\frac{1}{k}} \log^4 \Delta$,

(b) $|L_v \cap (\cap_{C \in \mathcal{C}_v} \mathrm{Save}_C)| \geq \frac{1}{2} \log^{4k} \Delta$*, and*

(c) $|T_{C,\gamma}| \leq 2\Delta^{\frac{1}{k}} \log^{4k-4} \Delta$.

Proof For each clique $C \in \mathcal{C}$ and colour γ s.t. $\gamma \in L_v$ for some $v \in C$, we place γ into Save_C with probability $p = \log^4 \Delta / \Delta^{\frac{1}{k}}$. For each vertex v, clique $C \in \mathcal{C}$ and colour γ we define E_v^1, E_v^2 and $E_{C,\gamma}^3$ to be the probability that v violates condition (a), v violates condition (b), and C, γ violate condition (c), respectively.

For any pair C, γ, we define $T_{C,\gamma}^{'}$ to be the set of vertices $u \in C$ such that $\gamma \in L_u \cap (\cap_{C_1 \in (\mathcal{C}_u - C)} \mathrm{Save}_{C_1})$, and we define the event $(E^3)_{C,\gamma}^{'}$ to be the event that $|T_{C,\gamma}^{'}| > 2\Delta^{\frac{1}{k}} \log^{4k-4} \Delta$. Note that if $(E^3)_{C,\gamma}^{'}$ does not hold then neither does $E_{C,\gamma}^3$.

Note further that

$$\begin{array}{rcccl} \mathbf{Exp}(|L_v \cap (\cup_{C \in \mathcal{C}_v} \mathrm{Save}_C)|) & \leq & |L_v| \times kp & \approx & k\Delta^{1-\frac{1}{k}} \log^4 \Delta \\ \mathbf{Exp}(|L_v \cap (\cap_{C \in \mathcal{C}_v} \mathrm{Save}_C)|) & \leq & |L_v| \times p^k & \approx & \log^{4k} \Delta \\ \mathbf{Exp}(|T_{C,\gamma}^{'}|) & \leq & \Delta \times p^{k-1} & \approx & \Delta^{\frac{1}{k}} \log^{4k-4} \Delta \end{array}$$

Thus, it is a straightforward application of the Chernoff Bound to show that the probability of any one event is (much) less than $e^{-\log^2 \Delta}$ (for the event $(E^3)_{C,\gamma}^{'}$ it is important to apply the linearity of the underlying hypergraph). Furthermore, each event, E, is mutually independent of all those events which do not have a subscript which is at distance less than 3 from some subscript of E. Observing that for any $C \in \mathcal{C}$, there are fewer than $k\Delta^2$ colours γ for which we define $C_{C,\gamma}$, it follows that each event is mutually independent of all but at most $2k\Delta^3$ other events, and so our lemma follows from the Local Lemma. $\Box$

For each vertex v we remove $\cup_{c \in \mathcal{C}_v} \mathrm{Save}_C$ from L_v, leaving at least $\Delta + \Delta^{1-\frac{1}{k}} \log^4 \Delta$ colours in L_v which is now equal to Avail_v. We then perform our standard iterative procedure, with one slight twist, we also need to keep track of the number of uncoloured vertices remaining in each $T_{C,\gamma}$ after each iteration. We recursively define T_i as follows:

- $T_0 = 2\Delta^{\frac{1}{k}} \log^{4k-4} \Delta$;
- $T_{i+1} = T_i(1 - \mathrm{Retain}_i) + \sqrt{T_i} \log^2 \Delta$.

We add the following condition to Lemma 6:
(d) For each $C \in \mathcal{C}$ *and* $\gamma \in \mathrm{Save}_C$*, we have* $|U_{T_{C,\gamma}}| \leq T_i$*.*

Claim: *The probability that a pair C, γ violates condition (d) is at most $\Delta^{-\log^2 \Delta}$.*

Proof Clearly, $\mathbf{Exp}(|U'_{T_{C,\gamma}}|) = (1 - \text{Retain}_i)|U_{T_{C,\gamma}}|$. That $|U'_{T_{C,\gamma}}|$ is sufficiently concentrated follows from Lemma 5(a). $\Box$

The strengthened version of Lemma 6 follows immediately.

Using the same techniques as in the proof of Lemma 7, we obtain $T_{i_0} \leq 3\log^{4k-4+7/k}\Delta$. We omit the details.

And this now yields the proof of our main theorem:

Proof of Theorem 1: We carry out our iterative procedure for i_0 iterations. Let H' be the graph induced by the uncoloured vertices. We will list colour H' using the lists Save_v. Since each list contains at least $\frac{1}{2}\log^{4k}\Delta$ colours, and since for each colour γ, no vertex has more than $kT_{i_0} = \mathrm{O}(\log^{4k-\frac{1}{2}}\Delta)$ neighbours with γ in their lists, we satisfy the conditions of the List Colouring Lemma, and so we can successfully complete the colouring. $\Box$

6 Hypergraphs of bounded codegree

In this section, we consider k-uniform hypergraphs of bounded codegree, i.e. those for which the maximum codegree of a pair of vertices is B. We prove Theorem 2 – that there is a c_k such that any such hypergraph has list chromatic index at most $\Delta + c_k(B^{\frac{1}{k}}\Delta^{1-\frac{1}{k}}(\log\frac{\Delta}{B})^4)$. Once again, we need only prove the theorem for large Δ as we can sweep the small Δ under the carpet by increasing the constant c_k. In fact, we can even assume that $\frac{\Delta}{B}$ is at least as big as a large constant, and so we assume that it is large enough to satisfy a number of inequalities scattered throughout the section.

The proof follows along the same lines as that of Theorem 1, with only two significant differences, which we outline below. Unless specified otherwise, all notation is as in the proof of Theorem 1.

The first difference arises in the proof of our analogue of Lemma 12:

Lemma 13 *We can choose a set of colours Save_C for each $C \in \mathcal{C}$, such that for each vertex v, clique C and colour $\gamma \in \text{Save}_C$:*

(a) $|L_v \cap (\cup_{C\in\mathcal{C}_v}\text{Save}_C)| \leq 2k\Delta^{1-\frac{1}{k}}B^{\frac{1}{k}}\log^4(\Delta/B)$,

(b) $|L_v \cap (\cap_{C\in\mathcal{C}_v}\text{Save}_C)| \geq \frac{1}{4}B\log^{4k}(\Delta/B)$, *and*

(c) $|T_{C,\gamma}| \leq 2\Delta^{\frac{1}{k}}B^{1-\frac{1}{k}}\log^{4k-4}(\Delta/B)$.

Proof For each colour γ and clique $C \in \mathcal{C}$, we place γ into Save_C with probability $(\log\frac{\Delta}{B})^4(\frac{B}{\Delta})^{\frac{1}{k}}$.

The proof is somewhat different than the proof of Lemma 12, due to the fact that if B is very large, say $\frac{\Delta}{\log \Delta}$, then $|T_{C,\gamma}|$ is not very highly concentrated. So we must settle for showing that for any C, $|T_{C,\gamma}|$ is too big for only a relatively small number of colours γ.

Set $\tau = \Delta^{\frac{1}{k}} B^{1-\frac{1}{k}} \log^{4k-4}(\Delta/B)$. Note that for each $C \in \mathcal{C}, \gamma \in L_C$, $\mathbf{Exp}(|T_{C,\gamma}|) \leq \tau$. For each clique $C \in \mathcal{C}$, we define Bad_C to be the set of colours $\gamma \in L_C$ such that $|T(C,\gamma)| > 2\tau$.

For each vertex v, we define E_v^1 to be the events that v violates condition (a) of our Lemma, and we define E_v^2 to be the event that $|L_v \cap (\cap_{C \in \mathcal{C}_v} \mathrm{Save}_C)| \geq \frac{1}{2} B \log^{4k}(\Delta/B)$. Finally, for each $C \in \mathcal{C}_v$ we define $E_{C,v}^3$ to be the event that $|\mathrm{Bad}_C \cap L_v| > B$.

We will show that with positive probability, none of these events hold. In this case, by deleting Bad_C from Save_C for each $C \in \mathcal{C}$, we satisfy the conditions of our Lemma.

It follows in the same manner as in the proof of Lemma 12 that $\mathbf{Pr}(E_v^1), \mathbf{Pr}(E_v^2) < \Delta^{-\ln \Delta}$. Now consider any B colours, $\gamma_1, ..., \gamma_B$. We will first bound the probability that all of these colours lie in Bad_C. To do so, we will apply our Concentration Lemma II.

Set $R = \sum_{i=1}^{B} |T_{C,\gamma_i}|$. Clearly $\mathbf{Exp}(R) = B\tau$. We will show that R is highly concentrated. To do so, observe that after choosing for each colour γ and clique $F \in \mathcal{C}$ whether γ goes into Save_F, we can find R triples (γ_i, F, u) where $F \neq C$, $\gamma_i \in \mathrm{Save}_F$ and $u \in C \cap F$. We define $b_{\gamma_i,F}$ to be the number of such triples the pair γ_i, F belongs to. Clearly, $\sum b_{\gamma_i,F} = R$, and $b_{\gamma_i,F} \leq B$. Therefore, by Concentration Lemma II,

$$\begin{aligned} \mathbf{Pr}(R - \mathbf{Exp}(R) > B\tau) &< 4\mathrm{e}^{-\frac{(B\tau)^2}{5B(B\tau)}} \\ &< \mathrm{e}^{-B\log^4(\Delta/B)}. \end{aligned}$$

Clearly then, the probability that $\gamma_1, ..., \gamma_B \in \mathrm{Bad}_C$ is at most $\mathbf{Pr}(R > 2B\tau) < \mathrm{e}^{-B\log^4(\Delta/B)}$. Therefore, the probability that $|L_v \cap \mathrm{Bad}_C| \geq B$ is at most

$$\binom{|L_v|}{B} \mathrm{e}^{-B\log^4(\Delta/B)} < \left(\frac{\mathrm{e}\Delta}{B}\right)^B \left(\frac{B}{\Delta}\right)^{B\log^3(\Delta/B)} < \left(\frac{B}{\Delta}\right)^{B\log^2(\Delta/B)},$$

which the reader can easily verify to be less than $\Delta^{-\log \Delta}$.

As in the proof of Lemma 12, each event is mutually independent of all but $O(\Delta^4)$ other events and so by the Local Lemma, with positive probability none of the events hold. □

We carry out the iterations of our procedure in the same manner as before, with only one change which will be specified later, and we keep track of the same parameters. We will continue until either L_i or S_i falls below $B\log^7(\Delta/B)$.

Our analogue of Lemma 4 is

Lemma 14 *If $L_i \geq S_i \geq \frac{1}{2}L_i \geq B\log^7(\Delta/B)$ then for each vertex v, clique $C \in \mathcal{C}$, subset $T \subseteq C$, and colour $\gamma \in L_C$:*

> *(a)* $\mathbf{Exp}(|U_T^{'}|) = |U_T|(1 - \mathrm{Retain}_i)$;
>
> *(b)* $\mathbf{Exp}(|L_v^{'}|) \geq |L_v| \times \mathrm{Keep}_i^k - \mathrm{O}(B)$;
>
> *(c)* $\mathbf{Exp}(n(C,\gamma)) \leq n(C,\gamma) \times \mathrm{Keep}_i^{k-1} \times (1 - \mathrm{Retain}_i) + \mathrm{O}(B)$.

The proof of Lemma 14 is virtually identical to that of Lemma 4, the only notable difference being that all $\mathrm{O}(\frac{1}{L_i}), \mathrm{O}(\frac{1}{S_i})$ terms become $\mathrm{O}(B/L_i), \mathrm{O}(B/S_i)$. We omit the details.

The only Lemma which must change significantly is Lemma 5. The difficulty here is that part (c) no longer holds for very large values of B, as $n(C,\gamma)$ is not neccessarily concentrated for every C,γ. Fortunately, we can apply the technique that we used in the proof of Lemma 13, and show that with high probability for each C, $n(C,\gamma)$ will be very close to its expected value for all but a small number of colours γ.

Set $\eta = S_i \times \mathrm{Keep}_i^{k-1} \times (1 - \mathrm{Retain}_i)$ and recall that for all $C \in \mathcal{C}, \gamma \in L_C$, we have $\mathbf{Exp}(n(C,\gamma)) \leq \eta$. For each clique $C \in \mathcal{C}$, we define Bad_C to be the set of colours $\gamma \in L_C$ such that $n(C,\gamma) > \eta + (BS_i)^{1/2}\log^2(\Delta/B)$.

Lemma 15 *For every vertex v, clique $C \in \mathcal{C}_v$, subset $T \subseteq C$, such that $|L_v|, |U_T| \geq B\log^7(\Delta/B)$, we have*

> *(a)* $\mathbf{Pr}(||U_T^{'}| - \mathbf{Exp}(|U_T^{'}|)| > (|U_T|^{1/2}\log^2\Delta) < \Delta^{-\log\Delta}$;
>
> *(b)* $\mathbf{Pr}(||L_v^{'}| - \mathbf{Exp}(|L_v^{'}|)| > L_i^{1/2}\log^2\Delta) < \Delta^{-\log\Delta}$;
>
> *(c)* $\mathbf{Pr}(|\mathrm{Bad}_C \cap L_v| > L_i^{1/2}\log^2\Delta) < \Delta^{-\log\Delta}$.

We will prove Lemma 15 at the end of this section. Only the proof of part (c) is different than the proof of Lemma 5.

At the end of each iteration, for each clique $C \in \mathcal{C}$, we delete Bad_C from L_C and from L_v for each $v \in C$. This is the only change that we make in our procedure. Note that if the property implicit in Lemma 15(c) holds for every vertex v, then the number of colours deleted from each L_v is negligible for our purposes.

As before, a straightforward application of the Local Lemma yields:

Lemma 16 *Define $U_0 = S_0 = \Delta, L_0 = \Delta + \Delta^{1-\frac{1}{k}}B^{1/k}\log^4(\Delta/B), T_0 = 2\Delta^{\frac{1}{k}}B^{1-\frac{1}{k}}\log^{4k-4}(\Delta/B)$ and recursively define*

- $U_{i+1} = U_i(1 - \mathrm{Retain}_i) + \sqrt{U_iB}\log^2(\Delta/B)$;
- $L_{i+1} = L_i \times \mathrm{Keep}_i^k - 2k\sqrt{L_iB}\log^2(\Delta/B)$;
- $S_{i+1} = S_i \times \mathrm{Keep}_i^{k-1} \times (1 - \mathrm{Retain}_i) + \sqrt{S_iB}\log^2(\Delta/B)$;

- $T_{i+1} = T_i(1 - \text{Retain}_i) + \sqrt{BT_i}\log^2(\Delta/B)$.

With positive probability, for every i such that $U_i, L_i, S_i, T_i \geq B\log^7(\Delta/B)$ and such that $L_j > S_j$ for each $j \leq i$, and for each vertex v, clique $C \in \mathcal{C}$ and colour γ, at the beginning of iteration i we have:

(a) $|U_C| \leq U_i$;

(b) $|L_v| \geq L_i$;

(c) $|S_{C,\gamma}| \leq S_i$;

(d) $|T_{C,\gamma}| \leq T_i$.

The rest of the proof of Theorem 2 now follows along the same lines as that of Theorem 1.

Proof of Theorem 2: We carry out our procedure until iteration $i_0 = \min\{i : L_i < B\log^7(\Delta/B)\}$. It follows as in the proof of Lemma 7 that $T_{i_0} < B\log^{4k-\frac{1}{2}}(\Delta/B)$. The details are very similar, and we omit them. Thus we can complete the colouring using the lists Save_v. □

All that remains is to prove Lemma 15, which we do now.

Proof of Lemma 15: Parts (a,b) follow exactly as in the corresponding parts of Lemma 5 (see the remark following that proof).

The proof of part (c) follows along the same lines as the similar part of the proof of Lemma 13. We first bound the probability that $\sqrt{B}$ particular colours $\gamma_1, ..., \gamma_{\sqrt{B}}$ all lie in Bad_C.

Set $R = \sum_{i=1}^{\sqrt{B}} n(C, \gamma_i)$. We apply our Concentration Lemma in the same manner as in the proof of Lemma 5(c), the only difference being that we must now take $\beta = 2k^2B$. We obtain that

$$\mathbf{Pr}(|R - \mathbf{Exp}(R)| > t) \leq \mathrm{e}^{-\theta(\frac{t^2}{BS_i})}.$$

Noting that $\mathbf{Exp}(R) \leq \sqrt{B}\eta$, and setting $t = \sqrt{B}(\sqrt{BS_i}\log^2(\Delta/B))$, yields that the probability of $\gamma_1, .., \gamma_{\sqrt{B}} \in \text{Bad}_C$ is at most $\mathrm{e}^{-\theta(B\log^2(\Delta/B))}$. Therefore,

$$\mathbf{Pr}(|L_v \cap \text{Bad}_C| \geq \sqrt{B}) < \binom{L_i}{\sqrt{B}}\mathrm{e}^{-\theta(B\log^2(\Delta/B))},$$

which can be shown to be at most $\Delta^{-\log\Delta}$ as in the proof of Lemma 13. □

7 An Algorithm

Beck[6] developed a polynomial time algorithm for two specific applications of the Local Lemma. Alon[2], amongst other things, showed that several known applications of the Local Lemma can be looked upon as special cases of the applications addressed by Beck, and thus also yield polytime algorithms. In [16], Molloy and Reed gave a Local Lemma-like set of conditions which if satisified not only guarantee the existence of the desired structure but also ensure that Beck's technique will produce a polynomial time (deterministic) algorithm for finding such a structure. Furthermore, they developed polytime algorithms for some problems to which Beck's technique did not apply, and put their techniques in a general setting which seems to apply to most of the known applications of the Local Lemma.

The following is a special case of the first result from [16] mentioned above, and it is straightforward to verify that it follows from that result.

Lemma 17 (Local Lemma Like Lemma) *Let F be a hypergraph such that each edge of F intersects at most d other edges. Suppose that each vertex v of F has a set $L(v)$ of permissable colours. We consider the set of satisfactory colourings of F according to those lists. For each edge e, we are explicitly given a boolean function* Bad *on the set of colourings of $V(e)$. Suppose that if we assign to each vertex v a colour uniformily from L_v then for each e, we have* $\mathbf{Pr}(\mathrm{Bad}(e) = 1) < \frac{1}{4d^9}$. *Then there exists a colouring such that* $\mathrm{Bad}(e) = 0$ *for every edge e, and furthermore we can find such a colouring in polynomial time.*

Note that the existence of the colouring follows easily from the Local Lemma. The interesting part of our new lemma is that the colouring can be found efficiently. This Lemma can be applied to every application of the Local Lemma in this paper, thus yielding the following:

Theorem 18 *For every k there is a c_k and a polynomial time algorithm which given a k-uniform linear hypergraph H of maximum degree Δ finds a $\Delta + c_k(\Delta^{1-\frac{1}{k}}(\log\ \Delta)^4)$ colouring of the edges of H.*

Proof We let H be a such a hypergraph. We can assume that $\Delta(H)$ is at least Δ_k as otherwise we can simply colour $E(H)$ greedily. Now, consider an arbitrary iteration of our colouring procedure.

We define the vertex set of F to be the set of uncoloured vertices of $G = L(H)$, and we create an edge of F for each bad event in our proof of Lemma 6, which contains all of the vertices whose colours determine whether that event holds. For example, for each clique $C \in \mathcal{C}$ and subset $T \subset C$ with $|U_T| > \log^7 \Delta$, we create an edge containing all the uncoloured vertices in $\cup_{u \in U_T} \cup_{C' \in C_v} C'$. For any edge e, we define $\mathrm{Bad}(e)$ to be 1 if the corresponding bad event holds, and 0 otherwise.

By construction, $\mathrm{Bad}(e)$ depends only on the colouring of the vertices of e. Furthermore, it is implicit in our proof of Lemma 6 that each edge of F intersects at most $O(\Delta^5)$ other edges. By Lemma 5, for each e the probability that $\mathrm{Bad}(e) = 1$ is at most $\Delta^{-\log \Delta}$. Therefore, our Local Lemma Like Lemma applies to produce an efficient algorithm for finding a colouring as guaranteed by Lemma 6, i.e. a colouring for which $|U_C^{'}| \leq \mathbf{Exp}(|U_C^{'}|)^{\frac{1}{2}} \log^2 \Delta$, etc.

The Local Lemma Like Lemma can be applied to make the proof of Lemma 12 constructive in the same manner. We omit the straightforward details.

All other aspects of our procedure are implicitly constructive. The only other step is to show that the proof of the List Colouring Lemma can be made constructive; unfortunately the Local Lemma Like Lemma does not apply here. There are two ways in which we can overcome this problem. The first is to use a more complicated argument based on the new techniques introduced in [16]. The details are rather involved, but they are relatively straightforward, and they follow along the general framework provided in [16]; for this reason, we omit them. The second (simpler) approach is to use a weaker version of the List Colouring Lemma, in which the constant "2e" is replaced by a much larger constant. This weaker version, which is clearly sufficient for our purposes, follows from the main result of [3]. That result is an application of the Local Lemma which can be made constructive using Lemma 17, thus completing the proof of Theorem 18. □

In the same way, we can obtain:

Theorem 19 *For every k there is a c_k and a polynomial time algorithm which given a k-uniform linear hypergraph H of maximum degree Δ and maximum codegree B finds a $\Delta + c_k(B^{\frac{1}{k}}\Delta^{1-\frac{1}{k}}(\log \frac{\Delta}{B})^6)$ colouring of the edges of H.*

8 Talagrand's Inequality

Finally, we prove our Concentration Lemmas from Section 4. As mentioned there, this is a straightforward corollary of a recent, very powerful concentration inequality due to Talagrand. Our corollary is very similar in flavour to many other applications of Talagrand's Inequality, see for example, [22] or [21].

First we present Talagrand's Inequality.

Let $X_1, ..., X_m$ be independent random variables. For any subset A of the possible outcomes, and any real ℓ, we define A_ℓ to be the set of possible outcomes which is within a distance ℓ of A with regards to an unusual measure. In particular, we say that $x = (x_1, ..., x_m) \in A_\ell$ iff for any real $b_1, ..., b_m$, there exists at least one point $y = (y_1, ..., y_m) \in A$ such that $\sum_{x_i \neq y_i} b_i < \ell \left(\sum_{i=1}^m b_i^2\right)^{1/2}$.

Talagrand's Inequality:

$$\mathbf{Pr}(A)(1 - \mathbf{Pr}(A_\ell)) \leq \mathrm{e}^{-\ell^2/4}.$$

Proof of Concentration Lemma: First, we will show that R is strongly concentrated near its median.

Define $A = \{y : R(y) \geq \mathbf{Med}(R)\}$, and $C = \{x : R(x) < \mathbf{Med}(R) - t\}$. We will prove that $C \subseteq \overline{A_\ell}$, where $\ell = t/\sqrt{\beta D}$. Consider some $x \in A_\ell$, and choose $b_1, ..., b_m$ satisfying the conditions of our Concentration Lemma. Since $x \in A_\ell$, there is some $y \in A$ such that

$$\begin{aligned} \sum_{x_i \neq y_i} b_i &< \ell \left(\sum_{i=1}^{m} b_i^2 \right)^{1/2} \\ &\leq \ell \left(\beta \sum_{i=1}^{m} b_i \right)^{1/2} \\ &\leq \ell \sqrt{\beta D} = t \end{aligned}$$

Therefore, by the conditions of our Concentration Lemma, $R(x) \leq R(y) + t$, and so since $y \in A$, we have $x \notin C$. Therefore $C \subseteq \overline{A_t}$, and so by Talagrand's Inequality.

$$\mathbf{Pr}(C) \leq \frac{1}{\mathbf{Pr}(A)} \mathrm{e}^{-\ell^2/4} = 2\mathrm{e}^{-\frac{t^2}{4\beta D}}.$$

Next, set $A^{'} = \{y : R(y) > \mathbf{Med}(R) + t\}$, and $C^{'} = \{x : R(x) \leq \mathbf{Med}(R)\}$. By a nearly identical argument, we obtain:

$$\mathbf{Pr}(A^{'}) \leq 2\mathrm{e}^{-\frac{t^2}{4\beta D}},$$

and so

11 $\mathbf{Pr}(|R - \mathbf{Med}(R)| > t) \leq 4\mathrm{e}^{-\frac{t^2}{4BD}}$.

To complete our proof, we will show that $|\mathbf{Exp}(R) - \mathbf{Med}(R)| < T^{'}\sqrt{D}$ where $T^{'}$ is a constant depending only on B. Clearly this will suffice.

First, observe that $\mathbf{Exp}(R) - \mathbf{Med}(R) = \mathbf{Exp}(R - \mathbf{Med}(R))$. From (11) we have

$$\begin{aligned} \mathbf{Exp}(R - \mathbf{Med}(R)) &\leq \sum_{i \geq 0} (i+1)\sqrt{D} \times \mathbf{Pr}(R - \mathbf{Med}(R) > i\sqrt{D}) \\ &< \sqrt{D} \times \sum_{i \geq 0} (i+1) 4\mathrm{e}^{-\frac{i^2}{4B}} \\ &= T_1 \sqrt{D} \end{aligned}$$

where T_i is a function of B. A lower bound on $\mathbf{Exp} - \mathbf{Med}(R)$ of $T_1\sqrt{D}$ is established in the same way, and we set $T^{'} = 2T_1$. $\square$

The proof of Concentration Lemma II is nearly identical. We omit the straightforward details.

Acknowledgments

We thank two referees for their helpful comments.

References

[1] M. Ajtai, J, Komlós and E. Szemerédi, *A dense infinite Sidon sequence*, Eur. J. Comb. **2** (1981), 1 - 11.

[2] N. Alon, *A parallel algorithmic version of the Local Lemma*, Random Structures and Algorithms, **2** (1991), 367 - 378.

[3] N. Alon, *The strong chromatic number of a graph.* Random Structures and Algorithms, **3** (1992), 1 - 7.

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

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

[6] J. Beck, *An algorithmic approach to the Lovász Local Lemma*, Random Structures and Algorithms, **2** (1991), 343 - 365.

[7] B. Bollobás and A. Harris, *List-colourings of graphs*, Graphs and Combinatorics **1** (1985), 115 - 127.

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

[9] 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.

[10] P. Erdős, A. Rubin and H. Taylor, *Choosability in graphs*, Congr. Num. **26** (1979), 125 - 157.

[11] R. Häggkvist and J. Janssen, *New bounds on the list chromatic index of the complete graph and other simple graphs*, Combin., Prob. and Comp. **6** (1997), 273 - 295.

[12] T. Jensen and B. Toft, *Graph Colouring Problems*, Wiley, 1995.

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

[14] A. Kostochka, V. Rödl and L. Talysheva, *On systems of small sets with no large Δ-subsystems*, Comb. Prob. and Comp. **8** (1999), 265 - 268.

[15] C. McDiarmid, *On the method of bounded differences.* Surveys in Combinatorics, Proceedings of the Twelfth British Combinatorial Conference (1989), 148 - 188.

[16] M. Molloy and B. Reed, *Further algorithmic aspects of the Local Lemma*, Proceedings of STOC 1998.

[17] M. Molloy and B. Reed, *Graph colouring via the probabilistic method*, in *Graph Theory and Computational Biology*, A. Gyrafas and L. Lovasz editors, J. Bolyai Math. Soc, 1999, 125 - 155.

[18] N. Pippenger and J. Spencer, *Asymptotic behavior of the chromatic index for hypergraphs*, J. Comb. Th. (A) **51** (1989), 24 - 42.

[19] B. Reed, *The list colouring constants*, J. Graph Th. **31** (1999), 149 - 153.

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

[21] J. Spencer, *Applications of Talagrand's Inequality*, manuscript.

[22] M. Steele, *Probability Theory and Combinatorial Optimization*, SIAM Regional Conference Series in Applied Mathematics (1997).

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

[24] V. Vizing, *Some unsolved problems in graph theory*, Russian Math Surveys **23**, pp. 125 - 141 (1968).

[25] V. Vizing, *Colouring the vertices of a graph with prescribed colours*, Diskret. Analiz. **29** (1976), 3 - 10.