A proof of the Erdős–Faber–Lovász conjecture: Algorithmic aspects

Dong Yeap Kang, Tom Kelly, Daniela Kühn, Abhishek Methuku, Deryk Osthus
School of Mathematics
University of Birmingham
Birmingham B15 2TT, United Kingdom.
Emails: [d.y.kang.1,t.j.kelly,d.kuhn,d.osthus]@bham.ac.uk, abhishekmethuku@gmail.com

***Abstract*—The Erdős-Faber-Lovász conjecture (posed in 1972) states that the chromatic index of any linear hypergraph on n vertices is at most n. Erdős considered this to be one of his three most favorite combinatorial problems and offered a $500 reward for a proof of this conjecture. We prove this conjecture for every large n. Here, we also provide a randomised algorithm to find such a colouring in polynomial time with high probability.**

***Keywords*-Hypergraph colouring; Rödl nibble; Erdős–Faber–Lovász;**

I. Introduction

Graph and hypergraph colouring problems are central to combinatorics, with applications and connections to many other areas, such as geometry, algorithm design, and information theory. A *proper edge-colouring* of a hypergraph $\mathcal{H}$ is an assignment of colours to the edges of $\mathcal{H}$ such that no two edges of the same colour share a vertex, and a *proper vertex-colouring* of $\mathcal{H}$ is an assignment of colours to the vertices of $\mathcal{H}$ such that every edge contains vertices of at least two different colours. The *chromatic index* of $\mathcal{H}$, denoted $\chi'(\mathcal{H})$, is the minimum number of colours used in a proper edge-colouring of $\mathcal{H}$, and the *chromatic number*, denoted $\chi(\mathcal{H})$, is the minimum number of colours used in a proper vertex-colouring of $\mathcal{H}$. For a graph G, Vizing's theorem [55] implies that $\chi'(G) \in \{\Delta(G), \Delta(G)+1\}$, where $\Delta(G)$ is the maximum degree of a vertex in G. A graph G is of *Class 1* if $\chi'(G) = \Delta(G)$ and it is of *Class 2* if $\chi'(G) = \Delta(G)+1$. Thus Vizing's theorem shows every graph is of Class 1 or Class 2. However, Holyer [23] showed that it is actually NP-hard to decide whether a graph is of Class 1 or Class 2 (and thus determining the chromatic index of a linear hypergraph is NP-hard as well. Here a hypergraph $\mathcal{H}$ is *linear* if every two distinct edges of $\mathcal{H}$ intersect in at most one vertex). In fact, Leven and Galil [42] showed that it is NP-hard to determine whether a graph is of Class 1 or Class 2 even if we restrict ourselves to d-regular graphs (i.e., graphs with all the degrees equal to d). However, Ferber and Jain [18] showed that d-regular graphs on an even number of vertices which are 'good' spectral expanders are of Class 1, and their proof yields a polynomial-time randomised colouring algorithm.

For vertex colourings the situation is similar – for a graph G it is also NP-hard to determine $\chi(G)$ exactly [37], [21]. In fact, for sufficiently large q, even colouring a q-colourable graph with $2^{\Omega(q^{1/3})}$ colours is NP-hard [24]. For $k \geq 3$, it is already NP-hard to determine whether a k-uniform hypergraph has a 2-coloring. See [7], [25], [49] for more results on the hardness of edge colouring and vertex colouring hypergraphs.

Thus it is natural to seek general bounds (and algorithms which attain these bounds) for these parameters. This is still challenging in particular for hypergraphs: for example, for $k \geq 3$ it is even non-trivial to determine the chromatic index of a complete k-uniform hypergraph $K_n^{(k)}$, a problem famously resolved by Baranyai's theorem [5].

A. The Erdős-Faber-Lovász conjecture

In 1972, Erdős, Faber, and Lovász conjectured (see [15]) the following equivalent bounds on colouring set systems, graphs, and hypergraphs. Let $n \in \mathbb{N}$.

(i) If $A_1, \dots, A_n$ are sets of size n such that every pair of them shares at most one element, then the elements of $\bigcup_{i=1}^n A_i$ can be coloured by n colours so that all colours appear in each A_i.

(ii) If G is a graph that is the union of n cliques, each having at most n vertices, such that every pair of cliques shares at most one vertex, then the chromatic number of G is at most n.

(iii) If $\mathcal{H}$ is a linear hypergraph with n vertices, then the chromatic index of $\mathcal{H}$ is at most n.

The formulation (iii) is the one that we will consider throughout the paper. For simplicity, we will refer to this conjecture as the EFL conjecture.

Erdős considered this to be 'one of his three most favorite combinatorial problems' (see e.g., [32]). The simplicity and elegance of its formulation initially led the authors to believe it to be easily solved (see e.g., the discussion in [9] and [15]). It was initially designed as a simple test case for a more general theory of hypergraph colourings. However, as the difficulty became apparent Erdős offered successively increasing rewards for a proof of the conjecture, which eventually reached $500.

Previous progress towards the conjecture includes the following results. Seymour [54] proved that every n-vertex linear hypergraph $\mathcal{H}$ has a matching of size at least $e(\mathcal{H})/n$, where $e(\mathcal{H})$ is the number of edges in $\mathcal{H}$. (Note that this

2575-8454/21/$31.00 ©2021 IEEE
DOI 10.1109/FOCS52979.2021.00107

immediately follows from the validity of the EFL conjecture, but it is already difficult to prove.) Kahn and Seymour [33] proved that every n-vertex linear hypergraph has fractional chromatic index at most n. Chang and Lawler [8] showed that every n-vertex linear hypergraph has chromatic index at most $\lceil 3n/2 - 2 \rceil$. A breakthrough of Kahn [28] yielded an approximate version of the conjecture, by showing that every n-vertex linear hypergraph has chromatic index at most $n + o(n)$. Recently Faber and Harris [17] proved the conjecture for linear hypergraphs whose edge sizes range between 3 and $cn^{1/2}$ for a small absolute constant $c > 0$. More background and earlier developments related to the EFL conjecture are detailed in the surveys of Kahn [29], [32]. See also the recent survey by the authors [34].

B. Main result

We prove that the EFL conjecture is true for every large n. Our proof also yields a randomised colouring algorithm. A full proof of this result can be found in [35]. In this extended abstract, we show how the proof of [35] can be modified to obtain a randomised polynomial-time algorithm.

Theorem I.1. *For every sufficiently large n, every linear hypergraph $\mathcal{H}$ on n vertices has chromatic index at most n. Moreover, there is a randomised polynomial-time algorithm which finds a proper edge n-colouring of $\mathcal{H}$ with high probability.*

There are three constructions for which Theorem I.1 is known to be tight: a complete graph K_n for any odd integer n (and minor modifications thereof), a finite projective plane of order k on $n = k^2 + k + 1$ points, and a degenerate plane $\{\{1, 2\}, \dots, \{1, n\}, \{2, \dots, n\}\}$. Note that the first example has bounded edge size (two), while the other two examples have unbounded edge size as n tends to infinity.

Kahn's proof [28] is based on a powerful method known as the Rödl nibble. Roughly speaking, this method builds a large matching using an iterative probabilistic procedure. It was originally developed by Rödl [51] to prove the Erdős–Hanani conjecture [14] on combinatorial designs. Another famous result based on this method is the Pippenger–Spencer theorem [50], which implies that the chromatic index of any uniform hypergraph $\mathcal{H}$ of maximum degree D and codegree $o(D)$ is $D+o(D)$. (Note that this in turn implies that the EFL conjecture holds for all large r-uniform linear hypergraphs of bounded uniformity $r \geq 3$.) In a seminal paper, Kahn [30] later developed the approach further to show that the same bound $D + o(D)$ even holds for the list chromatic index (an intermediate result in this direction, which also strengthens the Pippenger–Spencer theorem, was the main ingredient of his proof in [28]). The best bound on the $o(D)$ error term for the list chromatic index of such hypergraphs was obtained by Molloy and Reed [46], and for the chromatic index, the best bound was proved in [36]. Our proof will also rely on certain properties of the Rödl nibble.

In addition, our proof makes use of powerful colouring results for locally sparse graphs. This line of research goes back to Ajtai, Komlós, and Szemerédi [2] who (preceding Rödl [51]) developed a very similar semi-random nibble approach to give an upper bound $O(k^2/\log k)$ on the Ramsey number $R(3, k)$ by finding large independent sets in triangle-free graphs (the matching lower bound $R(3, k) = \Omega(k^2/\log k)$ was later established by Kim [39], also using a semi-random approach). The above Ramsey bound by Ajtai, Komlós and Szemerédi was subsequently strengthened by a highly influential result of Johansson [27], who showed that triangle-free graphs of maximum degree Δ have chromatic number $O(\Delta/\log \Delta)$ (and a related result was proved independently by Kim [38]). There are many generalisations and analogues of Johansson's theorem, in particular Frieze and Mubayi [19] proved a version of this for hypergraphs. It also turns out that the condition of being triangle-free can be relaxed (in various ways) to being 'locally sparse' [1], [3], [10], [56]. We will be able to apply such results to suitable parts of the line graph of our given linear hypergraph $\mathcal{H}$.

One step in our proof involves what may be considered a 'vertex absorption' argument; here certain vertices not covered by a matching produced by the Rödl nibble are 'absorbed' into the matching to form a colour class. (Vertex) absorption as a systematic approach was introduced by Rödl, Ruciński, and Szemerédi [52] to find spanning structures in hypergraphs (with precursors including [16], [40]). Absorption ideas were first used for edge decomposition problems in [41] to solve Kelly's conjecture on tournament decompositions. We will make use of an application of the main result of [41] to the overfull subgraph conjecture (which was derived in [22]).

II. Overview of the proof of Theorem I.1

In this section, we provide an overview of the proof of Theorem I.1. A more detailed overview for Sections II-A and II-B is given in the recent survey of the authors [34].

A. Colouring linear hypergraphs with bounded edge sizes

Here, we discuss the proof of Theorem I.1 in the special case when all edges of $\mathcal{H}$ have bounded size. In this subsection, we fix constants satisfying the hierarchy

$$0 < 1/n_0 \ll \xi \ll 1/r \ll \gamma \ll \varepsilon \ll \rho \ll 1,$$

we let $n \geq n_0$, and we let $\mathcal{H}$ be an n-vertex linear hypergraph such that every $e \in \mathcal{H}$ satisfies $2 \leq |e| \leq r$. We first describe the ideas which already lead to the near-optimal bound $\chi'(\mathcal{H}) \leq n + 1$.

Let G be the graph with $V(G) \coloneqq V(\mathcal{H})$ and $E(G) \coloneqq \{e \in \mathcal{H} : |e| = 2\}$. The first step of the proof is to include every edge of G in a 'reservoir' R independently with probability $1/2$ that we will use for 'absorption'. With high probability, each $v \in V(\mathcal{H})$ satisfies $d_R(v) = d_G(v)/2 \pm \xi n$. Since $\mathcal{H}$ is linear, this easily implies that $\Delta(\mathcal{H} \setminus R) \leq$

$(1/2+\xi)n$. So by the Pippenger–Spencer theorem [50], we obtain the nearly optimal bound $\chi'(\mathcal{H} \setminus R) \le (1/2+\gamma)n$. Now using R as a 'vertex-absorber', we would like to extend the colour classes of $\mathcal{H} \setminus R$ to cover as many vertices of U as possible, where $U := \{u \in V(\mathcal{H}) : d_G(u) \ge (1-\varepsilon)n\}$. This would allow us to control the maximum degree in the hypergraph consisting of uncoloured edges, so that it can then be coloured with few colours. To that end, we need the following important definition.

Definition II.1 (Perfect and nearly-perfect coverage)**.** Let $\mathcal{H}$ be a linear multi-hypergraph, let $\mathcal{N}$ be a set of edge-disjoint matchings in $\mathcal{H}$, and let $S \subseteq U \subseteq V(\mathcal{H})$.

- We say $\mathcal{N}$ has *perfect coverage* of U if each $N \in \mathcal{N}$ covers U.
- We say $\mathcal{N}$ has *nearly-perfect coverage of U with defects in S* if
  - (i) each $u \in U$ is covered by at least $|\mathcal{N}| - 1$ matchings in $\mathcal{N}$ and
  - (ii) each $N \in \mathcal{N}$ covers all but at most one vertex in U such that $U \setminus V(N) \subseteq S$.

We will construct some $\mathcal{H}' \subseteq \mathcal{H}$ and a proper edge-colouring $\psi : \mathcal{H}' \to C$ such that $\mathcal{H}' \supseteq \mathcal{H} \setminus R$, $|C| = (1/2+\gamma)n$ and the set of colour classes $\{\psi^{-1}(c) : c \in C\}$ has nearly-perfect coverage of U (with defects in U). Crucially, this means that $\mathcal{H} \setminus \mathcal{H}'$ is a graph and satisfies $\Delta(\mathcal{H} \setminus \mathcal{H}') \le n - |C|$. (Indeed, every vertex $u \in U$ satisfies $d_{\mathcal{H}}(u) \le n-1$ and is covered by all but at most one of the colour classes of ψ, and every vertex $v \notin U$ satisfies $d_{\mathcal{H}\setminus\mathcal{H}'}(v) \le d_R(v) \le ((1-\varepsilon)/2+\xi)n < n - |C|$.) Therefore, Vizing's theorem [55] implies that $\chi'(\mathcal{H} \setminus \mathcal{H}') \le \Delta(\mathcal{H} \setminus \mathcal{H}') + 1 \le n - |C| + 1$, so altogether we have $\chi'(\mathcal{H}) \le \chi'(\mathcal{H}') + \chi'(\mathcal{H} \setminus \mathcal{H}') \le n+1$, as claimed.

To construct $\mathcal{H}'$ and ψ we iteratively apply the Rödl nibble to (the leftover of) $\mathcal{H} \setminus R$ to successively construct large matchings N_i which are then removed from $\mathcal{H} \setminus R$ and form part of the colour classes of ψ. (The Rödl nibble is applied implicitly via [35, Corollary 4.3], which guarantees a large matching in a suitable hypergraph.) Crucially, each matching N_i exhibits pseudorandom properties, which allow us to use some edges of R to extend N_i into a matching M_i (which will form a colour class of ψ) with nearly-perfect coverage of U, as desired. (This is why we apply the Rödl nibble in our proof rather than the Pippenger-Spencer theorem.) Thus, R acts as a 'vertex-absorber' for $U \setminus V(N_i)$ and the final edge decomposition of the unused edges of R into matchings is achieved by Vizing's theorem. (Actually, this only works if $\mathcal{H} \setminus R$ is nearly regular, which is not necessarily the case. Thus, we first embed $\mathcal{H} \setminus R$ in a suitable nearly regular hypergraph $\mathcal{H}^*$ and prove that the respective matchings in $\mathcal{H}^*$ have nearly-perfect coverage of U, which suffices for our purposes.)

Let us now discuss how to improve the bound $\chi'(\mathcal{H}) \le n+1$ to $\chi'(\mathcal{H}) \le n$. Let $S := \{u \in U : d_G(u) < n-1\}$, and note that if $\{\psi^{-1}(c) : c \in C\}$ has either perfect coverage of U, or nearly-perfect coverage of U with defects in S, then $\Delta(\mathcal{H} \setminus \mathcal{H}') \le n - 1 - |C|$. In this case, we may use the same argument as before with Vizing's theorem to obtain $\chi'(\mathcal{H}) \le n$. However, it is not always possible to find such a colouring. For example, if $\mathcal{H}$ is a complete graph K_n for odd n (which is one of the extremal examples for Theorem I.1), then $U = V(\mathcal{H})$ and $S = \varnothing$, so it is not possible for even a single colour class to have nearly-perfect coverage of U with defects in S. However, we can adapt the above nibble-absorption-Vizing approach to work whenever $\mathcal{H}$ is not 'close' to K_n in the following sense.

Definition II.2 ((ρ,ε)-full)**.** Let $\mathcal{H}$ be an n-vertex linear hypergraph, and let G be the graph with $V(G) := V(\mathcal{H})$ and $E(G) := \{e \in \mathcal{H} : |e| = 2\}$. For $\varepsilon, \rho \in (0,1)$, $\mathcal{H}$ is *(ρ,ε)-full* if

- $|\{u \in V(\mathcal{H}) : d_G(u) \ge (1-\varepsilon)n\}| \ge (1-10\varepsilon)n$, and
- $|\{v \in V(\mathcal{H}) : d_G(v) = n-1\}| \ge (\rho - 15\varepsilon)n$.

As mentioned above, when $\mathcal{H}$ is not (ρ,ε)-full we can adapt the nibble-absorption-Vizing approach to show that $\chi'(\mathcal{H}) \le n$ (with a reservoir of density ρ rather than $1/2$). If $\mathcal{H}$ is (ρ,ε)-full then we will ensure that the leftover $\mathcal{H} \setminus \mathcal{H}' \subseteq R$ is a quasirandom almost regular graph (which involves a more careful choice of R – again it will have density close to ρ rather than $1/2$ but now it consists of a 'random' part and a 'regularising' part). This allows us to apply a result [22] on the overfull subgraph conjecture which implies that $\chi'(\mathcal{H} \setminus \mathcal{H}') \le \Delta(\mathcal{H} \setminus \mathcal{H}')$. (The result in [22] is obtained as a straightforward consequence of the result in [41] that robustly expanding regular graphs have a Hamilton decomposition, and thus, a 1-factorisation if they have even order.)

B. Colouring linear hypergraphs where all edges are large

Now we discuss how to prove Theorems I.1 when all edges of $\mathcal{H}$ have size at least some large constant. In this step it is often very useful to consider the line graph $L(\mathcal{H})$ of $\mathcal{H}$ and use the fact that $\chi(L(\mathcal{H})) = \chi'(\mathcal{H})$. In this subsection, we fix constants satisfying the hierarchy

$$0 < 1/n_0 \ll 1/r \ll \sigma \ll \delta,$$

we let $n \ge n_0$, and we let $\mathcal{H}$ be an n-vertex linear hypergraph such that every $e \in \mathcal{H}$ satisfies $|e| > r$. Now we sketch a proof that $\chi'(\mathcal{H}) \le n$ for such $\mathcal{H}$. If $\mathcal{H}$ is a finite projective plane of order k, where $k^2+k+1 = n$, then the line graph $L(\mathcal{H})$ is a clique K_n. Thus, $\chi'(\mathcal{H}) = \chi(L(\mathcal{H})) = n$, so the bound $\chi'(\mathcal{H}) \le n$ is best possible. Thus, we refer to the case where $\mathcal{H}$ has approximately n edges of size $(1\pm\delta)\sqrt{n}$ as the 'FPP-extremal' case. We also sketch how to prove the improved bound $\chi'(\mathcal{H}) \le (1-\sigma)n$ if $\mathcal{H}$ is not in the FPP-extremal case. As we discuss in the next subsection, we will need this result in the proof of Theorem I.1.

Consider an ordering $\preceq$ of the edges $e_1, e_2, \dots, e_m$ of $\mathcal{H}$ according to their size, i.e., $e_i \preceq e_j$ if $|e_i| > |e_j|$ for every $i, j \in [m]$. For an edge $e \in \mathcal{H}$, let $d^{\preceq}_{\mathcal{H}}(e)$ denote the number of edges in $\mathcal{H}$ which intersect e and precede e in $\preceq$. Clearly, a greedy colouring following this size-monotone ordering achieves a bound of $\chi'(\mathcal{H}) \leq \max_i d^{\preceq}_{\mathcal{H}}(e_i) + 1$ (this bound was also used in [8], [28]). Moreover, it is easy to see that if this greedy colouring algorithm fails to produce a colouring with at most $(1-\sigma)n$ colours, i.e., if an edge e satisfies $d^{\preceq}_{\mathcal{H}}(e) \geq (1-\sigma)n$, then almost all of the corresponding edges that intersect e and precede e must have size close to $|e|$.

Surprisingly, if one allows some flexibility in the ordering (in particular, if we allow it to be size-monotone only up to some edge e^* such that $d^{\preceq}_{\mathcal{H}}(e^*) \geq (1-\sigma)n$ while every edge f with $e^* \preceq f$ satisfies $d^{\preceq}_{\mathcal{H}}(f) < (1-\sigma)n$), then one can show much more: Either we can modify the ordering to reduce the number of edges which come before e^*, or there is a set $W \subseteq \mathcal{H}$ (where e^* is the last edge of W) such that

(W1) $|e^*| \approx |e|$ for every $e \in W$, and

(W2) the edges of W cover almost all pairs of vertices of $\mathcal{H}$.

If $|e^*| \leq (1-\delta)\sqrt{n}$, then one can show that $L(W)$ induces a 'locally sparse' graph (as $\mathcal{H}$ is linear). Moreover, (W1) implies that the maximum degree of $L(W)$ is not too large, and thus one can show that $\chi(L(W))$ is much smaller than $(1-\sigma)n$ (leaving enough room to colour the edges preceding W with a new set of colours). This together with (W2) allows us to extend the colouring of W to all of $\mathcal{H}$ using a suitable modification of the above greedy colouring procedure for the remaining edges in $\mathcal{H}$ to obtain that $\chi'(\mathcal{H}) \leq (1-\sigma)n$, as desired.

If $|e^*| \geq (1-\delta)\sqrt{n}$, then we first colour the edges of size at least $(1-\delta)\sqrt{n}$ (in particular, the edges of W) as follows. Let $\mathcal{H}' \subseteq \mathcal{H}$ be the hypergraph consisting of these edges. If $e(\mathcal{H}') \leq n$, then, of course, we may colour the edges of $\mathcal{H}'$ with different colours. Otherwise, if $t := e(\mathcal{H}') - n > 0$, the main idea is to find a matching of size t in the complement of $L(\mathcal{H}')$ (where $L(\mathcal{H}')$ will be close to being a clique of order not much more than n). By assigning the same colour to the edges of $\mathcal{H}'$ that are adjacent in this matching, we obtain $\chi'(\mathcal{H}') = \chi(L(\mathcal{H}')) \leq n$. Now we extend the colouring to all of $\mathcal{H}$ using a suitable modification of the above greedy colouring procedure again to obtain that $\chi'(\mathcal{H}) \leq n$, as desired.

C. Combining colourings of the large and small edges

We now describe how one can prove Theorem I.1 by building on the ideas described in Sections II-A and II-B. In this subsection and throughout the rest of the paper we work with constants satisfying the following hierarchy:

$$0 < 1/n_0 \ll 1/r_0 \ll \xi \ll 1/r_1 \ll \beta \ll \kappa \ll \gamma_1 \ll \varepsilon_1 \ll \rho_1 \ll \sigma \ll \delta \ll \gamma_2 \ll \rho_2 \ll \varepsilon_2 \ll 1. \tag{II.1}$$

Some of these constants are used to characterize the edges of a hypergraph by their size, as follows.

Definition II.3 (Edge sizes)**.** Let $\mathcal{H}$ be an n-vertex linear hypergraph with $n \geq n_0$.

- Let $\mathcal{H}_{\mathrm{small}} := \{e \in \mathcal{H} : |e| \leq r_1\}$. An edge $e \in \mathcal{H}$ is *small* if $e \in \mathcal{H}_{\mathrm{small}}$.
- Let $\mathcal{H}_{\mathrm{med}} := \{e \in \mathcal{H} : r_1 < |e| \leq r_0\}$. An edge $e \in \mathcal{H}$ is *medium* if $e \in \mathcal{H}_{\mathrm{med}}$.
- Let $\mathcal{H}_{\mathrm{large}} := \{e \in \mathcal{H} : |e| > r_0\}$. An edge $e \in \mathcal{H}$ is *large* if $e \in \mathcal{H}_{\mathrm{large}}$.
- Let $\mathcal{H}_{\mathrm{ex}} := \{e \in \mathcal{H} : |e| = (1 \pm \delta)\sqrt{n}\}$. An edge $e \in \mathcal{H}$ is *FPP-extremal* if $e \in \mathcal{H}_{\mathrm{ex}}$.
- Let $\mathcal{H}_{\mathrm{huge}} := \{e \in \mathcal{H} : |e| \geq \beta n/4\}$. An edge $e \in \mathcal{H}$ is *huge* if $e \in \mathcal{H}_{\mathrm{huge}}$.

Note that $\mathcal{H}_{\mathrm{small}}, \mathcal{H}_{\mathrm{med}}, \mathcal{H}_{\mathrm{large}}$ form a partition of the edges of $\mathcal{H}$. Also note that if $\mathcal{H}$ is an n-vertex linear hypergraph and $1/n \ll \alpha < 1$, then

$$|\{e \in \mathcal{H} : |e| \geq \alpha n\}| \leq 2/\alpha. \tag{II.2}$$

In the proof of Theorem I.1, given an n-vertex linear hypergraph $\mathcal{H}$ with $n \geq n_0$ (where we assume $\mathcal{H}$ has no singleton edges), we first find a proper edge-colouring $\psi_1 : \mathcal{H}_{\mathrm{med}} \cup \mathcal{H}_{\mathrm{large}} \to C_1$ as discussed in Section II-B, and then we extend it to a proper n-edge-colouring of $\mathcal{H}_{\mathrm{small}}$ by adapting the argument presented in Section II-A. The proof proceeds slightly differently depending on whether we are in the FPP-extremal case. As discussed in the previous subsection, in the non-FPP-extremal case, $\chi'(\mathcal{H}_{\mathrm{med}} \cup \mathcal{H}_{\mathrm{large}}) \leq (1-\sigma)n$, so we may assume $|C_1| = (1-\sigma)n$. In this case, we let $\gamma := \gamma_1$, $\varepsilon := \varepsilon_1$, and $\rho := \rho_1$; in the FPP-extremal case, we let $\gamma := \gamma_2$, $\rho := \rho_2$, and $\varepsilon := \varepsilon_2$. We define G and U as in Section II-A, and we define a suitable 'defect' set $S \subseteq U$ (whose choice now depends on the structure of $\mathcal{H}$). In order to extend the colouring ψ_1 of $\mathcal{H}_{\mathrm{med}} \cup \mathcal{H}_{\mathrm{large}}$ to $\mathcal{H}$, we need it to satisfy a few additional properties, which are provided by [35, Theorem 6.1]. Roughly, we need that

(1) each colour class of ψ_1 covers at most βn vertices, with exceptions for colour classes containing huge or medium edges, and

(2) at most γn colours are assigned by ψ_1 to colour medium edges.

We choose a 'reservoir' R from $E(G)$; how we choose it depends on whether we are in the FPP-extremal case. In the non-FPP-extremal case, we choose it as described in Section II-A, and in the FPP-extremal case, we include every edge of G incident to a vertex of U to be in R independently with probability ρ.

Let $C_{\mathrm{hm}} \subseteq C_1$ be the set of colours assigned to a huge or medium edge by ψ_1. Note that $e(\mathcal{H}_{\mathrm{huge}}) \leq 8/\beta$ by (II.2), so consequently, by (2), $|C_{\mathrm{hm}}| \leq 3\gamma n/2$. For each $c \in C_{\mathrm{hm}}$, we use [35, Lemma 7.11] to extend $\psi_1^{-1}(c)$ (in the sense of Section II-A) using edges of R, so that $\{\psi_1^{-1}(c) : c \in C_{\mathrm{hm}}\}$

has nearly perfect coverage of U with defects in S. There is possibly an exceptional colour class, which we call *difficult* (see [35, Definition 7.10]), that we need to consider in this step. This situation arises if $\mathcal{H}$ is close to being a degenerate plane. If $\mathcal{H}$ is the degenerate plane, then there is a huge edge e of size $n-1$, and U consists of a single vertex of degree $n-1$. Even though $\mathcal{H}$ is not (ρ,ε)-full, if c is assigned to the edge e, it is clearly impossible to extend $\psi^{-1}(c)$ to have perfect coverage of U, which would be necessary in order to finish the colouring with Vizing's theorem in the final step. However, if there is a difficult colour class that we cannot absorb, then we show that we can colour $\mathcal{H}$ directly (see [35, Lemma 7.12]).

We now construct some $\mathcal{H}'$ with $\mathcal{H}_{\mathrm{small}} \setminus R \subseteq \mathcal{H}' \subseteq \mathcal{H}_{\mathrm{small}}$ and a proper edge-colouring $\psi_2 : \mathcal{H}' \to C_2$ such that ψ_2 is compatible with ψ_1, $|C_2|$ is slightly larger than $(1-\rho+\gamma)n$, $C_2 \cap C_{\mathrm{hm}} = \varnothing$, and $\{\psi_1^{-1}(c) \cup \psi_2^{-1}(c) : c \in C_{\mathrm{hm}} \cup C_2\}$ has nearly-perfect coverage of U with defects in S. (Actually, as in Section II-A we obtain this coverage property only for a suitable auxiliary hypergraph $\mathcal{H}^* \supseteq \mathcal{H}'$, but we again ignore this here for simplicity.) In the non-FPP-extremal case, since $\rho = \rho_1 \ll \sigma$, this means we can reserve a set C_{final} of colours (of size close to ρn) which are used neither by ψ_1 nor by ψ_2. Then in the final step of the proof, we can colour the leftover graph $\mathcal{H}_{\mathrm{small}} \setminus \mathcal{H}' \subseteq R$ (with colours from C_{final}) as described in Section II-A. In the FPP-extremal case, we may have $|C_1| = n$, so we need to find a proper edge-colouring of $\mathcal{H}_{\mathrm{small}} \setminus \mathcal{H}'$ using colours from $C_1 \setminus C_2$ while avoiding conflicts with ψ_1. But in this case most pairs of vertices are contained in an edge of $\mathcal{H}_{\mathrm{ex}}$, which implies that $|U|$ is small. Moreover, every edge of the leftover graph $\mathcal{H}_{\mathrm{small}} \setminus \mathcal{H}' \subseteq R$ is incident to a vertex of U. These two properties allow us to colour the leftover graph $\mathcal{H}_{\mathrm{small}} \setminus \mathcal{H}'$ with $\Delta(\mathcal{H}_{\mathrm{small}} \setminus \mathcal{H}')$ colours while using (1) and (2) to avoid conflicts with ψ_1, as desired.

We conclude by discussing how to construct $\mathcal{H}'$ and ψ_2. Using the colours in C_2, we colour all of the edges of $\mathcal{H}_{\mathrm{small}} \setminus R$ and some of the remaining uncoloured (by ψ_1) edges of R based on the nibble and the absorption strategy outlined in Section II-A. For this, the following properties are crucial (which follow from (1) and the definition of $\mathcal{H}_{\mathrm{med}}$ respectively).

(a) $\psi_1^{-1}(c)$ covers at most βn vertices for each $c \in C_2$, and
(b) every vertex $v \in V(\mathcal{H})$ is contained in at most $n/(r_0-1)$ edges that are assigned a colour in C_2 by ψ_1 (since for any $c \in C_2$, either $\psi_1^{-1}(c)$ is empty or all the edges in $\psi_1^{-1}(c)$ are large). Thus, each edge in $\mathcal{H}_{\mathrm{small}}$ still has slightly more than $(1-\rho)n$ colours available in C_2 that do not conflict with ψ_1 (since any edge of $\mathcal{H}_{\mathrm{small}}$ intersects at most $r_1 n/(r_0-1)$ large edges and $r_1/(r_0-1) \ll \gamma$).

We will use (a) and (b) to show that the effect of the previously coloured edges (by ψ_1) on the Rödl nibble argument is negligible, i.e., we can adapt the arguments of Section II-A, so that the colouring ψ_2 of $\mathcal{H}'$ is compatible with ψ_1.

III. Algorithmic aspects of the proof

In this section, we show how the arguments in Sections 4–11 of [35] can be modified to obtain a randomised polynomial-time algorithm. In particular, this section is intended to be read in conjunction with Sections 4–11 of [35].

A. Embedding lemma

[35, Lemma 4.4] shows how an arbitrary bounded degree linear hypergraph $\mathcal{H}$ can be embedded into a regular uniform hypergraph $\mathcal{H}_{\mathrm{unif}}$. The proof of [35, Lemma 4.4] uses the existence of Steiner systems for which no randomised polynomial-time algorithm is provided in the literature. Here we describe how to use an explicit construction (given in [31]) instead, to prove essentially the same result.

Lemma III.1. *Let $0 < 1/N_0, 1/D_0, 1/C_0 \ll 1/r \leq 1/3$, where $r \in \mathbb{N}$. Let $N \geq N_0$, let $C \geq C_0$, let $D \geq D_0$, and let $\mathcal{H}$ be an N-vertex linear multi-hypergraph with $\Delta(\mathcal{H}) \leq D$. If every $e \in \mathcal{H}$ satisfies $|e| \leq r$, then there is a polynomial-time algorithm to find an r-uniform linear hypergraph $\mathcal{H}_{\mathrm{unif}}$ satisfying the following.*

(a) *$\mathcal{H} \subseteq \mathcal{H}_{\mathrm{unif}}|_{V(\mathcal{H})}$, and $\mathcal{H}_{\mathrm{unif}}|_{V(\mathcal{H})} \setminus \mathcal{H}$ only contains singleton edges.*
(b) *For any $v \in V(\mathcal{H}_{\mathrm{unif}})$, $D - C \leq d_{\mathcal{H}_{\mathrm{unif}}}(v) \leq D$. Moreover, if $d_{\mathcal{H}}(v) \geq D - C$ for $v \in V(\mathcal{H})$, then $d_{\mathcal{H}_{\mathrm{unif}}}(v) = d_{\mathcal{H}}(v)$.*
(c) *$v(\mathcal{H}_{\mathrm{unif}}) \leq r^5 D^3 N$.*

The bound in III.1(c) is slightly worse than the corresponding bound in [35, Lemma 4.4]. Nevertheless, it is sufficient since we only use III.1(c) in a weaker form, i.e., the bound $v(\mathcal{H}_{\mathrm{unif}}) \leq N^5$ suffices for the proof of Theorem I.1.

Proof of Lemma III.1 (sketch): Let $\mathcal{H}^*$ be an r-uniform linear hypergraph obtained from $\mathcal{H}$ by adding $r - |e|$ new vertices to each $e \in \mathcal{H}$.

Let $T := r^4 D^2$. For every integer $1 \leq d \leq D$, we construct a T-vertex r-uniform linear hypergraph $\mathcal{H}_d$ such that every vertex has degree either $d-1$ or d, as follows. Let $\mathcal{H}'_d$ be the r-partite r-uniform d-regular linear hypergraph with parts $X_1, \dots, X_r$ such that $X_i := \mathbb{Z}_{rd}$ for each $i \in [r]$, where $\mathcal{H}'_d := \{e_{x,y} : x \in \{0,\dots,rd-1\},\ y \in \{0,\dots,d-1\}, e_{x,y} \cap X_i = \{x + (i-1)y\}\ \forall i \in [r]\}$. Let $t := \lceil T/|V(\mathcal{H}'_d)| \rceil$. Since $t \geq |V(\mathcal{H}'_d)| - 1$, there exist integers $a_1 \geq \dots \geq a_t$, where each a_i is equal to either $|V(\mathcal{H}'_d)| - 1$ or $|V(\mathcal{H}'_d)|$ and $\sum_{i=1}^t a_i = T$. For $i \in [t]$, if $a_i = |V(\mathcal{H}'_d)|$ then let $\mathcal{G}_i := \mathcal{H}'_d$, and otherwise let $\mathcal{G}_i$ be a hypergraph obtained from $\mathcal{H}'_d$ by deleting exactly one

vertex. Finally, let $\mathcal{H}_d$ be the vertex-disjoint union of all $\mathcal{G}_i$ for $i \in [t]$.

We define our desired multi-hypergraph $\mathcal{H}_{\mathrm{unif}}$ by taking the union of T vertex-disjoint copies of $\mathcal{H}^*$, where the first copy is identified with $\mathcal{H}^*$. Then, for each $v \in V(\mathcal{H}^*)$ with $d_{\mathcal{H}^*}(v) < D - C$, let $v^1, \dots, v^T$ be the T clone vertices of $v \in V(\mathcal{H}^*)$ in $\mathcal{H}_{\mathrm{unif}}$, and extend $\mathcal{H}_{\mathrm{unif}}$ by making $\mathcal{H}_{\mathrm{unif}}[\{v^1, \dots, v^T\}]$ induce a copy of $\mathcal{H}_{D-d_{\mathcal{H}^*}(v)}$, which implies that $D - 1 \le d_{\mathcal{H}_{\mathrm{unif}}}(v^i) \le D$ for $1 \le i \le T$. ■

B. Local sparsity theorems

[35, Corollary 6.5] shows that if $\mathcal{H}$ is a linear hypergraph which is close to being r-uniform for some $r \le (1-\zeta)\sqrt{n}$, then $\chi'(\mathcal{H})$ is significantly less than n. [35, Corollary 6.5] is a straightforward consequence of [35, Theorem 6.4] by Molloy and Reed [47] which bounds the chromatic number of locally sparse graphs. Thus an algorithmic version of [35, Corollary 6.5] can be obtained from an algorithmic version of [35, Theorem 6.4]. This was obtained by Hurley, de Joannis de Verclos, and Kang [26] who improved the estimate on $\chi(G)$ in [35, Theorem 6.4] and remarked [26, Section 4] that their proof gives a randomised algorithm of time complexity linear in the number of vertices and polynomial in the degree using the algorithmic version of the local lemma of Moser and Tardos [48]. We can use this instead of [35, Theorem 6.4] to prove the following corollary which gives an algorithmic version of [35, Corollary 6.5].

Corollary III.2. *Let $0 < 1/n_0, 1/r \ll \alpha \ll \zeta < 1$, let $n \ge n_0$, and suppose $r \le (1-\zeta)\sqrt{n}$. If $\mathcal{H}$ is an n-vertex linear hypergraph such that every $e \in \mathcal{H}$ satisfies $|e| \in [r, (1+\alpha)r]$, then $\chi'(\mathcal{H}) \le (1-\zeta/500)n$. Moreover, there is a randomised polynomial-time algorithm which finds a proper edge-colouring of $\mathcal{H}$ using at most $(1-\zeta/500)n$ colours with high probability.*

C. Matchings and Colourings

Here we discuss how to obtain an algorithmic version of [35, Theorem 4.6] on edge colourings with lists (namely Corollary III.5) and of [35, Corollary 4.3] on pseudorandom matchings (namely Corollary III.7). For this, we can use (an algorithmic version of) the following result of Molloy and Reed [46].

Theorem III.3 (Molloy and Reed [46])**.** *Let $0 < 1/D_0 \ll 1/r < 1$, and let $D \ge D_0$. If $\mathcal{H}$ is a linear hypergraph such that $\Delta(\mathcal{H}) \le D$ and every $e \in \mathcal{H}$ satisfies $|e| \le r$, then $\chi'_\ell(\mathcal{H}) \le D + D^{1-1/r}\log^5 D$.*

We remark that Theorem III.3 is slightly stronger than the statement in [46, Theorem 1], since it allows the edges of $\mathcal{H}$ to be of size at most r rather than exactly r. However, Theorem III.3 can be easily deduced from [46, Theorem 1] by first embedding the hypergraph $\mathcal{H}$ in an r-uniform hypergraph $\mathcal{H}_{\mathrm{unif}}$ using Lemma III.1 and then applying [46, Theorem 1].

Although Molloy and Reed's proof of [46, Theorem 1] is non-constructive, they provided a polynomial-time algorithm in [46, Section 7] which finds a proper edge-colouring of $\mathcal{H}$ under the additional assumption $\Delta(\mathcal{H}) = O(1)$ using a 'local lemma like lemma' [46, Lemma 17]. Here we indicate how to get a randomised polynomial-time algorithm that finds a proper colouring of $\mathcal{H}$ with high probability without this additional assumption. First, using the algorithmic local lemma of Moser and Tardos [48, Theorem 1.2], we immediately obtain the following analogue of [46, Lemma 17]. In particular, Lemma III.4 does not require the assumption $\Delta(\mathcal{H}) = O(1)$.

Lemma III.4. *Let $\mathcal{F}$ be a hypergraph such that each edge of $\mathcal{F}$ intersects at most d other edges. Suppose that each vertex v of $\mathcal{F}$ has a list $L(v)$ of colours. For each edge e, we are explicitly given a boolean function $\mathit{Bad}_e : \prod_{v\in e} L(v) \to \{0,1\}$.*

Consider a probability space where each vertex $v \in V(\mathcal{F})$ is equipped with a random variable that chooses a colour from $L(v)$ independently and uniformly at random. If $\mathbb{P}(\mathrm{Bad}_e = 1) \le 1/e(d+1)$ for each $e \in \mathcal{F}$, then there exists an assignment of colours to all vertices of $\mathcal{F}$ such that $\mathrm{Bad}_e = 0$ for every edge e. Moreover, one can find such an assignment in polynomial time in $e(\mathcal{F})$ with high probability, provided that there is a polynomial-time algorithm to evaluate the Boolean function Bad_e for each $e \in \mathcal{F}$.

We can then use Lemma III.4 instead of [46, Lemma 17] in the proof of [46, Theorem 18] to obtain the following algorithmic variant of Theorem III.3 (and in particular an algorithmic version of [35, Theorem 4.6]).

Corollary III.5. *Let $0 < 1/D_0 \ll \alpha, 1/r < 1$, and let $D \ge D_0$. Let $\mathcal{H}$ be a linear hypergraph such that $\Delta(\mathcal{H}) \le D$ and every $e \in \mathcal{H}$ satisfies $|e| \le r$. For any $e \in \mathcal{H}$, let $C(e)$ be a list of colours such that $|C(e)| \ge D + D^{1-r}\log^5 D$. Then there is a randomised polynomial-time algorithm which finds a proper edge-colouring ϕ of $\mathcal{H}$ such that $\phi(e) \in C(e)$ with high probability.*

In particular, if C is a set of colours with $|C| \ge (1+\alpha)D$, and for each $e \in \mathcal{H}$, $C'(e) \subseteq C$ such that $|C'(e)| \le \alpha D/2$, then there is a randomised polynomial-time algorithm to find a proper edge-colouring $\phi : \mathcal{H} \to C$ such that $\phi(e) \notin C'(e)$ for each $e \in \mathcal{H}$ with high probability.

Now we discuss how to obtain Corollary III.7, which is an algorithmic version of [35, Corollary 4.3] for finding a pseudorandom matching in a hypergraph. To that end, we need an algorithmic version of a theorem of Ehard, Glock, and Joos [12, Theorem 1.2] that provides a matching $\mathcal{M}$ in a hypergraph $\mathcal{H}$ covering almost all vertices of every set in a given collection of sets. (This quantitatively improves an earlier result of Alon and Yuster [4].) Their proof in [12] actually provides a randomised algorithm except that

it uses Theorem III.3. However, fortunately, we can instead use its algorithmic version, i.e., Corollary III.5. This yields a randomised polynomial-time algorithm to construct the matching $\mathcal{M}$ with high probability.

Theorem III.6 (Ehard, Glock, and Joos [12])**.** *Let $r \geq 2$ be an integer, and let $\varepsilon := 1/(1500r^2)$. There exists Δ_0 such that the following holds for all $\Delta \geq \Delta_0$.*

Given an r-uniform linear hypergraph $\mathcal{H}$ which satisfies $\Delta(\mathcal{H}) \leq \Delta$ and $e(\mathcal{H}) \leq \exp(\Delta^{\varepsilon^2})$, and given a family $\mathcal{F}$ of subsets of $V(\mathcal{H})$ such that $|\mathcal{F}| \leq \exp(\Delta^{\varepsilon^2})$ and $\sum_{v\in S} d_{\mathcal{H}}(v) \geq \Delta^{26/25}$, there is a randomised polynomial-time algorithm that constructs a matching $\mathcal{M}$ in $\mathcal{H}$ with high probability such that for any $S \in \mathcal{F}$, $|S \cap V(\mathcal{M})| = (1 \pm \Delta^{-\varepsilon^2}) \sum_{v\in S} d_{\mathcal{H}}(v)/\Delta$.

Theorem III.6 implies the following desired algorithmic variant of [35, Corollary 4.3].

Corollary III.7. *Let $0 < 1/n_0 \ll 1/r, \kappa, \gamma < 1$. For any integer $n \geq n_0$, let $\mathcal{H}$ be an r-uniform linear n-vertex hypergraph such that every vertex has degree $(1 \pm \kappa)D$, where $D \geq n^{1/100}$. Let $\mathcal{F}$ be a set of subsets of $V(\mathcal{H})$ such that $|\mathcal{F}| \leq n^{2\log n}$. Then there is a randomised algorithm that constructs a matching $\mathcal{M}$ of $\mathcal{H}$ such that for any $S \in \mathcal{F}$ with $|S| \geq D^{1/20}$, we have $|S \setminus V(\mathcal{M})| = (\gamma \pm 4\kappa)|S|$.*

One can deduce Corollary III.7 from Theorem III.6 by applying Theorem III.6 to obtain a matching $\mathcal{M}_0$ (in $\mathcal{H}$) and then randomly removing each edge of $\mathcal{M}_0$ with probability γ to obtain a matching $\mathcal{M}$ which satisfies the assertion of Corollary III.7 with high probability by Chernoff's bound (see the paragraph following [35, Theorem 4.2] for more details).

D. Colouring large and medium edges

We need the following algorithmic version of [35, Theorem 6.1].

Theorem III.8. *Let $0 < 1/n_0 \ll 1/r_0 \ll 1/r_1, \beta \ll \gamma_1 \ll \sigma \ll \delta \ll \gamma_2 \ll 1$, and let $n \geq n_0$. Given an n-vertex linear hypergraph $\mathcal{H}$ where every $e \in \mathcal{H}$ satisfies $|e| > r_1$, there exists a randomised polynomial-time algorithm which finds, with high probability, either*

(III.8:a) *a proper edge-colouring of $\mathcal{H}$ using at most $(1-\sigma)n$ colours such that*
- (i) *every colour assigned to a huge edge is assigned to no other edge,*
- (ii) *every medium edge is assigned a colour from a set C_{med} of size at most $\gamma_1 n$ such that for every $c \in C_{\mathrm{med}}$, at most $\gamma_1 n$ vertices are incident to an edge coloured c, and*
- (iii) *for every colour $c \notin C_{\mathrm{med}}$ not assigned to a huge edge, at most βn vertices are incident to an edge coloured c,*

(III.8:b) *or a set of FPP-extremal edges of volume at least $1-\delta$ and a proper edge-colouring of $\mathcal{H}$ using at most n colours such that*
- (i) *for every colour c assigned to a huge edge, at most δn vertices are incident to an edge coloured c,*
- (ii) *every medium edge is assigned a colour from a set C_{med} of size at most $\gamma_2 n$ such that for every $c \in C_{\mathrm{med}}$, at most $\gamma_1 n$ vertices are incident to an edge coloured c, and*
- (iii) *for every colour $c \notin C_{\mathrm{med}}$ not assigned to a huge edge, at most βn vertices are incident to an edge coloured c.*

To prove Theorem III.8, we need algorithmic versions of [35, Lemma 5.1] and [35, Lemma 6.2]. These algorithmic versions are stated below as Lemmas III.9 and III.10 respectively.

Lemma III.9. *Let $0 < 1/n_0 \ll \delta \ll 1$, and let $n \geq n_0$. Given an n-vertex linear hypergraph $\mathcal{H}$ where every $e \in \mathcal{H}$ satisfies $|e| \geq (1-\delta)\sqrt{n}$, there is a polynomial-time algorithm that finds a proper edge-colouring of $\mathcal{H}$ with n colours, where each colour is assigned to at most two edges.*

The proof of Lemma III.9 in [35, Lemma 5.1] is clearly algorithmic, provided that one can find a maximum matching in the graph $\overline{L(\mathcal{H})}$ in polynomial time. (Here $L(\mathcal{H})$ denotes the *line graph* of the hypergraph $\mathcal{H}$, and $\overline{L(\mathcal{H})}$ denotes the complement of $L(\mathcal{H})$.) For this, there are well-known algorithms (e.g., [6], [11], [20], [45]), which can find a maximum matching in arbitrary graph in polynomial time.

We now turn to Lemma III.10. For any hypergraph $\mathcal{H}$, any ordering $\preceq$ of the edges of $\mathcal{H}$, and any $e \in \mathcal{H}$, let $\mathcal{H}^{\preceq e} := \{f \in \mathcal{H} : f \preceq e\}$, let $N^{\preceq}(e) := \{f \in \mathcal{H} : f \in N(e) \text{ and } f \preceq e\}$, and let $d^{\preceq}(e) := |N^{\preceq}(e)|$.

Lemma III.10 (Reordering lemma)**.** *Let $0 < 1/r_1 \ll \tau, 1/K$ where $\tau < 1$, $K \geq 1$, and $1 - \tau - 7\tau^{1/4}/K > 0$. Given an n-vertex linear hypergraph $\mathcal{H}$ where every $e \in \mathcal{H}$ satisfies $|e| \geq r_1$, there is a polynomial-time algorithm that finds a linear ordering $\preceq$ of the edges of $\mathcal{H}$ such that at least one of the following holds.*

- (a) *Every $e \in \mathcal{H}$ satisfies $d^{\preceq}(e) \leq (1-\tau)n$.*
- (b) *There is a set $W \subseteq \mathcal{H}$ such that*
  - (W1) $\max_{e\in W} |e| \leq (1 + 3\tau^{1/4}K^4) \min_{e\in W} |e|$ *and*
  - (W2) $\mathrm{vol}_{\mathcal{H}}(W) \geq \frac{(1-\tau-7\tau^{1/4}/K)^2}{1+3\tau^{1/4}K^4}$.

  Moreover, if e^* is the last edge of W, then
  - (O1) *for all $f \in \mathcal{H}$ such that $e^* \preceq f$ and $f \neq e^*$, we have $d^{\preceq}(f) \leq (1-\tau)n$ and*
  - (O2) *for all $e, f \in \mathcal{H}$ such that $f \preceq e \preceq e^*$, we have $|f| \geq |e|$.*

In the proof of [35, Lemma 6.2], the ordering $\preceq$ of the edges of $\mathcal{H}$ and $e^* \in \mathcal{H}$ were not chosen explicitly (i.e.,

they were chosen so that $e(\mathcal{H}^{\preceq e^*})$ is minimised subject to satisfying both (O1) and (O2). However, the proof of [35, Lemma 6.2] shows that to find an ordering $\preceq$ satisfying Lemma III.10, it suffices to find an ordering $\preceq$ such that either $\preceq$ satisfies III.10(a) or

(O3) there is an edge e^* with $d^{\preceq}(e^*) > (1-\tau)n$, satisfying (O1) and (O2), such that $|N(e) \cap \mathcal{H}^{\preceq e^*}| > (1-\tau)n$ for any $e \in N^{\preceq}(e^*)$.

It can be easily seen that the desired ordering $\preceq$ satisfying either III.10(a) or (O3) can be obtained by starting with an ordering $\preceq_0$ where $|e| \geq |f|$ for all $e \preceq_0 f$, and successively moving at most $e(\mathcal{H})$ edges to new positions. This yields a polynomial-time algorithm.

Algorithmic aspects of the proof of Theorem III.8: To prove Theorem III.8 we follow the proof of [35, Theorem 6.1] with some changes. The proof of [35, Theorem 6.1] relies on [35, Lemmas 5.1 and 6.2], [35, Corollary 6.5], and [35, Propositions 6.6, 6.7 and 6.8]. Below we list the required modifications.

We replace [35, Lemma 5.1] by Lemma III.9, [35, Lemma 6.2] by Lemma III.10, and [35, Corollary 6.5] by Lemma III.2, respectively. The proof of [35, Proposition 6.6] immediately gives a polynomial-time algorithm which finds an α-bounded edge-colouring using $b+2n/(\alpha^2 r)$ colours in polynomial-time given an arbitrary proper edge-colouring of $\mathcal{H}$ with b colours. In the proof of [35, Proposition 6.7], we proceeded by contradiction but the argument immediately shows that one can colour $\mathcal{H}$ greedily in the order prescribed by the ordering $\preceq$. Lastly, in the proof of [35, Proposition 6.8], we use Corollary III.5 instead of [35, Theorem 4.6] to find an edge-colouring of $\mathcal{H}$ with at most $\gamma n/2$ colours. ■

E. Vizing's theorem

We need the following algorithmic version of a theorem of Vizing [55].

Theorem III.11. *Let G be a graph with maximum degree D. Then there is a polynomial-time algorithm that finds a proper $(D+1)$-edge-colouring of G. Moreover, if the vertices of maximum degree in G induce a forest, then there is a polynomial-time algorithm that finds a proper D-edge-colouring of G.*

Note that Theorem III.11 immediately implies [35, Theorem 4.5]. For a proof of Theorem III.11, see the proof of [53, Theorem 28.1] (based on the ideas of Ehrenfeucht, Faber, and Kierstead [13]) and the remarks following the proof.

F. Algorithmic aspects of the proof of Theorem I.1

Below we follow the proof of Theorem I.1 shown in [35, Section 11] and (sequentially) present an exhaustive list of all the modifications required to the proof in order to produce a polynomial-time randomised algorithm that finds the desired colouring of $\mathcal{H}$ with high probability.

We replace [35, Theorem 6.1] by its algorithmic variant Theorem III.8, which is used to colour all medium and large edges of $\mathcal{H}$. Then to define a reservoir set R of edges (of size two) we use [35, Proposition 10.3 and Lemma 10.4] – the proof of [35, Proposition 10.3] immediately yields a polynomial-time randomised algorithm, and [35, Lemma 10.4] uses the Lovász (g, f)-factor theorem [43] which has a polynomial-time algorithm (see [44, Section 10.2]).

To extend the colour classes obtained from huge edges using edges in R, we use [35, Lemmas 7.11 and 7.12]. The proof of [35, Lemmas 7.12] uses the fact that any graph G with exactly one vertex of maximum degree can be properly edge-coloured with $\Delta(G)$ colours, and by Theorem III.11, there is a polynomial-time algorithm to find such a colouring. The proof of [35, Lemmas 7.11] requires finding a maximum matching in a certain graph, but it is well-known that there is a polynomial-time algorithm for doing so (see e.g., [6], [11], [20], [45]).

To extend the colour classes obtained from other medium and large edges using edges in R, we use [35, Lemmas 8.2 and 8.3]. The proof of [35, Lemmas 8.2] uses [35, Lemmas 8.1, 7.6 and 7.7].

In the proof of [35, Lemma 8.1] we use Lemma III.1 and Corollary III.7 instead of [35, Lemma 4.4] and [35, Corollary 4.3], respectively.

The proofs of [35, Lemmas 7.6 and 7.7] require finding a maximum matching in a certain graph, which can be done in polynomial time as mentioned above. In the proof of [35, Lemma 8.3] we use Corollary III.5 instead of [35, Theorem 4.6].

Finally, to colour the graph G_{final} consisting of remaining edges of R, we use either [35, Theorem 4.5], [35, Corollary 9.6], or [35, Lemma 9.2]. We replace [35, Theorem 4.5] by Theorem III.11 which has a polynomial-time algorithm. [35, Corollary 9.6] is derived from [22, Theorem 1.6] whose proof in [22] yields a polynomial-time algorithm, and the proof of [35, Lemma 9.2] requires finding a maximum matching in a certain graph for which, again, there is a polynomial-time algorithm (e.g., [6], [11], [20], [45]).

REFERENCES

[1] D. Achlioptas, F. Iliopoulos, and A. Sinclair, "Beyond the Lovász Local Lemma: Point to Set Correlations and Their Algorithmic Applications," in *2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS)*, 2019, pp. 725–744.

[2] M. Ajtai, J. Komlós, and E. Szemerédi, "A note on Ramsey numbers," *J. Combin. Theory Ser. A*, vol. 29, no. 3, pp. 354–360, 1980.

[3] N. Alon, M. Krivelevich, and B. Sudakov, "Coloring graphs with sparse neighborhoods," *J. Combin. Theory Ser. B*, vol. 77, no. 1, pp. 73–82, 1999.

[4] N. Alon and R. Yuster, "On a hypergraph matching problem," *Graphs Combin.*, vol. 21, no. 4, pp. 377–384, 2005.

[5] Z. Baranyai, "On the factorization of the complete uniform hypergraph," in *Infinite and finite sets (Colloq., Keszthely, 1973; dedicated to P. Erdős on his 60th birthday), Vol. I*, 1975, pp. 91–108. Colloq. Math. Soc. János Bolyai, Vol. 10.

[6] N. Blum, "A new approach to maximum matching in general graphs," in *International Colloquium on Automata, Languages, and Programming*. Springer, 1990, pp. 586–597.

[7] J. Brakensiek and V. Guruswami, "New hardness results for graph and hypergraph colorings," in *31st Conference on Computational Complexity*, ser. LIPIcs. Leibniz Int. Proc. Inform. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2016, vol. 50, pp. Art. No. 14, 27.

[8] W. I. Chang and E. L. Lawler, "Edge coloring of hypergraphs and a conjecture of Erdős, Faber, Lovász," *Combinatorica*, vol. 8, no. 3, pp. 293–295, 1988.

[9] F. Chung and R. Graham, *Erdős on graphs: His legacy of unsolved problems*. A K Peters, Ltd., Wellesley, MA, 1998.

[10] E. Davies, R. J. Kang, F. Pirot, and J.-S. Sereni, "Graph structure via local occupancy," *arXiv preprint arXiv:2003.14361*, 2020.

[11] J. Edmonds, "Paths, trees, and flowers," *Canadian J. Math.*, vol. 17, pp. 449–467, 1965.

[12] S. Ehard, S. Glock, and F. Joos, "Pseudorandom hypergraph matchings," *Combin. Probab. Comput.*, pp. 1–18, 2020.

[13] A. Ehrenfeucht, V. Faber, and H. A. Kierstead, "A new method of proving theorems on chromatic index," *Discrete Math.*, vol. 52, no. 2-3, pp. 159–164, 1984.

[14] P. Erdős and H. Hanani, "On a limit theorem in combinatorial analysis," *Publ. Math. Debrecen*, vol. 10, pp. 10–13, 1963.

[15] P. Erdős, "On the combinatorial problems which I would most like to see solved," *Combinatorica*, vol. 1, no. 1, pp. 25–42, 1981.

[16] P. Erdős, A. Gyárfás, and L. Pyber, "Vertex coverings by monochromatic cycles and trees," *J. Combin. Theory Ser. B*, vol. 51, no. 1, pp. 90–95, 1991.

[17] V. Faber and D. G. Harris, "Edge-coloring linear hypergraphs with medium sized edges," *Random Structures Algorithms*, vol. 55, no. 1, pp. 153–159, 2019.

[18] A. Ferber and V. Jain, "1-factorizations of pseudorandom graphs," in *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 2018, pp. 698–708.

[19] A. Frieze and D. Mubayi, "Coloring simple hypergraphs," *J. Combin. Theory Ser. B*, vol. 103, no. 6, pp. 767–794, 2013.

[20] H. N. Gabow and R. E. Tarjan, "Faster scaling algorithms for general graph-matching problems," *J. Assoc. Comput. Mach.*, vol. 38, no. 4, pp. 815–853, 1991.

[21] M. R. Garey and D. S. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*. USA: W. H. Freeman & Co., 1979.

[22] S. Glock, D. Kühn, and D. Osthus, "Optimal path and cycle decompositions of dense quasirandom graphs," *J. Combin. Theory Ser. B*, vol. 118, pp. 88–108, 2016.

[23] I. Holyer, "The NP-completeness of edge-coloring," *SIAM J. Comput.*, vol. 10, no. 4, pp. 718–720, 1981.

[24] S. Huang, "Improved hardness of approximating chromatic number," in *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*. Springer, 2013, pp. 233–243.

[25] ——, "$2^{(\log N)^{1/10-o(1)}}$ hardness for hypergraph coloring," 2015, arXiv preprint arXiv:1504.03923.

[26] E. Hurley, R. de Joannis de Verclos, and R. J. Kang, "An improved procedure for colouring graphs of bounded local density," in *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 2021, pp. 135–148.

[27] A. Johansson, "Asymptotic choice number for triangle free graphs," DIMACS technical report, Tech. Rep., 1996.

[28] J. Kahn, "Coloring nearly-disjoint hypergraphs with $n+o(n)$ colors," *J. Combin. Theory Ser. A*, vol. 59, no. 1, pp. 31–39, 1992.

[29] ——, "Asymptotics of Hypergraph Matching, Covering and Coloring Problems," in *Proceedings of the International Congress of Mathematicians*. Springer, 1995, pp. 1353–1362.

[30] ——, "Asymptotically good list-colorings," *J. Combin. Theory Ser. A*, vol. 73, no. 1, pp. 1–59, 1996.

[31] ——, "A linear programming perspective on the Frankl-Rödl-Pippenger theorem," *Random Structures Algorithms*, vol. 8, no. 2, pp. 149–157, 1996.

[32] ——, "On some hypergraph problems of Paul Erdős and the asymptotics of matchings, covers and colorings," in *The Mathematics of Paul Erdös I*. Springer, 1997, pp. 345–371.

[33] J. Kahn and P. D. Seymour, "A fractional version of the Erdős-Faber-Lovász conjecture," *Combinatorica*, vol. 12, no. 2, pp. 155–160, 1992.

[34] D. Kang, T. Kelly, D. Kühn, A. Methuku, and D. Osthus, "Graph and hypergraph colouring via nibble methods: A survey," 2021, arXiv preprint arXiv:2106.13733.

[35] ——, "A proof of the Erdős-Faber-Lovász conjecture," 2021, arXiv preprint arXiv:2101.04698.

[36] D. Kang, D. Kühn, A. Methuku, and D. Osthus, "New bounds on the size of nearly perfect matchings in almost regular hypergraphs," 2020, arXiv preprint arXiv:2010.04183.

[37] R. M. Karp, "Reducibility among combinatorial problems," in *Complexity of computer computations*. Springer, 1972, pp. 85–103.

[38] J. H. Kim, "On Brooks' theorem for sparse graphs," *Combin. Probab. Comput.*, vol. 4, no. 2, pp. 97–132, 1995.

[39] ——, "The Ramsey number $R(3,t)$ has order of magnitude $t^2/\log t$," *Random Structures Algorithms*, vol. 7, no. 3, pp. 173–207, 1995.

[40] M. Krivelevich, "Triangle factors in random graphs," *Combin. Probab. Comput.*, vol. 6, no. 3, pp. 337–347, 1997.

[41] D. Kühn and D. Osthus, "Hamilton decompositions of regular expanders: a proof of Kelly's conjecture for large tournaments," *Adv. Math.*, vol. 237, pp. 62–146, 2013.

[42] D. Leven and Z. Galil, "NP completeness of finding the chromatic index of regular graphs," *Journal of Algorithms*, vol. 4, no. 1, pp. 35–44, 1983.

[43] L. Lovász, "Subgraphs with prescribed valencies," *J. Combin. Theory*, vol. 8, no. 4, pp. 391–416, 1970.

[44] L. Lovász and M. D. Plummer, *Matching theory*. American Mathematical Soc., 2009, vol. 367.

[45] S. Micali and V. V. Vazirani, "An $O(\sqrt{|V|}|E|)$ algoithm for finding maximum matching in general graphs," in *21st Annual Symposium on Foundations of Computer Science (sfcs 1980)*. IEEE, 1980, pp. 17–27.

[46] M. Molloy and B. Reed, "Near-optimal list colorings," *Random Structures Algorithms*, vol. 17, no. 3-4, pp. 376–402, 2000.

[47] ——, *Graph colouring and the probabilistic method*, ser. Algorithms and Combinatorics. Springer-Verlag, Berlin, 2002, vol. 23.

[48] R. A. Moser and G. Tardos, "A constructive proof of the general Lovász local lemma," *J. ACM*, vol. 57, no. 2, pp. 1–15, 2010.

[49] P. Obszarski and A. Jastrzębski, "Edge-coloring of 3-uniform hypergraphs," *Discrete Appl. Math.*, vol. 217, no. part 1, pp. 48–52, 2017.

[50] N. Pippenger and J. Spencer, "Asymptotic behavior of the chromatic index for hypergraphs," *J. Combin. Theory Ser. A*, vol. 51, no. 1, pp. 24–42, 1989.

[51] V. Rödl, "On a packing and covering problem," *European J. Combin.*, vol. 6, no. 1, pp. 69–78, 1985.

[52] V. Rödl, A. Ruciński, and E. Szemerédi, "A Dirac-type theorem for 3-uniform hypergraphs," *Combin. Probab. Comput.*, vol. 15, no. 1-2, pp. 229–251, 2006.

[53] A. Schrijver, *Combinatorial optimization. Polyhedra and efficiency. Vol. A*, ser. Algorithms and Combinatorics. Springer-Verlag, Berlin, 2003, vol. 24, paths, flows, matchings, Chapters 1–38.

[54] P. D. Seymour, "Packing nearly-disjoint sets," *Combinatorica*, vol. 2, no. 1, pp. 91–97, 1982.

[55] V. G. Vizing, "The chromatic class of a multigraph," *Cybernetics*, vol. 1, no. 3, pp. 32–41, 1965.

[56] V. H. Vu, "A general upper bound on the list chromatic number of locally sparse graphs," *Combin. Probab. Comput.*, vol. 11, no. 1, pp. 103–111, 2002.