

Approximation Algorithms for Matroid-Intersection Coloring with Applications to Rota’s Basis Conjecture

Stephen Arndt
Carnegie Mellon University

Benjamin Moseley
Carnegie Mellon University

Kirk Pruhs
University of Pittsburgh

Chaitanya Swamy
University of Waterloo

Michael Zlatin
Pomona College

Abstract

We study algorithmic matroid intersection coloring. Given k matroids on a common ground set U of n elements, the goal is to partition U into the fewest number of color classes, where each color class is independent in all matroids. There is a rich literature examining the chromatic number of independence systems arising from the intersection of matroids, but a fundamental gap persists between existential bounds and efficient algorithms. It is known that $2\chi_{\max}$ colors suffice to color the intersection of two matroids, $(2k - 1)\chi_{\max}$ colors suffice for general k , where $\chi_{\max}$ is the maximum chromatic number of the individual matroids. However, these results are non-constructive, leveraging techniques such as topological Hall’s theorem and Sperner’s Lemma. On the algorithmic front, prior polynomial-time algorithms have typically relied on reducing classes of combinatorial matroids to partition matroids. This approach provably fails for general matroids, leaving a greedy $O(k \log n)$ -approximation via Set Cover as the best previously known bound, even when $k = 2$.

We provide the first polynomial-time algorithms to color two or more general matroids where the approximation ratio depends only on k and, in particular, is independent of n . For two matroids, we constructively match the $2\chi_{\max}$ existential bound, yielding a 2-approximation for the Matroid Intersection Coloring problem. For k matroids we achieve a $(k^2 - k)\chi_{\max}$ coloring, which is the first $O(1)$ -approximation for constant k . Our approach introduces a novel matroidal structure we call a *flexible decomposition*. We use this to formally reduce general matroid intersection coloring to graph coloring while avoiding the limitations of partition reduction techniques, and without relying on non-constructive topological machinery.

Furthermore, we give a *fully polynomial randomized approximation scheme* (FPRAS) for coloring the intersection of two matroids when $\chi_{\max}$ is large. This yields the first polynomial-time constructive algorithm for an asymptotic variant of Rota’s Basis Conjecture. This constructivizes Montgomery and Sauermann’s recent asymptotic breakthrough and generalizes it to arbitrary matroids.

1 Introduction

Matroids are fundamental discrete objects that arise in a variety of settings in combinatorics and optimization. Consequently, matroid-optimization problems have been extensively studied in the combinatorics, combinatorial-optimization, and Theoretical Computer Science (TCS) literature.

We investigate a natural and well-studied matroid-optimization problem called matroid-intersection coloring. In an instance of *k-matroid-intersection coloring*, we are given k matroids¹ $M_1 = (U, \mathcal{I}_1), \dots, M_k = (U, \mathcal{I}_k)$ on a common ground set, and the goal is to cover U using the minimum number of *common independent sets*. That is, a feasible solution consists of sets $S_1, \dots, S_q$, where $\bigcup_{c=1}^q S_c = U$, and each

¹A matroid is a tuple $M = (U, \mathcal{I})$, where U is a finite set and $\mathcal{I} \subseteq 2^U$ satisfies: (i) $\emptyset \in \mathcal{I}$; (ii) $B \in \mathcal{I}, A \subseteq B \Rightarrow A \in \mathcal{I}$; and (iii) if $A, B \in \mathcal{I}$, $|A| < |B|$, then $\exists e \in B \setminus A$ such that $A \cup \{e\} \in \mathcal{I}$. If M satisfies (i), (ii), we call M an *independence system*.

S_c is independent in the intersection of the k matroids, which is defined as the independence system $M_{\text{int}} = \bigcap_{i=1}^k M_i = (U, \mathcal{I}_{\text{int}} := \bigcap_{i=1}^k \mathcal{I}_i)$. Such a solution is called a q -coloring of $M_1, \dots, M_k$, where we interpret elements in each S_c as being “colored” with color c ; we seek a solution using the fewest number of colors. The optimal value is called the *chromatic number* of M_{int} and denoted $\chi(M_{\text{int}})$. While we have phrased matroid-intersection coloring as a covering problem, observe that one can always transform a solution to have pairwise-disjoint color-sets, so equivalently the goal is to partition U into the smallest number of common independent sets. As is standard in matroid optimization, we assume that each input matroid is specified via an independence oracle or a rank oracle.

The chromatic number of a single matroid $M = (U, \mathcal{I})$ with rank function r_M can be computed efficiently via a reduction to matroid intersection [Edm68, Sch03], which also yields the formula $\chi(M) = \max_{S \subseteq U} \lceil \frac{|S|}{r_M(S)} \rceil$, but matroid-intersection coloring quickly becomes intractable for $k \geq 2$. It is known that even with two matroids, the problem is *NP*-hard [BS21, BCK23, HKK24], and computing $\chi(M_{\text{int}})$ exactly requires an exponential number of independence-oracle calls [BS21]. It is therefore natural to consider approximation algorithms, and approximation results in the literature are typically stated relative to the natural lower bound $\chi_{\max} := \max_{i=1, \dots, k} \chi(M_i)$.

Matroid-intersection coloring has been studied both in the combinatorics literature [AB06, ABGK25, MS25] and the combinatorial-optimization and TCS literature [IMP21, BSY21, LMP21, AMPZ26]. The current state of affairs can be summarized as follows. The combinatorics literature has obtained various strong existential results via *nonconstructive*² means. The optimization and TCS literature obtains constructive results (i.e., polytime algorithms), but these apply only to certain “combinatorial” matroids (not general matroids). For general matroids, the only known constructive result prior to our work (*even for $k = 2$*) is an $O(k \log n)$ -approximation algorithm (i.e., an $O(k \log n) \cdot \chi(M_{\text{int}})$ -coloring) using the greedy algorithm for set cover.³

As is evident from this discussion, there is a significant gap between the nonconstructive results and constructive guarantees, which persists even for the case of two matroids. *Bridging this gap is the chief motivating goal of our work*, and we approach this research goal by considering the following two research directions.

- **Constructive Guarantees for Matroid-Intersection Coloring.** The state-of-the-art guarantees for matroid-intersection coloring are given by the following two nonconstructive results.⁴

Theorem 1.1 ([AB06]). *For any two matroids M_1 and M_2 , we have $\chi(M_1 \cap M_2) \leq 2\chi_{\max}$.*

Theorem 1.2 ([ABGK25]). *For k (general) matroids $M_1, \dots, M_k$, we have $\chi(M_{\text{int}}) \leq (2k - 1)\chi_{\max}$.*

The proofs of Theorem 1.1 and Theorem 1.2 are topological, using the notion of homotopical connectivity and relying on Sperner’s lemma to deduce the existence of a desired coloring. As Sperner’s lemma is essentially complete for the complexity class PPAD, making this approach constructive would involve surmounting well-known complexity-theoretic hurdles.

Prior algorithmic results focus on the setting where all matroids [IMP21, BSY21, LMP21], or all but one of the matroids [AMPZ26], are partition matroids. These approaches extend to standard “combinatorial” matroids (i.e., transversal matroids, gammoids, laminar matroids, graphic matroids), as one can reduce to partition matroids at the loss of a small factor in the coloring number [IMP21, BSY21, LMP21]. But this

²Since we are considering problems with a finite solution space, by “nonconstructive” we mean not readily amenable to conversion to a polynomial-time algorithm.

³The greedy step translates to solving a k -matroid-intersection problem, which admits an $O(k)$ -approximation.

⁴We state here the guarantees in terms of $\chi_{\max}$; with two matroids, a more refined, still nonconstructive, bound $\chi(M_1 \cap M_2) \leq \chi(M_1) + \chi(M_2)$ was recently obtained [BG25].

partition-reduction approach cannot be extended to general matroids because there are matroids where one cannot “simplify the independence-structure” and move to a partition matroid without incurring a *significant* blow-up in the chromatic number [LMP22, AKKG23]. In particular, the chromatic number of the partition matroid must necessarily increase as a function of $n = |U|$. In a sense, the work of [AMPZ26], who handle the case where one of the matroids is a general matroid can be seen as a limit of this partition-reduction approach, and as the authors of [AMPZ26] indicate, this approach is not amenable to handle the setting of even two general matroids.

This preface leads to the first question that we explore in this work.

Question 1. *Can we devise polynomial-time algorithms for k -matroid intersection coloring (with general matroids) whose guarantees match, or nearly match, those in Theorems 1.1 and 1.2?*

• **Asymptotic Rota’s Basis Conjecture.** The second research direction we consider involves one of the most-prominent problems in matroid-intersection coloring, namely *Rota’s Basis Conjecture*, which is the following famous conjecture regarding rearranging the bases of a matroid.

Conjecture 1 (Rota’s Basis Conjecture). Let $M_1 = (U, \mathcal{I}_1)$ be a rank- r matroid where U can be partitioned into r disjoint bases $B^1, B^2, \dots, B^r$. Then, there exist r disjoint bases $T_1, \dots, T_r$ such that $|T_i \cap B^j| = 1$ for all $i, j = 1, \dots, r$. The T_i bases are sometimes called rainbow bases, or transversal bases.

In the language of matroid-intersection coloring, taking $M_2 = (U, \mathcal{I}_2)$ to be the partition matroid with parts $B^1, \dots, B^r$ and capacity 1 on each part, Rota’s Basis Conjecture asserts that $\chi(M_1 \cap M_2) = r$. Note that in this setting $\chi_{\max} = \chi(M_1) = \chi(M_2) = r = \sqrt{n}$, where $n = |U|$.

Since Rota made this conjecture in 1989, it has received significant attention (see, e.g., [HR94, Dri97, Gly10, MS25] and the references therein). It was the subject of the 12th Polymath project [Cho17], which improved a $2\chi_{\max}$ upper bound on $\chi(M_1 \cap M_2)$ to $2\chi_{\max} - 2$. The conjecture holds for some special cases [MS25], perhaps most notably for real-representable matroids when the rank is within one of an odd prime [HR94, Dri97, Gly10]. Very recently, Montgomery and Sauermann [MS25] proved that Rota’s basis conjecture holds asymptotically, giving the current-best bounds on $\chi(M_1 \cap M_2)$. They prove the following.

Theorem 1.3 ([MS25]). *For any $\varepsilon > 0$, and any instance of Rota’s Basis Conjecture where $\chi(M_1)$ (or equivalently the rank r) is sufficiently large relative to ε , we have $\chi(M_1 \cap M_2) \leq (1 + \varepsilon)\chi_{\max}$.*

The proof of Theorem 1.3 is quite involved and also nonconstructive,⁵ which leads to the second question that we explore.

Question 2. *Is there a polynomial-time algorithm that obtains similar asymptotic guarantees for Rota’s Basis Conjecture as Theorem 1.3?*

1.1 Our Contributions

We make progress on the above questions, advancing the constructive frontier in the area of matroid-intersection coloring. We fully resolve Question 2, and resolve Question 1 for the case of two matroids, thus essentially eliminating the constructive-nonconstructive gap in the setting of two matroids.

Theorem 1.4 states our results for k -matroid-intersection coloring, which addresses Question 1.

Theorem 1.4 (Approximation results for matroid-intersection coloring).

⁵One key step in the proof involves finding roughly $\sqrt{n}$ rainbow independent sets having certain lexicographic-optimality properties. This seems quite challenging to obtain via an efficient algorithm.

- (a) *There is a polytime algorithm that returns a $2\chi_{\max}$ -coloring for 2-matroid-intersection coloring.*
- (b) *There is a polytime algorithm that returns a $k(k-1)\chi_{\max}$ -coloring for k -matroid-intersection coloring.*

Theorem 1.4 (a), which is an immediate corollary of Theorem 1.4 (b), resolves Question 1 for the case of two matroids. For $k > 2$, while we do not match the guarantee in Theorem 1.2, the $k(k-1)\chi_{\max}$ bound in Theorem 1.4 (b) significantly improves upon the previously-known approximation ratio of $O(k \log n)$ (via a reduction to set cover) for k -matroid-intersection coloring. In particular, this yields the first $O(1)$ -approximation algorithm for any fixed k .

In contrast to prior combinatorial approaches, which (as noted earlier) do not seem amenable to handling general matroids, our results are obtained by moving to an LP-relaxation of a matroid-intersection view of the problem, and leveraging an LP-rounding algorithm for matroid intersection (with multiple matroids). The rounding algorithm does not quite yield a coloring, but we identify suitable structure in the output of the rounding algorithm that enables us to convert the “approximate” coloring returned by the algorithm into a valid coloring. We elaborate further on these ideas in Section 2.1.

Our second main result is a constructive version of Theorem 1.3, thus fully resolving Question 2. Recall that in Rota’s Basis Conjecture (Conjecture 1), we have a matroid $M_1 = (U, \mathcal{I}_1)$ and r disjoint bases $B^1, \dots, B^r$, where $r = r_{M_1}(U)$, and to cast this in the setup of matroid-intersection coloring, we take M_2 to be the partition matroid encoding that at most one element is selected from each B^j basis. Let $n = |U|$.

Theorem 1.5. *For any $\varepsilon > 0$ and any instance of Rota’s Basis Conjecture where $\chi(M_1)$ (or equivalently the rank r) is sufficiently large relative to ε , there is a randomized algorithm with running time $\text{poly}(n, \frac{1}{\varepsilon})$ that returns a $(1 + \varepsilon)\chi_{\max}$ -coloring of $M_1 \cap M_2$ with high probability.⁶*

Theorem 1.5 actually follows from a much-more general result that we obtain. We develop a *fully polynomial randomized approximation scheme* (FPRAS) for two-matroid-intersection coloring with *any* two matroids (as opposed to M_2 being a partition matroid), when $\chi_{\max}$ is sufficiently large.

Theorem 1.6. *For any $\varepsilon > 0$ and any instance M_1, M_2 of two-matroid-intersection coloring with $\chi_{\max} \geq \frac{C \log n}{\varepsilon^5}$, where C is an absolute constant, there is a randomized algorithm with running time $\text{poly}(n, \frac{1}{\varepsilon})$ that returns a $(1 + \varepsilon)\chi_{\max}$ -coloring with high probability.*

Corollary 1.7. *Taking $\varepsilon = (\frac{C \log n}{\chi_{\max}})^{1/5}$ in Theorem 1.6, we obtain a coloring using $\chi_{\max} + O(\log^{1/5} n) \cdot \chi_{\max}^{4/5}$ colors with high probability; this translates to $(1 + o(1))\chi_{\max}$ colors when $\chi_{\max} = \omega(\log n)$.*

Theorem 1.5 is an immediate corollary of Theorem 1.6 since in Rota’s Basis Conjecture, we have $\chi_{\max} = \sqrt{n}$. Theorem 1.6 improves upon Theorem 1.3 in various ways: (i) first, it yields an *efficient* algorithm; (ii) second, it generalizes Theorem 1.3 by virtue of handling *arbitrary* instances of two-matroid-intersection coloring; (iii) third, the asymptotic dependence on $\chi_{\max}$, i.e., the $o(1)$ term in $(1 + o(1))\chi_{\max}$, is explicitly spelled out; (iv) our proof is substantially simpler than the proof of Theorem 1.3. As we discuss in Section 2.2, our algorithm and its analysis are based on a surprisingly clean approach that exploits polyhedral results on matroid intersection, in particular, the powerful *swap-rounding* technique in [CVZ11], to iteratively “peel” out common independent sets. In essence, we gain a striking amount of leverage from the use of this polyhedral rounding technique, and this is directly responsible for the simplicity of our proof.

One interesting insight to emerge from our work (specifically Theorem 1.6) is that asymptotically-optimal coloring is possible in the setting of Rota’s Basis Conjecture *not because one of the matroids is a partition matroid, but because $\chi_{\max}$ is sufficiently large*. This is again a direct consequence of the polyhedral approach we utilize, which is essentially opaque to the structure of the underlying matroids.

⁶It suffices to ensure $\frac{1}{\text{poly}(n)}$ success probability, since we can then boost the success probability to $1 - \frac{1}{n^\alpha}$ for any $\alpha > 0$ by running the algorithm $\text{poly}(n) \cdot O(\log n)$ times.

2 Technical Overview

At a high level, the main challenge in proving the existence of good colorings for matroid-intersection coloring, is that whether two elements of the ground set can be colored the same color is quite context-dependent, and depends on which other elements share that color. Taking the covering perspective, it is not clear how to compute common independent sets which reduce the chromatic number on the remaining uncovered elements. The proof of Theorem 1.4 surmounts the coloring challenge, and the proof of Theorem 1.6 surmounts the covering challenge. We proceed to highlight the main insights and ideas in these proofs.

2.1 Coloring k Matroids: Theorem 1.4 (see Section 4)

Recall that prior algorithmic results focus on the setting where all, or almost all, matroids [IMP21, BSY21, LMP21, AMPZ26] are partition matroids or are reducible to partition matroids. These techniques exploit the fact that, in partition matroids, it is substantially easier (compared to general matroids) to understand (and hence control) when a set is independent. This partition-reduction approach can handle certain combinatorial classes of matroids, but there are known classes where reduction to a partition matroid drastically increases the chromatic number [LMP22, AKKG23].

Our approach is different, but fairly natural. Let $q = \chi_{\max}$. Our algorithm consists of three main steps. First, we utilize a standard observation to cast the problem as an instance of $(k + 1)$ -matroid intersection. Second, we utilize an LP-rounding algorithm for this matroid-intersection problem. This algorithm does not, however, produce a feasible solution. While its output can be viewed as a kind of “approximate coloring”, this by itself is too weak a guarantee for our purposes. The third key step, which is our main conceptual contribution, is that we identify suitable structure in the output of the LP-rounding algorithm, and show that one can exploit this structure to set up a *graph-coloring* problem to resolve conflicts in the approximate coloring and obtain a valid coloring.

We now elaborate on these steps. The reduction to matroid intersection involves moving to a new ground set U' consisting of q disjoint copies $U_1, \dots, U_q$ of the original ground set U , and considering the disjoint union of q copies of each input matroid M_i , which we denote M'_i . We introduce an additional matroid M'_0 which encodes for each element in U , that at most one copy of each element $e \in U$ is picked. It is easy to see now that a valid q -coloring maps to a subset $R \subseteq U'$ which is a basis of M'_0 and independent in M'_i for each i , and also vice versa.

We consider the natural LP-relaxation of the problem and utilize the LP-rounding algorithm of [LOSZ20]. Their algorithm, which we refer to as the LOSZ-algorithm, outputs a basis R of M'_0 such that each “color class” $R \cap U_c$ is k -colorable in M_i for all $i \in [k]$, $c \in [q]$. However, this guarantee alone is too coarse for our purposes: the issue is that being k -colorable in each individual matroid does not mean that we can efficiently obtain a good coloring in their intersection. In general, all one can glean from this guarantee, by way of efficiently producing a coloring, is a $(k^k \cdot q)$ -coloring by partitioning $R \cap U_c$ into k independent sets of M_i , for all $i \in [k]$ and taking all the resulting intersections.⁷ Thus, this only yields a k^k -approximation algorithm; while this is a constant for fixed k , it is quite far from the $k(k - 1)\chi_{\max}$ -coloring that we are aiming for. Moreover, for $k = 2$, there are instances where $\chi(M_1) = \chi(M_2) = 2$, but $\chi(M_1 \cap M_2) > 2$, so 2-colorability in each individual matroid would not by itself yield the $2\chi_{\max}$ -coloring for 2-matroid-intersection coloring (Theorem 1.4(a)). The upshot is that the guarantee of [LOSZ20] by itself is insufficient for obtaining our approximation results.

We make progress via two chief insights. Our first insight is that one can identify much more structure in the set output by the LOSZ algorithm, encapsulated by a property that we call *flexible decomposition* (see Definition 4.1). We say that a set S admits a p -flexible decomposition in a matroid M if it has a partition

⁷As noted earlier, while it is true here that $R \cap U_c$ can be covered using $O(k^2)$ common independent sets (Theorem 1.2), the proof of this is quite nonconstructive and does not lead to an efficient algorithm.

$T_1, \dots, T_\ell$ such that every T_j part contains an M -independent set of size at least $|T_j| - (p - 1)$, and (*) the union of *any* combination of M -independent sets from these T_j parts yields an M -independent set. Observe that taking $p = 1$ corresponds precisely to saying that S is independent in M ; on the other hand, one can also infer that if S has a p -flexible decomposition in M , then S is p -colorable in M . Thus, p -flexible decomposition relaxes independence, but is a *strengthening* of p -colorability.

We prove that the set R output by the LOSZ algorithm is such that $R \cap U_c$ admits a k -flexible decomposition in M_i for all $i \in [k]$, $c \in [q]$, thereby strengthening the guarantee of [LOSZ20]. We refer to this compactly by saying that R is a (k, q) -pseudocoloring⁸ (see Definition 4.2 and Theorem 4.3).

Our second key insight is that the stronger guarantee provided by a (k, q) -pseudocoloring can be leveraged to convert this pseudocoloring into a $k(k - 1)q$ -coloring of $M_1, \dots, M_k$ (Theorem 4.6). We argue that, for each color class $c \in [q]$, one can create a suitable conflict graph G_c with vertex set $R \cap U_c$ such that: (a) G_c admits an efficiently-computable $k(k - 1)$ (vertex) coloring; and (b) such a coloring translates to a cover of $R \cap U_c$ by $k(k - 1)$ common independent sets. Thus, we obtain a $k(k - 1)q$ -coloring.

Roughly speaking, for a color class $c \in [q]$ and matroid M_i where $i \in [k]$, in order to ensure independence in M_i , due to property (*) it suffices to focus on a single part T_j of the k -flexible decomposition of $R \cap U_c$ in M_i and cover T_j by independent sets of M_i . Using matroid-exchange properties, this can be captured by defining suitable conflict pairs, encoded by edges involving elements in T_j . The union of these (vertex-disjoint) graphs for all parts T_j yields a graph H_i with maximum degree at most $k - 1$, and since we need to ensure independence in all matroids, we take the union of H_i over all $i \in [k]$ to obtain the conflict graph G_c . Thus, G_c has maximum degree at most $k(k - 1)$, and a simple greedy algorithm yields a $k(k - 1) + 1$ coloring of G_c . We improve this using (an algorithmic proof of) Brooks' theorem (Theorem 4.8) to a $k(k - 1)$ -coloring by proving that G_c does not contain a clique on $k(k - 1) + 1$ nodes. (For $k = 2$, there is a simpler argument showing that G_c is bipartite, and hence 2-colorable.)

It is instructive to compare, at a high level, our approach, with the partition-reduction approaches in prior work which achieve good approximations for combinatorial matroids. These approaches reduce to partition matroids at the very outset, yielding a simple sufficient condition for independence and allowing the resulting instance to be handled by combinatorial means. In contrast, our approach for general matroids leverages a reduction to matroid intersection and the LP-relaxation for the latter problem. The flexible-decomposition structure that we identify results from the output of the LP rounding algorithm and is not at all apparent at the beginning. Thus, we rely heavily on, and are guided by, the LP and the LP-rounding algorithm to isolate the desired structure, which we finally capitalize on to “clean up” the pseudocoloring and obtain a genuine coloring. We note that the use of graph coloring in the final clean-up step, that is, encoding independence via conflict-pair avoidance, is also a component of the partition-reduction approach.

2.2 Coloring Two Matroids when $\chi_{\max}$ is Large: Theorem 1.6 (see Section 5)

Here we give a technical overview of the proofs underlying the constructive result in Theorem 1.6.

From Edmonds' Condition (see ?? 3.1), we know that $\chi_{\max}$ is the maximum over all subsets S of U and over all matroids M_i , of the density of S in M_i , which is the ratio $|S|/r_i(S)$ of its cardinality to its rank in M_i . So, a natural approach to show the existence of a good coloring is to iteratively decide on some subsets of elements to be color classes, in such a way as to reduce the densities of the high density subsets of the remaining elements. This is the approach taken in the nonconstructive proof of Theorem 1.3 in [MS25]. However, the authors note that keeping track of these densities is very difficult, and this is one of the main reasons that their proof is so complex. We avoid this combinatorial difficulty using a polyhedral approach combined with randomization.

⁸Note that a $(1, q)$ -pseudocoloring is a q -coloring, so a (p, q) -pseudocoloring is a relaxation of q -coloring (as the name suggests).

It is helpful to first consider a simple, but flawed, randomized strategy. Suppose we sample each element independently and uniformly at random with probability p . It is not too difficult to show, via standard concentration bounds and a union bound, that if we delete these sampled elements, the density of *every* high density set (those with density close to $\chi_{\max}$) decreases to $(1-p+o(1))\chi_{\max}$ with high probability if $\chi_{\max} = \omega(\log n)$. In particular, if we choose $p = N/\chi_{\max}$ for some sufficiently large N , the maximum density drops by approximately N . The core insight here is that we do not need to identify or specifically target high density sets S . Because this “blind” uniform deletion effectively hits all sets with high concentration, it universally reduces the maximum density. This would be a perfect strategy if the randomly sampled set using $p = N/\chi_{\max}$ were a collection of N common independent sets in the matroid intersection. Indeed, we could simply color the sampled elements with N colors, and recurse on a subproblem with a chromatic number that is N smaller, implying that only $\chi_{\max} + o(\chi_{\max})$ iterations suffice. Of course, this method of independent sampling makes no structural progress because the resulting sampled set need not be a collection of independent sets in either matroid, let alone their intersection. However, this uniform hitting property suggests the utility of random sampling. So, we need a random sampling method to generate common independent sets that has (nearly) uniform marginals and the strong concentration of independent sampling.

Fortunately, the swap rounding technique of Chekuri, Vondrák, and Zenklusen [CVZ11] gives us exactly this type of guarantee. Swap rounding rounds a fractional point x in the matroid intersection polytope P to an integer point x_R corresponding to a common independent set R . Further, $\mathbb{E}[x_R] \approx x$, and *every* linear function on the coordinates of x_R strongly concentrates around its expectation. Our algorithm will take x to be the uniform fractional point in which each component is $1/\chi_{\max}$. It is easy to see that this point is in the matroid intersection polytope P (see ?? 3.2). It might appear that at first glance this point x is combinatorially “featureless”, revealing no information about the structure of a feasible integer coloring. But when swap rounding is applied N times to the uniform fractional point x , we obtain our desired properties, namely that every element is included with marginal probability approximately $N/\chi_{\max}$, *and* concentration bounds on every subset of elements comparable to independent random sampling.

Given the complexity of the nonconstructive proof of Theorem 1.3, our $(1+\varepsilon)$ -approximation algorithm is surprisingly clean. For $N = \lceil \varepsilon \chi_{\max} \rceil$, we just apply swap rounding N times to the uniform fractional point x , in which each component is $1/\chi_{\max}$, to obtain common independent sets $R_1, R_2, \dots, R_N$. The elements in each R_i are then colored the same color. We then “peel” these colored elements away, re-evaluate the maximum density of the remaining uncolored elements, and recurse on the still uncolored elements. Because swap rounding guarantees tight concentration, we are able to argue that the maximum density drops reliably by roughly a $1 - \varepsilon$ factor in each step.

3 Preliminaries and notation

For an integer $k \geq 1$, we use $[k]$ to denote the set $\{1, 2, \dots, k\}$. We use $\mathbb{R}_+$ to denote the nonnegative reals; similarly $\mathbb{Z}_+$ denotes the nonnegative integers. Recall that in the k -matroid-intersection coloring problem, we are given k matroids $M_1 = (U, \mathcal{I}_1), \dots, M_k = (U, \mathcal{I}_k)$ over a common ground set U , and we seek to cover U using the fewest number of common independent sets of these k matroids; a covering using q common independent sets is called a q -coloring of $M_1, \dots, M_k$.

More generally, given an *independence system* $N = (U, \mathcal{I})$ (i.e., $\mathcal{I}$ is a downwards-closed collection of sets), and a set $S \subseteq U$, we say that S is q -colorable in $(U, \mathcal{I})$, or admits a q -coloring in $(U, \mathcal{I})$, if S can be covered by q sets in $\mathcal{I}$. Thus, in the k -matroid-intersection coloring problem, we seek a q -coloring of U in the independence system $M_{\text{int}} = (U, \bigcap_{i \in [k]} \mathcal{I}_i)$ for the smallest value of q . The chromatic number, or coloring number, of N , denoted $\chi(N)$, is the smallest q for which there is a q -coloring of U in N .

We assume throughout that every matroid M we work with is loopless, that is, $\{e\}$ is independent in the

matroid, for every element e in its ground set. (Note that otherwise, no coloring of M exists.)

Matroid preliminaries. We collect a few relevant facts about matroids. Let $M = (U, \mathcal{I})$ be a matroid with rank function $r : 2^U \mapsto \mathbb{Z}_+$. Two polytopes commonly associated with M are:

- (a) its *matroid polytope*, $\mathcal{P}(M)$, which is the convex hull of indicator vectors of independent sets, which admits the inequality description $\mathcal{P}(M) = \{x \in \mathbb{R}_+^U : x(S) \leq r(S) \ \forall S \subseteq U\}$; and
- (b) its *matroid-base polytope*, $\mathcal{P}_B(M)$, which is the convex hull of indicator vectors of bases of M , and is given by $\{x \in \mathcal{P}(M) : x(U) = r(U)\}$.

Fact 3.1 (Edmonds' Condition [Edm68, Sch03]). *Let $M = (U, \mathcal{I})$ be a matroid with rank function $r : 2^U \rightarrow \mathbb{Z}_+$. Then $\chi(M) = \max_{S \subseteq U} \lceil \frac{|S|}{r(S)} \rceil$.*

Claim 3.2. *For any $q \geq \chi(M)$, we have $x = \mathbb{1}/q \in \mathcal{P}(M)$, where $\mathbb{1}$ is the all 1's vector in $\mathbb{R}^U$.*

Proof. By Fact 3.1, $q \geq \chi(M)$ implies that $|S| \leq q \cdot r(S)$ for all $S \subseteq U$, i.e., $\mathbb{1}/q \in \mathcal{P}(M)$. $\square$

Let $S \subseteq U$. The *refinement of M along S* yields the matroids $M_1 = M|_S = (S, \mathcal{I} \cap 2^S)$, which is the restriction of M to S , and $M_2 = M/S$, which is the matroid obtained from M by contracting S . We have $M_2 = (U - S, \{A \subseteq U - S : A \cup J \in \mathcal{I}\})$, where J is a maximal independent set contained in S ; it is well known that the definition of M/S does not depend on the choice of which maximal independent subset of S is chosen as J . Further, the rank function of the matroid contraction M/S is given by $r_{M/S}(A) = r(A \cup S) - r(S)$ for all $A \subseteq U - S$.

4 k -Matroid-Intersection Coloring: Proof of Theorem 1.4

We now present the $k(k-1)$ -approximation algorithm for k -matroid-intersection coloring, establishing Theorem 1.4. Recall that $M_i = (U, \mathcal{I}_i)$, $i \in [k]$, are the k input matroids, and $\chi_{\max} := \max_{i \in [k]} \chi(M_i)$. Our approximation guarantees are relative to $\chi_{\max}$, and as we discuss in Section 4.3, also yield improved integrality-gap bounds for the natural set-covering LP-formulation of the problem.

As outlined in Section 2.1, we leverage the reduction from q -coloring to a $(k+1)$ -matroid-intersection problem. Specifically, we consider a ground set U' consisting of q disjoint copies of U . We define M'_0 as the partition matroid on U' that ensures at most one copy of each original element $e \in U$ is selected, and for $i \in [k]$, we let M'_i be the disjoint union of q copies of the original matroid M_i . As established previously, finding a q -coloring of $M_1, \dots, M_k$ is equivalent to finding a basis R of M'_0 that is independent in $M'_1, \dots, M'_k$. We take $q = \chi_{\max}$, which ensures that the LP-relaxation of this $(k+1)$ -matroid-intersection problem is feasible (by Claim 3.2).

We apply the iterative refinement algorithm of [LOSZ20] (the LOSZ-algorithm) to this instance. The set R output by the LOSZ algorithm does not yield a coloring. While the analysis in [LOSZ20] ensures that each resulting “color class” $R \cap U_c$ is k -colorable in each M_i individually, as discussed earlier, this is insufficient for obtaining even a $\text{poly}(k)$ -approximation factor for k -matroid-intersection coloring. We instead prove that the LOSZ-algorithm has much more structure than previously known. We call this structure *flexible decomposition*, and exploit this structure to resolve conflicts in each color class via graph coloring.

Definition 4.1 (Flexible decomposition). Let $M = (U, \mathcal{I})$ be a matroid and $S \subseteq U$. Let $p \geq 1$ be an integer. We say that S admits a p -flexible decomposition $T_1, T_2, \dots, T_\ell$ in M , or that $T_1, \dots, T_\ell$ is a p -flexible decomposition of S in M , if: (a) $T_1, T_2, \dots, T_\ell$ is a partition of S ; (b) $r(T_j) \geq |T_j| - p + 1$ for all $j \in [\ell]$; and (c) for any collection of independent sets $I_j \subseteq T_j$ for $j \in [\ell]$, we have that $\bigcup_{j \in [\ell]} I_j \in \mathcal{I}$.

Unless otherwise stated, we will also require that the partition $T_1, \dots, T_\ell$ can be *efficiently* computed. We will often say “ S admits a p -flexible decomposition in M ” to refer to the fact there is an underlying (efficiently-computable) partition $T_1, \dots, T_\ell$ of S satisfying properties (b), (c).

Note that $S \in \mathcal{I}$ iff it admits a 1-flexible decomposition: if $S \in \mathcal{I}$, then the trivial partition comprising just the set S is a 1-flexible decomposition; conversely, if $T_1, \dots, T_\ell$ is a 1-flexible decomposition of S , then properties (b), (c) imply that $\bigcup_{j \in [\ell]} T_j \in \mathcal{I}$. Note also that if S has a p -flexible decomposition in M , then S is p -colorable in M : the combination of the $r(T_j)$ -size independent sets in the T_j parts, along with any combination of (the at most $p - 1$) elements excluded from these independent sets where we take (at most) one element from each part, yields a p -coloring in M . Thus, p -flexible decomposition strengthens the notion of p -colorability.

In Section 4.1, we show that the set R output by the LOSZ algorithm is such that $R \cap U_c$ admits a k -flexible decomposition in M_i for all $i \in [k]$, $c \in [q]$ (Theorem 4.3), strengthening the guarantee in [LOSZ20]. We call this notion of “approximate coloring” based on flexible decomposition a *pseudocoloring*.

Definition 4.2 (Pseudocoloring). Let $M_i = (U, \mathcal{I}_i)$, $i = 1, \dots, k$ be k matroids. Let $p, q \geq 1$ be integers. A (p, q) -pseudocoloring of $M_1, \dots, M_k$ is a partition $R_1, \dots, R_q$ of U such that R_c admits a p -flexible decomposition in M_i , for all $i \in [k]$, $c \in [q]$. We also sometimes say that $R_1, \dots, R_q$ is a (p, q) -pseudocoloring of $M_1, \dots, M_k$.

We next show in Section 4.2 that one can exploit this pseudocoloring and turn it into a valid coloring incurring a $k(k - 1)$ multiplicative-factor blow-up in the number of colors. We solve a *graph coloring problem* to achieve this end. For each color class $c \in [q]$, we create a suitable conflict graph G_c with vertex set $R \cap U_c$ and argue that: (a) we can efficiently compute a $k(k - 1)$ (vertex) coloring of G_c ; and (b) such a coloring translates to a cover of $R \cap U_c$ by $k(k - 1)$ common independent sets. Thus, since we start off with $q = \chi_{\max}$ color classes, we obtain a $k(k - 1)\chi_{\max}$ -coloring of $M_1, \dots, M_k$.

4.1 Flexible Decomposition via LP-Rounding and Iterative Refinement

As noted earlier, $S \subseteq U$ admits a 1-flexible decomposition in a matroid $M = (U, \mathcal{I})$ iff it is independent in M . Thus, a $(1, q)$ -coloring corresponds precisely to a q -coloring of $M_1, \dots, M_k$, and so a pseudocoloring is a relaxation of coloring.

Theorem 4.3. *Given matroids $M_i = (U, \mathcal{I}_i)$ for $i = 1, \dots, k$, we can efficiently find a $(k, \chi_{\max})$ -pseudocoloring of $M_1, \dots, M_k$.*

Theorem 4.3 will follow readily from the following theorem, which strengthens Theorem 2 in [LOSZ20] and is the main technical result of this section.

Theorem 4.4. *Let $M'_0 = (U'_0, \mathcal{I}'_0)$, and $M'_i = (U'_i, \mathcal{I}'_i)$, $i \in [k]$ be matroids, where $U'_i \subseteq U' := U'_0$ for all $i \in [k]$, and $w \in \mathbb{R}^{U'}$. Consider the following LP.*

$$\max \quad w^T x \quad \text{s.t.} \quad x \in \mathcal{P}_B(M'_0), \quad x|_{U'_i} \in \mathcal{P}(M'_i) \quad \forall i \in [k]. \quad (\text{LP}_{\text{mat}})$$

where $x|_S$ denotes the vector $(x_e)_{e \in S}$, for $S \subseteq U'$. Let $p_1, p_2, \dots, p_k$ be positive integers such that $\sum_{i \in [k]: e \in U'_i} \frac{1}{p_i} \leq 1$ for all $e \in U'$.

If (LP_{mat}) is feasible, then one can efficiently compute $R \subseteq U'$ such that (a) R is a basis of M'_0 ; (b) $w(R) \geq \text{OPT}_{\text{LP}_{\text{mat}}}$; and (c) $R \cap U'_i$ has a p_i -flexible decomposition in M'_i for all $i \in [k]$.

We prove Theorem 4.4 shortly, but first we show how this easily yields Theorem 4.3 as a corollary.

Claim 4.5. *Let $M = (U, \mathcal{I})$ be a matroid. Suppose $S \subseteq U$ has a p -flexible decomposition in M , and $A \subseteq S$. Then, A also has a p -flexible decomposition in M .*

Proof. Let $T_1, \dots, T_\ell$ be a p -flexible decomposition of S in M . Then it is easy to see that $\{T_j \cap A\}_{j \in [\ell]}$ (after discarding any empty sets) yields a p -flexible decomposition of A in M . $\square$

Proof of Theorem 4.3. We reduce the coloring problem to a matroid-intersection problem as discussed at the beginning of the section, and invoke Theorem 4.4 on this input. Let $q = \chi_{\max}$. Let U_c be a disjoint copy of U for all $c \in [q]$ and $U' = U'_0 = \bigcup_{c \in [q]} U_c$. Let $M'_0 = (U', \mathcal{I}'_0)$ be the partition matroid encoding that for every $e \in U$, at most one copy of e from $U_1, \dots, U_q$ is picked. For each $i \in [k]$, let $M'_i = (U', \mathcal{I}'_i)$ be the disjoint union of q copies of M_i , where the c -th copy resides over the ground set U_c for all $c \in [q]$.

The weight vector w will be irrelevant for our purposes, and we can take w to be arbitrary.

Observe that $x = \mathbb{1}/q$ is a feasible solution to (LP_{mat}) . Clearly, we have $x \in \mathcal{P}_B(M'_0)$ for the partition matroid M'_0 . For a matroid M'_i , $i \in [k]$, since $\chi(M_i) \leq q$, we also have that $\chi(M'_i) \leq q$, due to the definition of disjoint union. It follows from Claim 3.2 that $x \in \mathcal{P}(M'_i)$.

Let $R \subseteq U'$ be the output of Theorem 4.4 for the $M'_0, M'_1, \dots, M'_k$ matroids, taking $p_i = k$ for $i \in [k]$. We argue that $\{R \cap U_c\}_{c \in [q]}$, where we interpret each $R \cap U_c$ set as a subset of U , is a $(k, \chi_{\max})$ -pseudocoloring of $M_1, \dots, M_k$. Since R is a basis of M'_0 , it is immediate that $\{R \cap U_c\}_{c \in [q]}$ is a partition of U . Fix an index $i \in [k]$. Part (c) of Theorem 4.4 yields that $R \cap U'_i$ has a k -flexible decomposition in M'_i . By Claim 4.5, this implies that $R \cap U_c$ has a k -flexible decomposition in M'_i for all $c \in [q]$. Since M'_i restricted to U_c is simply a copy of M_i , this amounts to saying that $R \cap U_c$ has a k -flexible decomposition in M_i for all $c \in [q]$. $\square$

Proof of Theorem 4.4

As mentioned earlier, the algorithm for producing R is the same as the iterative-refinement based LP-rounding algorithm in [LOSZ20], which we describe below for completeness.

Algorithm LOSZ	// Iterative-refinement algorithm in [LOSZ20]
1: Initialize $\mathcal{M} = \{M'_1, \dots, M'_k\}$, $\widetilde{M} = (\widetilde{U}, \widetilde{\mathcal{I}}) \leftarrow M'_0$, $p_{M'_i} = p_i$ for all $i \in [k]$, and $R \leftarrow \emptyset$. 2: Compute an extreme-point optimal solution x^* to (LP_{mat}) for the matroids $\{\widetilde{M}\} \cup \mathcal{M}$. 3: For every $e \in \widetilde{U}$ with $x_e^* = 0$, delete e from $\widetilde{M}$ and all matroids in $\mathcal{M}$ containing e . 4: For every $e \in \widetilde{U}$ with $x_e^* = 1$, contract e in $\widetilde{M}$ and all matroids in $\mathcal{M}$ containing e ; also add e to R . 5: If the ground-set of some matroid in $M \in \mathcal{M}$ becomes empty as a result of these deletions and contractions, remove M from $\mathcal{M}$. $\triangleright$ Note that deletions and contractions remove elements from $\widetilde{U}$. 6: if $\widetilde{U} = \emptyset$ then return R . 7: while there is a matroid $M = (U, \mathcal{I}) \in \mathcal{M}$, set $S \subsetneq U$, $S \neq \emptyset$ with $x^*(S) = r_M(S)$ do $\triangleright r_M$ is rank f'n. of M 8: Update $\mathcal{M} \leftarrow \mathcal{M} \setminus \{M\} \cup \{M _S, M/S\}$ and set $p_{M _S} = p_{M/S} = p_M$. 9: Find a matroid $M = (U, \mathcal{I}) \in \mathcal{M}$ with $ U \leq r_M(U) + p_M - 1$ and remove M from $\mathcal{M}$. goto step 2.	

[LOSZ20] show that the algorithm is well-defined, terminates in polynomial time, and that parts (a) and (b) of the theorem statement hold. We do not repeat their arguments here and only note that (a) and (b) follow because we only pick elements with $x_e^* = 1$, and x^* yields a feasible solution to the new LP that is solved when we go to step 2. We focus on proving part (c), which strengthens the claim in [LOSZ20] that $R \cap U'_i$ is p_i -colorable in M'_i for all $i \in [k]$.

Note that each matroid in $\mathcal{M}$ arises from a specific input matroid. Fix an input matroid M'_i . Note that any matroid M that arises from M'_i via refinement in step 8 or contraction in step 4 has $p_M = p_i$. We can describe the sequence of matroids that arise from M'_i via a tree. The nodes of this tree are essentially the matroids arising from M'_i that are present in $\mathcal{M}$ at some point in the algorithm. The root node is M'_i . Consider a node $M = (U, \mathcal{I})$ of the tree. If M and $S \subseteq U$ are chosen in step 7 for refinement, then the children of M are the matroids $M|_S$ and M/S . It will be convenient to also include matroids obtained via

contraction as part of this tree. If we contract $S \subsetneq U$ in step 4, the children of M are the free matroid $(S, 2^S)$ (which is the same as $M|_S$, since $S \in \mathcal{I}$) and M/S . (In essence, contracting S can also be viewed as refining M along S .) The free matroid on S is not part of $\mathcal{M}$, so does not create any children and is a leaf of the tree; also if $S = U$ is contracted, then M gets removed from $\mathcal{M}$ in step 5 and is a leaf of the tree. Observe that every leaf of this tree is a free matroid, or a matroid removed from $\mathcal{M}$ in step 9. Clearly, this tree can be constructed efficiently as the algorithm unfolds.

We now argue that for every node $M = (U, \mathcal{I})$ of the tree, $R \cap U$ admits a p_i -flexible decomposition in M , by induction starting from the leaves of the tree. Applying this to the root node shows that $R \cap U'_i$ has a p_i -flexible decomposition in M'_i . For the base case, when M is a leaf, this is certainly true if M is a free matroid since then $R \cap U \in \mathcal{I}$. Otherwise, M must have been dropped from $\mathcal{M}$ in step 9 and we have $p_M = p_i$, in which case the trivial partition comprising the set $R \cap U$ yields a p_i -flexible decomposition.

Next, suppose that we have a non-leaf node $M = (U, \mathcal{I})$ with children $\overline{M} = M|_S$ (with ground-set S), $M'' = M/S$ (with ground-set $U - S$) for some $S \subsetneq U$, $S \neq \emptyset$. Inductively, we have that $R \cap S$ admits a p_i -flexible decomposition $A_1, \dots, A_\ell$ in $\overline{M}$, and $R \cap (U - S)$ admits a p_i -flexible decomposition $B_1, \dots, B_m$ in M'' . We claim that $A_1, \dots, A_\ell, B_1, \dots, B_m$ is a p_i -flexible decomposition of $R \cap U$ in M .

It is clear that $A_1, \dots, A_\ell, B_1, \dots, B_m$ partition $R \cap U$. We have $r_{\overline{M}}(T) = r_M(T)$ for all $T \subseteq S$ and $r_{M''}(T) = r_M(T \cup S) - r_M(S) \leq r_M(T)$ for all $T \subseteq U - S$. By property (b) of a p_i -flexible decomposition, we have $|A_j| \leq r_{\overline{M}}(A_j) + p_i - 1 = r_M(A_j) + p_i - 1$ for all $j \in [\ell]$. Similarly, we have $|B_j| \leq r_{M''}(B_j) + p_i - 1 \leq r_M(B_j) + p_i - 1$ for all $j \in [m]$. So $A_1, \dots, A_\ell, B_1, \dots, B_m$ satisfy property (b) as well. Finally, let $I_1, \dots, I_\ell$ be $\overline{M}$ -independent sets such that $I_j \subseteq A_j$ for all $j \in [\ell]$, and $I'_1, \dots, I'_m$ be M'' -independent sets such that $I'_j \subseteq B_j$ for all $j \in [m]$. Then, by property (c) of p_i -flexible decomposition, $I = \bigcup_{j \in [\ell]} I_j$ is independent in $\overline{M}$, and hence $I \in \mathcal{I}$, and $I'' = \bigcup_{j \in [m]} I'_j$ is independent in M'' . Since $M'' = M/S$, it follows that $I \cup I'' \in \mathcal{I}$. Thus, $A_1, \dots, A_\ell, B_1, \dots, B_m$ satisfy property (c), completing the induction step.

The proof above also shows that the flexible decomposition can be efficiently constructed. Indeed, the p_i -flexible decomposition of $R \cap U'_i$ is simply given by the $R \cap U$ sets corresponding to the leaf matroids $M = (U, \mathcal{I})$ of the tree constructed for M'_i . $\square$

4.2 Converting a Pseudocoloring into a Valid Coloring

We now show how to convert the pseudocoloring given by Theorem 4.3 to a valid coloring.

Theorem 4.6. *Given a (k, q) -pseudocoloring $R_1, \dots, R_q$ of $M_1 = (U, \mathcal{I}_1), \dots, M_k = (U, \mathcal{I}_k)$, one can efficiently obtain a $k(k-1)q$ -coloring of $M_1, \dots, M_k$.*

Combining this with Theorem 4.3, we immediately obtain a $k(k-1)\chi_{\max}$ -coloring of $M_1, \dots, M_k$ in polynomial time, thereby proving Theorems 1.4 (a) and 1.4 (b). The remainder of this section is devoted to the proof of Theorem 4.6.

Recall that $R_1, \dots, R_q$ being a (k, q) -pseudocoloring means that R_c has a polynomial time-computable k -flexible decomposition in M_i for all $i \in [k]$, $c \in [q]$. For each $c \in [q]$, we will define a conflict graph G_c with vertex-set R_c such that any vertex coloring of G_c using some s colors will translate to a cover of R_c by s common independent sets of $M_1, \dots, M_k$. Moreover, we will argue that G_c can be efficiently colored using at most $k(k-1)$ colors. This yields a $k(k-1)q$ -coloring of $M_1, \dots, M_k$, proving Theorem 4.6. This translation from graph coloring to matroid-intersection coloring is *enabled* by the structure imposed by flexible decompositions, which we crucially exploit.

In the sequel, we fix a color class R_c , where $c \in [q]$, and define G_c and prove the above two properties for G_c . Let $T_1^i, \dots, T_{\ell(i)}^i$ be a polynomial time-computable k -flexible decomposition of R_c in M_i , for $i \in [k]$. For each matroid $i \in [k]$, define the following graph $H_i = (R_c, E_i)$ with vertex-set R_c . For each $j \in [\ell(i)]$,

let $A_j^i \subseteq T_j^i$ be an independent set of M_i with $|A_j^i| \geq |T_j^i| - k + 1$. Such a set is guaranteed by property (b) of flexible decomposition (see Definition 4.1), and can be easily found in polynomial time. By the matroid-exchange property, for every $e \in T_j^i - A_j^i$, we can identify some $e' \in A_j^i$ such that $A_j^i \cup \{e\} - \{e'\} \in \mathcal{I}_i$; we include (e, e') as an edge in E_i . We also include edges between all pairs of elements in $T_j^i - A_j^i$ in E_i . We do this for all $j \in [\ell(i)]$ to obtain the graph H_i . Thus, H_i is the union of vertex-disjoint subgraphs, one for each part T_j^i of the k -flexible decomposition in M_i .

The conflict graph G_c for color class R_c is the union $(R_c, \bigcup_{i \in [k]} E_i)$ of all H_i 's. (Note that we include only one copy of an edge even if the edge is present in multiple H_i graphs.)

Lemma 4.7 argues that a vertex coloring of G_c yields a cover of R_c by common independent sets, and Lemma 4.9 shows that G_c can be colored using at most $k(k-1)$ colors.

Lemma 4.7. *Let $C_1, \dots, C_s$ be a vertex coloring of G_c using s colors. Then, $C_1, \dots, C_s$ is an s -coloring of $M_1, \dots, M_k$.*

Proof. Consider any $h \in [s]$. We argue that C_h is a common independent set of $M_1, \dots, M_k$, which will prove the lemma. We have that C_h is a stable set of G_c , i.e., there are no edges in G_c between any two nodes in C_h . We claim that $C_h \cap T_j^i$ is independent in M_i for all $i \in [k]$ and $j \in [\ell(i)]$. By property (c) of flexible decomposition, this implies that $C_h = C_h \cap R_c = \bigcup_{j \in [\ell(i)]} (C_h \cap T_j^i)$ is independent in M_i , so it suffices to prove the claim.

To see this, recall that $A_j^i \subseteq T_j^i$ is the independent set in M_i that was used to define the graph H_i (and hence, G_c). If $C_h \cap T_j^i \subseteq A_j^i$, then we are done. Otherwise, $C_h \cap T_j^i$ contains exactly one element $e \in T_j^i - A_j^i$, since H_i contains a clique on the elements in $T_j^i - A_j^i$. In this case, if (e, e') is an edge in H_i , where $e' \in A_j^i$, then we know that $e' \notin C_h$ and $A_j^i \cup \{e\} - \{e'\} \in \mathcal{I}_i$. It follows that $C_h \cap T_j^i \subseteq A_j^i \cup \{e\} - \{e'\}$ and so is independent in M_i . $\square$

To show that G_c can be efficiently colored using $k(k-1)$ colors, we utilize Brooks' theorem from graph theory, which has a simple proof via a clever greedy algorithm due to Lovasz [Lov75] that we describe in Appendix A for completeness.

Theorem 4.8 (Brooks, 1941). *Let G be a connected graph with maximum degree $\Delta(G) > 2$. Then $\chi(G) \leq \Delta(G)$ unless $G = K_{\Delta+1}$, and a coloring of G using at most $\Delta(G)$ colors can be computed in polynomial time.*

Lemma 4.9. *The conflict graph G_c can be efficiently colored using at most $k(k-1)$ colors.*

Proof. When $k = 2$, there is a particularly simple and crisp argument that is worth noting separately. In this case, note that each H_i graph is a *matching* (as $|T_j^i - A_j^i| \leq 1$ for all $j \in [\ell(i)]$) and G_c is a union of two matchings. So every path or cycle in G_c is an alternating path or alternating cycle with respect to these matchings. Therefore, G_c is bipartite, and hence 2-colorable.

Now consider general $k \geq 3$. Note that each H_i graph has maximum degree at most $k-1$, and so $\Delta = \Delta(G_c) \leq k(k-1)$. If $\Delta < k(k-1)$, then a straightforward greedy algorithm yields a $(\Delta+1)$ -coloring. So suppose $\Delta = k(k-1)$. We argue that G_c does not contain a clique K on $\Delta+1$ nodes. Then by Brooks' Theorem (Theorem 4.8), G_c can be colored efficiently using $\Delta = k(k-1)$ colors.

For a vertex v and graph (or edge-set) F , we use $\delta_F(v)$ to denote the edges of F incident to v . Suppose, for a contradiction, that G_c contains a clique K on $\Delta+1$ nodes. Observe that the clique K is a connected component of G_c , because the maximum degree of G_c is Δ . Further, for any node $v \in K$, we have $|\delta_G(v)| = k(k-1) \leq \sum_{i \in [k]} |\delta_{H_i}(v)|$, which implies that $|\delta_{H_i}(v)| = k-1$ for all $i \in [k]$. This has to hold for all $v \in K$. Thus, H_i is a $(k-1)$ -regular graph on K for all $i \in [k]$.

Further, recall that each H_i is the union of vertex-disjoint subgraphs, one subgraph for each part T_j^i of the k -flexible decomposition of R_c in M_i . Each of these subgraphs contains a clique on $|T_j^i - A_j^i| \leq k-1$

nodes along with one extra edge (e, e') for each $e \in T_j^i - A_j^i$ to some $e' \in A_j^i$. In order for this subgraph to be $(k-1)$ -regular on K , it must be a k -clique. Thus each H_i is the union of vertex-disjoint k -cliques on K .

However, K is a clique on $\Delta + 1 = k(k-1) + 1$ nodes. Because $k \nmid k(k-1) + 1$, H_i cannot be $(k-1)$ -regular on K for any $i \in [k]$, a contradiction. $\square$

Remarks. If the input is a (p, q) -pseudocoloring (where p need not be k), then we can still create the conflict graphs as above, and Lemma 4.7 continues to hold. We have $\Delta(G_c) \leq k(p-1)$, so a trivial greedy algorithm shows that G_c can be colored with at most $k(p-1) + 1$ colors, and we obtain a $(k(p-1) + 1)q$ coloring of $M_1, \dots, M_k$.

When $k = 3$, our algorithm produces a $6\chi_{\max}$ matroid-intersection coloring. In this case, each H_i graph (constructed for a given $c \in [q]$) consists of a collection of vertex-disjoint triangles or paths of length at most 3, and so G_c is a union of three such graphs. The hardest case here seems to be when each H_i consists of vertex-disjoint triangles. It is conjectured that in this case that a graph such as G_c can be 5-colored (see Question 1.8 in [AAA⁺18]). If this holds and a 5-coloring can be found efficiently, then this would likely yield an improved efficiently-computable $5\chi_{\max}$ -coloring for 3 matroids. This is notable as it would match the $5\chi_{\max}$ bound for 3-matroid-intersection coloring obtained by [ABGK25] via nonconstructive means.

4.3 Integrality-Gap Bounds

We show that our $k(k-1)\chi_{\max}$ -coloring result for k -matroid-intersection coloring also implies essentially a constructive $k(k-1)$ upper bound on the integrality gap of the natural covering-LP formulation of the problem.

Theorem 4.10. *Given matroids $M_1 = (U, \mathcal{I}_1), \dots, M_k = (U, \mathcal{I}_k)$, consider the following LP-relaxation for k -matroid-intersection coloring. Recall that $\mathcal{I}_{\text{int}} := \bigcap_{i \in [k]} \mathcal{I}_i$.*

$$\min \sum_{I \in \mathcal{I}_{\text{int}}} x_I \quad \text{s.t.} \quad \sum_{I \in \mathcal{I}_{\text{int}}: e \in I} x_I \geq 1 \quad \forall e \in U. \quad (\text{CovLP})$$

The algorithm described in Sections 4.1 and 4.2 returns an integer solution to (CovLP) of value at most $k(k-1) \cdot \lceil \text{OPT}_{\text{CovLP}} \rceil$.

Proof. Recall that $\chi_{\max} := \max_{i \in [k]} \chi(M_i)$ and $\chi(M_i) = \max_{S \subseteq U} \lceil \frac{|S|}{r_i(S)} \rceil$. It suffices to show that $\chi_{\max} \leq \lceil \text{OPT}_{\text{CovLP}} \rceil$. Let x^* be an optimal solution to (CovLP). Consider any $S \subseteq U$. Adding the inequalities of (CovLP) for all $e \in S$ gives $\sum_{I \in \mathcal{I}_{\text{int}}} x_I^* |I \cap S| \geq |S|$. Since $I \cap S$ is a common independent set contained in S , we have $|I \cap S| \leq r_i(S)$ for all $i \in [k]$. It follows that $(\sum_{I \in \mathcal{I}_{\text{int}}} x_I^*) \cdot (\min_{i \in [k]} r_i(S)) \geq |S|$, or equivalently, $\text{OPT}_{\text{CovLP}} \geq \max_{i \in [k]} \frac{|S|}{r_i(S)}$. This lower bound holds for any set $S \subseteq U$, so we have

$$\text{OPT}_{\text{CovLP}} \geq \max_{S \subseteq U} \max_{i \in [k]} \frac{|S|}{r_i(S)} = \max_{i \in [k]} \max_{S \subseteq U} \frac{|S|}{r_i(S)} \implies \lceil \text{OPT}_{\text{CovLP}} \rceil \geq \max_{i \in [k]} \chi(M_i) = \chi_{\max}. \quad \square$$

5 Two-Matroid-Intersection Coloring: Proof of Theorem 1.6

In this section, we present a fully polynomial randomized approximation scheme (FPRAS) for coloring the intersection of two matroids when the maximum of their chromatic numbers is large. In particular, we prove Theorem 1.6. For simplicity of analysis, we actually prove the following modified version of Theorem 1.6, which states a seemingly weaker guarantee without the use of a constant $C > 0$ in the bound. We show in Appendix B that Theorem 5.1 readily implies Theorem 1.6.

Stephen: Swamy and I also chatted about feasibility of iterative rounding LP occurring immediately at $\chi_{\max}$, explaining why $\chi_{\max}$ is a natural lower bound instead of just a trivial one. Possibly include in journal version

Theorem 5.1. *For any $\varepsilon \in (0, 0.001)$ and any instance M_1, M_2 of two-matroid-intersection coloring with $\chi_{\max} \geq \frac{\log n}{\varepsilon^5}$, there is a randomized algorithm with running time $\text{poly}(n, \frac{1}{\varepsilon})$ that returns a $(1 + 400\varepsilon)\chi_{\max}$ -coloring with probability at least $1 - 1/n^2$.*

Our algorithm for Theorem 5.1 will in fact produce a *covering* of $M_{\text{int}} = M_1 \cap M_2$, which is synonymous with coloring, because a covering can be trivially converted into a coloring by removing duplicate elements.

In Section 5.1, we give background on a central component of our algorithm, the swap-rounding algorithm for matroid intersection [CVZ11]. In Section 5.3, we state our algorithm. In Section 5.4, we analyze our algorithm and prove Theorem 5.1.

5.1 Background on Swap-Rounding Algorithm [CVZ11]

A central component to our FPRAS for coloring the intersection of two matroids when $\chi_{\max}$ is large is the swap-rounding algorithm for matroid intersection [CVZ11]. This allows us to round arbitrary fractional points in the matroid intersection polytope to common independent sets with strong expectation and concentration properties. In particular, the main theorem of [CVZ11] we leverage is the following:

Theorem 5.2 ([CVZ11]). *Let M_1 and M_2 be matroids on ground set U and $P = \mathcal{P}(M_1) \cap \mathcal{P}(M_2)$. Fix a constant $0 < \gamma \leq 1/2$. Then there is an efficient randomized rounding procedure *SwapRound* which, given a point $\mathbf{x} \in P$, outputs a common independent set $R \subseteq U$ such that $\mathbb{E}[\mathbb{1}_R] = (1 - \gamma)\mathbf{x}$. In addition, for any linear function $a(R) = \sum_{i \in R} a_i$ with $a_i \in [0, 1]$, and for any $\varepsilon \in [0, 1]$, we have:*

- *If $\mu \leq \mathbb{E}[a(R)]$, then $\Pr[a(R) \leq (1 - \varepsilon)\mu] \leq e^{-\mu\gamma\varepsilon^2/20}$.*
- *If $\mu \geq \mathbb{E}[a(R)]$, then $\Pr[a(R) \geq (1 + \varepsilon)\mu] \leq e^{-\mu\gamma\varepsilon^2/20}$.*

The algorithm for Theorem 5.2 starts by decomposing $\mathbf{x}$ into a convex combination of common independent sets in $\mathcal{I}_1 \cap \mathcal{I}_2$, which can be computed in polynomial time (see [Sch03]). The common independent sets are sequentially merged until only one remains, which is the output set R . Given two common independent sets $I, J \in \mathcal{I}_1 \cap \mathcal{I}_2$, the merge operation involves constructing an exchange graph between I, J in $M_1 \cap M_2$, finding a collection of paths and cycles representing feasible swaps between the two sets, and picking one at random to perform, which decreases the size of the symmetric difference $I \Delta J$. This is repeated until the two sets are the same, giving the merged set of the original I, J .

Runtime: A key property of *SwapRound*, which is not explicitly stated in [CVZ11], is that its runtime is polynomial in $n = |U|$ and $1/\gamma$ [Che26]. The fact that the runtime is also polynomial in $1/\gamma$ allows us to show our algorithm for two-matroid intersection coloring when $\chi_{\max}$ is large is an FPRAS and not just a PRAS. This distinction allows us to show our $(1 + \varepsilon)$ -approximation algorithm is efficient even for ε as small as $\text{poly}(1/n)$, instead of only for constant $\varepsilon > 0$.

Utility of SwapRound: Theorem 5.2 implies powerful concentration properties and has found a wide variety of applications. In our setting, we apply Theorem 5.2 in a few specific ways, namely we always use a uniform point $\alpha \mathbb{1} \in P$; we only care about linear functions of the form $a_i \in \{0, 1\}$ for all i , which correspond to $a(R) = |R \cap S|$ for some subset of elements S ; and we only care about the lower tail bound. We now state a corollary of Theorem 5.2 which will be all that we need for our setting.

Corollary 5.3. *Let M_1 and M_2 be matroids on ground set U and $P = \mathcal{P}(M_1) \cap \mathcal{P}(M_2)$. Fix a constant $0 < \gamma \leq 1/2$. Then there is an efficient randomized rounding procedure, *SwapRound*(M_1, M_2, α, γ), which, given a point $\alpha \mathbb{1} \in P$ outputs a common independent set $R \subseteq U$ satisfying:*

1. For all elements $e \in U$, we have $\Pr[e \in R] = (1 - \gamma)\alpha$.
2. For any subset $S \subseteq U$, if $\mu \leq \mathbb{E}[|R \cap S|]$, then for any $t > 0$ we have:

$$\Pr[|R \cap S| \leq \mu - t] \leq \exp\left(-\frac{\gamma t^2}{20\mu}\right)$$

Intuitively, Corollary 5.3 says that in the SwapRound procedure, the number of elements in the common independent set R which intersect any given set S has an exponential lower tail bound concentrating around the expectation.

Let $\text{SwapRound}(M_1, M_2, \alpha, \gamma)$ be the algorithm given by Corollary 5.3. When the matroids M_1, M_2 are implicit, we refer to the algorithm as $\text{SwapRound}(\alpha, \gamma)$. Notice that 3.2 shows that, for any matroid M , the vector $\mathbb{1}/q \in \mathcal{P}(M)$ for any $q \geq \chi(M)$. Hence, if $\chi_{\max}$ is the maximum chromatic number of M_1 and M_2 , then $\alpha = 1/\chi_{\max} \in \mathcal{P}(M_1) \cap \mathcal{P}(M_2)$ and can be used as a parameter in $\text{SwapRound}(\alpha, \gamma)$.

5.2 Stephen's Additional Preliminaries

5.2.1 Probability Theory Background

In Lemma 5.4, we state a standard fact about sums of conditionally subgaussian random variables, and we include the short proof for completeness.

Lemma 5.4 ([Fre75]). *Let $(\mathcal{F}_t)_{t=0}^m$ be a filtration, and let X_t be a real-valued $\mathcal{F}_t$ -measurable random variable for each $t \in [m]$. Let $v_1, \dots, v_m \geq 0$ and $\lambda \in \mathbb{R}$ be fixed constants. If*

$$\mathbb{E}\left[e^{\lambda X_i} \mid \mathcal{F}_{i-1}\right] \leq \exp\left(\frac{\lambda^2 v_i}{2}\right)$$

for all $i \in [m]$, then

$$\mathbb{E}\left[\exp\left(\lambda \sum_{i=1}^m X_i\right) \mid \mathcal{F}_0\right] \leq \exp\left(\frac{\lambda^2}{2} \sum_{i=1}^m v_i\right)$$

Proof. Define $S_t := \sum_{i=1}^t X_i$, with $S_0 := 0$. Since S_{t-1} is $\mathcal{F}_{t-1}$ -measurable, the tower property and the assumed conditional MGF bound give

$$\begin{aligned} \mathbb{E}\left[e^{\lambda S_t} \mid \mathcal{F}_0\right] &= \mathbb{E}\left[e^{\lambda S_{t-1}} \mathbb{E}\left[e^{\lambda X_t} \mid \mathcal{F}_{t-1}\right] \mid \mathcal{F}_0\right] \\ &\leq \exp\left(\frac{\lambda^2 v_t}{2}\right) \mathbb{E}\left[e^{\lambda S_{t-1}} \mid \mathcal{F}_0\right]. \end{aligned}$$

Iterating this inequality for $k = m, m-1, \dots, 1$ completes the proof. $\square$

Roughly speaking, $\mathcal{F}_t$ can represent the complete state of an algorithm after time step t , and X_t can be a random variable sampled at time t . The standard statement of the above lemma requires independence of the X_t 's. Conceptually, the lemma is saying that independence can be relaxed, and one only needs to ensure that the X_t 's are subgaussian *immediately prior to sampling*. This is very useful for proving concentration guarantees for randomized algorithms as generated sequences of random variables are often dependent.

EQUIVALENCE OF SUBGAUSSIAN MGF and CONCENTRATION.

5.2.2 Swap Rounding Background

Theorem 5.5 ([CVZ11]). *The input to $\text{SwapRound}(M_1, M_2, \mathbf{x}, \gamma)$ consists of two matroids $M_1 = (U, \mathcal{I}_1)$ and $M_2 = (U, \mathcal{I}_2)$ on common ground set U , a point $\mathbf{x} \in \mathcal{P}(M_1) \cap \mathcal{P}(M_2)$, and a constant $\gamma \in (0, 1/2]$. The output is a common independent set $R \in \mathcal{I}_1 \cap \mathcal{I}_2$ with the following two properties.*

- **Marginal Values** (Theorem 1.1 in [CVZ11]): $\mathbb{E}[\mathbb{1}_R] = (1 - \gamma)\mathbf{x}$.
- **Concentration** (Lemma 4.5 in [CVZ11]): For any constant $\lambda \in [-1/4, 1/4]$, integer $p \geq 2/\gamma$, and vector $\mathbf{a} \in [0, 1/p]^U$,

$$\mathbb{E} \left[e^{\lambda(a(R) - \mu_0)} \right] \leq e^{2\lambda^2 \mu_0}$$

where $\mu_0 = (1 - \gamma)\mathbf{a} \cdot \mathbf{x}$.

Note that the ‘‘Concentration’’ property is stated in terms of the stronger moment generating function (MGF) bound, since we will need this stronger bound in our analysis.

Utility of SwapRound: Theorem 5.5 implies powerful concentration properties and has found a wide variety of applications. In our setting, we apply Theorem 5.5 in a few specific ways, namely we always use a uniform point $\alpha \mathbb{1} \in P$; we only care about linear functions of the form $a_i \in \{0, 1\}$ for all i , which correspond to $a(R) = |R \cap S|$ for some subset of elements S ; and we only care about negative values of λ , which induce a lower tail concentration bound. Allowing vectors $\mathbf{a} \in [0, 1]^U$ requires a basic rescaling. We now state a corollary of Theorem 5.5 which will be all that we need for our setting.

Corollary 5.6. *The input to $\text{SwapRound}(M_1, M_2, \alpha \mathbb{1}, \gamma)$ consists of two matroids $M_1 = (U, \mathcal{I}_1)$ and $M_2 = (U, \mathcal{I}_2)$ on common ground set U , a constant $\alpha \in [0, 1]$ such that $\alpha \mathbb{1} \in \mathcal{P}(M_1) \cap \mathcal{P}(M_2)$, and a constant $\gamma \in (0, 1/2]$. The output is a common independent set $R \in \mathcal{I}_1 \cap \mathcal{I}_2$ with the following two properties.*

- **Marginal Values:** For all elements $e \in U$, we have $\Pr[e \in R] = (1 - \gamma)\alpha$.
- **Concentration:** For any constant $\lambda \in [0, \gamma/12]$ and subset of elements $S \subseteq U$,

$$\mathbb{E} \left[e^{-\lambda(|R \cap S| - \mu)} \right] \leq e^{6\lambda^2 \mu / \gamma}$$

where $\mu = (1 - \gamma)\alpha|S|$.

Proof. The ‘‘Marginal Values’’ condition follows immediately from Theorem 5.5. For the ‘‘Concentration’’ condition, given a subset of elements $S \subseteq U$, let $p = \lceil 2/\gamma \rceil$ and $\mathbf{a} = \mathbb{1}_S/p$. Note that $\mu_0 = (1 - \gamma)\mathbf{a} \cdot \mathbf{x} = (1 - \gamma)\alpha|S|/p = \mu/p$. Thus, Theorem 5.5 guarantees that

$$\mathbb{E} \left[e^{-\lambda(|R \cap S| - \mu)} \right] = \mathbb{E} \left[e^{-p\lambda(\mathbf{a}(R) - \mu_0)} \right] \tag{1}$$

$$\leq e^{2(-p\lambda)^2 \mu_0} \tag{2}$$

$$= e^{2p\lambda^2 \mu} \tag{3}$$

$$\leq e^{6\lambda^2 \mu / \gamma} \tag{4}$$

Line (1) follows by simplification. Line (2) follows by Theorem 5.5 and the fact that $p < 2/\gamma + 1 < 3/\gamma$, so $p\lambda \leq (3/\gamma)(\gamma/12) = 1/4$. Line (3) follows by simplification. Line (4) follows again by $p < 3/\gamma$. $\square$

5.3 Algorithm

In this section, we present our FPRAS for two-matroid intersection coloring when $\chi_{\max}$ is large. The algorithm proceeds in ℓ rounds. In each round, the algorithm samples several common independent sets using SwapRound. Each sampled common independent set will get its own color class in the final coloring. The elements in the sampled independent sets are removed from the ground set and the algorithm proceeds to the next round. Once the chromatic number of remaining elements becomes small enough (we choose ℓ so that this happens after ℓ rounds with high probability), the remaining elements are colored by applying the 2-approximation given by Theorem 1.4 (a). See Algorithm 2.

Algorithm 2 FPRAS for Two-Matroid Intersection Coloring for Large $\chi_{\max}$

Input: Matroids $M_1 = (U, \mathcal{I}_1)$ and $M_2 = (U, \mathcal{I}_2)$, $\varepsilon \in (0, 0.001)$

- 1: Initialize $\ell \leftarrow \left\lceil \frac{\log \varepsilon}{\log(1-\varepsilon+100\varepsilon^2)} \right\rceil$, $U^{(0)} \leftarrow U$
 - 2: **for** $i = 0$ to $\ell - 1$ **do** $\triangleright$ Phase 1: Iterative Peeling
 - 3: Let $M_1^{(i)}, M_2^{(i)}$ be M_1, M_2 restricted to $U^{(i)}$
 - 4: $\chi_{\max}^{(i)} \leftarrow \max \left\{ \chi(M_1^{(i)}), \chi(M_2^{(i)}) \right\}$
 - 5: $c_i \leftarrow \lceil \varepsilon \chi_{\max}^{(i)} \rceil$
 - 6: Sample c_i common independent sets $\{I_1, \dots, I_{c_i}\}$ using SwapRound $(M_1^{(i)}, M_2^{(i)}, 1/\chi_{\max}^{(i)}, \varepsilon)$
 - 7: $U^{(i+1)} \leftarrow U^{(i)} \setminus \bigcup_j I_j$ $\triangleright$ Remove all sampled elements
 - 8: Let $\mathcal{C}_{final}$ be a 2-approximate covering of $U^{(\ell)}$ (Theorem 1.4 (a)) $\triangleright$ Phase 2: Final Clean-up
 - 9: **Output:** All sampled sets from Phase 1 and the sets in $\mathcal{C}_{final}$
-

5.4 Analysis (Proof of Theorem 5.1)

In this section, we prove Theorem 5.1. Notice that the overall cost of our algorithm (i.e. the number of color classes it produces) is the sum of the costs in Phase 1 and Phase 2. In Phase 1, we pay $c_i = \lceil \varepsilon \chi_{\max}^{(i)} \rceil$ in each iteration i . The cost of Phase 2 depends on the chromatic number of the leftover elements after ℓ such Phase 1 iterations. Our key claim is that, with high probability, the maximum chromatic number of the matroids drops by at least a roughly $1 - \varepsilon$ factor in each iteration of Phase 1. This is formalized in Lemma 5.7, and is enough to prove Theorem 5.1, as we show below.

Lemma 5.7. *If $\chi_{\max} \geq \log n / \varepsilon^5$ and the chromatic number χ of a matroid at the start of a round is at least $\varepsilon \chi_{\max}$, then the chromatic number of the restricted matroid at the end of the round is at most $(1 - \varepsilon + 100\varepsilon^2) \chi$ with probability at least $1 - 1/n^{10}$.*

We prove Lemma 5.7 shortly, but first we show how this leads to Theorem 5.1.

Proof of Theorem 5.1. The chromatic number of each matroid after $i \leq \ell$ rounds is at most $(1 - \varepsilon + 100\varepsilon^2)^i \chi_{\max}$ with probability at least $1 - 2/n^9 \geq 1 - 1/n^2$ by union bounding over Lemma 5.7. Thus the total number of sampled sets is at most

Stephen: Mik comment on ℓ disappearing

$$\begin{aligned}
 \sum_{i=0}^{\ell-1} c_i &\leq \ell + \sum_{i=0}^{\ell-1} \varepsilon \chi_{\max}^{(i)} \leq \ell + \varepsilon \cdot \sum_{i=0}^{\ell-1} (1 - \varepsilon + 100\varepsilon^2)^i \chi_{\max} \leq \varepsilon \cdot \sum_{i=0}^{\infty} (1 - \varepsilon + 100\varepsilon^2)^i \chi_{\max} \\
 &\leq \frac{\varepsilon \chi_{\max}}{1 - (1 - \varepsilon + 100\varepsilon^2)} = \frac{\varepsilon \chi_{\max}}{\varepsilon - 100\varepsilon^2} < (1 + 200\varepsilon) \chi_{\max}
 \end{aligned}$$

where we use $\varepsilon < 0.001$ in the final inequality. After ℓ rounds, the chromatic number of the restricted matroids is at most

$$(1 - \varepsilon + 100\varepsilon^2)^\ell \chi_{\max} = \exp(\ell \log(1 - \varepsilon + 100\varepsilon^2)) \chi_{\max} \leq \exp(\log \varepsilon) \chi_{\max} = \varepsilon \chi_{\max}$$

Thus the 2-approximate covering will use $2\varepsilon \chi_{\max}$ sets, and so the algorithm will return a covering of $M_1 \cap M_2$ using at most $(1 + 200\varepsilon)\chi_{\max} + 2\varepsilon \chi_{\max} < (1 + 400\varepsilon)\chi_{\max}$ sets with probability at least $1 - 1/n^2$. $\square$

We now turn to proving Lemma 5.7. Recall that for a single matroid $M = (U, \mathcal{I})$, we have the explicit formula $\chi(M) = \max_{S \subseteq U} \lceil \frac{|S|}{r_M(S)} \rceil$ (Fact 3.1). It is easy to see that in the maximization expression, we can restrict to “flats” of the matroid, where a flat is a maximal subset of a given rank. Thus, to argue that the chromatic number of each matroid drops after an iteration of sampling, we need only prove that a large fraction of *all* high-density flats are covered by the sampled sets. There may be exponentially many such subsets, so care is required in order to make a union bound work.

Hence, our next step is to show that in each iteration of Phase 1, and for every subset S of the current ground set, the set of sampled elements T covers (approximately) an ε fraction of S . Importantly, we establish an exponential tail bound on the probability that this does not occur. This is formalized in Lemma 5.8. See Appendix C for background on subgaussian random variables necessary for Lemma 5.8.

Lemma 5.8. *Let $i \in [0, \ell - 1]$ be a given round. Let T_i be the union of the sampled common independent sets in the i 'th round. If $\chi_{\max}^{(i)} \geq \varepsilon \chi_{\max}$, then for all subsets $S \subseteq U^{(i)}$,*

$$\Pr[|T_i \cap S| \leq (\varepsilon - q\varepsilon^2)|S|] \leq \exp(-q^2\varepsilon^4|S|/100)$$

for all $q \geq 4$.

Proof. Let $S \subseteq U^{(i)}$ be arbitrary. Let $I_1, I_2, \dots, I_{c_i}$ be the sampled independent sets in the i 'th round. Let $\Delta_j = |I_j \cap (S \setminus \cup_{h < j} I_h)|$ be the number of new elements covered by I_j in S for $j = 1, 2, \dots, c_i$. Note that $T_i = \cup I_j$ and $|T_i \cap S| = \sum \Delta_j$.

Before the j 'th set is sampled, there are two possibilities:

1. **Many Elements Sampled:** $\sum_{h < j} \Delta_h \geq \varepsilon|S|$. In this case, we have already guaranteed $|T_i \cap S| \geq (\varepsilon - q\varepsilon^2)|S|$ for all $q \geq 4$.
2. **Few Elements Sampled:** $\sum_{h < j} \Delta_h < \varepsilon|S|$. In this case, at least $(1 - \varepsilon)|S|$ elements of S are still unsampled. Thus

$$\mathbb{E}[\Delta_j] \geq (1 - \varepsilon)(1 - \varepsilon) \frac{|S|}{\chi_{\max}^{(i)}} > (1 - 2\varepsilon) \frac{|S|}{\chi_{\max}^{(i)}}$$

and

$$\Pr[\Delta_j \leq \mathbb{E}[\Delta_j] - t] \leq \exp\left(-\frac{\varepsilon t^2}{20\mathbb{E}[\Delta_j]}\right) < \exp\left(-\frac{\varepsilon \chi_{\max}^{(i)} \cdot t^2}{20|S|}\right)$$

via Corollary 5.3 and using $\mathbb{E}[\Delta_j] < \frac{|S|}{\chi_{\max}^{(i)}}$. Thus Δ_j is a subgaussian random variable with mean $\mu \geq (1 - 2\varepsilon) \frac{|S|}{\chi_{\max}^{(i)}}$ and variance proxy $\sigma^2 = \frac{10|S|}{\varepsilon \chi_{\max}^{(i)}}$.

Stephen: Variance proxy equal or less than or equal?

The probability that $|T_i \cap S| \leq (\varepsilon - q\varepsilon^2) |S|$ is equal to the probability that $\sum \Delta_j \leq (\varepsilon - q\varepsilon^2) |S|$ where each Δ_j is a conditionally subgaussian random variable with mean $\mu \geq (1 - 2\varepsilon) \frac{|S|}{\chi_{\max}^{(i)}}$ and variance proxy $\sigma^2 = \frac{10|S|}{\varepsilon \chi_{\max}^{(i)}}$. Thus $\sum \Delta_j$ is a subgaussian random variable with mean $\bar{\mu} = \lceil \varepsilon \chi_{\max}^{(i)} \rceil \cdot \mu \geq (\varepsilon - 2\varepsilon^2) |S|$ and variance proxy $\bar{\sigma}^2 = \lceil \varepsilon \chi_{\max}^{(i)} \rceil \cdot \sigma^2 \geq 10|S|$. Thus

$$\Pr \left[\sum_j \Delta_j \leq (\varepsilon - 2\varepsilon^2) |S| - t \right] \leq \exp \left(-\frac{t^2}{20|S|} \right)$$

Plugging in $t = q/2 \cdot \varepsilon^2 |S|$ and using $q/2 + 2 \leq q$ gives the desired result. $\square$

Stephen: Add direct citation to conditionally subgaussian fact

With the concentration bound for individual subsets established, we now provide the global argument for the reduction of the chromatic number on each round. The strategy is to show that if the chromatic number does not drop sufficiently, there must exist a specific structural witness, a “high-density flat”, that was not well-covered by our sampled sets. By applying a union bound over the set of all such flats and utilizing the local guarantee from Lemma 5.8, we conclude that the chromatic number drops sufficiently with high probability.

Proof of Lemma 5.7. Fix one of the matroids at the start of a round let χ denote its chromatic number. Define a *high-density flat* as a flat S in the matroid s.t. $|S| > (1 - \varepsilon + 100\varepsilon^2) \chi r(S)$ where r is the rank function of the matroid. We first show that if the chromatic number of the restricted matroid at the end of the round is greater than $(1 - \varepsilon + 100\varepsilon^2) \chi$, then there exists a high-density flat S in the start matroid s.t. $|T \cap S| < (\varepsilon - 100\varepsilon^2) \chi r(S)$ where T is the union of the sampled common independent sets during the round. We then union bound this event over all high-density flats S and apply Lemma 5.8.

First, suppose the chromatic number of the restricted matroid at the end of the round is greater than $(1 - \varepsilon + 100\varepsilon^2) \chi$. Then the restricted matroid contains a subset S' s.t. $|S'| > (1 - \varepsilon + 100\varepsilon^2) \chi r(S')$ via Fact 3.1. Let S be the span of S' in the start matroid. Then S is a high-density flat in the start matroid, because S is a flat (since it is defined as the span of a collection of elements) and $|S| > (1 - \varepsilon + 100\varepsilon^2) \chi r(S)$ (since $|S| \geq |S'|$ and $r(S) = r(S')$). Further, $S' \subseteq S \setminus T$, giving $|T \cap S| = |S| - |S \setminus T| \leq |S| - |S'|$. Because $|S| \leq \chi r(S)$ via Fact 3.1, we have

$$|T \cap S| \leq |S| - |S'| < \chi r(S) - (1 - \varepsilon + 100\varepsilon^2) \chi r(S') = (\varepsilon - 100\varepsilon^2) \chi r(S)$$

Next, let $\mathcal{S}$ be the collection of all high-density flats in the start matroid, and let χ' be the chromatic number of the restricted matroid. Then

$$\Pr [\chi' > (1 - \varepsilon + 100\varepsilon^2) \chi] \leq \Pr [\exists S \in \mathcal{S} : |T \cap S| \leq (\varepsilon - 100\varepsilon^2) |S|] \quad (5)$$

$$\leq \sum_{S \in \mathcal{S}} \Pr [|T \cap S| \leq (\varepsilon - 100\varepsilon^2) |S|] \quad (6)$$

$$\leq \sum_{S \in \mathcal{S}} \exp(-100\varepsilon^4 |S|) \quad (7)$$

$$= \sum_{m=1}^n \sum_{S \in \mathcal{S} : r(S)=m} \exp(-100\varepsilon^4 |S|) \quad (8)$$

$$\leq \sum_{m=1}^n \sum_{S \in \mathcal{S} : r(S)=m} \exp(-50\varepsilon^4 \chi m) \quad (9)$$

$$\leq \sum_{m=1}^n n^m \exp(-50\varepsilon^4 \chi m) \quad (10)$$

$$\leq \sum_{m=1}^n \exp(m \log n - 50m \log n) \quad (11)$$

$$\leq \sum_{m=1}^n n^{-20m} \quad (12)$$

$$\leq n^{-10} \quad (13)$$

Line (5) follows from the previous paragraph. Line (6) follows by a union bound. Line (7) follows from Lemma 5.8. Line (8) follows by partitioning the high-density flats by rank. Line (9) follows by $|S| > (1 - \varepsilon + 100\varepsilon^2) \chi m > \chi m/2$. Line (10) follows from the fact that a flat of rank m is defined as the span of an independent set of size m , and so the number of flats of rank m in a matroid with n elements is at most $\binom{n}{m} \leq n^m$. Line (11) follows by $\chi \geq \varepsilon \chi_{\max} \geq \log n / \varepsilon^4$. Line (12) follows by computation. Line (13) follows from the fact that the sum contains n terms of size at most n^{-20} . $\square$

Stephen: Error on S and $\chi^r(S)$

Runtime of Algorithm 2. Lastly, we remark on the runtime of Algorithm 2. Phase 1 involves $\ell = O\left(\frac{\log(1/\varepsilon)}{\varepsilon}\right)$ rounds. In each round i , we call `SwapRound` $c_i = \lceil \varepsilon \chi_{\max}^{(i)} \rceil = O(\varepsilon n)$ times, and run for $\text{poly}(n)$ time to compute the chromatic numbers of $M_1^{(i)}, M_2^{(i)}$. The runtime of `SwapRound` $(M_1^{(i)}, M_2^{(i)}, 1/\chi_{\max}^{(i)}, \varepsilon)$ is polynomial in $n, 1/\varepsilon$ as discussed in the background section. Thus the runtime of Phase 1 is polynomial in $n, 1/\varepsilon$. The runtime of Phase 2 is polynomial in n , so the runtime of Algorithm 2 is polynomial in $n, 1/\varepsilon$.

A Algorithmic Proof of Brooks' Theorem (Theorem 4.8)

For completeness, we describe an algorithmic proof of Brooks' theorem due to Lovasz [Lov75].

The key to the proof is finding vertices $a, b, v \in V(G)$ such that v is adjacent to both a and b , the pair a and b are non-adjacent, and $G \setminus \{a, b\}$ is connected. Having found a, b , and v , we can construct an ordering of $V(G)$ such that the greedy coloring uses at most Δ colors.

Since $G \setminus \{a, b\}$ is connected, we may list its vertices as $x_1 = v, x_2, \dots, x_{n-2}$ so that each x_i with $i \geq 2$ is adjacent to some x_j with $j < i$. We now define a Δ -coloring of G . Assign a and b color 1. Then color $x_{n-2}, x_{n-3}, \dots, x_2$ greedily in reverse order with colors from $\{1, \dots, \Delta\}$. When coloring x_i for $i \geq 2$, at least one of its neighbors remains uncolored, hence x_i has at most $\Delta - 1$ previously colored neighbors, and a free color is always available. To color the final vertex $x_1 = v$, we notice that although v may have Δ neighbors, a and b share a color and hence v can be colored in $\{1, \dots, \Delta\}$ as well.

We turn to showing that there exists vertices $a, b, v \in V(G)$ such that v is adjacent to a and b , and a and b are non-adjacent, and $G \setminus \{a, b\}$ is connected. These vertices can be found in polynomial time by enumeration.

We may assume G is 2-vertex-connected. Otherwise, we color each 2-vertex-connected block separately and reconcile colors at cut vertices. If G is 3-vertex-connected, then since $G \neq K_{\Delta+1}$ some vertex v has two non-adjacent neighbors a and b , and 3-vertex-connectivity ensures $G \setminus \{a, b\}$ is connected. If G is 2-vertex-connected but not 3-vertex-connected, let x be a vertex that is not adjacent to every other vertex and has degree at least 3. This can be assumed to exist, see [BW14]. If $G \setminus \{x\}$ is 2-vertex-connected, set $a = x$, let b be any vertex at distance 2 from x , and let v be a common neighbor. Then $G \setminus \{a, b\}$ is connected since

$G \setminus \{x\}$ is 2-vertex-connected. Finally, if $G \setminus \{x\}$ has a cut vertex, take two leaves of the block-cut tree B_1 , B_2 with attachment vertices z_1, z_2 . Since G is 2-vertex-connected, x has a neighbor $a \in B_1 \setminus \{z_1\}$ and a neighbor $b \in B_2 \setminus \{z_2\}$. Now setting $v = x$, and taking a and b satisfy the properties as desired. $\square$

B Proof of Theorem 1.6

If $\varepsilon \geq 1$, then the $(1+\varepsilon)\chi_{\max}$ bound follows from Theorem 1.4(a), for any $\chi_{\max}$. So suppose $\varepsilon \in (0, 1)$. Let A_ε be the randomized algorithm given by Theorem 5.1. We take $C = 1000^5$, and the randomized algorithm in the statement of Theorem 1.6 to be $A_{\varepsilon/1000}$. (More precisely, the randomized algorithm is $A_{\varepsilon/1000}$, for $\varepsilon \in (0, 1)$, and the algorithm in Theorem 1.4(a) if $\varepsilon \geq 1$).

We now show that this randomized algorithm satisfies the guarantee of Theorem 1.6 for all $\varepsilon > 0$. We can focus on $\varepsilon \in (0, 1)$. When $\chi_{\max} \geq C \log n / \varepsilon^5 = \log n / (\varepsilon/1000)^5$, using Theorem 5.1, $A_{\varepsilon/1000}$ produces a feasible coloring of $M_{\text{int}} = M_1 \cap M_2$ using at most $(1 + 400(\varepsilon/1000))\chi_{\max} \leq (1 + \varepsilon)\chi_{\max}$ colors with probability at least $1 - 1/n^2$.

Lastly, we can convert the probability of $1 - 1/n^2$ into a high probability statement by running the algorithm $O(\log n)$ times and outputting the minimum coloring. $\square$

C Background on Subgaussian Random Variables for Lemma 5.8

Under construction In this section, we give background on subgaussian random variables for Lemma 5.8.

Definition C.1. A random variable X is *subgaussian* with *variance proxy* σ^2 if

$$\mathbb{E} \left[e^{\lambda(X - \mathbb{E}[X])} \right] \leq e^{\lambda^2 \sigma^2 / 2} \quad \forall \lambda \in \mathbb{R}$$

If the inequality only holds for $\lambda \geq 0$ then X is called *upper-tail subgaussian*, and if the inequality only holds for $\lambda \leq 0$ then X is called *lower-tail subgaussian*.

References

- [AAA⁺18] Ron Aharoni, Noga Alon, Michal Amir, Penny Haxell, Dan Hefetz, Zilin Jiang, Gal Kronenberg, and Alon Naor. Ramsey-nice families of graphs. *Electronic Journal of Combinatorics*, 72:29–44, 2018.
- [AB06] Ron Aharoni and Eli Berger. The intersection of a matroid and a simplicial complex. *Transactions of the American Mathematical Society*, 358(11):4895–4917, 2006.
- [ABGK25] Ron Aharoni, Eli Berger, He Guo, and Dani Kotlar. Coloring, list coloring, and fractional coloring in intersections of matroids. *Combinatorica*, 45:Article 46, 2025.
- [AKKG23] Dorna Abdolazimi, Anna R. Karlin, Nathan Klein, and Shayan Oveis Gharan. Matroid partition property and the secretary problem. In *14th Innovations in Theoretical Computer Science Conference (ITCS 2023)*, volume 251 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 2:1–2:9. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023.
- [AMPZ26] Stephen Arndt, Benjamin Moseley, Kirk Pruhs, and Michael Zlatin. Efficiently coloring the intersection of a general matroid and combinatorial matroids. In *Proceedings of the International Conference on Integer Programming and Combinatorial Optimization (IPCO 2026)*, 2026. To appear. Available at <https://arxiv.org/abs/2508.19473>.

- [BCK23] Kristóf Bérczi, Gergely Csáji, and Tamás Király. On the complexity of packing rainbow spanning trees. *Discrete Mathematics*, 346(113297), 2023.
- [BG25] Eli Berger and He Guo. Coloring the intersection of two matroids. *Proceedings of the American Mathematical Society*, 153(10):4145–4154, 2025.
- [BS21] Kristóf Bérczi and Tamás Schwarcz. Complexity of packing common bases in matroids. *Mathematical Programming*, 188(1):1–18, 2021.
- [BSY21] Kristóf Bérczi, Tamás Schwarcz, and Yutaro Yamaguchi. List coloring of two matroids through reduction to partition matroids. *SIAM Journal on Discrete Mathematics*, 35:2192–2209, 2021.
- [BW14] Bradley Baetz and David R Wood. Brooks’ vertex-colouring theorem in linear time. *arXiv preprint arXiv:1401.8023*, 2014.
- [Che26] Chandra Chekuri. Personal communication, 2026.
- [Cho17] Timothy Y. Chow. Rota’s conjecture. https://michaelnielsen.org/polymath1/index.php?title=Rota's_conjecture, 2017. Accessed: 2025-07-10.
- [CVZ11] Chandra Chekuri, Jan Vondrák, and Rico Zenklusen. Multi-budgeted matchings and matroid intersection via dependent rounding. In Dana Randall, editor, *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2011, San Francisco, California, USA, January 23-25, 2011*, pages 1080–1097. SIAM, 2011.
- [Dri97] Arthur A. Drisko. On the number of even and odd latin squares of order $p+1$. *Advances in Mathematics*, 128:20–35, 1997.
- [Edm68] Jack Edmonds. Matroid partition. In *Mathematics of the Decision Sciences, Part 1*, volume 11 of *Lectures in Applied Mathematics*, pages 335–345. American Mathematical Society, 1968.
- [Fre75] David A. Freedman. On tail probabilities for martingales. *The Annals of Probability*, 3(1):100–118, 1975.
- [Gly10] David G. Glynn. The conjectures of alon–tarsi and rota in dimension prime minus one. *SIAM Journal on Discrete Mathematics*, 24(2):394–399, 2010.
- [HKK24] Florian Hörsch, Tomáš Kaiser, and Matthias Kriesell. Rainbow bases in matroids. *SIAM Journal on Discrete Mathematics*, 38(2):1472–1491, 2024.
- [HR94] Rosa Huang and Gian-Carlo Rota. On the relations of various conjectures on latin squares and straightening coefficients. *Discrete Mathematics*, 128(1):225–236, 1994.
- [IMP21] Sungjin Im, Benjamin Moseley, and Kirk Pruhs. The matroid intersection cover problem. *Operations Research Letters*, 49:17–22, 2021.
- [LMP21] Marilena Leichter, Benjamin Moseley, and Kirk Pruhs. An efficient reduction of a gammoid to a partition matroid. In *29th Annual European Symposium on Algorithms (ESA 2021)*, volume 204 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 62:1–62:13. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021.
- [LMP22] Marilena Leichter, Benjamin Moseley, and Kirk Pruhs. On the impossibility of decomposing binary matroids. *Operations Research Letters*, 50:623–625, 2022.

- [LOSZ20] André Linhares, Neil K. Olver, Chaitanya Swamy, and Rico Zenklusen. Approximate multi-matroid intersection via iterative refinement. *Mathematical Programming*, 183:397–418, 2020.
- [Lov75] László Lovász. Three short proofs in graph theory. *Journal of Combinatorial Theory, Series B*, 19(3):269–271, 1975.
- [MS25] Richard Montgomery and Lisa Sauermann. Asymptotically-tight packing and covering with transversal bases in rota’s basis conjecture, 2025.
- [Sch03] Alexander Schrijver. *Combinatorial Optimization: Polyhedra and Efficiency*, volume 24 of *Algorithms and Combinatorics*. Springer, 2003.