Cover-Decomposition and Polychromatic Numbers

Béla Bollobás[1], David Pritchard[2], Thomas Rothvoß[3], and Alex Scott[4]

[1] University of Memphis, USA and University of Cambridge, UK
[2] EPFL, Lausanne, Switzerland
[3] MIT, Cambridge, USA
[4] Mathematical Institute, University of Oxford, UK

Abstract. A colouring of a hypergraph's vertices is *polychromatic* if every hyperedge contains at least one vertex of each colour; the *polychromatic number* is the maximum number of colours in such a colouring. Its dual, the *cover-decomposition number*, is the maximum number of disjoint hyperedge-covers. In geometric settings, there is extensive work on lower-bounding these numbers in terms of their trivial upper bounds (minimum hyperedge size & degree). Our goal is to get good lower bounds in natural hypergraph families not arising from geometry. We obtain algorithms yielding near-tight bounds for three hypergraph families: those with bounded hyperedge size, those representing paths in trees, and those with bounded VC-dimension. To do this, we link cover-decomposition to iterated relaxation of linear programs via discrepancy theory.

1 Introduction

In a set system on vertex set V, a subsystem is a *set cover* if each vertex of V appears in at least 1 set of the subsystem. Suppose each vertex appears in at least δ sets of the set system, for some large δ; does it follow that we can partition the system into 2 subsystems, such that each subsystem is a set cover?

Many natural families of set systems admit a universal constant δ for which this question has an affirmative answer. Such families are typically called *cover-decomposable*. But the family of *all* set systems is not cover-decomposable, as the following example shows. For any positive integer k, consider a set system which has $2k-1$ sets, and where every subfamily of k sets contain one mutually common vertex not contained by the other $k-1$ sets. This system satisfies the hypothesis of the question for $k=\delta$. But it has no set cover consisting of $\leq k-1$

[1]: supported in part by NSF grants CNS-0721983, CCF-0728928, DMS-0906634 and CCR-0225610, and ARO grant W911NF-06-1-0076. [2]: supported in part by an NSERC PDF. [3]: supported by the Alexander von Humboldt Foundation within the Feodor Lynen program, by ONR grant N00014-11-1-0053, and by NSF contract CCF-0829878.

sets, and it has only $2k-1$ sets in total; so no partition into two set covers is possible. This example above shows that some sort of restriction on the family is necessary to ensure cover-decomposability.

One positive example of cover-decomposition arises if every set has size 2: such hypergraphs are simply graphs. They are cover-decomposable with $\delta = 3$: any graph with minimum degree 3 can have its edges partitioned into two edge covers. More generally, Gupta [1] showed (also [2,3]) that we can partition the edges of any multigraph into $\lfloor \frac{3\delta+1}{4} \rfloor$ edge covers. This bound is tight, even for 3-vertex multigraphs.

Set systems in many geometric settings have been studied with respect to cover-decomposability; many positive and negative examples are known and there is no easy way to distinguish one from the other. In the affirmative case, as with Gupta's theorem, the next natural problem is to find for each $t \geq 2$ the smallest $\delta(t)$ such that when each vertex appears in at least $\delta(t)$ sets, a partition into t set covers is possible. The goal of this paper is to extend the study of cover-decomposition beyond geometric settings. A novel property of our studies is that we use iterated linear programming to find cover-decompositions.

1.1 Terminology and Notation

A *hypergraph* $H = (V, \mathcal{E})$ consists of a ground set V of vertices, together with a collection $\mathcal{E}$ of hyperedges, where each hyperedge $E \in \mathcal{E}$ is a subset of V. Hypergraphs are the same as *set systems*. We will sometimes call hyperedges just *edges* or *sets*. We permit $\mathcal{E}$ to contain multiple copies of the same hyperedge (e.g. to allow us to define "duals" and "shrinking" later), and we also allow hyperedges of cardinality 0 or 1. We only consider hypergraphs that are finite. (In many geometric cases, infinite versions of the problem can be reduced to finite ones, e.g. [4]; see also [5] for work on infinite versions of cover-decomposability.)

To *shrink* a hyperedge E in a hypergraph means to replace it with some $E' \subseteq E$. This operation is useful in several places.

A *polychromatic k-colouring* of a hypergraph is a function from V to a set of k colours so that for every edge, its image contains all colours. (Equivalently, the colour classes partition V into sets which each meet every edge, so-called *vertex covers/transversals*.) The maximum number of colours in a polychromatic colouring of H is called its *polychromatic number*, which we denote by $\mathsf{p}(H)$.

A *cover k-decomposition* of a hypergraph is a partition of $\mathcal{E}$ into k subfamilies $\mathcal{E} = \biguplus_{i=1}^{k} \{\mathcal{E}_i\}$ such that each $\bigcup_{E \in \mathcal{E}_i} E = V$. In other words, each $\mathcal{E}_i$ must be a set cover. The maximum k for which the hypergraph H admits a cover k-decomposition is called its *cover-decomposition number*, which we denote by $\mathsf{p}'(H)$.

The *dual* H^* of a hypergraph H is another hypergraph such that the vertex set of H^* corresponds to the edge set of H, and vice-versa, with incidences preserved. Thus the vertex-edge incidence matrices for H and H^* are transposes of one another. E.g., the standard notation for the example in the introduction is $\binom{[2k-1]}{k}^*$. From the definitions it is easy to see that the polychromatic and cover-decomposition numbers are dual to one another,

$$\mathsf{p}'(H) = \mathsf{p}(H^*).$$

The *degree* of a vertex v in a hypergraph is the number of hyperedges containing v; it is *d-regular* if all vertices have degree d. We denote the minimum degree by δ, and the maximum degree by Δ. We denote the minimum size of any hyperedge by r, and the maximum size of any hyperedge by R. Note that $\Delta(H) = R(H^*)$ and $\delta(H) = r(H^*)$. It is trivial to see that $\mathsf{p} \leq r$ in any hypergraph and dually that $\mathsf{p}' \leq \delta$. So the cover-decomposability question asks if there is a converse to this trivial bound: if δ is large enough, does p' also grow? To write this concisely, for a family $\mathcal{F}$ of hypergraphs, let its extremal *cover-decomposition function* $\overline{\mathsf{p}}'(\mathcal{F}, \delta)$ be

$$\overline{\mathsf{p}}'(\mathcal{F}, \delta) := \min\{\mathsf{p}'(H) \mid H \in \mathcal{F};\ \forall v \in V(H) : \text{degree}(v) \geq \delta\},$$

i.e. $\overline{\mathsf{p}}'(\mathcal{F}, \delta)$ is the best possible lower bound for p' among hypergraphs in $\mathcal{F}$ with min-degree $\geq \delta$. So to say that $\mathcal{F}$ is cover-decomposable means that $\overline{\mathsf{p}}'(\mathcal{F}, \delta) > 1$ for some constant δ. We also dually define

$$\overline{\mathsf{p}}(\mathcal{F}, r) := \min\{\mathsf{p}(H) \mid H \in \mathcal{F};\ \forall E \in \mathcal{E}(H) : |E| \geq r\}.$$

In the rest of the paper we focus on computing these functions. When the family $\mathcal{F}$ is clear from context, we write $\overline{\mathsf{p}}'(\delta)$ for $\overline{\mathsf{p}}'(\mathcal{F}, \delta)$ and $\overline{\mathsf{p}}(r)$ for $\overline{\mathsf{p}}(\mathcal{F}, r)$.

1.2 Results

In Section 2 we generalize Gupta's theorem to hypergraphs of bounded edge size. Let $\textsc{Hyp}(R)$ denote the family of hypergraphs with all edges of size at most R.

Theorem 1. *For all R, δ we have $\overline{\mathsf{p}}'(\textsc{Hyp}(R), \delta) \geq \max\{1, \delta/(\ln R + O(\ln \ln R))\}$.*

In proving Theorem 1, we first give a simple proof which is weaker by a constant factor, and then we refine the analysis. We use the Lovász Local Lemma (LLL) as well as discrepancy-theoretic results which permit us to partition a large hypergraph into two pieces with roughly-equal degrees. Next we show that Theorem 1 is essentially tight:

Theorem 2. *(a) For a constant C and all $R \geq 2, \delta \geq 1$ we have $\overline{\mathsf{p}}'(\textsc{Hyp}(R), \delta) \leq \max\{1, C\delta/\ln R\}$. (b) For any sequence $R, \delta \to \infty$ with $\delta = \omega(\ln R)$ we have $\overline{\mathsf{p}}'(\textsc{Hyp}(R), \delta) \leq (1 + o(1))\delta/\ln(R)$.*

Here (a) uses an explicit construction while (b) uses the probabilistic method.

By plugging Theorem 1 into an approach of [3], one obtains a good bound on the cover-decomposition number of *sparse* hypergraphs.

Corollary 3. *Suppose $H = (V, \mathcal{E})$ satisfies, for all $V' \subseteq V$ and $\mathcal{E}' \subseteq \mathcal{E}$, that the number of incidences between V' and $\mathcal{E}'$ is at most $\alpha|V'| + \beta|\mathcal{E}'|$. Then $\mathsf{p}'(H) \geq \frac{\delta(H) - \alpha}{\ln \beta + O(\ln \ln \beta)}$.*

(Duality yields a similar bound on the polychromatic number.) The proof is analogous to that in [3]: a max-flow min-cut argument shows that in this sparse hypergraph, we can shrink all edges to have size at most β, while keeping the minimum degree at least $\delta(H) - \alpha$.

In Section 3 we consider the following family of hypergraphs: the ground set is the edge set of an undirected tree, and each hyperedge must correspond to the edges lying in some path in the tree. We show that such systems are cover-decomposable:

Theorem 4. *For hypergraphs defined by edges of paths in trees,* $\overline{\mathsf{p}}'(\delta) \geq 1 + \lfloor(\delta - 1)/5\rfloor$.

To prove Theorem 4, we exploit the connection to discrepancy and iterated rounding, using an extreme point structure theorem for paths in trees from [6]. We also determine the extremal polychromatic number for such systems:

Theorem 5. *For hypergraphs defined by edges of paths in trees,* $\overline{\mathsf{p}}(r) = \lceil r/2 \rceil$.

This contrasts with a construction of Pach, Tardos and Tóth [7]: if we also allow hyperedges consisting of sets of "siblings," then $\overline{\mathsf{p}}(r) = 1$ for all r.

The *VC-dimension* is a prominent measure of set system complexity used frequently in geometry: it is the maximum cardinality of any $S \subseteq V$ such that $\{S \cap E \mid E \in \mathcal{E}\} = \mathbf{2}^S$. It is natural to ask what role the VC-dimension plays in cover-decomposability. We show the following — the proof is deferred to the full version[§].

Theorem 6. *For the family of hypergraphs with VC-dimension 1,* $\overline{\mathsf{p}}(r) = \lceil r/2 \rceil$ *and* $\overline{\mathsf{p}}'(\delta) = \lceil \delta/2 \rceil$.

By duality, the same holds for the family of hypergraphs whose duals have VC-dimension 1. We find Theorem 6 is best possible in a strong sense:

Theorem 7. *For the family of hypergraphs* $\{H \mid \text{VC-dim}(H), \text{VC-dim}(H^*) \leq 2\}$, *we have* $\overline{\mathsf{p}}(r) = 1$ *for all* r *and* $\overline{\mathsf{p}}'(\delta) = 1$ *for all* δ.

To prove this, we show the construction of [7] has primal and dual VC-dimension at most 2.

All of our lower bounds on $\overline{\mathsf{p}}$ and $\overline{\mathsf{p}}'$ can be implemented as polynomial-time algorithms. In the case of Theorem 1 this relies on the constructive LLL framework of Moser-Tardos [8]. In the tree setting (Theorem 4) the tree representing the hypergraph does not need to be explicitly given as input, since the structural property used in each iteration (Lemma 15) is easy to identify from the values of the extreme point LP solution. Note: since we also have the trivial bounds $\mathsf{p} \leq r, \mathsf{p}' \leq \delta$ these give *approximation algorithms* for p and p', e.g. Theorem 1 gives a $(\ln R + O(\ln \ln R))$-approximation for p'.

[§] http://arxiv.org/abs/1009.6144

1.3 Related Work

One practical motive to study cover-decomposition is that the hypergraph can model a collection of sensors [9,10], with each $E \in \mathcal{E}$ corresponding to a sensor which can monitor the set E of vertices; then monitoring all of V takes a set cover, and $\overline{\mathsf{p}}'$ is the maximum "coverage" of V possible if each sensor can only be turned on for a single time unit or monitor a single frequency. Another motive is that if $\overline{\mathsf{p}}'(\delta) = \Omega(\delta)$ holds for a family closed under vertex deletion, then the size of a *dual* ϵ*-net* is bounded by $O(1/\epsilon)$ [11].

A hypergraph is said to be *weakly* k*-colourable* if we can k-colour its vertex set so that no edge is monochromatic. Weak 2-colourability is also known as *Property B*, and these notions coincide with the property $\mathsf{p} \geq 2$. However, weak k-colourability does not imply $\mathsf{p} \geq k$ in general.

For a graph $G = (V, E)$, the *(closed) neighbourhood hypergraph* $\mathcal{N}(H)$ is defined to be a hypergraph on ground set V, with one hyperedge $\{v\} \cup \{u \mid \{u, v\} \in E\}$ for each $v \in V$. Then $\mathsf{p}(\mathcal{N}(G)) = \mathsf{p}'(\mathcal{N}(G))$ equals the *domatic number* of G, i.e. the maximum number of disjoint dominating sets. The paper of Feige, Halldórsson, Kortsarz & Srinivasan [12] obtains upper bounds for the domatic number and their bounds are essentially the same as what we get by applying Theorem 1 to the special case of neighbourhood hypergraphs; compared to our methods they use the LLL but not discrepancy or iterated LP rounding. They give a hardness-of-approximation result which implies that Theorem 1 is tight with respect to the approximation factor, namely for all $\epsilon > 0$, it is hard to approximate p' within a factor better than $(1 - \epsilon) \ln R$, under reasonable complexity assumptions. A generalization of results in [12] to packing polymatroid bases was given in [13]; this implies a weak version of Theorem 1 where the $\log R$ term is replaced by $\log |V|$.

A notable progenitor in geometric literature on cover-decomposition is the following question of Pach [14]. Take a convex set $A \subset \mathbb{R}^2$. Let $\mathbb{R}^2|\textsc{Translates}(A)$ denote the family of hypergraphs where the ground set V is a finite subset of $\mathbb{R}^2$, and each hyperedge is the intersection of V with some translate of A. Pach asked if such systems are cover-decomposable, and this question is still open. A state-of-the-art partial answer is due to Gibson & Varadarajan [10], who prove that $\overline{\mathsf{p}}(\mathbb{R}^2|\textsc{Translates}(A), \delta) = \Omega(\delta)$ when A is an open convex polygon.

The paper of Pach, Tardos and Tóth [7] obtains several negative results with a combinatorial method. They define a family of non-cover-decomposable hypergraphs based on trees and then they "embed" these hypergraphs into geometric settings. By doing this, they prove that the following families are not cover-decomposable: $\mathbb{R}^2|\textsc{Axis-Aligned-Rectangles}$; $\mathbb{R}^2|\textsc{Translates}(A)$ when A is a non-convex quadrilateral; and $\mathbb{R}^2|\textsc{Strips}$ and its dual. In contrast to the latter result, it is known that $\overline{\mathsf{p}}(\mathbb{R}^2|\textsc{Axis-Aligned-Strips}, r) \geq \lceil r/2 \rceil$ [15]. Recently it was shown [16] that $\mathbb{R}^3|\textsc{Translates}(\mathbb{R}^3_+)$ is cover-decomposable, giving cover-decomposability of $\mathbb{R}^2|\textsc{Homothets}(T)$ for any triangle T and a new proof (c.f. [17]) for $\mathbb{R}^2|\textsc{Bottomless-Axis-Aligned-Rectangles}$; the former contrasts with the non-cover-decomposability of $\mathbb{R}^2|\textsc{Homothets}(D)$ for D the unit disc [7].

Pálvölgyi [4] poses a fundamental combinatorial question: is there a function f so that in hypergraph families closed under edge deletion and duplication, $\overline{\mathsf{p}}'(\delta_0) \geq 2$ implies $\overline{\mathsf{p}}'(f(\delta_0)) \geq 3$? This is open for all $\delta_0 \geq 2$ and no counterexamples are known to the conjecture $f(\delta_0) = O(\delta_0)$.

Given a plane graph, define a hypergraph whose vertices are the graph's vertices, and whose hyperedges are the faces. For this family of hypergraphs, it was shown in [3] that $\overline{\mathsf{p}}(\delta) \leq \lfloor (3\delta - 5)/4 \rfloor$ using Gupta's theorem and a sparsity argument. This is the same approach which we exploit to prove Corollary 3.

Several different related colouring notions for paths in trees are considered in [18,19,20].

2 Hypergraphs of Bounded Edge Size

To get good upper bounds on $\overline{\mathsf{p}}'(\textsc{Hyp}(R), \delta)$, we will use the Lovász Local Lemma (LLL):

Lemma 8 (LLL, [21])**.** *Consider a collection of "bad" events such that each one has probability at most p, and such that each bad event is independent of the other bad events except at most D of them. (We call D the* dependence degree*.) If $p(D+1)\mathrm{e} \leq 1$ then with positive probability, no bad events occur.*

Our first proposition extends a standard argument about Property B [22, Theorem 5.2.1].

Proposition 9. $\overline{\mathsf{p}}'(\textsc{Hyp}(R), \delta) \geq \lfloor \delta / \ln(\mathrm{e}R\delta^2) \rfloor$.

I.e. given any hypergraph $H = (V, \mathcal{E})$ where every edge has size at most R and such that each $v \in V$ is covered at least δ times, we must show for $t = \lfloor \delta / \ln(\mathrm{e}R\delta^2) \rfloor$ that $\mathsf{p}'(H) \geq t$, i.e. that $\mathcal{E}$ can be decomposed into t disjoint set covers. It will be helpful here and later to make the degree of every vertex *exactly* δ, (this bounds the dependence degree). Indeed this is without loss of generality: else as long as $\deg(v) > \delta$ shrink some $E \ni v$ to $E \backslash \{v\}$ until $\deg(v)$ drops to δ; then observe that if we applying unshrinking to a vertex cover, it is still a vertex cover.

Proof of Proposition 9. Consider the following randomized experiment: for each hyperedge $E \in \mathcal{E}$, assign a random colour between 1 and t to E. If we can show that with positive probability, every vertex is incident with a hyperedge of each colour, then we will be done. In order to get this approach to go through,

For each vertex v define the *bad event* $\mathfrak{E}_v$ to be the event that v is not incident with a hyperedge of each colour. The probability of $\mathfrak{E}_v$ is at most $t(1 - \frac{1}{t})^\delta$, by using a union bound. The event $\mathfrak{E}_v$ only depends on the colours of the hyperedges containing v; therefore the events $\mathfrak{E}_v$ and $\mathfrak{E}_{v'}$ are independent unless v, v' are in a common hyperedge. In particular the dependence degree is less than $R\delta$. It follows by LLL that if

$$R\delta t(1 - \tfrac{1}{t})^\delta \leq 1/\mathrm{e},$$

then with positive probability, no bad events happen and we are done. We can verify that $t = \delta / \ln(\mathrm{e}R\delta^2)$ satisfies this bound. □

We will next show that the bound can be raised to $\Omega(\delta/\ln R)$. Intuitively, our strategy is the following. We have that $\delta/\ln(R\delta)$ is already $\Omega(\delta/\ln R)$ unless δ is superpolynomial in R. For hypergraphs where $\delta \gg R$ we will show that we can partition $\mathcal{E}$ into m parts $\mathcal{E} = \biguplus_{i=1}^{m} \mathcal{E}_i$ so that $\delta(V, \mathcal{E}_i)$ is at least a constant of δ/m, and such that δ/m is polynomial in R. Thus by Proposition 9 we can extract $\Omega((\delta/m)/\ln R)$ set covers from each $(V, \mathcal{E}_i)$, and their union proves $\bar{\mathsf{p}}' \geq \Omega(\delta/\ln R)$.

In fact, it will be enough to consider splitting $\mathcal{E}$ into two parts at a time, recursively. Then ensuring $\delta(V, \mathcal{E}_i) \gtrsim \delta/2$ $(i = 1, 2)$ amounts to a discrepancy-theoretic problem: given the incidence matrix with rows for edges and columns for vertices, we must 2-colour the rows by ± 1 so that for each column, the sum of the incident rows' colours is in $[-d, d]$, with the *discrepancy* d as small as possible. To get a short proof of a weaker version of Theorem 1, we can use an approach of Beck and Fiala [23]; moreover it is important to review their proof since we will extend it in Section 3.

Proposition 10 (Beck & Fiala [23])**.** *In a δ-regular hypergraph $H = (V, \mathcal{E})$ with all edges of size at most R, we can partition the edge set into $\mathcal{E} = \mathcal{E}_1 \uplus \mathcal{E}_2$ such that $\delta(V, \mathcal{E}_i) \geq \delta/2 - R$ for each $i \in \{1, 2\}$.*

Proof. Define a linear program with nonnegative variables $\{x_e, y_e\}_{e \in \mathcal{E}}$ subject to $x_e + y_e = 1$ and for all v, degree constraints $\sum_{e: v \in e} x_e \geq \delta/2$ and $\sum_{e: v \in e} y_e \geq \delta/2$. Note $x \equiv y \equiv \frac{1}{2}$ is a feasible solution. Let us abuse notation and when x or y is 0-1, use them interchangeably with the corresponding subsets of $\mathcal{E}$. So in the LP, a feasible integral x and y would correspond to a discrepancy-zero splitting of $\mathcal{E}$. We'll show that such a solution can be *nearly* found, allowing an additive R violation in the degree constraints. We use the following fact, which follows by double-counting. A constraint is *tight* if it holds with equality.

Lemma 11. *In any extreme-point solution (x, y) to the linear program, there is some tight degree constraint for whom at most R of the variables it involves are strictly between 0 and 1. This holds also if some variables are fixed at integer values and some of the degree constraints have been removed.*

Now we use the following iterated LP rounding algorithm. Each iteration starts with solving the LP and getting an extreme point solution. Then perform two steps: for each variable with an integral value in the solution, fix its value forever; and discard the constraint whose existence is guaranteed by the lemma. Eventually all variables are integral and we terminate.

For each degree constraint, either it was never discarded in which case the final integral solution satisfies it, or else it was discarded in some iteration. Now when the constraint was discarded it had at most R fractional variables, and was tight. So in the sum (say) $\sum_{e: v \in e} x_e = \delta$ there were at least $\delta - R$ variables fixed to 1 on the left-hand side. They ensure $\sum_{e: v \in e} x_e \geq \delta - R$ at termination, proving what we wanted. □

Here is how the Beck-Fiala theorem gives a near-optimal bound on $\bar{\mathsf{p}}'$.

Proposition 12. $\overline{\mathsf{p}}'(\text{HYP}(R), \delta) \geq \delta/O(\ln R)$.

Proof. If $\delta < 4R$ this already follows from Proposition 9. Otherwise apply Proposition 10 to the initial hypergraph, and then use shrinking to make both the resulting $(V, \mathcal{E}_i)$'s into regular hypergraphs. Iterate this process; stop splitting each hypergraph once its degree falls in the range $[R, 4R)$, which is possible since $\delta \geq 4R \Rightarrow \delta/2 - R \geq R$. Let M be the number of hypergraphs at the end.

Observe that in applying the splitting-and-shrinking operation to some $(V, \mathcal{E})$ to get $(V, \mathcal{E}_1)$ and $(V, \mathcal{E}_2)$, the sum of the degrees of $(V, \mathcal{E}_1)$ and $(V, \mathcal{E}_2)$ is at least the degree of $(V, \mathcal{E})$, minus $2R$ "waste". It follows that the total waste is at most $2R(M-1)$, and we have that $4RM + 2R(M-1) \geq \delta$. Consequently $M \geq \delta/6R$. As sketched earlier, applying Proposition 9 to the individual hypergraphs, and combining these vertex covers, shows that $\mathsf{p}' \geq M\lfloor R/\ln(\mathrm{e}R^3)\rfloor$ which gives the claimed bound. □

Now we get to the better bound with the correct multiplicative constant.

Proof of Theorem 1: $\forall R, \delta,\ \overline{\mathsf{p}}'(\text{HYP}(R), \delta) \geq \max\{1, \delta/(\ln R + O(\ln\ln R))\}$. Now Proposition 9 gives us the desired bound when δ is at most polylogarithmic in R, so we assume otherwise. Due to the crude bound in Proposition 12, we may assume R is sufficiently large when needed. We will need the following well-known discrepancy bound which follows using Chernoff bounds and the LLL; see also the full version.

Proposition 13. *For a constant C_1, in a d-regular hypergraph $H = (V, \mathcal{E})$ with all edges of size at most R, we can partition the edge set into $\mathcal{E} = \mathcal{E}_1 \uplus \mathcal{E}_2$ such that $\delta(V, \mathcal{E}_i) \geq d/2 - C_1\sqrt{d\ln(Rd)}$ $(i = 1, 2)$.*

Let $d_0 = \delta$ and $d_{i+1} = d_i/2 - C_1\sqrt{d_i \ln(Rd_i)}$. Thus each hypergraph obtained after i rounds of splitting has degree at least d_i; evidently $d_i \leq \delta/2^i$. We stop splitting after T rounds, where T will be fixed later to make d_T and $\delta/2^T$ polylogarithmic in R. The total degree loss due to splitting is

$$\delta - 2^T d_T = \sum_{i=0}^{T} 2^i(d_i - 2d_{i+1}) \leq \sum_{i=0}^{T-1} 2^i 2C_1\sqrt{d_i \ln Rd_i} \leq \sum_{i=0}^{T-1} 2^i 2C_1\sqrt{\frac{\delta}{2^i}\ln\frac{R\delta}{2^i}}$$
$$= 2C_1\sqrt{\delta}\sum_{i=0}^{T-1}\sqrt{2^i \ln\frac{R\delta}{2^i}}.$$

This sum is an arithmetic-geometric series dominated by the last term, so that we deduce $\delta - 2^T d_T = O(\sqrt{\delta 2^T \ln(R\delta/2^T)})$. Pick T such that $\delta/2^T$ is within a constant factor of $\ln^3 R$, then we deduce

$$d_T \geq \delta/2^T(1 - O(\sqrt{2^T/\delta\ln(R\delta/2^T)})) \geq \delta/2^T(1 - O(\ln^{-1}(R))).$$

Consequently with Proposition 9 we see that

$$\mathsf{p}' \geq 2^T d_T/(\ln R + O(\ln\ln R)) \geq \delta(1 - O(\ln^{-1}(R)))/(\ln R + O(\ln\ln R))$$

which gives the claimed bound. □

2.1 Lower Bounds

Now we show that the bounds obtained previously are tight.

Proof of Theorem 2(a). We want to show, for a constant C and all $R \geq 2, \delta \geq 1$ we have $\overline{\mathsf{p}}'(\textsc{Hyp}(R), \delta) \leq \max\{1, C\delta/\ln R\}$. Consider the hypergraph $H = \binom{[2k-1]}{k}^*$ in the introduction. It is k-regular, it has $\mathsf{p}'(H) = 1$, and $R(H) = \binom{2k-2}{k-1}$.

Since $\overline{\mathsf{p}}'(\textsc{Hyp}(R), \delta)$ is non-increasing in R, we may reduce R by a constant factor to assume that either $R = 2$, (in which case we are done by Gupta's theorem) or $R(H) = \binom{2k-2}{k-1}$ for some integer k. Note this gives $k = \Theta(\log R)$. Moreover, if $\delta \leq k$ then H proves the theorem, so assume $\delta \geq k$. Again by monotonicity, we may increase δ by a constant factor to make δ a multiple of k. Let $\mu = \delta/k$.

Consider the hypergraph μH obtained by copying each of its edges μ times, for an integer $\mu \geq 1$; note that it is δ-regular. The argument in the introduction shows that any set cover has size at least k and therefore average degree at least $k\binom{2k-2}{k-1}/\binom{2k-1}{k} = k^2/(2k-1) = \Theta(\ln R)$. Thus $\overline{\mathsf{p}}'(\mu H) = O(\delta/\ln R)$ which proves the theorem. □

Proof of Theorem 2(b). We want to show as $R, \delta \to \infty$ with $\delta = \omega(\ln R)$, we have $\overline{\mathsf{p}}'(\textsc{Hyp}(R), \delta) \leq (1 + o(1))\delta/\ln(R)$. We assume an additional hypothesis, that $R \geq \delta$; this will be without loss of generality as we can handle the case $\delta > R$ using the μ-replication trick from the proof of Theorem 2(a), since our argument is again based on lower-bounding the minimum size of an set cover.

Let $\delta' = \delta(1+o(1))$ and $R' = R(1-o(1))$ be parameters that will be specified shortly. We construct a random hypergraph with $n = R'^2\delta'$ vertices and $m = R'\delta'^2$ edges, where for each vertex v and each edge E, we have $v \in E$ with independent probability $p = 1/R'\delta'$. Thus each vertex has expected degree δ' and each edge has expected size R'. A standard Chernoff bound together with $np = \omega(\ln m)$ shows the maximum edge size is $(1+o(1))R'$ asymptotically almost surely (*a.a.s.*); pick R' such that this $(1 + o(1))R'$ equals R. Likewise, since $mp = \omega(\ln n)$ a.a.s. the actual minimum degree is at least $(1 - o(1))\delta'$ which we set equal to δ.

We will show that this random hypergraph has $\mathsf{p}' \geq (1+o(1))\delta/\ln R$ a.a.s. via the following bound, which is based off of an analogous bound for Erdős-Renyi random graphs in [12, §2.5]:

Claim 14. *A.a.s. the minimum set cover size is at least* $\frac{1}{p}\ln(pn)(1 - o(1))$.

The proof is given in the full version. This claim finishes the proof since it implies that the maximum number of disjoint set covers p' is at most $(1 + o(1))mp/\ln(pn) = (1 + o(1))\delta'/\ln(R') = (1 + o(1))\delta/\ln(R)$. □

Aside from the results above, not much else is known about specific values of $\overline{\mathsf{p}}'(\textsc{Hyp}(R), \delta)$ for small R, δ. The Fano plane gives $\overline{\mathsf{p}}(\textsc{Hyp}(3), 3) = 1$: if its seven sets are partitioned into two parts, one part has only three sets, and it is not hard to verify the only covers consisting of three sets are pencils through a point and therefore preclude the remaining sets from forming a cover. Moreover, Thomassen [24] showed that every 4-regular, 4-uniform hypergraph has Property B; together with monotonicity we deduce that $\overline{\mathsf{p}}(\textsc{Hyp}(3), 4) \geq \overline{\mathsf{p}}(\textsc{Hyp}(4), 4) \geq 2$.

3 Paths in Trees

Let TreeEdges|Paths denote the following family of hypergraphs: the ground set is the edge set of an undirected tree, and each hyperedge must correspond to the edges lying in some path in the tree. Such systems are cover-decomposable:

Theorem 4. $\overline{\mathsf{p}}'(\textsc{TreeEdges}|\textsc{Paths}, \delta) \geq 1 + \lfloor(\delta - 1)/5\rfloor$.

Proof. In other words, given a family of paths covering each edge at least $\delta = 5k + 1$ times, we can partition the family into $k + 1$ covers. We use induction on k; the case $k = 0$ is evidently true.

We will use an iterated LP relaxation algorithm similar to the one used in Proposition 10. However, it is more convenient to get rid of the y variables; it is helpful to think of them implicitly as $y = 1 - x$. Thus our linear program will have one variable $0 \leq x_P \leq 1$ for every path P. Fix integers A, B such that $A + B = \delta$, and the LP will aim to make x the indicator vector of an A-fold cover, and $1 - x$ the indicator vector of a B-fold cover. So for each edge e of the tree, we will have one *covering* constraint $\sum_{P:e\in P} x_e \geq A$ and one *packing* constraint $\sum_{P:e\in P} x_e \leq |P : e \in P| - B$ (corresponding to coverage for y). Note that the linear program has a feasible fractional solution $x \equiv A/\delta$.

As before, the iterated LP relaxation algorithm repeatedly finds an extreme point solution of the linear program, fixes the value of variables whenever they have integral values, and discards certain constraints. We will use the following analogue of Lemma 11, which is an easy adaptation of a similar result for packing in [6]; we give more details in the full version.

Lemma 15. *Suppose some x variables are fixed to 0 or 1, and some covering/packing constraints are discarded. Any extreme point solution x^* has the following property: there is a tight covering or packing constraint involving at most 3 variables which are fractional in x^*.*

When such a constraint arises, we discard it. As before, any non-discarded constraint is satisfied by the integral x at termination. Additionally, consider a discarded constraint, say a covering one $\sum_{P:e\in P} x_e \geq A$ for some P. When it is discarded, it holds with equality, and the left-hand side consists of 0's, 1's, and at most 3 fractional values. Since A is an integer, it follows that there are at least $A - 2$ 1's on the LHS. The final x still has these variables equal to 1; so overall, x is the characteristic vector of an $(A - 2)$-fold cover, and likewise $1 - x$ is the characteristic vector of a $(B - 2)$-fold cover.

Finally, fix $A = 3$ and $B = \delta - 3$. The final integral x covers every edge at least $3 - 2$ times — it is a cover. The final $1 - x$ covers every edge at least $\delta - 5 = 5(k - 1) + 1$ times. Hence we can use induction to continue splitting $1 - x$, giving the theorem. □

For the related settings where we have paths covering nodes, or dipaths covering arcs, more involved combinatorial lemmas [6, full version] give that $\overline{\mathsf{p}}'(\delta)$ is at least $1 + \lfloor(\delta - 1)/13\rfloor$. We think Theorem 4 is not tight; the best upper bound on $\overline{\mathsf{p}}'$ we know is $\lfloor(3\delta + 1)/4\rfloor$.

For polychromatic numbers and systems of paths in trees, we have:

Theorem 5. $\overline{\mathsf{p}}(\textsc{TreeEdges}|\textsc{Paths}, r) = \lceil r/2\rceil$.

Proof Sketch. For the lower bound, colour the edges of the tree by giving all edges at level i the colour $i \bmod \lceil r/2\rceil$.

In the upper bound, it is enough to consider even r. We use a Ramsey-like argument. Take a complete t-ary tree with $\frac{r}{2}$ levels of edges, so each leaf-leaf path has r edges. In any $(\frac{r}{2} + 1)$-colouring, $t/(\frac{r}{2} + 1)$ of the edges incident on the root have the same colour. Iterating the argument, for large enough t, two root-leaf paths have the same sequence of $\frac{r}{2}$ colours, and their union shows the colouring is not polychromatic. □

4 Future Work

In the *sensor cover* problem (e.g. [10]) each hyperedge has a given duration; we seek to schedule each hyperedge at an offset so that every item in the ground set is covered for the contiguous time interval $[0, T]$ with the *coverage* T as large as possible. Cover-decomposition is the special case where all durations are unit. Clearly T is at most the minimum of the duration-weighted degrees, which we denote by $\overline{\delta}$. Is there always a schedule with $T = \Omega(\overline{\delta}/\ln R)$ if all hyperedges have size at most R? The LLL is viable but splitting does not work and new ideas are needed. In the full version we get a positive result for graphs ($R = 2$):

Theorem 16. *Every instance of sensor cover in graphs has a schedule of coverage at least* $\overline{\delta}/8$.

Acknowledgments. We thank Jessica McDonald, Dömötör Pálvölgyi, and Oliver Schaud for helpful discussions on these topics.

References

1. Gupta, R.P.: On the chromatic index and the cover index of a mulltigraph. In Alavi, Y., Lick, D.R., eds.: Theory and Applications of Graphs: Int. Conf. Kalamazoo, May 11-15, 1976. Volume 642 of Lecture Notes in Mathematics. Springer Verlag (1978) 204–215
2. Andersen, L.D.: Lower bounds on the cover-index of a graph. Discrete Mathematics **25** (1979) 199–210
3. Alon, N., Berke, R., Buchin, K., Buchin, M., Csorba, P., Shannigrahi, S., Speckmann, B., Zumstein, P.: Polychromatic colorings of plane graphs. Discrete & Computational Geometry **42** (2009) 421–442 Preliminary version appeared in *Proc. 24th SOCG*, pages 338–345, 2008.

4. Pálvölgyi, D.: Decomposition of Geometric Set Systems and Graphs. PhD thesis, École Polytechnique Fédérale de Lausanne (2010) arXiv:1009.4641.
5. Elekes, M., Mátrai, T., Soukup, L.: On splitting infinite-fold covers. Fund. Math. **212** (2011) 95–127 arXiv:0911.2774.
6. Könemann, J., Parekh, O., Pritchard, D.: Max-weight integral multicommodity flow in spiders and high-capacity trees. In: Proc. 6th WAOA. (2008) 1–14
7. Pach, J., Tardos, G., Tóth, G.: Indecomposable coverings. Canadian Mathematical Bulletin **52** (2009) 451–463 Preliminary version in *Proc. 7th CJCDGCGT* (2005), pages 135-148, 2007.
8. Moser, R.A., Tardos, G.: A constructive proof of the general Lovász Local Lemma. J. ACM **57** (2010) Preliminary version in *Proc. 41st STOC*, 343–350, 2009.
9. Buchsbaum, A.L., Efrat, A., Jain, S., Venkatasubramanian, S., Yi, K.: Restricted strip covering and the sensor cover problem. In: Proc. 18th SODA. (2007) 1056–1063
10. Gibson, M., Varadarajan, K.: Decomposing coverings and the planar sensor cover problem. In: Proc. 50th FOCS. (2009) 159–168
11. Pach, J., Tardos, G.: Tight lower bounds for the size of epsilon-nets. In: Proc 27th SoCG. (2011) 458–463 arXiv:1012.1240.
12. Feige, U., Halldórsson, M.M., Kortsarz, G., Srinivasan, A.: Approximating the domatic number. SIAM J. Comput. **32** (2002) 172–195 Preliminary version appeared in *Proc. 32nd STOC*, pages 134-143, 2000.
13. Călinescu, G., Chekuri, C., Vondrák, J.: Disjoint bases in a polymatroid. Random Struct. Algorithms **35** (2009) 418–430
14. Pach, J.: Decomposition of multiple packing and covering. In: Kolloquium über Diskrete Geometrie, Salzburg, Inst. Math. U. Salzburg (1980) 169–178
15. Aloupis, G., Cardinal, J., Collette, S., Imahori, S., Korman, M., Langerman, S., Schwartz, O., Smorodinsky, S., Taslakian, P.: Colorful strips. In: Proc. 9th LATIN. (2010) 2–13
16. Keszegh, B., Pálvölgyi, D.: Octants are cover decomposable. arXiv:1101.3773 (2011)
17. Keszegh, B.: Weak conflict free colorings of point sets and simple regions. In: Proc. 19th CCCG. (2007) Extended version "Coloring half-planes and bottomless rectangles" at arXiv:1105.0169.
18. Cheriyan, J., Jordán, T., Ravi, R.: On 2-coverings and 2-packings of laminar families. In: Proc. 7th ESA. (1999) 510–520
19. Cheilaris, P., Keszegh, B., Pálvölgyi, D.: Unique-maximum and conflict-free colorings for hypergraphs and tree graphs. In: Proc. 7th Japanese-Hungarian Symp. Disc. Math. Appl. (2011) 207–216 arXiv:1002.4210.
20. Erlebach, T., Jansen, K.: The maximum edge-disjoint paths problem in bidirected trees. SIAM J. Discrete Math. **14** (2001) 326–355 Preliminary version in *Proc. 9th ISAAC*, 1998.
21. Erdős, P., Lovász, L.: Problems and results on 3-chromatic hypergraphs and some related questions. In Hajnal, A., Rado, R., Sós, V.T., eds.: Infinite and Finite Sets (Coll. Keszthely, 1973). Volume 10 of Colloq. Math. Soc. János Bolyai. North-Holland (1975) 609–627
22. Alon, N., Spencer, J.H.: The Probabilistic Method (3rd Edition). Wiley, New York (2008)
23. Beck, J., Fiala, T.: "Integer-making" theorems. Discrete Applied Mathematics **3** (1981) 1–8
24. Thomassen, C.: The even cycle problem for directed graphs. J. Amer. Math. Soc. **5** (1992) 217–229