Density-Sensitive Total Weight for Partite Decompositions of Uniform Hypergraphs

Bolzano team*

July 2026

Abstract

For a complete d-partite d-uniform hypergraph $K(A_1, \dots, A_d)$, define its weight as $\sum_i |A_i|$. Consider covers and edge-disjoint exact partitions of an n-vertex d-uniform hypergraph into such pieces. For every fixed $d \geq 3$ and $0 < \varepsilon < 1$, the worst-case minimum total weight among hypergraphs with $m \leq (1-\varepsilon)\binom{n}{d}$ edges is

$$\Theta_{d,\varepsilon}\left(m \min\left\{1, \frac{\log(1/\gamma)}{\log n}\right\}\right), \qquad \gamma = \frac{m}{\binom{n}{d}}.$$

The formula is the same for covers and exact partitions. The lower bound uses an exact-density random hypergraph; the upper bound repeatedly extracts density-sensitive balanced boxes and then partitions the sparse remainder. Espuña's extraction algorithm makes the construction deterministic and polynomial-time. Cone examples show that the average total-weight scale W/n need not control the maximum number of pieces through one vertex. Complete and near-complete endpoint bounds are also given, while the sharp transition as $\gamma \to 1$ remains open.

1 Introduction

1.1 Background and motivating problem

All logarithms are natural unless a base is displayed; changing the base only changes constants depending on d. Let H be a d-uniform hypergraph. A complete d-partite d-graph, or d-clique in the terminology of [10], is a hypergraph

$$K(A_1, \dots, A_d) = \big\{\{a_1, \dots, a_d\} : a_i \in A_i \text{ for all } i\big\},$$

where the parts $A_1, \dots, A_d$ are nonempty and pairwise disjoint. A family of such pieces is a cover of H if every edge of H is contained in at least one piece, and it is an exact partition if every edge of H is contained in exactly one piece. In both cases all pieces are required to be subhypergraphs of H.

Krapivin, Przybocki, Sanhueza-Matamala, and Subercaseaux [10] prove the optimal worst-case local partition theorem for uniform hypergraphs in Theorem 5 of their full version and, separately, the density-aware graph partition theorem in full-version Theorem 7. Their hypergraph framework continues the local-decomposition program of Csirmaz, Ligeti, and Tardos [8]. The motivating direction is the first item of their Section 10:

*Dissemination note. All results produced by Bolzano are publicly available at bolzano.app. When a result appears especially promising, we give it a preliminary review and contact the authors of the source paper, inviting them to use it in any way they consider appropriate. To avoid excessive unsolicited email, we reserve this outreach for a small number of cases. This note presents a result that either fell just below our threshold for direct outreach or about which we contacted the source-paper authors but received no response. We are disseminating it more broadly in the hope that it may still be useful to the community.

Problem 1 (Krapivin et al., full version, Section 10 [10])**.** Study analogous density-aware results for d-uniform hypergraphs, with edge density $\gamma_d = |E(H)|/\binom{n}{d}$.

This is a broad future-work direction rather than a numbered conjecture. The main theorem answers its total-weight reading for fixed d away from the complete endpoint. It does not determine the worst-case local-load function at each density. Instead, it identifies a link-concentration obstruction showing why the total-weight answer divided by n is not a valid maximum-load principle.

The balanced complete-partite subgraph used in the upper bound belongs to the classical hypergraph Zarankiewicz theory initiated by Erdős [3]. Espuña recently gave a deterministic polynomial-time algorithm attaining the same order of balanced part size, with explicit density dependence [4]. Exact partitions of random uniform hypergraphs were studied by Peng [6]; that work optimizes the number of pieces in a random model, whereas the present result concerns worst-case total weight at a prescribed edge count.

1.2 Contribution

For a piece $P = K(A_1, \dots, A_d)$ write

$$w(P) = \sum_{i=1}^{d} |A_i|, \qquad e(P) = \prod_{i=1}^{d} |A_i|.$$

For a cover or partition $\mathcal{P}$, set

$$W(\mathcal{P}) = \sum_{P \in \mathcal{P}} w(P), \qquad L(\mathcal{P}) = \max_{v \in V(H)} |\{P \in \mathcal{P} : v \in V(P)\}|.$$

For later comparison with the complete endpoint, define

$$\begin{aligned}
W_{\mathrm{cov}}(H) &= \min\{W(\mathcal{P}) : \mathcal{P} \text{ is a cover of } H\},\\
W_{\mathrm{part}}(H) &= \min\{W(\mathcal{P}) : \mathcal{P} \text{ is an exact partition of } H\},\\
L_{\mathrm{cov}}(H) &= \min\{L(\mathcal{P}) : \mathcal{P} \text{ is a cover of } H\},\\
L_{\mathrm{part}}(H) &= \min\{L(\mathcal{P}) : \mathcal{P} \text{ is an exact partition of } H\}.
\end{aligned}$$

The central contribution is the following.

(i) Theorem 2 determines, up to constants depending only on d and the distance from density one, the worst-case total weight of both covers and exact partitions at every edge count.

(ii) Theorem 3 exhibits link concentration that forces one vertex to have much larger load than the average W/n, so a local theorem requires information beyond global density.

(iii) Theorem 4 and Proposition 22 place the non-dense result against the complete and small-complement regimes and isolate the remaining dense transition.

The link example that follows quantifies a limitation already visible in Example 25 of the full version of Krapivin et al. [10]. The complete-endpoint section is included as context: it presents weighted reformulations of classical Graham–Pollak [9], Katona–Szemerédi [5], and Alon [1] results, rather than claiming new endpoint determinations.

Theorem 2 (Density-aware total weight away from the complete endpoint)**.** *Fix $d \geq 3$ and $0 < \varepsilon < 1$. Let $M = \binom{n}{d}$, let $1 \leq m \leq (1-\varepsilon)M$, and put $\gamma = m/M$. Among all n-vertex d-uniform hypergraphs with m edges, the worst-case minimum total weight of an exact partition into complete d-partite pieces is*

$$\Theta_{d,\varepsilon}\left(m \min\left\{1, \frac{\log(1/\gamma)}{\log n}\right\}\right).$$

The same asymptotic formula holds for covers.

The lower bound is a density-controlled random construction; the upper bound is a density-sensitive extraction and peeling argument. The theorem gives $\Theta_d(m)$ in the sparse range $m \lesssim n^{d-1}$ and

$$\Theta_{d,\varepsilon}\left(\gamma n^d \frac{\log(1/\gamma)}{\log n}\right)$$

in the intermediate range $1/n \lesssim \gamma \leq 1-\varepsilon$.

Theorem 2 supplies an average-load bound W/n for its constructed partition. The next theorem shows that this average need not control the busiest vertex.

Theorem 3 (A link-concentration example)**.** *Fix $d \geq 3$, let $N \geq d-1$, and let $n \geq N+1$. Let U have N vertices, add an apex x and $n-N-1$ isolated padding vertices, and let*

$$H = \left\{\{x\} \cup S : S \in \binom{U}{d-1}\right\}.$$

Then H is an n-vertex hypergraph of density $\Theta_d(N^{d-1}/n^d)$. Writing $k = \lfloor (d-1)/2 \rfloor$, every exact partition has apex load at least

$$\max\left\{N-d+2, \frac{\binom{N}{k}}{\binom{d-1}{k}}\right\} = \Omega_d\left(N^{\lfloor (d-1)/2 \rfloor}\right),$$

and every cover has apex load at least $\lceil \log_2(N-d+3) \rceil$.

For instance, if $N = n^a$ with $0 < a < 1/(d-2)$, then the apex lower bound $\Omega(N)$ for partitions is asymptotically larger than the average scale $|E(H)|/n = \Theta_d(N^{d-1}/n)$. Thus link data are indispensable for a per-instance local statement at the total-weight scale. Determining the density-constrained worst-case local parameter remains open.

The complete endpoint also behaves differently from the entropy-type non-complete formula.

Theorem 4 (Complete endpoint benchmarks)**.** *For fixed $d \geq 2$, covers of the complete d-uniform hypergraph $K_n^{(d)}$ have optimal total weight $\Theta_d(n \log n)$ and optimal maximum load $\Theta_d(\log n)$.*

For exact partitions and fixed $d \geq 3$,

$$W_{\mathrm{part}}(K_n^{(d)}) \geq \Omega_d(n^{\lceil d/2 \rceil}), \qquad L_{\mathrm{part}}(K_n^{(d)}) \geq \Omega_d(n^{\lceil d/2 \rceil - 1}),$$

and there is an exact partition with

$$W_{\mathrm{part}}(K_n^{(d)}) \leq O_d\left(n^{\lceil d/2 \rceil} (\log n)^{\mathbf{1}_{d \text{ even}}}\right), \quad L_{\mathrm{part}}(K_n^{(d)}) \leq O_d\left(n^{\lceil d/2 \rceil - 1} (\log n)^{\mathbf{1}_{d \text{ even}}}\right).$$

In particular, for odd $d = 2s+1$ the exact-partition values are $\Theta_d(n^{s+1})$ for total weight and $\Theta_d(n^s)$ for maximum load.

The remaining dense transition $\gamma \to 1$ is not settled here. Section 4.3 records repair bounds for $K_n^{(d)}$ with few missing edges, but these do not yet interpolate sharply between Theorem 2 and Theorem 4.

2 Total weight away from the complete endpoint

2.1 The random lower bound

Lemma 5 (Exact-density random lower bound)**.** *Let $M = \binom{n}{d}$, let $1 \le m < M$, set $\gamma = m/M$, and put*

$$R = \max\left\{1, \frac{\log n}{\log(1/\gamma)}\right\}.$$

For fixed d, there is an m-edge d-uniform hypergraph H on $[n]$ such that every complete d-partite subhypergraph $P \subseteq H$ satisfies

$$e(P) \le C_d R\, w(P).$$

Consequently every cover or exact partition $\mathcal{P}$ of H satisfies

$$W(\mathcal{P}) \ge c_d\, m \min\left\{1, \frac{\log(1/\gamma)}{\log n}\right\}.$$

Proof. Choose H uniformly among all m-edge subhypergraphs of $\binom{[n]}{d}$. Fix an ordered candidate piece $P = K(A_1, \dots, A_d)$ with total part size $s = w(P)$ and edge count $a = e(P)$. If $a > m$, then $P \subseteq H$ has probability zero. Otherwise

$$\Pr(P \subseteq H) = \frac{\binom{M-a}{m-a}}{\binom{M}{m}} \le \left(\frac{m}{M}\right)^a = \gamma^a.$$

For a fixed s, the number of ordered choices of pairwise disjoint parts with total size s is at most $\binom{n}{s} d^s \le (dn)^s$. Since $R \log(1/\gamma) \ge \log n$, if $a > K_d R s$ then

$$\Pr(P \subseteq H) \le \exp[-K_d R s \log(1/\gamma)] \le n^{-K_d s}.$$

The probability that some ordered candidate of total size s and edge-to-weight ratio above $K_d R$ is contained in H is therefore at most

$$(dn)^s n^{-K_d s} = d^s n^{-(K_d - 1)s}.$$

For K_d sufficiently large, summing over $s \ge 1$ gives $o(1)$. Hence with positive probability every piece contained in H has $e(P) \le C_d R w(P)$.

For such an H, if $\mathcal{P}$ covers H, then

$$m = |E(H)| \le \sum_{P \in \mathcal{P}} e(P) \le C_d R \sum_{P \in \mathcal{P}} w(P) = C_d R\, W(\mathcal{P}).$$

Since

$$R^{-1} = \min\left\{1, \frac{\log(1/\gamma)}{\log n}\right\},$$

the claimed lower bound follows. Exact partitions are a special case of covers for this argument. $\square$

2.2 Sparse ordered-star partitions

Lemma 6 (Sparse universal partition). *Every m-edge d-uniform hypergraph on an ordered vertex set has an exact partition into complete d-partite pieces with*

$$W \le m + (d-1)\min\left\{m, \binom{n}{d-1}\right\} \le dm.$$

Proof. For each $(d-1)$-set $S = \{s_1 < \cdots < s_{d-1}\}$, define

$$T(S) = \{v : s_{d-1} < v,\ S \cup \{v\} \in E(H)\}.$$

If $T(S)$ is nonempty, form

$$P_S = K(\{s_1\}, \dots, \{s_{d-1}\}, T(S)).$$

This piece is contained in H. Every edge $\{x_1 < \cdots < x_d\}$ belongs to the unique piece indexed by $S = \{x_1, \dots, x_{d-1}\}$. Thus these pieces form an exact partition.

Their total weight is

$$\sum_{S:T(S)\neq\emptyset} \big((d-1) + |T(S)|\big) = m + (d-1)|\{S : T(S) \neq \emptyset\}|.$$

The number of nonempty S is at most m and at most $\binom{n}{d-1}$, proving the bound. □

2.3 The extraction lemma

The upper bound in the intermediate-density range uses a complete partite piece with a density-sensitive edge-to-weight ratio. For constant density, the existence of a balanced box of this order is classical [3]; the lemmas below give the precise variable-density form needed by the peeling argument.

Lemma 7 (Homomorphic boxes). *Let G be an r-partite r-uniform hypergraph with vertex classes $X_1, \dots, X_r$ and partite density*

$$\alpha = \frac{e(G)}{|X_1|\cdots|X_r|}.$$

If one chooses independently and uniformly ordered s-tuples $\mathbf{x}_i \in X_i^s$ for $i = 1, \dots, r$, allowing repetitions, then the probability that all s^r transversals are edges of G is at least α^{s^r}.

Proof. The proof is by induction on r. The case $r = 1$ is immediate. For $r \ge 2$, first expose $\mathbf{x}_1, \dots, \mathbf{x}_{r-1}$. Let Z be the fraction of vertices $y \in X_r$ such that every transversal using y in the last coordinate is an edge. Conditional on the first $r-1$ tuples, the last tuple completes a homomorphic box with probability Z^s, so the desired probability is $\mathbb{E}Z^s$.

For $y \in X_r$, let α_y be the partite density of the link of y on $X_1, \dots, X_{r-1}$. By the induction hypothesis,

$$\mathbb{E}Z \ge \frac{1}{|X_r|}\sum_{y\in X_r} \alpha_y^{s^{r-1}} \ge \left(\frac{1}{|X_r|}\sum_{y\in X_r}\alpha_y\right)^{s^{r-1}} = \alpha^{s^{r-1}},$$

where Jensen's inequality is used and the average link density is α. A second application of Jensen gives

$$\mathbb{E}Z^s \ge (\mathbb{E}Z)^s \ge \alpha^{s^r}.$$

□

Lemma 8 (Balanced honest box). *Let G be a d-partite d-uniform hypergraph with vertex classes $X_1, \dots, X_d$, minimum part size N_*, and partite density ρ. If*

$$N_* \rho^{s^{d-1}} \geq 4ds^2,$$

then G contains a copy of $K^{(d)}_{s,\dots,s}$ with s distinct vertices in each class.

Proof. Choose random ordered s-tuples in $X_1, \dots, X_{d-1}$, allowing repetitions. Let Y be the number of vertices $y \in X_d$ that complete all transversals. By Lemma 7 applied to links and Jensen's inequality,

$$\mathbb{E}Y \geq |X_d| \rho^{s^{d-1}}.$$

Let $\mathcal{B}$ be the event that at least one of the first $d-1$ sampled tuples has a repeated entry. By the union bound,

$$\Pr(\mathcal{B}) \leq \sum_{i=1}^{d-1} \frac{\binom{s}{2}}{|X_i|} \leq \frac{ds^2}{2N_*}.$$

Writing $\beta = \rho^{s^{d-1}}$, the hypothesis gives $\Pr(\mathcal{B}) \leq \beta/8$. Since always $Y \leq |X_d|$,

$$\mathbb{E}(Y \mathbf{1}_{\neg \mathcal{B}}) \geq |X_d|\beta - |X_d| \Pr(\mathcal{B}) \geq \frac{7}{8}|X_d|\beta \geq \frac{7}{8} N_* \beta \geq s.$$

Thus some repetition-free choice of the first $d-1$ tuples has at least s common neighbors in X_d. Choosing any s of them gives the required honest complete d-partite box. □

Lemma 9 (Balanced slice). *Let H be a d-uniform hypergraph on n vertices with density $\rho = e(H)/\binom{n}{d}$. Put $N = \lfloor n/d \rfloor$. There are pairwise disjoint sets $X_1, \dots, X_d \subseteq V(H)$, each of size N, such that the transversal d-partite subhypergraph across them has partite density at least ρ.*

Proof. Choose the ordered disjoint N-sets uniformly at random. For a fixed edge $e \in E(H)$, the probability that it is transversal across the selected classes is

$$\frac{d! N^d}{(n)_d}.$$

The expected number of transversal edges is therefore $e(H) d! N^d/(n)_d$. Dividing by the total number N^d of possible transversals, the expected partite density is

$$\frac{d! e(H)}{(n)_d} = \frac{e(H)}{\binom{n}{d}} = \rho.$$

Some choice has density at least the expectation. □

Proposition 10 (Non-dense extraction). *Fix $d \geq 2$ and $q_0 > 0$. There is $c = c(d, q_0) > 0$ such that every nonempty n-vertex d-uniform hypergraph H of density ρ with $q = \log(1/\rho) \geq q_0$ contains a complete d-partite piece $P \subseteq H$ satisfying*

$$\frac{e(P)}{w(P)} \geq c \max\left\{1, \frac{\log n}{q}\right\}.$$

Proof. The finitely many small n may be absorbed into the constant, so assume n is large in terms of d and q_0. Put $T = \log n/q$.

If T is bounded by a constant T_0, a single edge gives $e(P)/w(P) = 1/d$, which is a constant multiple of $\max\{1, T\}$ after decreasing c.

Assume $T > T_0$. Choose a constant $a \in (0, 1)$ with $a^{d-1} < 1/2$ and set

$$s = \left\lfloor aT^{1/(d-1)} \right\rfloor.$$

For T_0 large, $s \geq (a/2)T^{1/(d-1)} \geq 1$. By Lemma 9, H has a balanced d-partite slice with all parts of size $N = \lfloor n/d \rfloor$ and partite density $\rho' \geq \rho$. Since $s^{d-1} \leq a^{d-1}T$,

$$(\rho')^{s^{d-1}} \geq \rho^{s^{d-1}} = \exp[-qs^{d-1}] \geq n^{-a^{d-1}}.$$

Also $s^2 \leq a^2(\log n/q_0)^{2/(d-1)}$. Because $1 - a^{d-1} > 0$, for all sufficiently large n we have

$$N(\rho')^{s^{d-1}} \geq 4ds^2.$$

Lemma 8 gives a copy of $K^{(d)}_{s,\dots,s}$ inside H. For this piece,

$$\frac{e(P)}{w(P)} = \frac{s^d}{ds} = \frac{s^{d-1}}{d} \geq c'(d, q_0)\frac{\log n}{q}.$$

This proves the proposition. □

2.4 Peeling proof of the upper bound

Proposition 11 (Non-dense total-weight upper bound)**.** *Fix $d \geq 2$ and $q_0 > 0$. There is $C = C(d, q_0)$ such that every n-vertex d-uniform hypergraph H with m edges, $M = \binom{n}{d}$, density $\gamma = m/M$, and $0 < \gamma \leq e^{-q_0}$ has an exact partition into complete d-partite pieces with*

$$W \leq C\, m \min\left\{1, \frac{\log(1/\gamma)}{\log n}\right\}.$$

Proof. Let $q = \log(1/\gamma)$. If $q \geq \log n$, then $\gamma \leq 1/n$ and Lemma 6 gives $W \leq dm$, which is the desired bound.

Assume $q < \log n$. Then $m > M/n$. Starting with $H_0 = H$, repeatedly apply Proposition 10 to the current hypergraph H_i while its edge count $x_i = e(H_i)$ is greater than M/n. Put

$$q_i = \log(M/x_i).$$

Since $x_i \leq m$ and $x_i > M/n$, we have $q_0 \leq q_i < \log n$. The extracted piece P_i satisfies

$$\frac{e(P_i)}{w(P_i)} \geq c\frac{\log n}{q_i}.$$

Let $a_i = e(P_i)$. Then

$$w(P_i) \leq \frac{a_i q_i}{c \log n}.$$

Delete the edges of P_i and continue. When the process stops, let H_* be the leftover hypergraph, with $x_* = e(H_*) \leq M/n$. Partition H_* by Lemma 6, at cost at most $dx_* \leq dM/n$.

It remains to estimate the peeled cost. During step i, the edge count decreases from x_i to $x_i - a_i$. Since $y \mapsto \log(M/y)$ is decreasing,

$$a_i q_i \le \int_{x_i - a_i}^{x_i} \log(M/y)\, dy.$$

Summing all steps gives

$$\sum_i a_i q_i \le \int_{x_*}^{m} \log(M/y)\, dy \le m \log(M/m) + m = m(q+1).$$

Thus the peeled pieces have total weight at most

$$\frac{m(q+1)}{c \log n} \le C_1(d, q_0) \frac{mq}{\log n}.$$

The leftover cost is absorbed in the same scale. Since $m > M/n$, we have $\gamma > 1/n$. There is $c_0 = c_0(q_0) > 0$ such that

$$\gamma \log(1/\gamma) \ge c_0 \frac{\log n}{n} \qquad (1/n < \gamma \le e^{-q_0}).$$

Indeed, if $\gamma \le e^{-1}$ the function $t \log(1/t)$ is increasing on $[1/n, e^{-1}]$, and if $\gamma > e^{-1}$ then $q_0 < 1$ and $\gamma \log(1/\gamma) \ge e^{-1} q_0$. Since $m = \gamma M$ and $q = \log(1/\gamma)$, rearranging this inequality yields

$$M/n \le C_2(q_0) \frac{mq}{\log n}.$$

Therefore the leftover ordered-star cost also has size $O_{d,q_0}(mq/\log n)$. Since $q < \log n$, this is exactly the asserted intermediate-range bound. □

Proof of Theorem 2. The lower bound is Lemma 5. For the upper bound, apply Proposition 11 with

$$q_0 = -\log(1 - \varepsilon),$$

because $m \le (1-\varepsilon)M$ means $\gamma \le e^{-q_0}$. The proposition gives an exact partition and therefore also a cover. The lower bound applies to both covers and partitions, so the two estimates match up to constants depending only on d and ε. □

Corollary 12 (Polynomial-time construction)**.** *For fixed $d \ge 3$ and $0 < \varepsilon < 1$, an exact partition attaining the upper bound in Theorem 2 can be found by a deterministic polynomial-time algorithm.*

Proof. For a current hypergraph of density ρ, Espuña's algorithm [4, Thm. 1; arXiv v2, Thm. 1.2] finds in polynomial time a balanced complete d-partite subgraph with side length

$$t = \left\lfloor \left(\frac{\log n}{\log(16/\rho)} \right)^{1/(d-1)} \right\rfloor,$$

returning a single edge when this quantity is smaller than two. Throughout the peeling regime, $q = \log(1/\rho) \ge q_0 = -\log(1-\varepsilon)$, and hence $\log(16/\rho) = q + \log 16 = \Theta_\varepsilon(q)$. The returned piece therefore has the same edge-to-weight order $\Omega_{d,\varepsilon}(\max\{1, \log n/q\})$ used in Proposition 11. Repeatedly call this algorithm, delete the edges of the returned piece, and finish with the ordered-star partition from Lemma 6. At least one edge is removed per call, so there are at most m calls; the construction and the output are polynomial in the input size. The cost analysis of Proposition 11 is unchanged. □

3 Maximum load and link obstructions

We use two elementary facts about complete graphs. They are included to make the link argument independent of external decomposition theorems.

Lemma 13 (Bicliques in a complete graph)**.** *An exact partition of K_r into complete bipartite graphs has at least $r-1$ pieces. A cover of K_r by complete bipartite graphs has at least $\lceil \log_2 r \rceil$ pieces.*

Proof. Suppose first that K_r is exactly partitioned by $K(A_i, B_i)$, $1 \le i \le k$, and that $k < r-1$. The homogeneous equations

$$\sum_{v=1}^{r} x_v = 0, \qquad \sum_{v \in A_i} x_v = 0 \quad (1 \le i \le k)$$

impose fewer than r linear constraints, so they have a nonzero real solution x. Exactness gives

$$\sum_{u<v} x_u x_v = \sum_{i=1}^{k} \left(\sum_{u \in A_i} x_u \right) \left(\sum_{v \in B_i} x_v \right) = 0.$$

On the other hand, the first constraint gives

$$\sum_{u<v} x_u x_v = \frac{1}{2} \left[\left(\sum_v x_v \right)^2 - \sum_v x_v^2 \right] = -\frac{1}{2} \sum_v x_v^2 < 0,$$

a contradiction. This is the Graham–Pollak argument [9].

For a cover, orient the two sides of each of its k bicliques and assign to each vertex v the set $C(v)$ of indices for which v lies on the first side. If $u \ne v$, some biclique covers uv; in that biclique one endpoint is on the first side and the other on the second, so $C(u) \ne C(v)$. Thus the r vertices receive distinct subsets of $[k]$, whence $r \le 2^k$. □

Lemma 14 (Link trace)**.** *Let H be a d-uniform hypergraph and $v \in V(H)$. If $\mathcal{P}$ is an exact partition of H into complete d-partite pieces, then the pieces containing v induce an exact partition of the link $L_H(v)$ into complete $(d-1)$-partite pieces. If $\mathcal{P}$ is a cover, they induce a cover of the link.*

Proof. Suppose $P = K(A_1, \dots, A_d)$ contains v, say $v \in A_j$. Its trace on the link of v is

$$K(A_1, \dots, A_{j-1}, A_{j+1}, \dots, A_d).$$

Every edge of this trace, together with v, is an edge of P and hence of H. Conversely, every edge $S \in L_H(v)$ corresponds to the d-edge $S \cup \{v\}$, which is covered by some piece containing v. In the partition case the covering piece is unique, so the induced link pieces are also an exact partition. □

Proof of Theorem 3. The link of the apex x is the complete $(d-1)$-uniform hypergraph on U. By Lemma 14, every piece containing x gives a complete $(d-1)$-partite piece in this link.

For exact partitions, fix an $(d-3)$-set $S \subseteq U$; when $d=3$, take $S = \emptyset$. Restrict attention to link edges containing S. After deleting the fixed vertices of S, these edges are precisely the edges of a complete graph on $N-d+3$ vertices. The induced traces of the $(d-1)$-partite link pieces give an exact biclique partition of this complete graph. Lemma 13 gives at least $(N-d+3)-1 = N-d+2$ bicliques. Hence the apex load is at least $N-d+2$.

There is also a stronger higher-rank estimate. The link traces themselves form an exact partition of $K_N^{(d-1)}$. Lemma 19, applied with $r = d-1$ and $k = \lfloor (d-1)/2 \rfloor$, shows that their number is at least $\binom{N}{k}/\binom{d-1}{k}$. Taking the larger of the two lower bounds gives the asserted exact-partition estimate.

For covers, the same restriction gives a biclique cover of K_{N-d+3}. The second assertion of Lemma 13 gives apex load at least $\lceil \log_2(N-d+3) \rceil$.

Finally,

$$|E(H)| = \binom{N}{d-1}, \qquad \frac{|E(H)|}{\binom{n}{d}} = \Theta_d(N^{d-1}/n^d).$$

□

4 Known endpoint benchmarks and weighted reformulations

4.1 Covers of the complete hypergraph

The minimum total order of complete partite pieces covering a complete uniform hypergraph was studied by Radhakrishnan [7] and, in the modern hypergraph-covering literature, by Babu and Vishwanathan [2]. In particular, the $\Theta_d(n \log n)$ total-weight scale below is known. A proof is included to relate that parameter to the maximum-load notation used here.

Lemma 15 (Weighted biclique-cover bound)**.** *If bicliques $K(A_i, B_i)$ cover K_n, then*

$$\sum_i \big(|A_i| + |B_i|\big) \geq n \log_2 n.$$

Proof. Orient the two sides of each biclique. For every vertex v, form a partial word $\sigma_v \in \{0, 1, *\}^k$: its ith coordinate is 0 when $v \in A_i$, is 1 when $v \in B_i$, and is $*$ otherwise. Let $c(v)$ be the number of specified coordinates, and let $Q_v \subseteq \{0,1\}^k$ be the subcube of binary words extending σ_v. Then $|Q_v| = 2^{k-c(v)}$.

If $u \neq v$, a biclique covering uv assigns opposite bits to u and v in some coordinate. Consequently $Q_u \cap Q_v = \varnothing$. The subcubes therefore satisfy

$$\sum_v 2^{k-c(v)} \leq 2^k, \qquad \text{or equivalently} \qquad \sum_v 2^{-c(v)} \leq 1.$$

Since $x \mapsto 2^{-x}$ is convex, Jensen's inequality yields

$$2^{-n^{-1}\sum_v c(v)} \leq \frac{1}{n} \sum_v 2^{-c(v)} \leq \frac{1}{n}.$$

Thus $\sum_v c(v) \geq n \log_2 n$. Counting specified coordinates by bicliques instead of vertices gives $\sum_v c(v) = \sum_i (|A_i| + |B_i|)$. □

Proposition 16 (Known complete-cover endpoint)**.** *For fixed $d \geq 2$, covers of $K_n^{(d)}$ by complete d-partite pieces have optimal total weight $\Theta_d(n \log n)$ and optimal maximum load $\Theta_d(\log n)$.*

Proof. For the upper bound, take $k = C_d \log n$ independent random colorings $f_j : V(K_n^{(d)}) \to [d]$. A fixed d-set is rainbow under one coloring with probability $p_d = d!/d^d > 0$, so for C_d large enough the probability that some d-set is rainbow under none of the colorings is below 1. Hence there are $O_d(\log n)$ colorings so that every d-set is rainbow in at least one. Each coloring gives one complete d-partite piece with parts $f_j^{-1}(1),\dots,f_j^{-1}(d)$, of weight n. These pieces cover $K_n^{(d)}$, giving $W = O_d(n\log n)$ and $L = O_d(\log n)$.

For the lower bound, let $\mathcal{P}$ be any complete d-partite cover. From each piece $P = K(A_1,\dots,A_d)$ form all $\binom{d}{2}$ bicliques $K(A_i,A_j)$ for $i<j$. Their total weight is $(d-1)w(P)$. These bicliques cover the complete graph K_n: any pair of vertices can be extended to a d-edge, and any piece covering that edge places the pair in two distinct parts.

Lemma 15 says that the induced biclique cover has total weight at least $n\log_2 n$. Therefore $(d-1)W(\mathcal{P}) \ge n\log_2 n$, so $W(\mathcal{P}) = \Omega_d(n\log n)$. Since $L(\mathcal{P}) \ge W(\mathcal{P})/n$, the load lower bound follows. □

4.2 Exact partitions of the complete hypergraph

Alon proved that the minimum number of complete r-partite pieces in an exact partition of $K_n^{(r)}$ is $\Theta_r(n^{\lfloor r/2\rfloor})$, and that the exact value for $r=3$ is $n-2$ [1]. The split-product construction and inclusion-matrix argument below are weighted versions of this classical Graham–Pollak line; they are retained to track total weight and local load.

Lemma 17 (Weighted split-product construction). *For fixed $d\ge 2$, $K_n^{(d)}$ has an exact partition into complete d-partite pieces with*

$$W \le O_d\Big(n^{\lceil d/2\rceil}(\log n)^{\mathbf{1}_{d\ \text{even}}}\Big), \qquad L \le O_d\Big(n^{\lceil d/2\rceil-1}(\log n)^{\mathbf{1}_{d\ \text{even}}}\Big).$$

Proof. Define recursively a family $D_r(V)$ of complete r-partite pieces partitioning $\binom{V}{r}$. If $|V|<r$, take the empty family. For $r=1$ and $V\neq\varnothing$, take the single 1-partite piece on V. In the remaining case $r\ge 2$ and $|V|\ge r$, split $V = A\sqcup B$ as evenly as possible. Include the pieces of $D_r(A)$ and $D_r(B)$, and for every $1\le k\le r-1$, every $P = K(P_1,\dots,P_k)\in D_k(A)$, and every $Q = K(Q_1,\dots,Q_{r-k})\in D_{r-k}(B)$, include the product piece

$$P\otimes Q = K(P_1,\dots,P_k,Q_1,\dots,Q_{r-k}).$$

Every r-edge e has a unique value $k=|e\cap A|$. If $k=0$ or $k=r$, induction partitions it inside one side. If $1\le k\le r-1$, induction gives unique pieces covering $e\cap A$ and $e\cap B$, and their product covers e. Thus $D_r(V)$ is an exact partition.

Let $C_r(n), W_r(n), L_r(n)$ denote the number of pieces, total weight, and maximum load in this recursive construction. The product rule gives, for $a=\lfloor n/2\rfloor$ and $b=\lceil n/2\rceil$,

$$C_r(n) \le C_r(a) + C_r(b) + \sum_{k=1}^{r-1} C_k(a)C_{r-k}(b),$$

$$W_r(n) \le W_r(a) + W_r(b) + \sum_{k=1}^{r-1}\big(C_{r-k}(b)W_k(a) + C_k(a)W_{r-k}(b)\big),$$

and, after bounding separately the load of vertices in A and in B,

$$L_r(n) \le \max\left\{\begin{aligned} &L_r(a) + O_r\left(\sum_{k=1}^{r-1} C_{r-k}(b)L_k(a)\right), \\ &L_r(b) + O_r\left(\sum_{k=1}^{r-1} C_k(a)L_{r-k}(b)\right)\end{aligned}\right\}.$$

Both alternatives are estimated below with the actual values $a = \lfloor n/2 \rfloor$ and $b = \lceil n/2 \rceil$. The inductive estimates apply separately to a and b; since each is at most $\lceil n/2 \rceil$ and $h(a), h(b) \le h(n)$, no monotonicity of C_r or L_r is being assumed. We verify the asymptotics rather than suppressing the parity bookkeeping. Put $h(n) = \log(2n)$. Unrolling a two-sided recurrence

$$F(n) \le F(\lfloor n/2 \rfloor) + F(\lceil n/2 \rceil) + Kn^p h(n)^q$$

by levels shows that the level-j toll is at most $Kn^p 2^{-j(p-1)} h(n)^q$. Hence $F(n) = O(n^p h(n)^q)$ for $p > 1$, while for $p = 1$ there are $O(h(n))$ nonzero levels and $F(n) = O(nh(n)^{q+1})$. Similarly, a one-sided recurrence $G(n) \le G(\lceil n/2 \rceil) + Kn^p h(n)^q$ has order $O(n^p h(n)^q)$ for $p > 0$, and $O(h(n)^{q+1})$ for $p = 0$.

We prove simultaneously, by induction on r, that

$$C_r(n) = O_r\left(n^{\lfloor r/2 \rfloor} h(n)^{\eta_r}\right), \tag{1}$$

$$W_r(n) = O_r\left(n^{\lceil r/2 \rceil} h(n)^{\theta_r}\right), \tag{2}$$

$$L_r(n) = O_r\left(n^{\lceil r/2 \rceil - 1} h(n)^{\theta_r}\right), \tag{3}$$

where $\eta_r = 1$ for odd $r \ge 3$ and is zero otherwise, while $\theta_r = 1$ for even r and is zero for odd r. The case $r = 1$ is $C_1 = 1$, $W_1 = n$, and $L_1 = 1$.

For the piece count, $r = 2$ has constant cross toll and hence $C_2(n) = O(n)$. For $r = 3$, the two cross terms are $C_1(a)C_2(b) + C_2(a)C_1(b) = O(n)$; the $p = 1$ recurrence therefore gives $C_3(n) = O(nh(n))$. If $r = 2s \ge 4$, every even–even product among the cross terms has order $O_r(n^s)$, while an odd–odd product has order $O_r(n^{s-1}h(n)^2) = O_r(n^s)$. Thus the cross toll is $O_r(n^s)$ and the $p = s > 1$ recurrence gives $C_{2s}(n) = O_r(n^s)$. If $r = 2s + 1 \ge 5$, one factor in every cross product has odd uniformity and the other even uniformity; the toll is $O_r(n^s h(n))$. Since $s > 1$, the two-sided recurrence preserves this order. This proves (1).

For total weight, $r = 2$ has toll $O(n)$ and hence $W_2(n) = O(nh(n))$. For even $r = 2s \ge 4$, each term $C_{r-k}W_k + C_k W_{r-k}$ has order $O_r(n^s h(n))$: an even weight factor supplies one logarithm, while an odd count factor supplies it in the opposite parity case. The two-sided recurrence with $s > 1$ gives $W_{2s}(n) = O_r(n^s h(n))$. For odd $r = 2s + 1$, the terms pairing an even count with an odd weight have order $O_r(n^{s+1})$; the remaining terms are $O_r(n^s h(n)^2) = O_r(n^{s+1})$. Since $s + 1 > 1$, the recurrence gives $W_{2s+1}(n) = O_r(n^{s+1})$. This proves (2).

Finally consider the one-sided load recurrence. For $r = 2$ its cross toll is constant, so $L_2(n) = O(h(n))$. For even $r = 2s \ge 4$, every cross term $C_{r-k}L_k$ is $O_r(n^{s-1}h(n))$; since $s - 1 > 0$, the one-sided recurrence preserves this order. For odd $r = 2s + 1$, the cross toll is $O_r(n^s)$: the potentially logarithmic terms have the smaller order $O_r(n^{s-1}h(n)^2)$ and are absorbed. The one-sided recurrence gives $L_{2s+1}(n) = O_r(n^s)$. This proves (3). Taking $r = d$ establishes the lemma. □

Lemma 18 (Full row rank of inclusion matrices)**.** *Let $0 \le k \le m \le N - k$. Form the matrix $I^N_{k,m}$ whose rows are the k-subsets $S \subseteq [N]$, whose columns are the m-subsets $U \subseteq [N]$, and whose (S, U) entry is 1 when $S \subseteq U$ and is 0 otherwise. Then $I^N_{k,m}$ has full row rank over $\mathbb{R}$.*

Proof. Suppose coefficients c_S satisfy

$$\sum_{\substack{S \subseteq U \\ |S| = k}} c_S = 0 \tag{4}$$

for every m-set U. We prove $c_S = 0$ by induction on k. For $k = 0$ this is immediate. Let $k \ge 1$, fix distinct $x, y \in [N]$, and write $\Omega = [N] \setminus \{x, y\}$. For every $(m-1)$-set $R \subseteq \Omega$, subtract (4) for

$R \cup \{x\}$ and for $R \cup \{y\}$. The terms indexed by subsets of R cancel, leaving

$$\sum_{\substack{A \subseteq R \\ |A|=k-1}} \left(c_{A\cup\{x\}} - c_{A\cup\{y\}}\right) = 0.$$

The induction hypothesis applies on the $(N-2)$-element ground set with parameters $k-1, m-1$, because $k-1 \le m-1 \le (N-2)-(k-1)$. Hence $c_{A\cup\{x\}} = c_{A\cup\{y\}}$ for every $(k-1)$-set $A \subseteq \Omega$.

Thus two k-sets that differ by one element have equal coefficients. The graph of k-sets joined by one-element exchanges is connected, so all coefficients equal a common value c. Substitution into (4) gives $\binom{m}{k}c = 0$, and therefore $c = 0$. □

Lemma 19 (Rank lower bound for complete partitions)**.** *Let $r \ge 2$, $n \ge r$, and $k = \lfloor r/2 \rfloor$. Every exact partition of $K_n^{(r)}$ into complete r-partite pieces has at least*

$$\frac{\binom{n}{k}}{\binom{r}{k}}$$

pieces.

Proof. Put $\ell = r - k$. Let D be the matrix whose rows are k-subsets $S \subseteq [n]$, whose columns are ℓ-subsets $T \subseteq [n]$, and whose entry is 1 if $S \cap T = \emptyset$ and 0 otherwise. Complementing a column identifies D with the inclusion matrix whose columns are the $(n-\ell)$-sets. Since $n \ge r = k + \ell$ and $\ell \ge k$, the inequalities $k \le n - \ell \le n - k$ hold. Lemma 18 therefore gives

$$\operatorname{rank}(D) = \binom{n}{k}.$$

For one complete r-partite piece $P = K(A_1, \dots, A_r)$, the corresponding submatrix D^P decomposes as

$$D^P = \sum_{\substack{I \subseteq [r] \\ |I|=k}} u_I v_I^{\mathsf{T}},$$

where u_I indicates k-sets choosing one vertex from each part A_i with $i \in I$, and v_I indicates ℓ-sets choosing one vertex from each complementary part. Hence $\operatorname{rank}(D^P) \le \binom{r}{k}$.

If $P_1, \dots, P_t$ exactly partition $K_n^{(r)}$, then $D = \sum_{i=1}^t D^{P_i}$. By rank subadditivity,

$$\binom{n}{k} = \operatorname{rank}(D) \le \sum_{i=1}^{t} \operatorname{rank}(D^{P_i}) \le t\binom{r}{k},$$

which proves the claim. □

Proof of Theorem 4. The cover statement is Proposition 16. The exact-partition upper bounds are Lemma 17. For the lower bounds, let $\mathcal{P}$ be an exact partition of $K_n^{(d)}$ and fix a vertex v. By Lemma 14, the pieces containing v induce an exact partition of the link of v, namely $K_{n-1}^{(d-1)}$. Lemma 19 applied with $r = d-1$ implies that at least $\Omega_d(n^{\lfloor (d-1)/2 \rfloor}) = \Omega_d(n^{\lceil d/2 \rceil - 1})$ pieces contain v. This is the load lower bound. Summing over all v gives the total-weight lower bound.

For odd $d = 2s+1$, the lower and upper exponents match exactly: $W = \Theta_d(n^{s+1})$ and $L = \Theta_d(n^s)$. For even $d = 2s$, the bounds differ by at most one logarithmic factor. □

Example 20 (The median partition of $K_n^{(3)}$; [1])**.** Order the vertices $1 < 2 < \cdots < n$. For each $2 \le b \le n-1$, let

$$P_b = K(\{1, \dots, b-1\}, \{b\}, \{b+1, \dots, n\}).$$

The edges of P_b are exactly the triples whose median is b. Every triple has a unique median, so the $n-2$ pieces P_b form an exact partition. Every piece has weight n, and every vertex belongs to every piece; hence

$$W = n(n-2), \qquad L = n-2.$$

These values are optimal. Indeed, for every vertex v, the link-trace lemma turns the pieces containing v into an exact biclique partition of K_{n-1}. Lemma 13 forces at least $n-2$ such pieces. Thus every vertex has load at least $n-2$, and summing the loads gives weight at least $n(n-2)$.

4.3 Near-complete repair

Lemma 21 (Deleting forbidden transversals from one box)**.** *Let $B = K(A_1, \dots, A_d)$ and let $F_B \subseteq E(B)$ with $|F_B| = h$. Then $E(B) \setminus F_B$ has an exact partition into complete d-partite boxes with total weight at most*

$$(1 + (d-1)h)w(B),$$

and every vertex of B has local load at most $1 + (d-1)h$ inside this repair.

Proof. Delete the forbidden edges one at a time. Suppose the current exact partition contains a box $Q = K(C_1, \dots, C_d)$ containing the next forbidden edge $f = \{x_1, \dots, x_d\}$ with $x_i \in C_i$. Replace Q by the boxes

$$K(\{x_1\}, \dots, \{x_{j-1}\}, C_j \setminus \{x_j\}, C_{j+1}, \dots, C_d), \qquad j = 1, \dots, d,$$

discarding empty boxes. Every edge of $Q \setminus \{f\}$ has a first coordinate in which it differs from f, and this first differing coordinate assigns it to exactly one replacement box. Thus the replacement is an exact partition of $E(Q) \setminus \{f\}$.

Each replacement box has weight at most $w(Q) \le w(B)$. Therefore one deletion step increases total weight by at most $(d-1)w(B)$ and increases the load of any fixed vertex by at most $d-1$. Starting from the single box B and performing h steps proves the lemma. □

Proposition 22 (Near-complete exact partitions)**.** *Let $H = K_n^{(d)} \setminus F$ with $|F| = t$ and fixed $d \ge 3$. Then H has an exact partition with*

$$W \le O_d\Big(n^{\lceil d/2 \rceil} (\log n)^{\mathbf{1}_{d \text{ even}}}\Big) + (d-1)nt$$

and

$$L \le O_d\Big(n^{\lceil d/2 \rceil - 1} (\log n)^{\mathbf{1}_{d \text{ even}}}\Big) + (d-1)t.$$

Conversely every exact partition satisfies

$$W \ge c_d n^{\lceil d/2 \rceil} - dt, \qquad L \ge c_d n^{\lceil d/2 \rceil - 1} - dt/n.$$

Proof. Start with the recursive exact partition of $K_n^{(d)}$ from Lemma 17. For each complete box P in that partition, let $F_P = F \cap E(P)$. The sets F_P partition F. Apply Lemma 21 to P with forbidden set F_P. Since every original box has weight at most n, the total added weight is at most

$$(d-1)n \sum_P |F_P| = (d-1)nt,$$

and the load at each vertex increases by at most $(d-1)t$. This gives the upper bounds.

For the lower bounds, let $\mathcal{P}$ be an exact partition of H. For each vertex v, let t_v be the number of missing edges containing v. The link $L_H(v)$ is obtained from $K_{n-1}^{(d-1)}$ by deleting t_v edges. The pieces of $\mathcal{P}$ containing v exactly partition this link. If we add each missing link edge as a singleton complete $(d-1)$-partite piece, we obtain an exact partition of $K_{n-1}^{(d-1)}$ using at most $\ell_{\mathcal{P}}(v)+t_v$ pieces. Lemma 19 gives

$$\ell_{\mathcal{P}}(v)+t_v \geq c_d n^{\lceil d/2\rceil-1}.$$

Thus $\ell_{\mathcal{P}}(v) \geq c_d n^{\lceil d/2\rceil-1} - t_v$. Summing over v and using $\sum_v t_v = dt$ gives the weight lower bound; choosing a vertex with $t_v \leq dt/n$ gives the load lower bound. □

Remark 23. Proposition 22 is useful only in small-complement regimes. It does not determine the full transition between the non-dense formula of Theorem 2 and the complete-endpoint behavior of Theorem 4. The corresponding transition for covers is also open.

5 Concluding remarks

Theorem 2 gives a sharp density-sensitive answer to the total-weight part of Problem 1 whenever the edge density stays a fixed distance below one. The proof is robust across the sparse/intermediate boundary: singleton-prefix pieces are optimal at the sparsest scales, while balanced boxes of side $\Theta((\log n/\log(1/\gamma))^{1/(d-1)})$ drive the intermediate peeling argument. The exact-density random model supplies matching obstructions for both covers and partitions.

Three issues remain outside the result. First, the transition $\gamma = 1-o(1)$ is not determined by either the non-dense extraction theorem or the linear-in-the-number-of-holes repair. Second, for even uniformity the weighted exact complete endpoint still has a logarithmic gap. Third, Corollary 12 supplies a polynomial-time implementation, but a near-linear or output-sensitive algorithm with a sharp running-time bound remains open. On the local side, the cone example shows that link concentration must be incorporated before one can expect a per-instance theorem comparable to the total-weight scale.

References

[1] N. Alon. Decomposition of the complete r-graph into complete r-partite r-graphs. *Graphs and Combinatorics* 2(1):95–100, 1986. doi:10.1007/BF01788083.

[2] A. Babu and S. Vishwanathan. Improved bounds for multicovering hypergraphs. arXiv:2208.12589v2, 2022; revised 2026.

[3] P. Erdős. On extremal problems of graphs and generalized graphs. *Israel Journal of Mathematics* 2(3):183–190, 1964. doi:10.1007/BF02759942.

[4] F. Espuña. Finding partite hypergraphs efficiently. *Information Processing Letters* 193:106624, 2026. doi:10.1016/j.ipl.2026.106624; arXiv:2508.10641v2.

[5] G. Katona and E. Szemerédi. On a problem of graph theory. *Studia Scientiarum Mathematicarum Hungarica* 2:23–28, 1967.

[6] X. Peng. On the decomposition of random hypergraphs. *Journal of Combinatorial Theory, Series B* 129:18–37, 2018. doi:10.1016/j.jctb.2017.09.002.

[7] J. Radhakrishnan. Improved bounds for covering complete uniform hypergraphs. *Information Processing Letters* 41(4):203–207, 1992. doi:10.1016/0020-0190(92)90181-T.

[8] L. Csirmaz, P. Ligeti, and G. Tardos. Erdős–Pyber theorem for hypergraphs and secret sharing. *Graphs and Combinatorics* 31(5):1335–1346, 2015. doi:10.1007/s00373-014-1448-7.

[9] R. L. Graham and H. O. Pollak. On the addressing problem for loop switching. *Bell System Technical Journal* 50(8):2495–2519, 1971. doi:10.1002/j.1538-7305.1971.tb02618.x.

[10] A. Krapivin, B. Przybocki, N. Sanhueza-Matamala, and B. Subercaseaux. Optimal and efficient partite decompositions of hypergraphs. In *Proceedings of the 58th Annual ACM Symposium on Theory of Computing (STOC 2026)*, pages 2187–2198, 2026; arXiv:2511.11855v1. doi:10.1145/3798129.3800921.