HAL open science

Equitable orientations of sparse uniform hypergraphs

Nathann Cohen, William Lochet

To cite this version:

Nathann Cohen, William Lochet. Equitable orientations of sparse uniform hypergraphs. The Electronic Journal of Combinatorics, 2016, 23 (4), pp.4-31. hal-01417832

HAL Id: hal-01417832

https://hal.science/hal-01417832

Submitted on 16 Dec 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Equitable orientations of sparse uniform hypergraphs

Nathann Cohen
CNRS and Université Paris-Sud
nathann.cohen@gmail.com

William Lochet*
Université Côte d'Azur, CNRS, Inria, I3S, France
LIP, ENS de Lyon, CNRS, Université de Lyon, France
william.lochet@gmail.com

Submitted: May 20, 2016; Accepted: Jan 2, 2012; Published: XX
Mathematics Subject Classifications: 05C65, 05D40

Abstract

Caro, West, and Yuster studied how r-uniform hypergraphs can be oriented in such a way that (generalizations of) indegree and outdegree are as close to each other as can be hoped. They conjectured an existence result of such orientations for sparse hypergraphs, of which we present a proof.

1 Introduction

In [1], Caro, West, and Yuster presented a generalization to hypergraphs of the notion of *orientation* defined for graphs. Their acknowledged purpose is to study how hypergraphs can be oriented in such a way that minimum and maximum degree are close to each other, knowing that reaching an additive difference of $\leqslant 1$ is always achievable in the case of graphs. Identifying an *orientation* of an edge with a total ordering of its elements, they define a notion of degree on oriented r-uniform hypergraphs.

Definition 1. Let $\mathcal{H}$ be a r-uniform hypergraph, and let every $S \in \mathcal{H}$ define a total order on its elements as a bijection $\sigma_S : S \mapsto [r]$. The degree $d_P(U)$ of a set of vertices $U \subseteq V(\mathcal{H})$ with respect to a set of positions $P \subseteq [r]$ (where $|U| = |P|$) is equal to:

$$d_P(U) = |\{S \in \mathcal{H} : U \subseteq S \text{ and } \sigma_S(U) = P\}|$$

From there they define *equitable* orientations:

Definition 2. The orientation of a r-uniform hypergraph $\mathcal{H}$ is said to be p-equitable if $|d_P(U) - d_{P'}(U)| \leqslant 1$ for any choice of $U \subseteq V(\mathcal{H})$ and $P, P' \subseteq [r]$ of cardinality p. It is said to be *nearly p-equitable* if the looser requirement $|d_P(U) - d_{P'}(U)| \leqslant 2$ holds.

*Supported by ANR under contract STINT ANR-13-BS02-0007

They gave proof that all hypergraphs admit a 1-equitable as well as a $(r-1)$-equitable orientation, and also proved that some hypergraphs do not admit a p-equitable orientation for all values of p. Additionally, they parameterized the notion of maximum degree in order to focus on hypergraphs which are *sparse* with respect to the problem at hand:

$$\Delta_p(\mathcal{H}) = \max_{\substack{U \subseteq V(\mathcal{H}) \\ |U|=p}} |\{S \in \mathcal{H} : U \subseteq S\}|$$

Thus, they proved that for any fixed value of p and k, and for every sufficiently large integer r, every r-uniform hypergraph $\mathcal{H}$ with $\Delta_p(\mathcal{H}) \leqslant k$ admits a nearly p-equitable orientation. They conjectured that this setting actually ensured the existence of a p-equitable orientation, which we prove here.

Theorem 3. *Let p, k be fixed integers. There exists r_0 such that for every $r \geqslant r_0$, every r-uniform hypergraph with $\Delta_p(\mathcal{H}) \leqslant k$ admits a p-equitable orientation.*

Note that, as r is big compared to $\Delta_p(\mathcal{H})$, a p-equitable orientation means that $d_P(U)$ is equal to 0 or 1 for every choice of set of positions P and set of vertices U.

In order to prove the existence of nearly p-equitable orientation, Caro, West, and Yuster [1] used the Lovász Local Lemma. In [3], Moser and Tardos presented an elegant algorithmic proof of it which developed the technique of entropy compression. Our proof uses that technique and the following Lemma (proved in Section 3) that counts what can be seen as a generalization of derangements.

Lemma 4. *Let $p, k \in \mathbb{N}$ and $\alpha < 1$ be fixed. Let X be a set of cardinality r and let $\mathcal{L}_S$ be, for every $S \in \binom{X}{p}$, a collection of p-subsets of X with $|\mathcal{L}_S| \leqslant k$. Then, if no p-subset occurs in more than r^α of the $\mathcal{L}_S$, a random permutation σ of X satisfies $\sigma(S) \notin \mathcal{L}_S$ for every S with probability $\geqslant (1 - 2k/\binom{r}{p})^{\binom{r}{p}} = e^{-2k} + o(1)$ when r grows large.*

2 Algorithm

In what follows, we assume that every finite set S has an implicit enumeration on its elements, and in particular that the edges of a hypergraph $\mathcal{H}$ are implicitly ordered. We will say that i *represents* an element $s \in S$ when s is the i-th element of S in this implicit ordering.

We will orient the edges of $\mathcal{H}$ one by one as a (partial) equitable orientation of $\mathcal{H}$, i.e. in such a way that any p-subset of $V(\mathcal{H})$ never appears more than once at the same position among the oriented edges. To do so, we require the partial orientation to enforce an additional property.

Definition 5. Let $\mathcal{H}$ be a partially oriented r-uniform hypergraph. We say that an edge $S \in \mathcal{H}$ is *pressured* by a family $\{S_1, \ldots, S_l\}$ of edges (oriented by $\sigma_{S_1}, \ldots, \sigma_{S_l}$) if there exists $P \in \binom{[r]}{p}$ such that $\sigma_{S_i}^{-1}(P) \subseteq S$ for every i.

Note that Lemma 11 ensures that a partial orientation of $\mathcal{H}$ can be extended to an unoriented edge S, provided that no family of more than r^α oriented edges pressures S. It asserts, for $c < e^{-2k}$ and r sufficiently large, that at least $cr!$ orientations of S are admissible for this extension: we name them *good* permutations of S. Algorithm 1 selects an ordering randomly among them, while ensuring that no other edge is pressured by a family of edges larger than $r_1 = \lfloor r^\alpha \rfloor$.

Data: A r-uniform hypergraph $\mathcal{H}$ with $\Delta_p(\mathcal{H}) \leqslant k$
Result: A p-equitable orientation of $\mathcal{H}$

while *not all edges are oriented* **do**
  $S_1 \leftarrow$ unoriented edge of smallest index
  Pick for S_1 the orientation indexed v_i (among $\geqslant cr!$ available)
  if *some edge S of $\mathcal{H}$ is pressured by a family* $\{S_1, \ldots, S_{r_1}\}$ **then**
    Cancel the orientation of all edges S_i.
end
Return the oriented $\mathcal{H}$

Algorithm 1: A non-deterministic algorithm

Algorithm 1 starts with every edge being unoriented. At each step it orients the unoriented edge of smallest index by choosing a random permutation amongst the $cr!$ first good permutations. We call *bad event* the event that an edge $S \in \mathcal{H}$ is pressured by a family $\{S_1, \ldots, S_{r_1}\}$ of cardinality r_1. If a bad event occurs after orienting S_1, then the algorithm erases the orientation of the $S_1, \ldots, S_{r_1}$.

It is trivial to see that Algorithm 1 only returns p-equitable orientations of $\mathcal{H}$. Moreover, every time the algorithm chooses a random permutation, it does so among at least $cr!$ good ones by Lemma 11. Note that we need to consider large families pressuring already oriented edges: indeed, we might have to cancel the orientation of such an edge to redefine it again later.

Theorem 6. *Let $p, k \in \mathbb{N}$, $\alpha, c \in \mathbb{R}_{>0}$ with $\alpha < 1$ and $c < e^{-2k}$. For every sufficiently large r, there is a set of random choices for which Algorithm 1 terminates.*

In order to prove this result we will analyse the possible executions of the M first steps of Algorithm 1. To this end we make it deterministic by defining a log (following the idea of [3]) and obtain Algorithm 2, in the following way:

- Take as input a vector $v \in [cr!]^M$ which simulates the random choices.
- Output a log when it is not able to orient all edges.

We define a *log of order M* to be a triple (R, X, F) where:

- R is a binary word whose length lies between M and $2M$.
- X is a sequence of h 7-tuples of integers $(x_1, x_2, x_3, x_4, x_5, x_6, x_7)$ where:

$$x_1 \leqslant \binom{r}{p} \qquad x_2 \leqslant k \qquad x_3 \leqslant \binom{\binom{r}{p}}{r_1-1} \qquad x_4 \leqslant k^{r_1-1}$$
$$x_5 \leqslant p!^{r_1-1} \qquad x_6 \leqslant (r-p)!^{r_1-1} \qquad x_7 \leqslant r!$$

- F is an integer smaller than $(r!+1)^{|\mathcal{H}|}$ representing a partial orientation of $\mathcal{H}$.

The log of order M (or just log) is actually a trace of the deterministic algorithm's execution after M steps. Its objective is to encode which orientations get canceled during the algorithm's execution. We will show later that Algorithm 2 cannot produce the same log from two different input vectors $v, v' \in [cr!]^M$. and that, for M big enough, that the set of possible log is smaller than $(cr!)^M$. We now describe the log and how Algorithm 2 produces it.

- R is initialized to the empty word. We append 1 to R whenever Algorithm 2 adds a new orientation; we append 0 whenever it cancels one.
- Consider the following bad event: after orienting S_1, an edge $S \in \mathcal{H}$ is pressured by a family $\{S_1, \ldots, S_{r_1}\}$ of cardinality r_1. We note s_i the set of vertices that S_i maps to P. We associate the following 7-tuple which identifies the sets S_i as well as their orientation:
  - $x_1 < \binom{r}{p}$ represents the set s_1 among the $\binom{r}{p}$ possible subsets of size p of S_1.
  - $x_2 < k$ identifies S as one of the (at most k) edges containing s_1.
  - $x_3 < \binom{\binom{r}{p}}{r_1-1}$ is an integer representing the set of subsets $s_2, \ldots s_{r_1}$ amongst the $\binom{r}{p}$ subsets of size p of S.
  - $x_4 < k^{r_1-1}$ is an integer representing the sequence $(y_2, \ldots, y_{r_1}) \in [k]^{r_1-1}$ such that the y_l-th edge containing s_l is S_l.
  - $x_5 < p!^{r_1-1}$ is an integer representing the sequence $(p_1, \ldots, p_{r_1})$, where $p_i \in [p!]$ represents the subpermutation of S_i onto s_i (we know it's a permutation of P).
  - $x_6 < (r-p)!^{r_1-1}$ is the integer representing the sequence $[p_2, \ldots, p_{r_1}]$, where $p_i \in [(r-p)!]$ represent the subpermutation of S_i onto $[r] \setminus s_i$.
  - $x_7 < r!$ is the integer representing the permutation chosen for S_1.

  X is the list of the 7-tuples describing the bad events, in the order in which they happen.
- F is the integer representing the partial orientation of $\mathcal{H}$ (i.e. a choice among $r!+1$ per edge of $\mathcal{H}$) after M steps.

This gives the following Algorithm 2:

Data:

1. A r-uniform hypergraph $\mathcal{H}$ with $\Delta_p(\mathcal{H}) \leqslant k$,
2. A vector $v \in [cr!]^M$

Result: A p-equitable orientation of $\mathcal{H}$, or a log of order M

```
R ← ∅, X ← ∅
for 1 ⩽ i ⩽ M do
    S_1 ← unoriented edge of smallest index
    Pick for S the orientation indexed v_i among ⩾ cr! available
    if some edge of H is pressured by {S_1, ..., S_{r_1}} then
        Append 1 to the end of R
        Append to X a 7-tuple describing the conflict
        Cancel the orientation of all r_1 + 1 edges involved in the conflict
    else if all edges are oriented then
        Return the oriented H
    else
        Append 0 to the end of R
    end
end
F ← the integer representing the partial orientation of H.
Return (R, X, F)
```

Algorithm 2: A deterministic algorithm

We will show the following claim.

Claim 7. *Let e be a vector in $[cr!]^M$ from which Algorithm 2 cannot produce a p-equitable orientation of $\mathcal{H}$ and outputs a log (R, X, F). We can reconstruct e from (R, X, F).*

Proof of the claim. First we show that we can find, for every $z \leqslant M$, the set $C(z)$ of edges which are oriented after z steps. We proceed by induction on z, starting from $C(0) = \emptyset$. At step $z+1$, Algorithm 2 chooses a orientation for the smallest index i not in $C(z)$. If, in R, the $(z+1)$-th 1 is not followed by a 0, then there is no bad event triggered by this step. In this case the set $C(z+1)$ is the set $C(z) \cup i$. Suppose now that the $(z+1)$-th 1 is followed by a sequence of 0: this means that the algorithm encountered a bad event. By looking at the number of sequences of 0 in R before the $z+1$-th 1 we can deduce the number of bad events before this one. This mean we can find, in X, the 7-tuple $(x_1, x_2, x_3, x_4, x_5, x_6, x_7)$ associated to this bad event. We take the following notations for the bad event : After orienting S_1, an edge S of H is pressured by a family $\{S_1, \dots S_{r_1}\}$ of cardinality r_1 We note s_i the subset of S_i that are sent to P. S_1 is the last edge we oriented (known by induction), x_1 indicates s_1 amongst the subset of S, x_2 indicates S amongst the set of edges containing s_1, x_3 indicates the s_d for $d \in [2..r_1]$, and x_4 indicates the S_d for $d \in [2..r_1]$. In this case the set $C(z+1)$ is the set $C(z)$ for which we removed all the ES_d for $d \in [2..r_1]$.

We can now deduce the set $S(z)$ of all chosen orientations after z steps. We also proceed by induction, this time starting from step M. By construction, F is exactly the integer representing the partial orientation of $\mathcal{H}$ at step M. If the last letter of R is a 1, this means the last step of the algorithm consisted only of choice of a orientation. We just showed that we know which orientation was chosen after $M-1$ steps, so we can deduce the state of all orientation after $M-1$ steps. If the last letter is a 0, Algorithm 2 encountered a bad event. Keeping the notation of the bad event, let $(x_1, x_2, x_3, x_4, x_5, x_6, x_7)$ be the 7-tuple associated to this bad event. Like before x_1, x_2, x_3, x_4 and the knowledge of $C(M-1)$ allow us to know which permutations Algorithm 2 erased at this step. Moreover x_7 tells us the random choice made by Algorithm 2 and from x_7 and x_1 we can deduce P. For each s_i we know the orientation chosen for S_i at the step $M-1$ sends P onto s_i, from x_5 we deduce exactly in which order and from x_6 we get the rest of the orientation. Therefore we can deduce the set of chosen orientations before the bad event occurred. With the sets $S(z)$ and $C(z)$ known for all $z \leqslant M$ we can easily deduce e. $\diamond$

The previous claim has the following corollary:

Corollary 8. *If $\mathcal{H}$ admits no p-equitable orientation, then Algorithm 2 defines an injection from the set of vectors $[cr!]^M$ into L^M.*

Let L_M be the set of all possible logs after M steps of Algorithm 2. To show Theorem 6 it suffices to show that, for M big enough, $|L_M|$ is strictly smaller than $(cr!)^M$.

Lemma 9. *For M big enough, $|L_M| < (cr!)^M$.*

Proof. We will compute a bound for $|L_M|$. R is a binary word of size $\leqslant 2M$, and there are at most 4^M such words. X is a list of 7-tuples. As Algorithm 2 made M choices and each bad event removes r_1 of those, there exist at most $\frac{M}{r_1}$ bad events. Moreover, for each 7-tuple, $(x_1, x_2, x_3, x_4, x_5, x_6, x_7)$ we have $x_1 \leqslant \binom{r}{p}$, $x_2 \leqslant k$, $x_3 \leqslant \binom{\binom{r}{p}}{r_1-1}$, $x_4 \leqslant k^{r_1-1}$, $x_5 \leqslant p!^{r_1-1}$, $x_6 \leqslant (r-p)!^{r_1-1}$, $x_7 \leqslant r!$. Using the bounds $\binom{n}{k} \leqslant (\frac{n \cdot e}{k})^k$ or $\binom{n}{k} \leqslant n^k$ we get the following bound.

$$
\begin{aligned}
|X| &\leqslant \left(r^p \cdot k \cdot \left(\frac{r^p \cdot e}{r_1 - 1} \right)^{r_1-1} \cdot (k \cdot p! \cdot (r-p)!)^{r_1-1} \cdot r! \right)^{M/r_1} \\
&\leqslant \frac{(r! \cdot (r^p)^{r_1} \cdot (r-p)!^{r_1-1})^{M/r_1} \cdot (k \cdot e \cdot p!)^M}{(r_1-1)^{M(r_1-1)/r_1}} \\
&\leqslant \left[r^p \cdot r!^{r_1} \cdot \left(\frac{r^p}{r(r-1)\dots(r-p+1)} \right)^{r_1-1} \right]^{M/r_1} \cdot \left(\frac{k \cdot e \cdot p!}{(r_1-1)^{(r_1-1)/r_1}} \right)^M
\end{aligned}
$$

We can assume $r > 2p$, and so $\frac{r}{r-p+1} < 2$:

$$
|X| \leqslant r!^M \cdot \left(r^{p/r_1} \cdot 2^p \cdot \frac{k \cdot e \cdot p!}{(r_1-1)^{(r_1-1)/r_1}} \right)^M
$$

As $|F| \leqslant (r!+1)^{|\mathcal{H}|}$ and $|L_M| \leqslant |F||X||R|$ we get the following bound on $|L_M|$:

$$|L_M| \leqslant r!^M \cdot \left(4 \cdot r^{p/r_1} \cdot 2^p \cdot \frac{k \cdot e \cdot p!}{(r_1-1)^{(r_1-1)/r_1}}\right)^M \cdot (r!+1)^{|\mathcal{H}|}$$

□

3 Derangements

The results of this section are based on a lemma from Erdős and Spencer [2]:

Lemma 10 (Lopsided Lovász Local Lemma)**.** *Let $A_1, \ldots, A_m$ be events in a probability space, each with probability at most p. Let G be a graph defined on those events such that for every A_i, and for every set S avoiding both A_i and its neighbours, the following relation holds:*

$$P[A_i | \bigwedge_{A_j \in S} \bar{A}_j] \leqslant P[A_i]$$

Then if $4dp \leqslant 1$, all the events can be avoided simultaneously:

$$P[\bar{A}_1 \wedge \cdots \wedge \bar{A}_m] \geqslant (1-2p)^m > 0$$

Thanks to this result we can prove the following, which can be seen as a generalization of the fact that a random permutation of n points is a derangement with asymptotic probability $1/e$.

Lemma 11. *Let $p, k \in \mathbb{N}$ and $\alpha < 1$ be fixed. Let X be a set of cardinality r and let $\mathcal{L}_S$ be, for every $S \in \binom{X}{p}$, a collection of p-subsets of X with $|\mathcal{L}_S| \leqslant k$. Then, if no p-subset occurs in more than r^α of the $\mathcal{L}_S$, a random permutation σ of X satisfies $\sigma(S) \notin \mathcal{L}_S$ for every S with probability $\geqslant (1 - 2k/\binom{r}{p})^{\binom{r}{p}} = e^{-2k} + o(1)$ when r grows large.*

Proof. For every $S \in \binom{X}{p}$, we define the *bad event* B_S with:

$$B_S = \bigvee_{S' \in \mathcal{L}_S} [\sigma(S) = S']$$

Each B_S has a probability $P[B_S] \leqslant k/\binom{r}{p}$. On these bad events we define a lopsidependency graph (see [2]) G_B with the following adjacencies:

$$\left\{(B_{S_1}, B_{S_2}) : S_1, S_2 \in \binom{X}{p} \text{ s.t. } \left[S_1 \bigcup \mathcal{L}_{S_1}\right] \bigcap \left[S_2 \bigcup \mathcal{L}_{S_2}\right] \neq \emptyset\right\}$$

As a p-subset of X intersects at most $O(r^{p-1})$ others, and noting that every p-subset can occur at most r^α times, we have that:

$$\Delta(G_B) \leqslant (k+1)r^\alpha \times O(r^{p-1}) = o(r^p)$$

In order to apply the Lopsided Lovász Local Lemma to the events B_S and graph G_B, we must ensure for every $S \in \binom{X}{p}$ and $S_B \subseteq V(G_B) \backslash N_{G_B}[B_S]$ that:

$$P(B_S | \bigwedge_{B_{S'} \in S_B} \bar{B}_{S'}) \leqslant P(B_S) \tag{1}$$

Indeed, if we denote by T (for *trace*) the number of elements of $\bigcup_{B_{S'} \in S_B} S'$ sent by the random permutation σ into $\bigcup \mathcal{L}_S$:

$$P(B_S) = \sum_t P(B_S \mid T = t) P(T = t)$$

$$P(B_S | \bigwedge_{B_{S'} \in S_B} \bar{B}_{S'}) = \sum_t P(B_S \mid T = t, \bigwedge_{B_{S'} \in S_B} \bar{B}_{S'}) P(T = t \mid \bigwedge_{B_{S'} \in S_B} \bar{B}_{S'})$$

As $\bigcup \mathcal{L}_S$ is disjoint from the $\bigcup \mathcal{L}_{S'}, \forall B_{S'} \in S_B$, we have:

$$P(B_S \mid T = t, \bigwedge_{B_{S'} \in S_B} \bar{B}_{S'}) = P(B_S \mid T = t)$$

And thus:

$$P(B_S | \bigwedge_{B_{S'} \in S_B} \bar{B}_{S'}) = \sum_t P(B_S \mid T = t) P(T = t \mid \bigwedge_{B_{S'} \in S_B} \bar{B}_{S'})$$

In order to prove (1), we will first need the following observation:

Claim 12. *$P(B_S \mid T = t)$ is a decreasing function of t.*

Proof of the claim. We compute the value of $P(B_S \mid T = t)$ exactly, denoting by $r' \leqslant r$ the cardinality of $\bigcup_{B_{S'} \in S_B} S'$. It is equal to 0 when $t > r' - p$, and is otherwise equal to:

$$\begin{aligned}
P(B_S \mid T = t) &= \sum_{S' \in \mathcal{L}_S} P(\sigma(S) = S' \mid T = t) \\
&= \frac{|\mathcal{L}_S|}{\binom{r-t}{p}} \frac{\binom{r'-p}{t}}{\binom{r'}{t}} \\
&= \left(|\mathcal{L}_S| \frac{(r'-p)!p!}{r'!} \right) \left(\frac{(r-p-t)!(r'-t)!}{(r'-p-t)!(r-t)!} \right) \\
&= P(B_S \mid T = t-1) \left(\frac{(r'-p-t+1)}{(r-p-t+1)} \frac{(r-t+1)}{(r'-t+1)} \right) \\
&\leqslant P(B_S \mid T = t-1)
\end{aligned}$$

◇

Additionally, we will prove a relationship on the members of $\sum_t P(T = t)$ and on those of $\sum_t P(T = t \mid \bigwedge_{B_{S'} \in S_B} \bar{B}_{S'})$, which both sum to 1.

Claim 13. *If $P(T = t \mid \bigwedge_{B_{S'} \in S_B} \bar{B}_{S'})$ is nonzero, then*

$$\frac{P(T = t+1)}{P(T = t)} \leqslant \frac{P(T = t+1 | \bigwedge_{B_{S'} \in S_B} \bar{B}_{S'})}{P(T = t \mid \bigwedge_{B_{S'} \in S_B} \bar{B}_{S'})}$$

Proof of the claim. According to Bayes' Theorem applied to the right side of the equation,

$$\frac{P(T = t+1 | \bigwedge_{B_{S'} \in S_B} \bar{B}_{S'})}{P(T = t \mid \bigwedge_{B_{S'} \in S_B} \bar{B}_{S'})} = \frac{P(\bigwedge_{B_{S'} \in S_B} \bar{B}_{S'} \mid T = t+1) P(T = t+1)}{P(\bigwedge_{B_{S'} \in S_B} \bar{B}_{S'} \mid T = t) P(T = t)}$$

We thus only need to ensure the following, which is a consequence of Lemma 14:

$$P(\bigwedge_{B_{S'} \in S_B} \bar{B}_{S'} \mid T = t+1) \geqslant P(\bigwedge_{B_{S'} \in S_B} \bar{B}_{S'} \mid T = t)$$

◇

We are now ready to prove (1), and we define d_t for every t where $P(T = t)$ is nonzero:

$$d_t = P(T = t) - P(T = t \mid \bigwedge_{B_{S'} \in S_B} \bar{B}_{S'})$$

Because d_t is a difference of probability distributions the sum $\sum_t d_t$ is null, and we can rewrite (1) using d_t:

$$\begin{aligned}
0 &\leqslant P(B_S) - P(B_S | \bigwedge_{B_{S'} \in S_B} \bar{B}_{S'}) \\
&\leqslant \sum_t P(B_S \mid T = t) P(T = t) - \sum_t P(B_S \mid T = t) P(T = t \mid \bigwedge_{B_{S'} \in S_B} \bar{B}_{S'}) \\
&\leqslant \sum_t d_t P(B_S \mid T = t)
\end{aligned}$$

We will thus prove that the sum $\sum_t d_t P(B_S \mid T = t)$ is nonnegative. It is a consequence of Claim 13 that all nonnegative values of d_t appear before all nonpositive ones, and so that there is a t_0 such that $d_t \geqslant 0$ iff $t \leqslant t_0$. As a result, $|\sum_{t \leqslant t_0} d_t| = |\sum_{t > t_0} d_t| = \frac{1}{2}$ and we can write:

$$\begin{aligned}
\sum_t d_t P(B_S \mid T = t) &= \sum_{t \leqslant t_0} d_t P(B_S \mid T = t) + \sum_{t > t_0} d_t P(B_S \mid T = t) \\
&\geqslant \frac{1}{2} P(B_S \mid T = t_0) - \frac{1}{2} P(B_S \mid T = t_0 + 1) \geqslant 0 \quad \text{(by Claim 12)}
\end{aligned}$$

The second hypothesis of Lemma 10 is that $4pd \leqslant 1$, which translates in our case to $4 \frac{k}{\binom{r}{p}} o(r^p) = o(1)$ and is thus satisfied when r grows large. Hence, we have that:

$$P[\bigwedge_S \bar{B}_S] \geqslant \left[1 - 2k / \binom{r}{p}\right]^{\binom{r}{p}} = e^{-2k} + o(1)$$

□

Lemma 14. *Let A, B be two sets of size r, and let $\sigma : A \mapsto B$ be a random bijection. For every $A_1, \ldots, A_k \subset A' \subset A$ and $B_1, \ldots, B_k \subset B' \subset B$, the following function increases with t.*

$$P\left[\bigwedge_i [\sigma(A_i) \neq B_i] \;\middle|\; \sigma(A')\backslash B' \text{ has cardinality } t\right] \tag{2}$$

Proof. We implicitly assume in this proof that the conditionning event has a nonzero probability for t and $t+1$. Let S_1, S_2 be two sets of cardinality $|A'|$ with symmetric difference $S_1 \Delta S_2 = \{x, y\}$ where $x \in S_2$ is an element of $B\backslash B'$. Let σ_{xy} be the permutation transposing x and y. Then,

$$\begin{aligned} P\left[\bigwedge_i [\sigma(A_i) \neq B_i] \;\middle|\; \sigma(A') = S_1\right] &\leqslant P\left[\bigwedge_i [\sigma_{xy}\sigma(A_i) \neq B_i] \;\middle|\; \sigma(A') = S_1\right] \\ &= P\left[\bigwedge_i [\sigma(A_i) \neq B_i] \;\middle|\; \sigma(A') = S_2\right] \end{aligned}$$

We are now ready to derive the result:

$$(2) = \frac{1}{\binom{|B\backslash B'|}{t}\binom{|B'|}{|A'|-t}} \sum_{\substack{S\subseteq B \\ |S|=|A'| \\ |S\backslash B'|=t}} P\left[\bigwedge_i [\sigma(A_i) \neq B_i] \;\middle|\; \sigma(A') = S\right]$$

Using our previous remark, we find an upper bound on the last term of the equation by averaging it over sets S' obtained from S by the exchange of two elements:

$$\begin{aligned} (2) &\leqslant \frac{1}{\binom{|B\backslash B'|}{t}\binom{|B'|}{|A'|-t}} \sum_{\substack{S\subseteq B \\ |S|=|A'| \\ |S\backslash B'|=t}} \frac{1}{(|B\backslash B'|-t)(|A'|-t)} \sum_{\substack{S'\subseteq B \\ |S'|=|A'| \\ |S'\backslash B'|=t+1 \\ |S\Delta S'|=2}} P\left[\bigwedge_i [\sigma(A_i) \neq B_i] \;\middle|\; \sigma(A') = S'\right] \\ &= \frac{1}{\binom{|B\backslash B'|}{t}\binom{|B'|}{|A'|-t}} \frac{(t+1)(|B'|-|A'|+t+1)}{(|B\backslash B'|-t)(|A'|-t)} \sum_{\substack{S'\subseteq B \\ |S'|=|A'| \\ |S'\backslash B'|=t+1}} P\left[\bigwedge_i [\sigma(A_i) \neq B_i] \;\middle|\; \sigma(A') = S'\right] \\ &= \frac{\binom{|B\backslash B'|}{t+1}\binom{|B'|}{|A'|-t-1}}{\binom{|B\backslash B'|}{t}\binom{|B'|}{|A'|-t}} \frac{(t+1)(|B'|-|A'|+t+1)}{(|B\backslash B'|-t)(|A'|-t)} P\left[\bigwedge_i [\sigma(A_i) \neq B_i] \;\middle|\; \sigma(A')\backslash B' \text{ has cardinality } t+1\right] \\ &= P\left[\bigwedge_i [\sigma(A_i) \neq B_i] \;\middle|\; \sigma(A')\backslash B' \text{ has cardinality } t+1\right] \end{aligned}$$

□

References

[1] Yair Caro, Douglas West, and Raphael Yuster. Equitable hypergraph orientations. *Electron. J. Combin.*, 18(1), 2011.

[2] Paul Erdős and Joel Spencer. Lopsided Lovász local lemma and Latin transversals. *Discrete Appl. Math.*, 30(2-3):151–154, 1991. doi:10.1016/0166-218X(91)90040-4.

[3] Robin A. Moser and Gábor Tardos. A constructive proof of the general Lovász Local Lemma. *J. ACM*, 57(2):1–11, 2010. doi:10.1145/1667053.1667060.